Chapter 2

Martingale Measures

2.1 A General Discrete-Time Market Model

Information Structure

Fix a time set $\mathbb{T} = \{0, 1, \dots, T\}$, where the *trading horizon* T is treated as the terminal date of the economic activity being modelled, and the points of $\mathbb{T}$ are the admissible *trading dates*. We assume as given a fixed probability space $(\Omega, \mathcal{F}, P)$ to model all 'possible states of the market'.

In most of the simple models discussed in Chapter 1, Ω is a *finite* probability space (i.e., has a finite number of points ω each with $P(\{\omega\}) > 0$). In this situation, the σ-field $\mathcal{F}$ is the power set of Ω, so that every subset of Ω is $\mathcal{F}$-measurable.

Note, however, that the finite models can equally well be treated by assuming that, on a general sample space Ω, the σ-field $\mathcal{F}$ in question is finitely generated. In other words, there is a finite partition $\mathcal{P}$ of Ω into mutually disjoint sets $A_1, A_2, \dots, A_n$ whose union is Ω and that generates $\mathcal{F}$ so that $\mathcal{F}$ also contains only finitely many events and consists precisely of those events that can be expressed in terms of $\mathcal{P}$. In this case, we further demand that the probability measure P on $\mathcal{F}$ satisfies $P(A_i) > 0$ for all i.

In both cases, the only role of P is to identify the events that investors agree are *possible*; they may disagree in their assignment of probabilities to these events. We refer to models in which either of the preceding additional assumptions applies as *finite market models*. Although most of our examples are of this type, the following definitions apply to general market models. Real-life markets are, of course, always finite; thus the additional 'generality' gained by considering arbitrary sample spaces and σ-fields is a question of mathematical convenience rather than wider applicability!

The *information structure* available to the investors is given by an increasing (finite) sequence of *sub-σ-fields* of $\mathcal{F}$: we assume that $\mathcal{F}_0$ is trivial; that is, it contains only sets of P-measure 0 or 1. We assume that $(\Omega, \mathcal{F}_0)$ is complete (so that any subset of a null set is itself null and $\mathcal{F}_0$ contains all

P-null sets) and that $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots \subset \mathcal{F}_T = \mathcal{F}$. An increasing family of σ-fields is called a *filtration* $\mathbb{F} = (\mathcal{F}_t)_{t \in \mathbb{T}}$ on $(\Omega, \mathcal{F}, P)$. We can think of $\mathcal{F}_t$ as containing the information available to our investors at time t: investors learn without forgetting, but we assume that they are not prescient-insider trading is not possible. Moreover, our investors think of themselves as 'small investors' in that their actions will not change the probabilities they assign to events in the market. Again, note that in a finite market model each σ-field $\mathcal{F}_t$ is generated by a minimal finite partition $\mathcal{P}_t$ of Ω and that $\mathcal{P}_0 = \{\Omega\} \subset \mathcal{P}_1 \subset \mathcal{P}_2 \subset \cdots \subset \mathcal{P}_T = \mathcal{P}$. At time t, all our investors know which cell of $\mathcal{P}_t$ contains the 'true state of the market', but none of them knows more.

Market Model and Numéraire

The definitions developed in this chapter will apply to general discrete market models, where the sample space need not be finite. Fix a probability space $(\Omega, \mathcal{F}, P)$, a natural number d, the *dimension* of the market model, and assume as given a $(d+1)$-dimensional stochastic process $S = \left\{S_t^i : t \in \mathbb{T}, i = 0, 1, \ldots, d\right\}$ to represent the time evolution of the *securities price process*. The security labelled 0 is taken as a riskless (nonrandom) *bond* (or *bank account*) with price process S^0, while the d risky (random) stocks labelled $1, 2, \ldots, d$ have price processes $S^1, S^2, \ldots, S^d$. The process S is assumed to be *adapted* to the filtration $\mathbb{F}$, so that for each $i \leq d$, S_t^i is $\mathcal{F}_t$-measurable; that is, the prices of the securities at all times up to t are known at time t. Most frequently, we in fact take the filtration $\mathbb{F}$ as that generated by the price process $S = \left(S^1, S^2, \ldots, S^d\right)$. Then $\mathcal{F}_t = \sigma\left(S_u : u \leq t\right)$ is the smallest σ-field such that all the $\mathbb{R}^{d+1}$-valued random variables $\left\{S_u = \left(S_u^0, S_u^1, \ldots, S_u^d\right), u \leq t\right\}$ are $\mathcal{F}_t$-measurable. In other words, at time t, the investors know the values of the price vectors $(S_u : u \leq t)$, but they have no information about later values of S.

The tuple $(\Omega, \mathcal{F}, P, \mathbb{T}, \mathbb{F}, S)$ is the *securities market model*. We require at least one of the price processes to be strictly positive throughout; that is, to act as a benchmark, known as the *numéraire*, in the model. As is customary, we generally assign this role to the bond price S^0, although in principle any strictly positive S^i could be used for this purpose.

Note on Terminology: The term 'bond' is the one traditionally used to describe the riskless security that we use here as numéraire, although 'bank account' and 'money market account' are popular alternatives. We continue to use 'bond' in this sense until Chapter 9, where we discuss models for the evolution of interest rates; in that context, the term 'bond' refers to a certain type of risky asset, as is made clear.

2.2 Trading Strategies

Value Processes

Throughout this section, we fix a securities market model $(\Omega, \mathcal{F}, P, \mathbb{T}, \mathbb{F}, S)$. We take S^0 as a strictly positive *bond* or riskless security, and without loss of generality we assume that $S^0(0) = 1$, so that the initial value of the bond S^0 yields the units relative to which all other quantities are expressed. The *discount factor* $\beta_t = \frac{1}{S^0_t}$ is then the sum of money we need to invest in bonds at time 0 in order to have 1 unit at time t. Note that we allow the discount *rate* - that is, the increments in β_t - to vary with t; this includes the case of a constant interest rate $r > 0$, where $\beta_t = (1+r)^{-t}$.

The securities $S^0, S^1, S^2, \ldots, S^d$ are traded at times $t \in \mathbb{T}$: an investor's *portfolio* at time $t \geq 1$ is given by the $\mathbb{R}^{d+1}$-valued random variable $\theta_t = (\theta^i_t)_{0 \leq i \leq d}$ with *value process* $V_t(\theta)$ given by

$$V_0(\theta) = \theta_1 \cdot S_0, \qquad V_t(\theta) = \theta_t \cdot S_t = \sum_{i=0}^{d} \theta^i_t S^i_t \text{ for } t \in \mathbb{T},\ t \geq 1.$$

The value $V_0(\theta)$ is the investor's initial endowment. The investors select their time t portfolio once the stock prices at time $t-1$ are known, and they hold this portfolio during the time interval $(t-1, t]$. At time t the investors can adjust their portfolios, taking into account their knowledge of the prices S^i_t for $i = 0, 1, \ldots, d$. They then hold the new portfolio θ_{t+1} throughout the time interval $(t, t+1]$.

Market Assumptions

We require that the *trading strategy* $\theta = \{\theta_t : t = 1, 2, \ldots, T\}$ consisting of these portfolios be a *predictable* vector-valued stochastic process: for each $t < T$, θ_{t+1} should be $\mathcal{F}_t$-measurable, so θ_1 is $\mathcal{F}_0$-measurable and hence constant, as $\mathcal{F}_0$ is assumed to be trivial. We also assume throughout that we are dealing with a 'frictionless' market; that is, there are no transaction costs, unlimited short sales and borrowing are allowed (the random variables θ^i_t can take any real values), and the securities are perfectly divisible (the S^i_t can take any positive real values).

Self-Financing Strategies

We call the trading strategy θ *self-financing* if any changes in the value $V_t(\theta)$ result entirely from net gains (or losses) realised on the investments; the value of the portfolio after trading has occurred at time t and before stock prices at time $t+1$ are known is given by $\theta_{t+1} \cdot S_t$. If the total value of the portfolio has been used for these adjustments (i.e., there are no withdrawals and no new funds are invested), then this means that

$$\theta_{t+1} \cdot S_t = \theta_t \cdot S_t \text{ for all } t = 1, 2, \ldots, T-1. \tag{2.1}$$

Writing $\Delta X_t = X_t - X_{t-1}$ for any function X on $\mathbb{T}$, we can rewrite (2.1) at once as

$$\Delta V_t(\theta) = \theta_t \cdot S_t - \theta_{t-1} \cdot S_{t-1} = \theta_t \cdot S_t - \theta_t \cdot S_{t-1} = \theta_t \cdot \Delta S_t; \tag{2.2}$$

that is, the gain in value of the portfolio in the time interval $(t-1, t]$ is the scalar product in $\mathbb{R}^d$ of the new portfolio vector θ_t with the vector ΔS_t of price increments. Thus, defining the *gains process* associated with θ by setting

$$G_0(\theta) = 0, \qquad G_t(\theta) = \theta_1 \cdot \Delta S_1 + \theta_2 \cdot \Delta S_2 + \cdots + \theta_t \cdot \Delta S_t,$$

we see at once that θ is self-financing if and only if

$$V_t(\theta) = V_0(\theta) + G_t(\theta) \text{ for all } t \in \mathbb{T}. \tag{2.3}$$

This means that θ is self-financing if and only if the value $V_t(\theta)$ arises solely as the sum of the initial *endowment* $V_0(\theta)$ and the gains process $G_t(\theta)$ associated with the strategy θ.

We can write this relationship in yet another useful form: since $V_t(\theta) = \theta_t \cdot S_t$ for any $t \in \mathbb{T}$ and *any* strategy θ, it follows that we can write

$$\begin{aligned}
\Delta V_t &= V_t - V_{t-1} \\
&= \theta_t \cdot S_t - \theta_{t-1} \cdot S_{t-1} \\
&= \theta_t \cdot (S_t - S_{t-1}) + (\theta_t - \theta_{t-1}) \cdot S_{t-1} \\
&= \theta_t \cdot \Delta S_t + (\Delta\theta_t) \cdot S_{t-1}.
\end{aligned} \tag{2.4}$$

Thus, the strategy θ is self-financing if and only if

$$(\Delta\theta_t) \cdot S_{t-1} = 0. \tag{2.5}$$

This means that, for a self-financing strategy, the vector of changes in the portfolio θ is orthogonal in $\mathbb{R}^{d+1}$ to the *prior* price vector S_{t-1}. This property is sometimes easier to verify than (2.1). It also serves to justify the terminology: the cumulative effect of the time t variations in the investor's holdings (which are made *before* the time t prices are known) should be to balance each other. For example, if $d = 1$, we need to balance $\Delta\theta_t^0 S_{t-1}^0$ against $\Delta\theta_t^1 S_{t-1}^1$ since by (2.5) their sum must be zero.

Numéraire Invariance

Trivially, (2.1) and (2.3) each have an equivalent 'discounted' form. In fact, given any numéraire (i.e., any process (Z_t) with $Z_t > 0$ for all $t \in \mathbb{T}$), it follows that a trading strategy θ is self-financing relative to S if and only if it is self-financing relative to ZS since

$$(\Delta\theta_t) \cdot S_{t-1} = 0 \text{ if and only if } (\Delta\theta_t) \cdot Z_{t-1}S_{t-1} = 0 \text{ for } t \in \mathbb{T} \setminus \{0\}\,.$$

Thus, changing the choice of 'benchmark' security will not alter the class of trading strategies under consideration and thus will not affect market behaviour. This simple fact is sometimes called the 'numéraire invariance theorem'; in continuous-time models it is not completely obvious (see Chapter 9 and [102]). We will also examine the numéraire invariance of other market entities. While the use of different discounting conventions has only limited mathematical significance, economically it amounts to understanding the way in which these entities are affected by a change of currency.

Writing $\overline{X}_t = \beta_t X_t$ for the discounted form of the vector X_t in $\mathbb{R}^{d+1}$, it follows (using $Z = \beta$ in the preceding equation) that θ is self-financing if and only if $(\Delta\theta_t) \cdot \overline{S}_{t-1} = 0$, that is, if and only if

$$\theta_{t+1} \cdot \overline{S}_t = \theta_t \cdot \overline{S}_t \text{ for all } t = 1, 2, \ldots, T-1, \tag{2.6}$$

or, equivalently, if and only if

$$\overline{V}_t(\theta) = V_0(\theta) + \overline{G}_t(\theta) \text{ for all } \ t \in \mathbb{T}. \tag{2.7}$$

To see the last equivalence, note first that (2.4) holds for any θ with $\overline{S}$ instead of S, so that for self-financing θ we have $\Delta\overline{V}_t = \theta_t \cdot \Delta\overline{S}_t$; hence (2.7) holds. Conversely, (2.7) implies that $\Delta\overline{V}_t = \theta_t \cdot \Delta\overline{S}_t$, so that $(\Delta\theta_t) \cdot \overline{S}_{t-1} = 0$ and so θ is self-financing.

We observe that the definition of $\overline{G}(\theta)$ does not involve the amount θ_t^0 held in bonds (i.e., in the security S^0) at time t. Hence, if θ is self-financing, the initial investment $V_0(\theta)$ and the predictable real-valued processes θ^i $(i = 1, 2, \ldots, d)$ completely determine θ^0, just as we have seen in the one-period model in Section 1.4.

Lemma 2.2.1. *Given an $\mathcal{F}_0$-measurable function V_0 and predictable real-valued processes $\theta^1, \theta^2, \ldots, \theta^d$, the unique predictable process θ^0 that turns*

$$\theta = \left(\theta^0, \theta^1, \theta^2, \cdots, \theta^d\right)$$

into a self-financing strategy with initial value $V_0(\theta) = V_0$ is given by

$$\theta_t^0 = V_0 + \sum_{u=1}^{t-1} \left(\theta_u^1 \Delta\overline{S}_u^1 + \cdots + \theta_u^d \Delta\overline{S}_u^d\right) - \left(\theta_t^1 \overline{S}_{t-1}^1 + \cdots + \theta_t^d \overline{S}_{t-1}^d\right). \tag{2.8}$$

Proof. The process θ^0 so defined is clearly predictable. To see that it produces a self-financing strategy, recall by (2.7) that we only need to observe that this value of θ^0 is the unique predictable solution of the equation

$$\begin{aligned}\overline{V}_t(\theta) &= \theta_t^0 + \theta_t^1 \overline{S}_t^1 + \theta_t^2 \overline{S}_t^2 + \cdots + \theta_t^d \overline{S}_t^d \\ &= V_0 + \sum_{u=1}^{t} \left(\theta_u^1 \Delta\overline{S}_u^1 + \theta_u^2 \overline{S}_u^2 + \cdots + \theta_u^d \Delta\overline{S}_u^d\right).\end{aligned}$$

□

Admissible Strategies

Let Θ be the class of all self-financing strategies. So far, we have not insisted that a self-financing strategy must at all times yield non-negative total wealth; that is, that $V_t(\theta) \geq 0$ for all $t \in \mathbb{T}$. From now on, when we impose this additional restriction, we call such self-financing strategies *admissible*; they define the class Θ_a.

Economically, this requirement has the effect of restricting certain types of short sales: although we can still borrow certain of our assets (i.e., have $\theta_t^i < 0$ for some values of i and t), the overall value process must remain non-negative for each t. But the additional restriction has little impact on the mathematical modelling, as we show shortly.

We use the class Θ_a to define our concept of 'free lunch'.

Definition 2.2.2. An *arbitrage opportunity* is an admissible strategy θ such that

$$V_0(\theta) = 0, \qquad V_t(\theta) \geq 0 \text{ for all } t \in \mathbb{T}, \qquad E\left(V_T(\theta)\right) > 0.$$

In other words, we require $\theta \in \Theta_a$ with initial value 0 but final value strictly positive with positive probability. Note, however, that the probability measure P enters into this definition only through its null sets: the condition $E\left(V_T(\theta)\right) > 0$ is equivalent to $P(V_T(\theta)) > 0) > 0$, justifying the following definition.

Definition 2.2.3. The market model is *viable* if it does not contain any arbitrage opportunities; that is, if $\theta \in \Theta_a$ has $V_0(\theta) = 0$, then $V_T(\theta) = 0$ a.s..

'Weak Arbitrage Implies Arbitrage'

To justify the assertion that restricting attention to admissible claims has little effect on the modelling, we call a self-financing strategy $\theta \in \Theta$ a *weak arbitrage* if

$$V_0(\theta) = 0, \qquad V_T(\theta) \geq 0, \qquad E\left(V_T(\theta)\right) > 0.$$

The following calculation shows that if a weak arbitrage exists then it can be adjusted to yield an admissible strategy - that is, an arbitrage as defined in Definition 2.2.2.

Note. If the price process is a martingale under some equivalent measure-as will be seen shortly-then any hedging strategy with zero initial value and positive final expectation will automatically yield a positive expectation at all intermediate times by the martingale property.

Suppose that θ is a weak arbitrage and that $V_t(\theta)$ is *not* non-negative a.s. for all t. Then there exists $t < T$, and $A \in \mathcal{F}_t$ with $P(A) > 0$ such that

$$(\theta_t \cdot S_t)(\omega) < 0 \text{ for } \omega \in A, \theta_u \cdot S_u \geq 0 \;\text{ a.s. for } u > t.$$

We amend θ to a new strategy ϕ by setting $\phi_u(\omega) = 0$ for all $u \in \mathbb{T}$ and $\omega \in \Omega \setminus A$, while on A we set $\phi_u(\omega) = 0$ if $u \leq t$, and for $u > t$ we define

$$\phi_u^0(\omega) = \theta_u^0(\omega) - \frac{\theta_t \cdot S_t}{S_t^0(\omega)}, \phi_u^i(\omega) = \theta_u^i(\omega) \text{ for } i = 1, 2, \ldots, d.$$

This strategy is obviously predictable. It is also self-financing: on $\Omega \setminus A$ we clearly have $V_u(\phi) \equiv 0$ for all $u \in \mathbb{T}$, while on A we need only check that $(\Delta\phi_{t+1}) \cdot S_t = 0$ by the preceding construction (in which $\Delta\theta_u$ and $\Delta\phi_u$ differ only when $u = t + 1$) and (2.5). We observe that $\phi_t^i = 0$ on A^c for $i \geq 0$ and that, on A,

$$\Delta\phi_{t+1}^0 = \phi_{t+1}^0 = \theta_{t+1}^0 - \frac{\theta_t \cdot S_t}{S_t^0}, \Delta\phi_{t+1}^i = \theta_{t+1}^i \text{ for } i = 1, 2, \ldots, d.$$

Hence

$$(\Delta\phi_{t+1}) \cdot S_t = \mathbf{1}_A(\theta_{t+1} \cdot S_t - \theta_t \cdot S_t) = \mathbf{1}_A(\theta_t \cdot S_t - \theta_t \cdot S_t) = 0$$

since θ is self-financing.

We show that $V_u(\phi) \geq 0$ for all $u \in \mathbb{T}$, and $P(V_T(\phi) > 0) > 0$. First note that $V_u(\phi) = 0$ on $\Omega \setminus A$ for all $u \in \mathbb{T}$. On A we also have $V_u(\phi) = 0$ when $u \leq t$, but for $u > t$ we obtain

$$V_u(\phi) = \phi_u \cdot S_u = \theta_u^0 S_u^0 - \frac{(\theta_t \cdot S_t) S_u^0}{S_t^0} + \sum_{i=1}^d \theta_u^i S_u^i = \theta_u \cdot S_u - (\theta_t \cdot S_t)\left(\frac{S_u^0}{S_t^0}\right).$$

Since, by our choice of t, $\theta_u \cdot S_u \geq 0$ for $u > t$, and $(\theta_t \cdot S_t) < 0$ while $S^0 \geq 0$, it follows that $V_u(\phi) \geq 0$ for all $u \in \mathbb{T}$. Moreover, since $S_t^0 > 0$, we also see that $V_T(\phi) > 0$ on A.

This construction shows that the existence of what we have called weak arbitrage immediately implies the existence of an arbitrage opportunity. This fact is useful in the fine structure analysis for finite market models we give in the next chapter.

Remark 2.2.4. Strictly speaking, we should deal separately with the possibility that the investor's initial capital is negative. This is of course ruled out if we demand that all trading strategies are admissible. We can relax this condition and consider a one-period model, where a trading strategy is just a portfolio θ, chosen at the outset with knowledge of time 0 prices and held throughout the period. In that case, an arbitrage is a portfolio that leads from a non-positive initial outlay to a non-negative value at time 1. Thus here we have two possible types of arbitrage since the portfolio θ leads to one of two conclusions:

a) $V_0(\theta) < 0$ and $V_1(\theta) \geq 0$ or

b) $V_0(\theta) = 0$ and $V_1(\theta) \geq 0$ and $P(V_1(\theta) > 0) > 0$.

In this setting, the assumption that there are no arbitrage opportunities leads to two conditions on the prices:

(i) $V_1(\theta) = 0$ implies $V_0(\theta) = 0$ or

(ii) $V_1(\theta) \geq 0$ and $P(V_1(\theta)) > 0$ implies $V_0(\theta) \geq 0$.

The reader will easily construct arbitrages if either of these conditions fails. In our treatment of multi-period models, we consistently use admissible strategies, so that Definition 2.2.3 is sufficient to define the viability of pricing models.

Uniqueness of the Arbitrage Price

Fix H as a contingent claim with maturity T so H is a non-negative $\mathcal{F}_T$-measurable random variable on $(\Omega, \mathcal{F}_T, P)$. The claim is said to be *attainable* if there is an admissible strategy θ that *generates* (or *replicates*) it, that is, such that

$$V_T(\theta) = H.$$

We should expect the value process associated with a generating strategy to be given *uniquely*: the existence of two admissible strategies θ and θ' with $V_t(\theta) \neq V_t(\theta')$ would violate the *Law of One Price*, and the market would therefore allow riskless profits and not be viable. A full discussion of these economic arguments is given in [241].

The next lemma shows, conversely, that in a viable market the *arbitrage price* of a contingent claim is indeed unique.

Lemma 2.2.5. *Suppose H is an attainable contingent claim in a viable market model. Then the value processes of all generating strategies for H are the same.*

Proof. If θ and ϕ are admissible strategies with

$$V_T(\theta) = H = V_T(\phi)$$

but $V(\theta) \neq V(\phi)$, then there exists $t < T$ such that

$$V_u(\theta) = V_u(\phi) \text{ for all } u < t, \qquad V_t(\theta) \neq V_t(\phi).$$

The set $A = \{V_t(\theta) > V_t(\phi)\}$ is in $\mathcal{F}_t$ and we can assume $P(A) > 0$ without loss of generality. The random variable $X = V_t(\theta) - V_t(\phi)$ is $\mathcal{F}_t$-measurable and defines a self-financing strategy ψ as by letting

$$\psi_u(\omega) = \theta_u(\omega) - \phi_u(\omega) \text{ for } u \leq t \text{ on } A, \text{ for } u \in \mathbb{T}, \text{ on } A^c,$$
$$\psi_u^0 = \beta_t X \text{ and } \psi_u^i = 0 \text{ for } i = 1, 2, \ldots, d \text{ for } u > t, \text{ on } A.$$

It is clear that ψ is predictable. Since both θ and ϕ are self-financing, it follows that (2.1) also holds with ψ for $u < t$, while if $u > t$, $\psi_{u+1} \cdot S_u =$

$\psi_u \cdot S_u$ on A^c similarly. On A, we have $\psi_{u+1} = \psi_u$. Thus we only need to compare $\psi_t \cdot S_t = V_t(\theta) - V_t(\phi)$ and $\psi_{t+1} \cdot S_t = \mathbf{1}_{A^c}(\theta_{t+1} - \phi_{t+1}) \cdot S_t + \mathbf{1}_A \beta_t X S_t^0$. Now note that $S_t^0 = \beta_t^{-1}$ and that $X = V_t(\theta) - V_t(\phi)$, while on A^c the first term becomes $(\theta_t - \phi_t) \cdot S_t = V_t(\theta) - V_t(\phi)$ and the latter vanishes. Thus $\psi_{t+1} \cdot S_t = V_t(\theta) - V_t(\phi) = \psi_t \cdot S_t$.

Since $V_0(\theta) = V_0(\phi)$, ψ is self-financing with initial value 0. But $V_T(\psi) = \mathbf{1}_A(\beta_t X S_t^0) = \mathbf{1}_A \beta_t \beta_T^{-1} X$ is non-negative a.s. and is strictly positive on A, which has positive probability. Hence ψ is a weak arbitrage, and by the previous section the market cannot be viable. $\square$

We have shown that in a viable market it is possible to associate a unique time t value (or *arbitrage price*) to any attainable contingent claim H. However, it is not yet clear how the generating strategy, and hence the price, are to be found in particular examples. In the next section, we characterise viable market models without having to construct explicit strategies and derive a general formula for the arbitrage price instead.

2.3 Martingales and Risk-Neutral Pricing

Martingales and Their Transforms

We wish to characterise viable market models in terms of the behaviour of the *increments* of the discounted price process $\overline{S}$. To set the scene, we first need to recall some simple properties of martingales. Only the most basic results needed for our purposes are described here; for more details consult, for example, [109], [199], [236], [299].

For these results, we take a general probability space $(\Omega, \mathcal{F}, P)$together with any filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in \mathbb{T}}$, where, as before, $\mathbb{T} = \{0, 1, \ldots, T\}$. Consider stochastic processes defined on this *filtered probability space* (also called *stochastic basis*) $(\Omega, \mathcal{F}, P, \mathbb{F}, \mathbb{T})$. Recall that a stochastic process $X = (X_t)$ is *adapted* to $\mathbb{F}$ if X_t is $\mathcal{F}_t$-measurable for each $t \in \mathbb{T}$.

Definition 2.3.1. An $\mathbb{F}$-adapted process $M = (M_t)_{t \in \mathbb{T}}$ is an $(\mathbb{F}, P)$-*martingale* if $E(|M_t|) < \infty$ for all $t \in \mathbb{T}$ and

$$E(M_{t+1} \,|\mathcal{F}_t) = M_t \text{ for all } t \in \mathbb{T} \setminus \{T\}. \tag{2.9}$$

If the equality in (2.9) is replaced by $\leq$ ($\geq$), we say that M is a *supermartingale* (*submartingale*).

Note that M is a martingale if and only if

$$E(\Delta M_{t+1} \,|\mathcal{F}_t) = 0 \text{ for all } t \in \mathbb{T} \setminus \{T\}.$$

Thus, in particular, $E(\Delta M_{t+1}) = 0$. Hence

$$E(M_{t+1}) = E(M_t) \text{ for all } t \in \mathbb{T} \setminus \{T\},$$

so that a martingale is 'constant on average'. Similarly, a submartingale increases, and a supermartingale decreases, on average. Thinking of M_t as representing the current capital of a gambler, a martingale therefore models a 'fair' game, while sub- and supermartingales model 'favourable' and 'unfavourable' games, respectively (as seen from the perspective of the gambler, of course!).

The linearity of the conditional expectation operator shows trivially that any linear combination of martingales is a martingale, and the tower property shows that M is a martingale if and only if

$$E\left(M_{s+t} \,|\mathcal{F}_s\right) = M_s \text{ for } t = 1, 2, \ldots, T - s.$$

Moreover, (M_t) is a martingale if and only if $(M_t - M_0)$ is a martingale, so we can assume $M_0 = 0$ without loss whenever convenient.

Many familiar stochastic processes are martingales. The simplest example is given by the successive conditional expectations of a single integrable random variable X. Set $M_t = E\left(X \,|\mathcal{F}_t\right)$ for $t \in \mathbb{T}$. By the tower property,

$$E\left(M_{t+1} \,|\mathcal{F}_t\right) = E\left(E\left(X \,|\mathcal{F}_{t+1}\right) |\mathcal{F}_t\right) = E\left(X \,|\mathcal{F}_t\right) = M_t.$$

The values of the martingale M_t are successive best mean-square estimates of X, as our 'knowledge' of X, represented by the σ-fields $\mathcal{F}_t$, increases with t.

More generally, if we model the price process of a stock by a martingale M, the conditional expectation (i.e., our best mean-square estimate at time s of the future value M_t of the stock) is given by its current value M_s. This generalises a well-known fact about processes with independent increments: if the zero-mean process W is adapted to the filtration $\mathbb{F}$ and $(W_{t+1} - W_t)$ is independent of $\mathcal{F}_t$, then $E\left(W_{t+1} - W_t \,|\mathcal{F}_t\right) = E\left(W_{t+1} - W_t\right) = 0$. Hence W is a martingale.

Exercise 2.3.2. Suppose that the centred (i.e., zero-mean) integrable random variables $(Y_t)_{t \in \mathbb{T}}$ are independent, and let $X_t = \sum_{u \leq t} Y_u$ for each $t \in \mathbb{T}$. Show that X is a martingale for the filtration it generates. What can we say when the Y_t have positive means?

Exercise 2.3.3. Let $(Z_n)_{n \geq 1}$ be independent identically distributed random variables, adapted to a given filtration $(\mathcal{F}_n)_{n \geq 0}$. Suppose further that each Z_n is non-negative and has mean 1. Show that $X(0) = 1$ and that $X_n = Z_1 Z_2 \cdots Z_n$ $(n \geq 1)$ defines a martingale for $(\mathcal{F}_n)$, provided all the products are integrable random variables, which holds, for example, if all $Z_n \in \mathcal{L}^\infty(\Omega, \mathcal{F}, P)$.

Note also that any *predictable* martingale is almost surely constant: if M_{t+1} is $\mathcal{F}_t$-measurable, we have $E\left(M_{t+1} \,|\mathcal{F}_t\right) = M_{t+1}$ and hence M_t and M_{t+1} are a.s. equal for all $t \in \mathbb{T}$. This is no surprise: if at time t we know the value of M_{t+1}, then our best estimate of that value will be perfect.

The construction of the gains process associated with a trading strategy now suggests the following further definition.

Definition 2.3.4. Let $M = (M_t)$ be a martingale and $\phi = (\phi_t)_{t\in\mathbb{T}}$ a predictable process defined on $(\Omega, \mathcal{F}, P, \mathbb{F}, \mathbb{T})$. The process $X = \phi \cdot M$ given for $t \geq 1$ by

$$X_t = \phi_1 \Delta M_1 + \phi_2 \Delta M_2 + \cdots + \phi_t \Delta M_t \tag{2.10}$$

and

$$X_0 = 0$$

is the *martingale transform* of M by ϕ.

Martingale transforms are the discrete analogues of the stochastic integrals in which the martingale M is used as the 'integrator'. The Itô calculus based upon this integration theory forms the mathematical backdrop to martingale pricing in continuous time, which comprises the bulk of this book. An understanding of the technically much simpler martingale transforms provides valuable insight into the essentials of stochastic calculus and its many applications in finance theory.

The Stability Property

If $\phi = (\phi_t)_{t\in\mathbb{T}}$ is bounded and predictable, then ϕ_{t+1} is $\mathcal{F}_t$-measurable and $\phi_{t+1}\Delta M_{t+1}$ remains integrable. Hence, for each $t \in \mathbb{T} \setminus \{T\}$, we have

$$E\left(\Delta X_{t+1} \left| \mathcal{F}_t \right.\right) = E\left(\phi_{t+1}\Delta M_{t+1} \left| \mathcal{F}_t \right.\right) = \phi_{t+1} E\left(\Delta M_{t+1} \left| \mathcal{F}_t \right.\right) = 0.$$

Therefore $X = \phi \cdot M$ is a martingale with $X_0 = 0$. Similarly, if ϕ is also non-negative and Y is a supermartingale, then $\phi \cdot Y$ is again a supermartingale.

This stability under transforms provides a simple, yet extremely useful, characterisation of martingales.

Theorem 2.3.5. *An adapted real-valued process M is a martingale if and only if*

$$E\left((\phi \cdot M)_t\right) = E\left(\sum_{u=1}^{t} \phi_u \Delta M_u\right) = 0 \textit{ for } t \in \mathbb{T} \setminus \{0\} \tag{2.11}$$

for each bounded predictable process ϕ.

Proof. If M is a martingale, then so is the transform $X = \phi \cdot M$, and $X_0 = 0$. Hence $E\left((\phi \cdot M)_t\right) = 0$ for all $t \geq 1$ in $\mathbb{T}$.

Conversely, if (2.11) holds for M and every predictable ϕ, take $s > 0$, let $A \in \mathcal{F}_s$ be given, and define a predictable process ϕ by setting $\phi_{s+1} = \mathbf{1}_A$, and $\phi_t = 0$ for all other $t \in \mathbb{T}$. Then, for $t > s$, we have

$$0 = E\left((\phi \cdot M)_t\right) = E(\mathbf{1}_A(M_{s+1} - M_s)).$$

Since this holds for all $A \in \mathcal{F}_s$, it follows that $E\left(\Delta M_{s+1} \left| \mathcal{F}_s \right.\right) = 0$, so M is a martingale. $\square$

2.4 Arbitrage Pricing: Martingale Measures

Equivalent Martingale Measures

We now return to our study of viable securities market models. Recall that we assume as given an arbitrary complete measurable space $(\Omega, \mathcal{F})$ on which we consider various probability measures. We also consider a filtration $\mathbb{F} = (\mathcal{F}_t)_{t\in\mathbb{T}}$ such that $(\Omega, \mathcal{F}_0)$ is complete, and $\mathcal{F}_T = \mathcal{F}$. Finally, we are given a $(d+1)$-dimensional stochastic process $S = \{S_t^i : t \in \mathbb{T}, 0 \leq i \leq d\}$ with $S_0^0 = 1$ and S^0 interpreted as a riskless bond providing a discount factor $\beta_t = \frac{1}{S_t^0}$ and with S^i $(i = 1, 2, \ldots, d)$ interpreted as risky stocks. Recall that we are working in a general securities market model: we do *not* assume that the resulting market model is finite or that the filtration $\mathbb{F}$ is generated by S.

Suppose that the discounted vector price process $\bar{S}$ happens to be a martingale under some probability measure Q; that is,

$$E_Q\left(\Delta \bar{S}_t^i | \mathcal{F}_{t-1}\right) = 0 \text{ for } t \in \mathbb{T} \setminus \{0\} \text{ and } i = 1, 2, \ldots, d.$$

Note that, in particular, this assumes that the discounted prices are integrable with respect to Q. Suppose that $\theta = \left\{\theta_t^i : i \leq d, t = 1, 2, \ldots, T\right\} \in \Theta_a$ is an admissible strategy whose discounted value process is also Q-integrable for each t. Recall from (2.7) that the discounted value process of θ has the form

$$\begin{aligned}
\bar{V}_t(\theta) &= V_0(\theta) + \overline{G}_t(\theta) \\
&= \theta_1 \cdot S_0 + \sum_{u=1}^{t} \theta_u \cdot \Delta \bar{S}_u \\
&= \sum_{i=1}^{d} \left(\theta_1^i S_0^i + \sum_{u=1}^{t} \theta_u \Delta \overline{S}_u^i \right).
\end{aligned}$$

Thus the discounted value process $\overline{V}(\theta)$ is a constant plus a finite sum of martingale transforms; and therefore it is a martingale with initial (constant) value $V_0(\theta)$. Hence we have $E\left(\overline{V}_t(\theta)\right) = E\left(V_0(\theta)\right) = V_0(\theta)$.

We want to show that this precludes the existence of arbitrage opportunities. If we know in advance that the value process of every admissible strategy is integrable with respect to Q, this is easy: if $V_0(\theta) = 0$ and $V_T(\theta) \geq 0$ a.s. (Q), but $E_Q\left(\overline{V}_t(\theta)\right) = 0$, it follows that $V_T(\theta) = 0$ a.s. (Q). This remains true a.s. (P), provided that the probability measure Q has the same null sets as P (we say that Q and P are *equivalent measures* and write $Q \sim P$). If such a measure can be found, then no self-financing strategy θ can lead to arbitrage; that is, the market is viable. This leads to an important definition.

Definition 2.4.1. A probability measure $Q \sim P$ is an *equivalent martingale measure* (EMM) for S if the discounted price process $\overline{S}$ is a (vector) martingale under Q for the filtration $\mathbb{F}$. That is, for each $i \leq d$ the discounted price process $\overline{S}^i$ is an $(\mathbb{F}, Q)$-martingale (recall that $\overline{S}^0 \equiv 1$).

To complete the argument, we need to justify the assumption that the value processes we have considered are Q-integrable. This follows from the following remarkable proposition (see also [132]).

Proposition 2.4.2. *Given a viable model $(\Omega, \mathcal{F}, P, \mathbb{T}, \mathbb{F}, S)$, suppose that Q is an equivalent martingale measure for S. Let H be an attainable claim. Then $\beta_T H$ is Q-integrable and the discounted value process for any generating strategy θ satisfies*

$$\overline{V}_t(\theta) = E_Q\left(\beta_T H \,|\mathcal{F}_t\right) \ \ a.s.\,(P) \ \ for\ all\ t \in \mathcal{F}. \tag{2.12}$$

Thus $\overline{V}(\theta)$ is a non-negative Q-martingale.

Proof. Choose a generating strategy θ for H and let $\overline{V} = \overline{V}(\theta)$ be its discounted value process. We show by backward induction that $\overline{V}_t \geq 0$ a.s. $(P)$for each t. This is clearly true for $t = T$ since $\overline{V}_T = \beta_T H \geq 0$ by definition. Hence suppose that $\overline{V}_t \geq 0$. If θ_t is unbounded, replace it by the bounded random vectors $\theta_t^n = \theta_t \mathbf{1}_{A_n}$, where $A_n = \{|\theta_t| \leq n\}$, so that $\overline{V}_{t-1}(\theta^n) = \overline{V}_{t-1}(\theta)\mathbf{1}_{A_n}$ is $\mathcal{F}_{t-1}$-measurable and Q-integrable. Then we can write

$$\overline{V}_{t-1}(\theta^n) = \overline{V}_t(\theta^n) - \sum_{i=1}^d \theta_t^{n,i} \Delta\overline{S}_t^i \geq - \sum_{i=1}^d \theta_t^{n,i} \Delta\overline{S}_t^i,$$

so that

$$\begin{aligned}
\overline{V}_{t-1}(\theta)\mathbf{1}_{A_n} &= \overline{V}_{t-1}(\theta^n) \\
&= E_Q\left(\overline{V}_{t-1}(\theta^n) \,|\mathcal{F}_{t-1}\right) \\
&\geq - \sum_{i=1}^d \theta_t^{n,i} E_Q\left(\Delta\overline{S}_t^i \,|\mathcal{F}_{t-1}\right) \\
&= 0.
\end{aligned}$$

Letting n increase to ∞, we see that $\overline{V}_{t-1}(\theta) \geq 0$.

Thus we have a.s. (P) on each A_n that

$$\begin{aligned}
E_Q\left(\overline{V}_t(\theta) \,|\mathcal{F}_{t-1}\right) - \overline{V}_{t-1}(\theta) &= E_Q\left(\sum_{i=1}^d \theta_t^{n,i} \Delta\overline{S}_t^i \,|\mathcal{F}_{t-1}\right) \\
&= \sum_{i=1}^d \theta_t^{n,i} E_Q\left(\Delta\overline{S}_t^i \,|\mathcal{F}_{t-1}\right) \\
&= 0.
\end{aligned}$$

Again letting n increase to ∞, we have the identity

$$E_Q\left(\overline{V}_t(\theta)\,|\mathcal{F}_{t-1}\right) = \overline{V}_{t-1}(\theta) \text{ a.s.}\,(P)\,. \tag{2.13}$$

Finally, as $V_0 = \theta_1 \cdot S_0$ is a non-negative constant, it follows that $E_Q\left(\overline{V}_1\right) = V_0$. But by the first part of the proof $\overline{V}_1 \geq 0$ a.s. (P) and hence a.s. (Q), so $\overline{V}_1 \in L^1(Q)$. We can therefore begin an induction, using (2.13) at the inductive step, to conclude that $\overline{V}_t \in L^1(Q)$ and $E_Q\left(\overline{V}_t(\theta)\right) = V_0$ for all $t \in \mathbb{T}$. Thus $\overline{V}(\theta)$ is a non-negative Q-martingale, and since its final value is $\beta_T H$, it follows that $\overline{V}_t(\theta) = E_Q\left(\beta_T H\,|\mathcal{F}_t\right)$ a.s. (P) for each $t \in \mathbb{T}$. □

Remark 2.4.3. The identity (2.12) not only provides an alternative proof of Lemma 2.2.5 by showing that the price of any attainable European claim is independent of the particular generating strategy, since the right-hand side does not depend on θ, but also provides a means of calculating that price without having to construct such a strategy. Moreover, the price does not depend on the choice of any particular equivalent martingale measure: the left-hand side does not depend on Q.

Exercise 2.4.4. Use Proposition 2.4.2 to show that if θ is a self-financing strategy whose final discounted value is bounded below a.s. $(P)$by a constant, then for any EMM Q the expected final value of θ is simply its initial value. What conclusion do you draw for trading only with strategies that have bounded risk?

We have proved that the existence of an equivalent martingale measure for S is *sufficient* for viability of the securities market model. In the next chapter, we discuss the *necessity* of this condition. Mathematically, the search for equivalent measures under which the given process $\overline{S}$ is a martingale is often much more convenient than having to show that no arbitrage opportunities exist for $\overline{S}$.

Economically, we can interpret the role of the martingale measure as follows. The probability assignments that investors make for various events do not enter into the derivation of the arbitrage price; the only criterion is that agents prefer more to less and would therefore become arbitrageurs if the market allowed arbitrage. The price we derive for the contingent claim H must thus be the same for all risk preferences (probability assignments) of the agents as long as they preclude arbitrage. In particular, an economy of risk-neutral agents will also produce the arbitrage price we derived previously. The equivalent measure Q, under which the discounted price process is a martingale represents the probability assignment made in this risk-neutral economy, and the price that this economy assigns to the claim will simply be the average (i.e., expectation under Q) discounted value of the payoff H.

Thus the existence of an equivalent martingale measure provides a general method for pricing contingent claims, which we now also formulate in terms of undiscounted value processes.

Martingale Pricing

We summarise the role played by martingale measures in pricing claims. Assume that we are given a viable market model $(\Omega, \mathcal{F}, P, \mathbb{F}, S)$ and some equivalent martingale measure Q. Recall that a *contingent claim* in this model is a non-negative ($\mathcal{F}$-measurable) random variable H representing a contract that pays out $H(\omega)$ dollars at time T if $\omega \in \Omega$ occurs. Its time 0 value or (current) *price* $\pi(H)$ is then the value that the parties to the contract would deem a 'fair price' for entering into this contract.

In a viable model, an investor could hope to evaluate $\pi(H)$ by constructing an admissible trading strategy $\theta \in \Theta_a$ that exactly replicates the returns (cash flow) yielded by H at time T. For such a strategy θ, the initial investment $V_0(\theta)$ would represent the price $\pi(H)$ of H. Recall that H is an *attainable claim* in the model if there exists a *generating strategy* $\theta \in \Theta_a$ such that $V_T(\theta) = H$, or, equivalently, $\overline{V}_t(\theta) = \beta_T H$. But as Q is a martingale measure for S, $\overline{V}(\theta)$ is, up to a constant, a martingale transform, and hence a martingale, under Q, it follows that for all $t \in \mathbb{T}$,

$$\overline{V}_t(\theta) = E_Q\left(\beta_T H \,|\mathcal{F}_t\right),$$

and thus

$$V_t(\theta) = \beta_t^{-1} E_Q\left(\beta_T H \,|\mathcal{F}_t\right) \tag{2.14}$$

for any $\theta \in \Theta_a$. In particular,

$$\pi(H) = \overline{V}_0(\theta) = E_Q\left(\beta_T H \,|\mathcal{F}_0\right) = E_Q\left(\beta_T H\right). \tag{2.15}$$

Market models in which all European contingent claims are attainable are called *complete*. These models provide the simplest class in terms of option pricing since any contingent claim can be priced simply by calculating its (discounted) expectation relative to an equivalent martingale measure for the model.

Uniqueness of the EMM

We have shown in Proposition 2.4.2 that for an attainable European claim H the identity $\bar{V}_0(\theta) = E_Q\left(\beta_T H\right)$ holds for *every* EMM Q in the model and for every replicating strategy θ.

This immediately implies that in a complete model the EMM must be unique. For if Q and R are EMMs in a complete pricing model, then any European claim is attainable. It follows that $E_Q\left(\beta_T H\right) = E_R\left(\beta_T H\right)$ and hence also

$$E_Q\left(H\right) = E_R\left(H\right), \tag{2.16}$$

upon multiplying both sides by β_T, which is non-random. In particular, equation (2.16) holds when the claim is the indicator function of an arbitrary set $F \in \mathcal{F}_T = \mathcal{F}$. This means that $Q = R$; hence the EMM is unique. Moreover, our argument again verifies that the *Law of One Price* (see Lemma 2.2.5) must hold in a viable model; that is, we cannot have two admissible trading strategies θ, θ' that satisfy $V_T(\theta) = V_T(\theta')$ but $V_0(\theta) \neq V_0(\theta')$. Our modelling assumptions are thus sufficient to guarantee consistent pricing mechanisms (in fact, this consistency criterion is strictly weaker than viability; see [241] for simple examples).

The Law of One Price permits valuation of an attainable claim H through the initial value of a self-financing strategy that generates H; the valuation technique using risk-neutral expectations gives the price $\pi(H)$ *without* prior determination of such a generating strategy. In particular, consider a single-period model and a claim H (an Arrow-Debreu security) defined by

$$H(\omega) = \begin{cases} 1 & \text{if } \omega = \omega' \\ 0 & \text{otherwise,} \end{cases}$$

where $\omega' \in \Omega$ is some specified state. If H is attainable, then

$$\pi(H) = E_Q(\beta_T H) = \frac{1}{\beta_T} Q(\{\omega'\}).$$

This holds even when β is random. The ratio $\frac{Q(\{\omega'\})}{\beta_T(\omega')}$ is known as the *state price* of ω'. In a finite market model, we can similarly define the change of measure density $\Lambda = \Lambda(\{\omega\})_{\omega\in\Omega}$, where $\Lambda(\{\omega\}) = \frac{Q(\{\omega\})}{P(\{\omega\})}$) as the *state price density*. See [241] for details of the role of these concepts.

Superhedging

We adopt a slightly more general approach (which we shall develop further in Chapter 5 and exploit more fully for continuous-time models in Chapters 7 to 10) to give an explicit justification of the 'fairness' of the option price when viewed from the different perspectives of the buyer and the seller (option writer), respectively.

Definition 2.4.5. Given a European claim $H = f(S_T)$, an $(x, H)-$*hedge* is an initial investment x in an admissible strategy θ such that $V_T(\theta) \geq H$ a.s.

This approach to hedging is often referred to as defining a *superhedging strategy.* This clearly makes good sense from the seller's point of view, particularly for claims of American type, where the potential liability may not always be covered exactly by replication. By investing x in the strategy θ at time 0, an investor can cover his potential liabilities whatever the stock price movements in $[0, T]$. When there is an admissible strategy θ exactly replicating H, the initial investment $x = \pi(H)$ is an example of an $(x, H)-$hedge. Since the strategy θ exactly covers the final liabilities, (i.e., $V_T(\theta) = H$), we call this a *minimal hedge.*

All prices acceptable to the option seller must clearly ensure that the initial receipts for the option enable him to invest in a hedge (i.e., must ensure that there is an admissible strategy whose final value is at least H). The *seller's price* can thus be defined as

$$\pi_s = \inf \{z \geq 0 : \text{there exists } \theta \in \Theta_a \text{ with } V_T(\theta) = z + G_T(\theta) \geq H \text{ a.s.}\}.$$

The buyer, on the other hand, wants to pay no more than is needed to ensure that his final wealth suffices to cover the initial outlay, or borrowings. So his price will be the maximum he is willing to borrow, $y = -V_0$, at time 0 to invest in an admissible strategy θ, so that the sum of the option payoff and the gains from following θ cover his borrowings. The *buyer's price* is therefore

$$\pi_b = \sup \{y \geq 0 : \text{there exists } \theta \in \Theta_a \text{ with } -y + G_T(\theta) \geq -H \text{ a.s.}\}.$$

In particular, θ must be self-financing, so that $\beta_T V_T(\theta) = V_0 + \beta_T G_T(\theta)$, and since βS is a Q-martingale, we have $E_Q(\beta_T G_T(\theta)) = 0$. So the seller's price requires that $z \geq E_Q(\beta_T H)$ for each z in (2.21)and hence $\pi_s \geq E_Q(\beta_T H)$.

Similarly, for the buyer's price, we require that $-y + E_Q(\beta_T H) \geq 0$ and hence also $\pi_b \leq E_Q(\beta_T H)$. We have proved the following proposition.

Proposition 2.4.6. *For any integrable European claim H in a viable pricing model,*

$$\pi_b \leq E_Q(\beta_T H) \leq \pi_s. \tag{2.17}$$

If the claim H is attained by an admissible strategy θ, the minimal initial investment z in the strategy θ that will yield final wealth $V_T(\theta) = H$ is given by $E_Q(\beta_T H)$, and conversely this represents the maximal initial borrowing y required to ensure that $-y + G_T(\theta) + H \geq 0$. This proves the following corollary.

Corollary 2.4.7. *If the European claim H is attainable, then the buyer's price and seller's price are both equal to $E_Q(\beta_T H)$. Thus, in a complete model, every European claim H has a unique price, given by $\pi = E_Q(\beta_T H)$, and the generating strategy θ for the claim is a minimal hedge.*

2.5 Strategies Using Contingent Claims

Our definition of arbitrage involves trading strategies that include only primary securities (i.e., a riskless bank account which acts as numéraire and a collection of risky assets, which we called 'stocks' for simplicity). Our analysis assumes that these assets are traded independently of other assets. In real markets, however, investors also have access to derivative (or secondary) securities, whose prices depend on those of some underlying assets. We have grouped these under the term 'contingent claim' and we have considered how such assets should be priced. Now we need to consider an extended concept of arbitrage since it is possible for an investor to build a trading strategy including both primary securities and contingent claims, and we use this combination to seek to secure a riskless profit. We must therefore identify circumstances under which the market will preclude such profits.

Thus our concept of a trading strategy should be extended to include such combinations of primary and secondary securities, and we shall show that the market remains viable precisely when the contingent claims are priced according to the martingale pricing techniques for European contingent claims that we have developed. To achieve this, we need to restrict attention to trading strategies involving a bank account, stocks, and *attainable European* contingent claims.

Assume that a securities market model $(\Omega, \mathcal{F}, P, \mathbb{T}, \mathbb{F}, S)$ is given. We allow trading strategies to include attainable European claims, so that the value of the investor's portfolio at time $t \in \mathbb{T}$ will have the form

$$V_t = \theta_t \cdot S_t + \gamma_t \cdot Z_t = \sum_{i=0}^{d} \theta_t^i S_t^i + \sum_{j=1}^{m} \gamma_t^j Z_t^j, \tag{2.18}$$

where S^0 is the bank account, $\left\{S_t^i : i = 1, 2, \ldots, d\right\}$ are the prices of d risky stocks, and $Z_t = (Z_t^j)_{j \leq m}$ are the values of m attainable European contingent claims with time T payoff functions given by $(Z^j)_{j \leq m}$. We write $S = (S^i)_{0 \leq i \leq d}$. Recall that an attainable claim Z^j can be replicated exactly by a self-financing strategy involving only the process S. The holdings of each asset are assumed to be predictable processes, so that for $t = 1, 2, \ldots, T$, θ_t^i and γ_t^j are $\mathcal{F}_{t-1}$-measurable for $i = 0, 1, \ldots, d$ and $j = 1, 2, \ldots, m$. We call our model an *extended securities market model.*

The trading strategy $\phi = (\theta, \gamma)$ *is* self-financing if its initial value is

$$V_0(\phi) = \theta_1 \cdot S_0 + \gamma_1 \cdot Z_0$$

and for $t = 1, 2, \ldots, T-1$ we have

$$\theta_t \cdot S_t + \gamma_t \cdot Z_t = \theta_{t+1} \cdot S_t + \gamma_{t+1} \cdot Z_t. \tag{2.19}$$

Note that $\cdot$ denotes the inner product in $\mathbb{R}^{d+1}$ and $\mathbb{R}^m$, respectively. A new feature of the extended concept of a trading strategy is that the final values of some of its components are known in advance since the final portfolio has value

$$V_T(\phi) = \theta_T \cdot S_T + \gamma_T \cdot Z,$$

as $Z = (Z_T^j)_{j \leq m}$ represents the m payoff functions of the European claims. Moreover, unlike stocks, we have to allow for the possibility that the values Z_t^j can be zero or negative (as can be the case with forward contracts). However, with these minor adjustments we can regard the model simply as a securities market model with one riskless bank account and $d+m$ risky assets. With this in mind, we extend the concept of arbitrage to this model.

Definition 2.5.1. An arbitrage opportunity in the extended securities market model is a self-financing trading strategy ϕ such that $V_0(\phi)=0$, $V_T(\phi)\geq 0$, and $E_P\left(V_T(\phi)\right)>0$. We call the model *arbitrage-free* if no such strategy exists.

As in the case of weak arbitrage in Section 2.2, we do not demand that the value process remain non-negative throughout $\mathbb{T}$. That this has no effect on the pricing of the contingent claims can be seen from the following result.

Theorem 2.5.2. *Suppose that $(\Omega,\mathcal{F},P,\mathbb{T},\mathbb{F},S)$ is an extended securities market model admitting an equivalent martingale measure Q. The model is arbitrage-free if and only if every attainable European contingent claim with payoff Z has value process given by* $\left\{S_t^0 E_Q\left(\frac{Z}{S_T^0}|\mathcal{F}_t\right): t\in\mathbb{T}\right\}$.

Proof. Let $\theta=(\theta^i)_{i\leq d}$ be a generating strategy for Z. The value process of θ is then given as in equation (2.14) by $V_t(\theta)=S_t^0 E_Q\left(\frac{Z}{S_T^0}|\mathcal{F}_t\right)$ since the discount process is $\beta_t=\frac{1}{S_t^0}$ when S^0 is the numeraire.

We need to show that the model is arbitrage-free precisely when the value process $(Z_t)_{t\in\mathbb{T}}$ of the claim Z is equal to $(V_t(\theta))_{t\in\mathbb{T}}$. Suppose therefore that for some $u\in\mathbb{T}$ these processes differ on a set D of positive P-measure. We first assume that $D=\{Z_u>V_u(\theta)\}$, which belongs to $\mathcal{F}_u$. To construct an arbitrage, we argue as follows: do nothing for $\omega\notin D$, and for $\omega\in D$ wait until time u. At time u, sell short one unit of the claim Z for $Z_u(\omega)$, invest $V_u(\omega)$ of this in the portfolio of stocks and bank account according to the prescriptions given by strategy θ, and bank the remainder $(Z_u(\omega)-V_u(\omega))$ until time T. This produces a strategy ϕ, where

$$\phi_t=\begin{cases}0 & \text{if } t\leq u\\ \left(\theta_t^0+\frac{Z_u-V_u(\theta)}{S_u^0},\theta_t^1,\ldots,\theta_t^d,-1\right)\mathbf{1}_D & \text{if } t>u.\end{cases}$$

It is not hard to show that this strategy is self-financing; it is evidently predictable. Its value process $V\left(\phi\right)$ has $V_0(\phi)=0$ since in fact $V_t(\phi)=0$ for all $t\leq u$, while $V_T(\phi)(\omega)=0$ for $\omega\notin D$. For $\omega\in D$, we have

$$(\theta_T\cdot S_T)(\omega)=V_T(\theta)(\omega)=Z(\omega)$$

since θ replicates Z. Hence

$$\begin{aligned}V_T(\phi)(\omega)&=\left(\theta_T\cdot S_T+(Z_u-V_u(\theta))\frac{S_T^0}{S_u^0}-Z\right)(\omega)\\&=\left((Z_u-V_u(\theta))\frac{S_T^0}{S_u^0}\right)(\omega)\end{aligned}$$

$> 0.$

This shows that ϕ is an arbitrage opportunity in the extended model since $V_T(\phi) \geq 0$ and $P(V_T(\phi) > 0) = P(D) > 0$.

To construct an arbitrage when $Z_u < V_u(\theta)$ for some $u \leq T$ on a set E with $P(E) > 0$, we simply reverse the positions described above. On E at time u, shortsell the amount $V_u(\theta)$ according to the strategy θ, buy one unit of the claim Z for Z_u, place the difference in the bank, and do nothing else. Hence, if the claim Z does not have the value process $V(\theta)$ determined by the replicating strategy θ, the extended model is not arbitrage-free.

Conversely, suppose that every attainable European claim Z has its value function given via the EMM Q as $Z_t = S_t^0 E_Q\left(\frac{Z}{S_T^0}\,|\mathcal{F}_t\right)$ for each $t \leq T$, and let $\psi = (\phi, \gamma)$ be a self-financing strategy, involving S and m attainable European claims $(Z^j)_{j \leq m}$, with $V_0(\psi) = 0$ and $V_T(\psi) \geq 0$. We show that $P(V_T(\psi) = 0) = 1$, so that ψ cannot be an arbitrage opportunity in the extended model. Indeed, consider the discounted value process $\overline{V}(\psi) = \frac{V(\psi)}{S^0}$ at time $t > 0$:

$$\begin{aligned} E_Q\left(\overline{V}_t(\psi)\,|\mathcal{F}_{t-1}\right) &= E_Q\left(\sum_{i=0}^{d} \phi_t^i \overline{S}_t^i + \sum_{j=1}^{m} \gamma_t^j \frac{Z_t^j}{S_t^0}\,|\mathcal{F}_{t-1}\right) \\ &= \sum_{i=0}^{d} \phi_t^i E_Q\left(\overline{S}_t^i\,|\mathcal{F}_{t-1}\right) + \sum_{j=1}^{m} \gamma_t^j E_Q\left(\overline{V}_t^j(\theta^j)\,|\mathcal{F}_{t-1}\right). \end{aligned}$$

Here we use the fact that $\overline{S}^i = \frac{S^i}{S^0}$ is a martingale under Q and, defining $\overline{V}_t^j(\theta^j)$ as the discounted value process of the replicating strategy for the claim Z^j, we see that $\overline{V}_t^j(\theta^j) = E_Q\left(\frac{Z^j}{S_T^0}\,|\mathcal{F}_t\right) = \frac{Z_t^j}{S_t^0}$. Since each process $\overline{V}^j(\theta^j)$, $j \leq m$, is a Q-martingale, it follows that

$$E_Q\left(\overline{V}_t(\psi)\,|\mathcal{F}_{t-1}\right) = \sum_{i=0}^{d} \phi_t^i \overline{S}_{t-1}^i + \sum_{j=1}^{m} \gamma_t^j \overline{V}_{t-1}^j(\theta^j) = \overline{V}_{t-1}(\psi)$$

since the strategy $\psi = (\phi, \gamma)$ is self-financing, so that $\overline{V}(\psi)$ is also a Q-martingale. Consequently, $E_Q\left(\overline{V}_t(\psi)\right) = E_Q\left(V_0(\psi)\right) = 0$. Therefore $Q(\overline{V}_T(\psi) = 0) = 1$, and since $Q \sim P$ it follows that $P(V_T(\psi) = 0) = 1$. Therefore the extended securities market model is arbitrage-free. □

This result should not come as a surprise. It remains the case that the only independent sources of randomness in the model are the stock prices $S_1, S_2, \ldots, S_d$, since the contingent claims used to construct trading strategies are priced via an equivalent measure for which their discounted versions are martingales. However, it does show that the methodology is consistent. We return to extended market models when examining possible arbitrage-free prices for claims in incomplete models in Chapter 4.

Some Consequences of Call-Put parity

In the call-put parity relation (1.3), the discount rate is given by $\beta_{t,T} = \beta^{T-t}$, where $\beta = (1+r)$. Write (1.3) in the form

$$S_t = C_t - P_t + \beta^{T-t}K. \tag{2.20}$$

With the price of each contingent claim expressed at the expectation under the risk-neutral measure Q of its discounted final value, we show that the right-hand side of (2.20) is independent of K. Indeed,

$$\begin{aligned}
S_t &= \beta^{T-t}[E_Q\left((S_T-K)^+\right) - E_Q\left((K-S_T)^+\right) + K] \\
&= \beta^{T-t}\left(\int_{\{S_T\geq K\}}(S_T-K)dQ - \int_{\{S_T<K\}}(K-S_T)dQ + K\right) \\
&= \beta^{T-t}\left(\int_{\Omega}(S_T-K)dQ + K\right) \\
&= \beta^{T-t}E_Q\left(S_T\right) \\
&= \beta_t^{-1}E_Q\left(\beta^T S_T\right).
\end{aligned}$$

This shows that call-put parity is a consequence of the martingale property of the discounted price under Q in any market model that allows pricing of contingent claims by expectation under an equivalent martingale measure.

Remark 2.5.3. The identity also leads to the following interesting observation due to Marek Capinski, which first appeared in [35]. Recall the Modigliani-Miller theorem (see [20]), which states that the value of a firm is independent of the way in which it is financed. Since its value is represented by the sum of its equity (stock) and debt, the theorem states that the level of debt has no impact on the value of the firm. This can be interpreted in terms of options, as follows.

If the firm's borrowings at time 0 are represented by $\beta_T K$, so that it faces repayment of debt at K by time T, the stockholders have the option to buy back this debt at that time, in order to avert bankruptcy of the firm. They will only do so if the value S_T of the firm at time T is at least K. The firm's stock can therefore be represented as a European call option on S with payoff K at time T, and thus the current (time 0) value of the stock is the call option price C_0. The total current value of the firm is $S_0 = C_0 - P_0 + \beta^T K$, where P_0 is a put option on S with the same strike and horizon as the call. The calculation above shows that S_0 is independent of K, as the Modigliani-Miller theorem claims. Moreover, the current value of the debt is given via the call-put parity relation as $(\beta^T K - P_0)$. This is lower than the present value $\beta_T K$ of K, so that P_0 reflects the default risk (i.e., risk that the debt may not be recovered in full at time T).

2.6 Example: The Binomial Model

We now take another look at the Cox-Ross-Rubinstein binomial model, which provides a very simple, yet striking, example of the strength of the martingale methods developed so far.

The CRR Market Model

The Cox-Ross-Rubinstein binomial market model was described in Chapter 1. Recall that we assumed that $d = 1$. There is a single stock S^1 and a riskless bond S^0, which accrues interest at a fixed rate $r > 0$. Taking $S_0^0 = 1$, we have $S_t^0 = (1+r)^t$ for $t \in \mathbb{T}$, and hence $\beta_t = (1+r)^{-t}$. The ratios of successive stock values are Bernoulli random variables; that is, for all $t < T$, either $S_t^1 = S_{t-1}^1(1+a)$ or $S_t^1 = S_{t-1}^1(1+b)$, where $b > a > -1$ are fixed throughout, while S_0^1 is constant. We can thus conveniently choose the sample space

$$\Omega = \{1+a, 1+b\}^T$$

together with the natural filtration $\mathbb{F}$ generated by the stock price values; that is, $\mathcal{F}_0 = \{\emptyset, \Omega\}$, and $\mathcal{F}_t = \sigma(S_u^1 : u \leq t)$ for $t > 0$. Note that $\mathcal{F}_T = \mathcal{F} = 2^\Omega$ is the σ-field of all subsets of Ω.

The measure P on Ω is the measure induced by the ratios of the stock values. More explicitly, we write S for S^1 for the rest of this section to simplify the notation, and set $R_t = \frac{S_t}{S_{t-1}}$ for $t > 0$. For $\omega = (\omega_1, \omega_2, \ldots, \omega_T)$ in Ω, define

$$P(\{\omega\}) = P(R_t = \omega_t, t = 1, 2, \ldots, T). \tag{2.21}$$

For any probability measure Q on $(\Omega, \mathcal{F})$, the relation $E_Q\left(\overline{S}_t \,|\mathcal{F}_{t-1}\right) = \overline{S}_{t-1}$ is equivalent to

$$E_Q\left(R_t \,|\mathcal{F}_{t-1}\right) = 1 + r$$

since $\frac{\beta_t}{\beta_{t-1}} = 1 + r$. Hence, if Q is an equivalent martingale measure for S, it follows that $E_Q(R_t) = 1 + r$. On the other hand, R_t only takes the values $1+a$ and $1+b$; hence its average value can equal $1+r$ only if $a < r < b$. We have yet again verified the following result.

Lemma 2.6.1. *For the binomial model to have an EMM, we must have*

$$a < r < b.$$

When the binomial model is viable, there is a *unique* equivalent martingale measure Q for S. We construct this measure in the following lemma.

Lemma 2.6.2. *The discounted price process $\overline{S}$ is a Q-martingale if and only if the random variables (R_t) are independent, identically distributed, and $Q(R_1 = 1+b) = q$ and $Q(R_1 = 1+a) = 1-q$, where $q = \frac{r-a}{b-a}$.*

Proof. Under independence, the (R_t) satisfy

$$E_Q\left(R_t\,|\mathcal{F}_{t-1}\right) = E_Q\left(R_t\right) = q(1+b)+(1-q)(1+a) = q(b-a)+1+a = 1+r.$$

Hence, by our earlier discussion, $\overline{S}$ is a Q-martingale.

Conversely, if $E_Q\left(R_t\,|\mathcal{F}_{t-1}\right) = 1+r$, then, since R_t takes only the values $1+a$ and $1+b$, we have

$$(1+a)Q(R_t = 1+a\,|\mathcal{F}_{t-1}) + (1+b)Q(R_t = 1+b\,|\mathcal{F}_{t-1}) = 1+r,$$

while

$$Q(R_t = 1+a\,|\mathcal{F}_{t-1}) + Q(R_t = 1+b\,|\mathcal{F}_{t-1}) = 1.$$

Letting $q = Q(R_t = 1+b\,|\mathcal{F}_{t-1})$, we obtain

$$(1+a)(1-q) + (1+b)q = 1+r.$$

Hence $q = \frac{r-a}{b-a}$. The independence of the R_t follows by induction on $t > 0$. For $\omega = (\omega_1, \omega_2, \ldots, \omega_T) \in \Omega$, we see inductively that

$$Q\left(R_1 = \omega_1, R_2 = \omega_2, \ldots, R_t = \omega_t\right) = \prod_{i=1}^{t} q_i,$$

where $q_i = q$ when $\omega_i = 1+b$ and equals $1-q$ when $\omega_i = 1+a$. Thus the (R_t) are independent and identically distributed as claimed. □

Remark 2.6.3. Note that $q \in (0,1)$ if and only if $a < r < b$. Thus a viable binomial market model admits a *unique* EMM given by Q as in Lemma 2.6.2.

The CRR Pricing Formula

The CRR pricing formula, obtained in Chapter 1 by an explicit hedging argument, can now be deduced from our general martingale formulation by calculating the Q-expectation of a European call option on the stock. More generally, the value of the call $C_T = (S_T - K)^+$ at time $t \in \mathbb{T}$ is given by (2.14); that is,

$$V_t(C_T) = \frac{1}{\beta_t} E_Q\left(\beta_T C_T\,|\mathcal{F}_t\right).$$

Since $S_T = S_t \prod_{u=t+1}^{T} R_u$ (by the definition of (R_u)), we can calculate this expectation quite easily since S_t is $\mathcal{F}_t$-measurable and each R_u $(u > t)$ is independent of $\mathcal{F}_t$. Indeed,

$$V_t(C_T) = \beta_t^{-1}\beta_T E_Q\left(\left[S_t \prod_{u=t+1}^{T} R_u - K\right]^+ |\mathcal{F}_t\right)$$

$$
\begin{aligned}
&= (1+r)^{t-T} E_Q \left(\left[S_t \prod_{u=t+1}^{T} R_u - K \right]^+ |\mathcal{F}_t \right) \\
&= v(t, S_t). \qquad (2.22)
\end{aligned}
$$

Here

$$
\begin{aligned}
v(t,x) &= (1+r)^{t-T} E_Q \left(\left[x \prod_{u=t+1}^{T} R_u - K \right]^+ \right) \\
&= (1+r)^{t-T} \sum_{u=0}^{T-t} \binom{T-t}{u} q^u (1-q)^{T-t-u} \left[x(1+b)^u (1+a)^{T-t-u} - K \right]^+
\end{aligned}
$$

and, in particular, the price at time 0 of the European call option C with payoff $C_T = (S_T - K)^+$ is given by

$$
v(0, S_0) = (1+r)^{-T} \sum_{u=A}^{T} \binom{T}{u} q^u (1-q)^{T-u} \left[S_0 (1+b)^u (1+a)^{T-u} - K \right], \qquad (2.23)
$$

where A is the first integer k for which $S_0(1+b)^k(1+a)^{T-k} > K$. The CRR option pricing formula (1.5.3) now follows exactly as in Chapter 1.

Exercise 2.6.4. Show that for the replicating strategy $\theta = (\theta^0, \theta^1)$ describing the value process of the European call C, the stock portfolio θ^1 can be expressed in terms of the differences of the value function as $\theta^1_t = \theta(t, S_{t-1})$, where

$$
\theta(t,x) = \frac{v(t, x(1+b)) \ - v(t, x(1+a))}{x(b-a)}.
$$

Exercise 2.6.5. Derive the call-put parity relation (2.20) by describing the values of the contingent claims involved as expectations relative to Q.

2.7 From CRR to Black-Scholes

Construction of Approximating Binomial Models

The binomial model contains all the information necessary to deduce the famous Black-Scholes formula for the price of a European call option in a continuous-time market driven by Brownian motion. A detailed discussion of the mathematical tools used in that model is deferred until Chapter 6, but we now describe how the random walks performed by the steps in the binomial tree lead to Brownian motion as a limiting process when we reduce the step sizes continually while performing an ever larger number of steps within a fixed time interval $[0, T]$. From this we will see how the Black-Scholes price arises as a limit of CRR prices.

Consider a one-dimensional stock price process $S = (S_t)$ on the finite time interval $[0, T]$ on the real line, together with a European put option with payoff function $f_T = (K - S_T)^+$ on this stock. We use put options here because the payoff function f is bounded, thus allowing us to deduce that the relevant expectations (using EMMs) converge once we have shown via a central limit theorem that certain random variables converge weakly. The corresponding result for call options can then be derived using call-put parity.

We wish to construct a discrete-time binomial model beginning with the same constant stock price S_0 and with N steps in $[0, T]$. Thus we let $h_N = \frac{T}{N}$ and define the discrete timeline $\mathbb{T}_N = \{0, h_N, 2h_N, \ldots, Nh_N\}$. The European put P^N with strike K and horizon T is then defined on $\mathbb{T}_N$. By (2.14), (exactly as in the derivation of (2.22)), P^N has CRR price P_0^N given by

$$P_0^N = (1 + \rho_N)^{-N} E_{Q_N}\left(\left[K - S_0 \prod_{k=1}^{N} R_k^N\right]^+\right), \tag{2.24}$$

where, writing S_k^N for the stock price at time kh_N, the ratios $R_k^N = \frac{S_k^N}{S_{k-1}^N}$ take values $1 + b_N$ or $1 + a_N$ at each discrete time point kh_N $(k \leq N)$.

The values of a_N, b_N and the riskless interest rate ρ_N have yet to be chosen. Once they are fixed, with $a_N < \rho_N < b_N$, they will uniquely determine the risk-neutral probability measure Q_N for the Nth binomial model since by Lemma 2.6.2 the binomial random variables $(R_k^N)_{k \leq N}$ are then an independent and identically distributed sequence. We obtain, as before, that

$$Q_N(R_1^N = 1 + b_N) = q_N = \frac{\rho_N - a_N}{b_N - a_N}. \tag{2.25}$$

We treat the parameters from the Black-Scholes model as given and adjust their counterparts in our CRR models in order to obtain convergence. To this end, we fix $r \geq 0$ and set $\rho_N = rh_N$, so that the discrete-time riskless rate satisfies $\lim_{N\to\infty}(1 + \rho_N)^N = e^{rT}$, so that r acts as the 'instantaneous' rate of return.

Fix $\sigma > 0$, which will act as the volatility per unit time of the Black-Scholes stock price, and for each fixed N we now fix a_N, b_N by demanding that the discounted logarithmic returns are given by

$$\log\left(\frac{1 + b_N}{1 + \rho_N}\right) = \sigma\sqrt{h_N} = \sigma\sqrt{\frac{T}{N}}, \quad \log\left(\frac{1 + a_N}{1 + \rho_N}\right) = -\sigma\sqrt{h_N} = -\sigma\sqrt{\frac{T}{N}},$$

so that

$$u_N = 1 + b_N = \left(1 + \frac{rT}{N}\right) e^{\sigma\sqrt{\frac{T}{N}}}, \quad d_N = 1 + a_N = \left(1 + \frac{rT}{N}\right) e^{-\sigma\sqrt{\frac{T}{N}}}.$$

Note that the discount factor at each step is $1+\rho_N = 1+\frac{rT}{N}$ for each $k \le N$. The random variables

$$\left\{ Y_k^N = \log\left(\frac{R_k^N}{1+\rho_N}\right) : k \le N \right\}$$

are independent and identically distributed. We shall consider their sum

$$Z_N = \sum_{k=1}^{N} Y_k^N = \sum_{k=1}^{N} R_k^N - N\log(1+\rho_N)$$

for each N. The discounted stock price is thus

$$\overline{S}_N^N = (1+\rho_N)^{-N} \prod_{k=1}^{n} R_k^N = \exp\left\{ \sum_{k=1}^{N} Y_k^N \right\} = e^{Z_N},$$

so that the N^{th} put option price becomes

$$P_0^N = E_{Q_N}\left(\left[\left(1+\frac{rt}{N}\right)^{-N} K - S_0 e^{Z_N} \right]^+ \right). \tag{2.26}$$

Convergence in Distribution

The values taken by Y_1^N are $\pm\sigma\sqrt{h_N}$, so its second moment is $\sigma^2 h_N = \sigma^2 \frac{T}{N}$, while its mean is given by

$$\mu_N = (2q_N - 1)\sigma\sqrt{h_N} = (2q_N-1)\sigma\sqrt{\frac{T}{N}}.$$

Our choices will imply that q_N converges to $\frac{1}{2}$ as $N \to \infty$. We show this by checking the rate of convergence. First recall some notation: $a_N = a + o\left(\frac{1}{N}\right)$ means that $N(a_N - a) \to 0$ as $N \to \infty$.

Since $1 - q_N = \frac{u_N - \rho_N}{u_N - d_N}$, we see that $2q_N - 1$ is of order $\frac{1}{\sqrt{N}}$:

$$\begin{aligned} 2q_N - 1 = 1 - 2(1-q_N) &= 1 - 2\left(\frac{e^{\sigma\sqrt{h_N}} - 1}{e^{\sigma\sqrt{h_N}} - e^{-\sigma\sqrt{h_N}}}\right) \\ &= 1 - \frac{e^{\sigma\sqrt{h_N}} - 1}{\sinh(\sigma\sqrt{h_N})}. \end{aligned}$$

Expanding into Taylor series the right-hand side has the form

$$1 - \frac{x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots}{x + \frac{x^3}{3!} + \cdots} = \frac{-\frac{x^2}{2} - \frac{x^4}{4!} + \cdots}{x + \frac{x^3}{3!} + \cdots},$$

so that $2q_N - 1 = -\frac{1}{2}\sigma\sqrt{h_N} + o\left(\frac{1}{N}\right)$. Thus $\mu_N = -\frac{1}{2}\frac{\sigma^2 T}{N} + o\left(\frac{1}{N}\right)$, so that $N\mu_N \to -\frac{1}{2}\sigma^2 T$ as $N \to \infty$.

Since the second moment of Y_1^N is $\sigma^2 \frac{T}{N}$, its variance σ_N^2 therefore satisfies

$$\sigma_N^2 = \sigma^2 \frac{T}{N} + o\left(\frac{1}{N}\right). \tag{2.27}$$

We apply the central limit theorem for triangular arrays (see, e.g. , [168, VII.5.4] or [45, Corollary to Theorem 3.1.2]) in the following form to the independent and identically distributed random variables (Y_k^N) for $k \le N$ and $N \in \mathbb{N}$.

Theorem 2.7.1 (Central Limit Theorem). *For $N \ge 1$, let $(Y_k^N)_{k\le N}$ be an independent and identically distributed sequence of random variables, each with mean μ_N and variance σ_N^2. Suppose that there exist real μ and $\Sigma^2 > 0$ such that $N\mu_N \to \mu$ and $\sigma_N^2 = \Sigma^2 + o\left(\frac{1}{N}\right)$ as $N \to \infty$. Then the sums $Z_N = \sum_{k=1}^N Y_k^N$ converge in distribution to a random variable $Z \sim \mathcal{N}(\mu, \Sigma^2)$.*

We prove this by verifying the Lindeberg-Feller condition for the Y_k^N, namely that for all $\varepsilon > 0$

$$\sum_{k=1}^N E_{Q_N}\left((Y_k^N)^2 \mathbf{1}_{\{|Y_k^N|>\varepsilon\}}\right) \to 0 \text{ as } N \to \infty. \tag{2.28}$$

We have seen that for fixed N and all $k \le N$, $(Y_k^N)^2$ is constant on Ω and takes the value $\frac{\sigma^2 T}{N}$. Therefore, using Chebychev's inequality with $|Y_k^N| = \sigma\sqrt{\frac{T}{N}}$, we see that for each $k \le N$

$$E_{Q_N}\left((Y_k^N)^2 \mathbf{1}_{\{|Y_k^N|>\varepsilon\}}\right) = \frac{\sigma^2 T}{N} P(|Y_k^N| > \varepsilon) \le \frac{\sigma^2 T}{N} \frac{E\left(|Y_k^N|\right)}{\varepsilon},$$

and since the right-hand side equals $\frac{\sigma^3}{\varepsilon}\left(\frac{T}{N}\right)^{\frac{3}{2}}$, the Lindeberg condition is satisfied. The Lindeberg-Feller Theorem completes the proof.

For the sequence (Y_k^N) defined above, the conditions of the theorem are satisfied with $\mu = -\frac{1}{2}\sigma^2 T$ and $\Sigma = \frac{1}{2}\sigma^2 T$ with σ as fixed above. Thus (Z_N) converges in distribution to $Z \sim \mathcal{N}(-\frac{1}{2}\sigma^2 T, \sigma^2 T)$, while $(1+\rho_N)^{-N} \to e^{-rT}$ as $N \to \infty$. It follows that the limit of the CRR put option prices (P_0^N) is given by

$$E\left((e^{-rT}K - S_0 e^Z)^+\right), \tag{2.29}$$

where the expectation is now taken with respect to the distribution of Z.

The Black-Scholes Formula

Standardising Z, we see that the random variable $X = \frac{1}{\sigma\sqrt{T}}(Z + \frac{1}{2}\sigma^2 T)$ has distribution $N(0,1)$; that is, $Z = \sigma\sqrt{T}X - \frac{1}{2}\sigma^2 T$. The limiting value of P_0^N can be found by evaluating the integral

$$\int_{-\infty}^{\infty} \left[e^{-rT}K - S_0 e^{-\frac{1}{2}\sigma^2 T + \sigma\sqrt{T}x}\right]^+ \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} dx. \tag{2.30}$$

Observe that the integrand is non-zero only when

$$\sigma\sqrt{T}x + \left(r - \frac{1}{2}\sigma^2\right)T < \log\left(\frac{K}{S_0}\right),$$

that is, on the interval $(-\infty, \gamma)$, where

$$\gamma = \frac{\log\left(\frac{K}{S_0}\right) - (r - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}.$$

Thus the put option price for the limiting pricing model reduces to

$$\begin{aligned} P_0 &= Ke^{-rT}(\Phi(\gamma)) - S_0 \int_{-\infty}^{\gamma} e^{-\frac{\sigma^2 T}{2}} e^{\sigma\sqrt{T}x - \frac{x^2}{2}} \frac{dx}{\sqrt{2\pi}} \\ &= Ke^{-rT}(\Phi(\gamma)) - S_0 \int_{-\infty}^{\gamma} e^{-\frac{1}{2}(x-\sigma\sqrt{T})^2} \frac{dx}{\sqrt{2\pi}} \\ &= Ke^{-rT}(\Phi(\gamma)) - S_0 \left(\Phi(\gamma - \sigma\sqrt{T})\right). \end{aligned}$$

Here Φ denotes the cumulative normal distribution function.

Setting $d_- = -\gamma$ and $d_+ = d_- + \sigma\sqrt{T}$, and using the symmetry of Φ, we obtain $1 - \Phi(\gamma) = \Phi(-\gamma) = \Phi(d_-)$ and $1 - \Phi(\gamma - \sigma\sqrt{T}) = \Phi(d_+)$, where

$$d_\pm = \frac{\log\left(\frac{S_0}{K}\right) + (r \pm \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}. \tag{2.31}$$

By call-put parity, this gives the familiar *Black-Scholes formula* for the call option: the time 0 price of the call option $f_T = (S_T - K)^+$ is given by

$$V_0(C) = C_0 = S_0\Phi(d_+) - e^{-rT}K\Phi(d_-). \tag{2.32}$$

Remark 2.7.2. An alternative derivation of this approximating procedure, using binomial models where for each n the probabilities of the 'up' and 'down' steps are equal to $\frac{1}{2}$, can be found in [35].

By replacing T by $T - t$ and S_0 by S_t, we can read off the value process V_t for the option similarly; in effect this treats the option as a contract written at time t with time to expiry $T - t$,

$$V_t(C) = S_t\Phi(d_{t+}) - e^{-r(T-t)}K\Phi(d_{t-}), \tag{2.33}$$

where

$$d_{t\pm} = \frac{\log\left(\frac{S_t}{K}\right) + (r \pm \frac{1}{2}\sigma^2)(T - t)}{\sigma\sqrt{T - t}}.$$

The preceding derivation has not required us to study the dynamics of the 'limit stock price' S; it is shown in Chapter 7 that this takes the form

$$dS_t = S_t\mu dt + \sigma S_t dW_t, \tag{2.34}$$

where W is a Brownian motion. The stochastic calculus necessary for the solution of such stochastic differential equations is developed in Chapter 6. However, we can already note one remarkable property of the Black-Scholes formula: it does not involve the mean return μ of the stock but depends on the riskless interest rate r and the volatility σ. The mathematical reason for this lies in the change to a risk-neutral measure (which underlies the martingale pricing techniques described in this chapter), which eliminates the drift term from the dynamics.

Dependence of the Option Price on the Parameters

Write $C_t = V_t(C)$ for the Black-Scholes value process of the call option; i.e.,

$$C_t = S_t\Phi(d_{t+}) - e^{-r(T-t)}K\Phi(d_{t-}),$$

where $d_{t\pm}$ is given as in (2.33). As we have calculated for the case $t = 0$, the European put option with the same parameters in the Black-Scholes pricing model is given by

$$P_t = Ke^{-r(T-t)}\Phi(-d_{t-}) - S_t\Phi(-d_{t+}).$$

We examine the behaviour of the prices C_t at extreme values of the parameters. (The reader may consider the put prices P_t similarly.)

When S_t increases, $d_{t\pm}$ grows indefinitely, so that $\Phi(d_{t\pm})$ tends to 1, and so C_t has limiting value $S_t - Ke^{-r(T-t)}$. In effect, the option becomes a forward contract with delivery price K since it is 'certain' to be exercised at time T. Similar behaviour is observed when the volatility σ shrinks to 0 since again $d_{t\pm}$ become infinite, and the riskless stock behaves like a bond (or money in the bank).

When $t \to T$ (i.e., the time to expiry decreases to 0) and $S_t > K$, then $d_{t\pm}$ becomes ∞ and $e^{-r(T-t)} \to 1$, so that C_t tends to $S_t - K$. On the other hand, if $S_t < K$, $\log\left(\frac{S_t}{K}\right) < 0$ so that $d_{t\pm} = -\infty$ and $C_t \to 0$. Thus, as expected, $C_t \to (S_T - K)^+$ when $t \to T$.

Remark 2.7.3. Note finally that there is a natural 'replicating strategy' given by (2.33) since this value process is expressed as a linear combination of units of stocks S_t and bonds S_t^0 with $S_0^0 = 1$ and $S_t^0 = \beta_t^{-1}S_0^0 = e^{rt}$. Writing the value process $V_t = \theta_t \cdot S_t$ (where by abuse of notation $S = (S^0, S)$), we obtain

$$\theta_t^0 = -Ke^{-rT}\Phi(d_{t-}), \qquad \theta_t^1 = \Phi(d_{t+}). \tag{2.35}$$

In Chapter 7, we consider various derivatives of the Black-Scholes option price, known collectively as 'the Greeks', with respect to its different parameters. This provides a sensitivity analysis with parameters that are widely used in practice.

http://www.springer.com/978-0-387-21292-0

Mathematics of Financial Markets
Elliott, R.J.; Kopp, P.E.
2005, XII, 354 p., Hardcover
ISBN: 978-0-387-21292-0