FINANCIAL STOCHASTICS

Notes for the Course by

A.W. van der Vaart

UPDATE, November 2007

Reading and Believing at Own Risk

CONTENTS

LITERATURE

These notes are based on interpretations and unifications of the following works, where sometimes the interpretations are loose. The core terminology for these notes is taken from Hunt and Kennedy.

DISCLAIMER: These notes are meant to be of help, but contrary to appearance, are not a solid text, but always in progress. Unfortunately there is no reliable textbook covering the complete material.

[1] Baxter, M. and Rennie, A., (1996). *Financial calculus*. Cambridge University Press, Cambridge.

[2] Chung, K.L. and Williams, R.J., (1990). *Introduction to stochastic integration, second edition*. Birkhäuser, London.

[3] Hunt, P.J. and Kennedy, J.E., (1998). *Financial Engineering*. Wiley.

[4] Jacod, J. and Shiryaev, A.N., (1987). *Limit theorems for stochastic processes*. Springer-Verlag, Berlin.

[5] Kramkov, D.O., (1996). Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets. *Probability Theory and Related Fields* **105**, 459–479.

[6] Kopp, P.E. and Elliott, R.J., (1999). *Mathematics and financial markets*. Springer-Verlag, New York.

[7] Karatzas, I. and Shreve, S.E., (1988). *Brownian motion and stochastic calculus*. Springer-Verlag, Berlin.

[8] Karatzas, I. and Shreve, S.E., (1998). *Methods of mathematical finance*. Springer-Verlag, Berlin.

[9] Musiela, M. and Rutkowski, M., (1997). *Martingale Methods in Financial Modelling*. Springer-Verlag, Berlin.

[10] Revuz, D. and Yor, M., (1994). *Continuous martingales and Brownian motion*. Springer, New York.

[11] Rogers, L.C.G. and Williams, D., (2000). *Diffusions, Markov Processes and Martingales, volumes 1 and 2*. Cambridge University Press, Cambridge.

[12] Schoutens, W., (2003). *Lévy Processes in Finance*. Wiley.

1
Black-Scholes

A *financial derivative* is a contract that is based on the price of an underlying asset, such as a stock price or bond price. An "option", of which there are many different types, is an important example. A main focus of "financial engineering" is on finding the "fair price" of such a derivative. Following the work by Black and Scholes in the 1970s the prices of derivatives are found through the principle of "no arbitrage". The concepts of "completeness" and "arbitrage-freeness" of an economy are central to make this work.

In this chapter we discuss, as an introduction, the pricing of the European call option in the original model used by Black and Scholes. Even though this model is unrealistically simple, the ensuing "Black-Scholes formula" is still the most frequently used tool in practice.

We denote by S_t the price of a stock at time $t \geq 0$, and assume that S satisfies the stochastic differential equation

$$dS_t = \mu S_t \, dt + \sigma S_t \, dW_t. \tag{1.1}$$

Here W is a Brownian motion process on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, and μ and $\sigma > 0$ are given numbers. The filtration $\{\mathcal{F}_t\}$ is the completed natural filtration generated by W, and it is assumed that S is continuous and adapted to this filtration.

This simple stochastic differential equation satisfies the conditions of the Itô theorem and hence this theorem guarantees the existence of a unique solution S. Using Itô's formula it is straightforward to verify directly that the solution takes the form

$$S_t = S_0 e^{(\mu - \frac{1}{2}\sigma^2)\, t + \sigma \, W_t}.$$

In particular, the stock price is strictly positive if $S_0 > 0$, as we shall assume. The number σ is called the *volatility* of the stock. It determines the variability of the stock over time. The number μ gives the drift of the stock. It is responsible for the exponential growth of a typical stock price.

We assume that, besides in the stock S, we can invest in a "riskless asset" with a predetermined yield, much as putting money in a savings account with a fixed interest rate. We assume that the price R_t of this investment at time t satisfies the differential equation

$$dR_t = rR_t\,dt, \qquad\qquad R_0 = 1.$$

Here the number r is called the *interest rate*. Of course, the differential equation can be solved explicitly to give

$$R_t = e^{rt}.$$

This is the "continuously compounded interest" over the interval $[0,t]$. An amount of money V invested in the asset R at time 0 grows with certainty to an amount Ve^{rt} at time t. Thus the interest rate gives the "time value of money": one unit of money that we are promised to receive at time t is worth e^{-rt} units of money at time 0. Dividing an amount V by the number e^{rt} to get the amount Ve^{-rt} is called *discounting*.

The interest rate r should be interpreted as being corrected for inflation or other effects, so that it gives the "true" economic rate of growth of a risk-free asset.

A *portfolio* or *strategy* (ϕ,ψ) is defined to be a pair of predictable processes ϕ and ψ. The pair (ϕ_t,ψ_t) gives the numbers of bonds and stocks owned at time t, giving the *portfolio value*

$$V_t = \phi_t R_t + \psi_t S_t. \tag{1.2}$$

The predictable processes ϕ and ψ can depend on the past until "just before t" and we may think of changes in the content of the portfolio as a reallocation of bonds and stocks that take place just before the new stock price is known. A portfolio is "self-financing" if such reshuffling can be carried out without import or export of money, whence changes in the value of the portfolio are due only to changes in the values of the underlying assets. More precisely, we call the strategy (ϕ,ψ) *self-financing* if

$$dV_t = \phi_t\,dR_t + \psi_t\,dS_t. \tag{1.3}$$

This is to be interpreted in the sense that V must be a semimartingale satisfying $V = V_0 + \phi\cdot R + \psi\cdot S$. It is implicitly required that ϕ and ψ are suitable integrands relative to R and S. In finance the reshuffling of a portfolio to attain a certain aim is called "hedging". Originally hedging was understood to be a strategy to control or limit risk, but following several crashes "hedging" now also has a connotation of being "risky" and "greedy" in everyday language.

A *contingent claim* with expiry time $T>0$ is defined to be an $\mathcal{F}_T$-measurable random variable. It is interpreted as the "pay-off" at the expiry time of a "derivative", a contract based on the stock. In this section we consider contracts that can only be exercised, or "paid out", at expiry time.

The European call option is an important example. It is a contract that gives the right, but not the obligation, to buy a stock at a predetermined price K at a predetermined time T in the future. A main question we like to answer is: what is the fair price of a European option at time 0? Because an option gives a right and not an obligation, it must cost some money to acquire one.

It is easy to determine the value of the European option at time T. If the stock price S_T is higher than the "strike price" K (the "option is in the money"), then we will exercise our right, and buy the stock for the price K. If we wished we could sell the stock immediately after acquiring it, and thus make a net gain of $S_T - K$. On the other hand, if the stock price S_T would be below the strike price K ("the option is out of the money"), then the option is worthless, and we would not use our right. We can summarize the two cases by saying that the value of the option at the time T is equal to $(S_T - K)^+$, where $x^+ = x \vee 0$ is the positive part of a number.

At time 0 this value is still unknown and hence cannot be used as the price of the option. A first, naive idea would be to set the price of an option equal to the expected value $\mathrm{E}(S_T - K)^+$. This does not take the time value of money into account, and hence a better idea would be to set the option price equal to the expected discounted value $\mathrm{E}e^{-rT}(S_T - K)^+$. This is still too naive, as Black and Scholes first argued in the beginning 1970s. It turns out that the correct price is an expectation, but the expectation must be computed under a different probability measure.

Before giving Black and Scholes' argument (in a martingale form introduced by Harrison and Pliska in the 1980s), we note that the same type of reasoning applies to many other contracts as well. The main points are that the contract can pay out at expiry time only and that the value of the contract at expiry time T can be expressed as an $\mathcal{F}_T$-measurable random variable, called a *contingent claim*. Some examples of contingent claims are:

(i) *European call option*: $(S_T - K)^+$.
(ii) *European put option*: $(K - S_T)^+$.
(iii) *Asian call option*: $\bigl(\int_0^T S_t\,dt - K\bigr)^+$.
(iv) *lookback call option*: $S_T - \min_{0\le t\le T} S_t$,
(v) *down and out barrier option*: $(S_T - K)^+ 1\{\min_{0\le t\le T} S_t \ge L\}$.

The constants K and L and the expiry time T are fixed in the contract. There are many more possibilities; the more complicated contracts are referred to as *exotic options*. Note that in (iii)–(v) the claim depends on the history of the stock price throughout the period $[0, T]$. All contingent claims can be priced following the same no-arbitrage approach that we outline below.

Because the claims we wish to evaluate always have a finite term T, all the processes in our model matter only on the interval $[0,T]$. We may or must understand the assumptions and assertions accordingly.

The process $\tilde{S}$ defined by $\tilde{S}_t = R_t^{-1}S_t$ gives the discounted price of the stock. It turns out that claims can be priced by reference to a "martingale measure", defined as the (unique) measure that turns the process $\tilde{S}$ into a martingale. By Itô's formula and (1.1),

$$
\begin{aligned}
d\tilde{S}_t &= e^{-rt}\,dS_t + S_t(-re^{-rt})\,dt \\
&= \frac{\mu - r}{\sigma}\sigma e^{-rt}S_t\,dt + \sigma e^{-rt}S_t\,dW_t.
\end{aligned}
\tag{1.4}
$$

Under the measure $\mathbb{P}$ governing the Black-Scholes stochastic differential equation (1.1) the process W is a Brownian motion by assumption, and hence $\tilde{S}$ is a local martingale if and only if its drift component vanishes, i.e. if $\mu \equiv r$. This will rarely be the case in the real world.

On the other hand, under a different probability measure the discounted stock process $\tilde{S}$ may well be a martingale, i.e. it may be a martingale if viewed as a process on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \tilde{\mathbb{P}})$, equipped with a different probability measure $\tilde{\mathbb{P}}$. Because the drift term in (1.4) causes the problem, we could first rewrite the equation as

$$
d\tilde{S}_t = \sigma e^{-rt}S_t\,d\tilde{W}_t, \tag{1.5}
$$

for the process $\tilde{W}$ defined by

$$
\tilde{W}_t = W_t - \theta t, \qquad \theta = \frac{r-\mu}{\sigma}.
$$

(The number θ is called the *market price of risk*. If it is zero, then the real world is already "risk-neutral"; if not, then the number θ measures the deviation from a risk-neutral market relative to the volatility process.) If we could find a probability measure $\tilde{\mathbb{P}}$ such that the process $\tilde{W}$ is a $\tilde{\mathbb{P}}$-martingale, then the process $\tilde{S}$ would be a $\tilde{\mathbb{P}}$-local martingale. *Girsanov's theorem*, which we shall obtain in Chapter 2, permits us to do this: it implies that the process $\tilde{W}$ is a Brownian motion under the measure $\tilde{\mathbb{P}}$ with density $d\tilde{\mathbb{P}}/d\mathbb{P} = \exp(\theta \cdot W_T - \frac{1}{2}\theta^2 T)$ relative to $\mathbb{P}$. Presently, we do need to know the nature of $\tilde{\mathbb{P}}$; all that is important is the existence of a probability measure $\tilde{\mathbb{P}}$ such that $\tilde{W}$ is a Brownian motion.

We claim that the "reasonable price" at time 0 for a contingent claim with pay-off X at time T is the expectation under the measure $\tilde{\mathbb{P}}$ of the discounted value of the claim at time T, i.e.

$$
V_0 = \tilde{\mathrm{E}} R_T^{-1} X,
$$

where $\tilde{\mathrm{E}}$ denotes the expectation under $\tilde{\mathbb{P}}$. This is a consequence of economic, no-arbitrage reasoning, based on the following theorem.

1.6 Theorem. *Let X be a nonnegative contingent claim with $\tilde{\mathrm{E}}R_T^{-1}|X| < \infty$. Then there exists a self-financing strategy with value process V as in (1.2) such that:*
(i) $V \geq 0$ up to indistinguishability.
(ii) $V_T = X$ almost surely.
(iii) $V_0 = \tilde{\mathrm{E}}R_T^{-1}X$.

Proof. The process $\tilde{S} = R^{-1}S$ is a continuous semimartingale under $\mathbb{P}$ and a continuous local martingale under $\tilde{\mathbb{P}}$, in view of (1.5). Let $\tilde{V}$ be a cadlag version of the $\tilde{\mathbb{P}}$-martingale

$$\tilde{V}_t = \tilde{\mathrm{E}}\big(R_T^{-1}X \,|\, \mathcal{F}_t\big).$$

Suppose that there exists a predictable process ψ such that

$$d\tilde{V}_t = \psi_t \, d\tilde{S}_t.$$

Then $\tilde{V}$ is continuous, because $\tilde{S}$ is continuous, and hence predictable. Define

$$\phi = \tilde{V} - \psi\tilde{S}.$$

Then ϕ is predictable, because $\tilde{V}$, ψ and $\tilde{S}$ are predictable. The value of the portfolio (ϕ, ψ) is given by $V = \phi R + \psi S = (\tilde{V} - \psi\tilde{S})R + \psi S = \tilde{V}R$ and hence, by Itô's formula and (1.4),

$$\begin{aligned} dV_t &= \tilde{V}_t \, dR_t + R_t \, d\tilde{V}_t = (\phi_t + \psi_t\tilde{S}_t)\, dR_t + R_t\psi_t \, d\tilde{S}_t \\ &= (\phi_t + \psi_t R_t^{-1}S_t)\, dR_t + R_t\psi_t\big(-S_tR_t^{-2}\, dR_t + R_t^{-1}\, dS_t\big) \\ &= \phi_t \, dR_t + \psi_t \, dS_t. \end{aligned}$$

Thus the portfolio (ϕ, ψ) is self-financing. Statements (i)–(iii) of the theorem are clear from the definition of $\tilde{V}$ and the relation $V = R\tilde{V}$.

We must still prove the existence of the process ψ. In view of (1.5) we need to determine this process ψ such that

$$d\tilde{V}_t = \psi_t \sigma S_t e^{-rt} \, d\tilde{W}_t.$$

Because the process $\sigma S_t e^{-rt}$ is strictly positive, it suffices to determine a predictable process H such that $d\tilde{V}_t = H_t \, d\tilde{W}_t$.

The process $\tilde{W}$ is a $\tilde{\mathbb{P}}$-Brownian motion and $\tilde{V}$ is a $\tilde{\mathbb{P}}$-martingale. Furthermore, the filtrations generated by $\tilde{W}$ and W are identical. Therefore, the existence of an appropriate process H follows from the representing theorem for Brownian local martingales. According to this theorem, proved in Chapter 3, any local martingale relative to the filtration $\mathcal{F}_t$, which by assumption is the filtration generated by $\tilde{W}$, can be represented as a stochastic integral relative to $\tilde{W}$. ∎

We interpret the preceding theorem economically as saying that $V_0 = \tilde{\mathrm{E}} R_T^{-1} X$ is the just price at time 0 for the contingent claim X. We argue this by exhibiting "winning" strategies if the price is higher or lower than V_0.

Suppose that the price P_0 of the claim at time 0 would be higher than V_0. Then rather than buying the option we could buy the portfolio (ϕ_0, ψ_0) as in the theorem. This would save us the amount $P_0 - V_0$. We could next reshuffle our portfolio during the time interval $[0, T]$ according to the strategy (ϕ, ψ) and hence, by the theorem, end up with a portfolio (ϕ_T, ψ_T) with value V_T exactly equal to $V_T = (S_T - K)^+$, the value of the option at expiry time. Thus we have made a certain profit and we would never buy the option.

On the other hand, if the price P_0 of the option were lower than V_0, then anybody in the possession of a portfolio (ϕ_0, ψ_0), worth V_0, might sell this at time 0, buy an option for P_0 and put money $V_0 - P_0$ aside. During the term of the option the portfolio $(-\phi_0, -\psi_0)$, where the minus indicates "sold", could be reshuffled according to the inverse strategy $(-\phi_t, -\psi_t)$, yielding a capital at time T of $-V_T = -(S_T - K)^+$. Besides we also have an option worth $(S_T - K)^+$, and the money $V_0 - P_0$ set aside at time 0. Again this leads to a certain profit and hence nobody would keep stocks and savings.

For a general claim it may or may not be easy to evaluate the expectation $\tilde{\mathrm{E}} R_T^{-1} X$ explicitly, as the claim X may depend on the full history of the stock process. For pricing claims that depend on the final value S_T of the stock only, it is straightforward calculus to obtain a concrete formula.

For example, consider the price of a European call option. First we write the stock price in terms of the $\tilde{\mathbb{P}}$-Brownian motion $\tilde{W}$ as

$$S_t = S_0 e^{(r - \frac{1}{2}\sigma^2)\, t + \sigma\, \tilde{W}_t}.$$

In particular, under $\tilde{\mathbb{P}}$, the variable $\log(S_t/S_0)$ is normally distributed with mean $(r - \frac{1}{2}\sigma^2)t$ and variance $\sigma^2 t$. The price of a European call option can be written as, with Z a standard normal variable,

$$e^{-rT} \tilde{\mathrm{E}} (S_T - K)^+ = e^{-rT} \mathrm{E} \Big(S_0 e^{(r - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T} Z} - K \Big)^+.$$

It is straightforward calculus to evaluate this explicitly, giving the classical *Black-Scholes formula*

$$S_0 \Phi\Big(\frac{\log(S_0/K) + (r + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}} \Big) - K e^{-rT} \Phi\Big(\frac{\log(S_0/K) + (r - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}} \Big).$$

Remarkably, the drift coefficient μ does not make part of this equation: it plays no role in the pricing formula. Apparently, the systematic part of the stock price diffusion can be completely hedged away.

That $\tilde{\mathrm{E}} R_T^{-1} X$ is the fair price of the claim X is not a mathematical theorem, but the outcome of economic reasoning. One may or may not find this reasoning completely convincing. A hole in the pricing argument concerns the uniqueness of the strategy (ϕ, ψ) in the theorem. If there were another self-financing strategy (ϕ', ψ') with the same value V_T' at expiry time, but a different value V_0' at time 0, then clearly the argument would not be tenable, as it would lead to two "fair" prices.

If there existed two strategies with the same outcome at time T, but different values at time 0, then the difference strategy defined as $(\phi - \phi', \psi - \psi')$ would start with a nonzero value $V_0 - V_0'$, but end with certainty with a zero value $V_T - V_T'$. Thus this strategy, or its negative, would make money with certainty. This is referred to as an *arbitrage*. If an economy allows arbitrage, then the preceding economic argument does not make sense. We shall see later that the Black-Scholes economy is "arbitrage-free", provided that this concept is defined appropriately.

The following example shows that this issue must be treated with care. It exhibits, within the Black-Scholes model, a self-financing strategy that can be initiated with value 0 at time 0, and leads with certainty to an arbitrarily large positive value at time T. The construction is comparable to the well-known "doubling scheme" in a fair betting game that pays out twice the stake, or nothing, both with probability $1/2$. The doubling scheme consists of playing until we win the bet for the first time, doubling the stake at each time we loose. If we win after $n + 1$ bets, then our total gain is $-1 - 2 - 4 - \cdots - 2^{n-1} + 2^n = 1 > 0$. Because it is certain that we win eventually, we gain with certainty.

The reason that the doubling scheme does not work in practice, and it is the same with the investment strategy in the following example, is that we may need to play arbitrarily many times and our expected loss before we finally win is infinite. Thus we need an infinite capital to be certain that we can bring our strategy to an end as planned. From a practical point of view such a strategy is not real, and it is reasonable to exclude it from consideration when considering a "best" betting strategy.

Finance theory similarly excludes certain strategies, such as the one in the following example, from consideration. Among the set of remaining "admissible strategies", suitably defined, no strategy allows for arbitrage.

1.7 Example (Arbitrage). Consider the special case of the Black-Scholes model, where $\mu = r = 0$, $\sigma = 1$, and $S_0 = 1$. We claim that for every constant $\alpha > 0$ there exists a self-financing strategy (ϕ, ψ) with value process V satisfying $V_0 = 0$ and $V_T \geq \alpha$. In other words, the economy permits arbitrage of arbitrarily large size if we allow all self-financing trading strategies.

We construct the strategy by stopping another strategy (ϕ', ψ') as soon as the latter strategy's value process reaches the level α. Define

$$(\phi'_t, \psi'_t) = \Bigl(\frac{-1}{\sqrt{T-t}} + \int_0^t \frac{1}{S_s\sqrt{T-s}}\, dS_s, \frac{1}{S_t\sqrt{T-t}}\Bigr).$$

In view of the fact that $R_t = 1$ by definition, for every t, the value process of this strategy is given by

$$V'_t = \phi'_t R_t + \psi'_t S_t = \int_0^t \frac{1}{S_s\sqrt{T-s}}\, dS_s.$$

The equation shows that $V'_0 = 0$. Furthermore, the strategy (ϕ', ψ') is self-financing, as $\phi'_t\, dR_t + \psi'_t\, dS_t = \psi'_t\, dS_t = dV'_t$.

For U equal to the stopping time $U = \inf\{t > 0\colon V'_t = \alpha\}$, define $(\phi, \psi) = (\phi', \psi')1_{[0,U]} + (\alpha, 0)1_{(U,T]}$. This corresponds to waiting until the value of the portfolio under strategy (ϕ', ψ') reaches the level α and next investing all the money (i.e. the value α) in the risk-free asset R. It is intuitively clear that the new strategy is self-financing, as it is self-financing before U, at U, and clearly also after U. This can also be proved rigorously. (Cf. Exercise 4.3.)

The value process of the strategy (ϕ, ψ) is equal to α on the time interval $[U, T]$ (as the interest rate is zero), and hence has value α at time T whenever $U < T$. This is the case with probability one, as we shall now show.

The process $Y_t = V'_{T(1-e^{-t})}$ can be shown to be a Brownian motion, and $V'_t = Y_{-\log(1-t/T)}$. If t increases from 0 to T, then $-\log(1-t/T)$ increases from 0 to ∞. The properties of Brownian motion yield that $\limsup_{t\to\infty} Y_t = \infty$ almost surely. Therefore, the value process V'_t reaches any level α on the interval $[0, T)$ with certainty. $\square$

1.8 EXERCISE. Verify that the process Y in the preceding example is a Brownian motion process. [Hint: compute its quadratic variation process and use Lévy's theorem, Theorem 2.6.]

2
Change of Measure

Girsanov's theorem concerns the martingale property under a change of the probability measure on the underlying filtered space. Because densities often come as "exponential processes", we first recall the definition of such processes.

2.1 Exponential Processes

The *exponential process* corresponding to a continuous semimartingale X is the process $\mathcal{E}(X)$ defined by

$$\mathcal{E}(X)_t = e^{X_t - \frac{1}{2}[X]_t}.$$

The name "exponential process" would perhaps suggest the process e^X rather than the process $\mathcal{E}(X)$ as defined here. The additional term $\frac{1}{2}[X]$ in the exponent of $\mathcal{E}(X)$ is motivated by the extra term in the Itô formula. An application of this formula to the right side of the preceding display yields

$$d\mathcal{E}(X)_t = \mathcal{E}(X)_t \, dX_t. \tag{2.1}$$

(Cf. the proof of the following theorem.) If we consider the differential equation $df(x) = f(x)\,dx$ as the true definition of the exponential function $f(x) = e^x$, then $\mathcal{E}(X)$ is the "true" exponential process of X, not e^X.

Besides that, the exponentiation as defined here has the nice property of turning local martingales into local martingales.

2.2 Theorem. *The exponential process $\mathcal{E}(X)$ of a continuous local martingale X with $X_0 = 0$ is a local martingale. Furthermore,*

(i) If $\mathrm{E}e^{\frac{1}{2}[X]_t} < \infty$ for every $t \geq 0$, then $\mathcal{E}(X)$ is a martingale.

(ii) If X is an L_2-martingale and $\mathrm{E}\int_0^t \mathcal{E}(X)_s^2\, d[X]_s < \infty$ for every $t \geq 0$, then $\mathcal{E}(X)$ is an L_2-martingale.

Proof. By Itô's formula applied to the function $f(X_t, [X]_t) = \mathcal{E}(X)_t$, we find that

$$d\mathcal{E}(X)_t = \mathcal{E}(X)_t\, dX_t + \tfrac{1}{2}\mathcal{E}(X)_t\, d[X]_t + \mathcal{E}(X)_t\, (-\tfrac{1}{2})\, d[X]_t.$$

This simplifies to (2.1) and hence $\mathcal{E}(X) = 1+\mathcal{E}(X)\cdot X$ is a stochastic integral relative to X. If X is a local martingale, then so is $\mathcal{E}(X)$. Furthermore, if X is an L_2-martingale and $\int 1_{[0,t]}\mathcal{E}(X)^2\, d\mu_X < \infty$ for every $t \geq 0$, then $\mathcal{E}(X)$ is an L_2-martingale. This condition reduces to the condition in (ii).

The proof of (i) should be skipped at first reading. If $0 \leq T_n \uparrow \infty$ is a localizing sequence for $\mathcal{E}(X)$, then Fatou's lemma gives

$$\mathrm{E}\bigl(\mathcal{E}(X)_t \,|\, \mathcal{F}_s\bigr) \leq \liminf_{n\to\infty} \mathrm{E}\bigl(\mathcal{E}(X)_{t\wedge T_n} \,|\, \mathcal{F}_s\bigr) = \liminf_{n\to\infty} \mathcal{E}(X)_{s\wedge T_n} = \mathcal{E}(X)_s.$$

Therefore, the process $\mathcal{E}(X)$ is a supermartingale. It is a martingale if and only if its mean is constant, where the constant must be $\mathrm{E}\mathcal{E}(X)_0 = 1$.

By the representation theorem for Brownian martingales we may assume that the local martingale X takes the form $X_t = B_{[X]_t}$ for a process B that is a Brownian motion relative to a certain filtration. For every fixed t the random variable $[X]_t$ is a stopping time relative to this filtration. We conclude that it suffices to prove: if B is a Brownian motion and T a stopping time with $\mathrm{E}\exp(\frac{1}{2}T) < \infty$, then $\mathrm{E}\exp(B_T - \frac{1}{2}T) = 1$.

Because $2B_s$ is normally distributed with mean zero and variance $4s$,

$$\mathrm{E}\int_0^t \mathcal{E}(B)_s^2\, ds = \int_0^t \mathrm{E}e^{2B_s}e^{-s}\, ds = \int_0^t e^s\, ds < \infty$$

By (ii) it follows that $\mathcal{E}(B)$ is an L_2-martingale. For given $a < 0$ define $S_a = \inf\{t \geq 0\colon B_t - t = a\}$. Then S_a is a stopping time, so that $\mathcal{E}(B)^{S_a}$ is a martingale, whence $\mathrm{E}\mathcal{E}(B)_{S_a\wedge t} = 1$ for every t. It can be shown that S_a is finite almost surely and

$$\mathrm{E}\mathcal{E}(B)^{S_a} = \mathrm{E}e^{B_{S_a} - \frac{1}{2}S_a} = 1.$$

(The distribution of S_a is known in closed form. See e.g. Rogers and Williams I.9, p18-19; because $B_{S_a} = S_a + a$, the right side is the expectation of $\exp(a + \frac{1}{2}S_a)$.) With the help of an integration lemma we conclude that $\mathcal{E}(B)_{S_a\wedge t} \to \mathcal{E}(B)_{S_a}$ in L_1 as $t \to \infty$, and hence $\mathcal{E}(B)^{S_a}$ is uniformly integrable. By the optional stopping theorem, for any stopping time T,

$$1 = \mathrm{E}\mathcal{E}(B)_T^{S_a} = \mathrm{E}1_{T<S_a}e^{B_T - \frac{1}{2}T} + \mathrm{E}1_{T\geq S_a}e^{B_{S_a} - \frac{1}{2}S_a}.$$

Because the sample paths of the process $t \mapsto B_t - t$ are bounded on compact time intervals, $S_a \uparrow \infty$ if $a \downarrow -\infty$. Therefore, the first term on the right converges to $\mathrm{E}\exp(B_T - \frac{1}{2}T)$, by the monotone convergence theorem. The second term is equal to

$$\mathrm{E}1_{T\geq S_a} e^{S_a + a - \frac{1}{2}S_a} \leq e^a \mathrm{E} e^{\frac{1}{2}T}.$$

If $\mathrm{E}\exp(\frac{1}{2}T) < \infty$, then this converges to zero as $a \to -\infty$. ∎

In applications it is important to determine whether the process $\mathcal{E}(X)$ is a martingale, rather than just a local martingale. No simple necessary and sufficient condition appears to be possible, although the condition in (i), which is known as *Novikov's condition*, is optimal in the sense that the factor $\frac{1}{2}$ in the exponent cannot be replaced by a smaller constant, in general.

2.3 EXERCISE. Let X be a continuous semimartingale with $X_0 = 0$. Show that $Y = \mathcal{E}(X)$ is the *unique* solution to the pair of equations $dY = Y\,dX$ and $Y_0 = 1$. [Hint: use Itô's formula to show that $d\big(Y\mathcal{E}(X)^{-1}\big) = 0$ for every solution Y, so that $Y\mathcal{E}(X)^{-1} \equiv Y_0\mathcal{E}(X)_0^{-1} = 1$.]

2.4 EXERCISE. Show that $\mathcal{E}(X)^T = \mathcal{E}(X^T)$ for every stopping time T.

2.5 EXERCISE. Show that $\mathcal{E}(X)\mathcal{E}(Y) = \mathcal{E}(X+Y)$ if $[X,Y] = 0$.

The differential equation $dY_t = Y_{t-}\,dX_t$ makes perfect sense also if X is a general semimartingale, and allows to extend the definition of the exponential process in a sensible way. We define the *exponential process* corresponding to a general semimartingale X as the solution Y of this equation. An application of Itô's formula (for possibly discontinuous semimartingales) shows that a solution exists and is given by

$$\mathcal{E}(X)_t = e^{X_t - \frac{1}{2}[X^c]_t} \prod_{s\leq t} \big(1 + \Delta X_s\big) e^{-\Delta X_s}.$$

This solution is unique and, of course, reduces to the earlier definition of the exponential process if X is continuous.

2.2 Lévy's Theorem

The (predictable) quadratic variation process of a Brownian motion is the identity function. Lévy's theorem asserts that Brownian motion is the only continuous local martingale with this quadratic variation process. It is a useful tool to show that a given process is a Brownian motion. The continuity is essential, because the compensated Poisson process is another example of a martingale with predictable quadratic variation process equal to the identity.

2.6 Theorem (Lévy). *Let M be a continuous local martingale, 0 at 0, such that $[M]$ is the identity function. Then M is a Brownian motion process.*

Proof. For a fixed real number θ consider the complex-valued stochastic process

$$X_t = e^{i\theta M_t + \frac{1}{2}\theta^2 t}.$$

By application of Itô's formula to $X_t = f(M_t, t)$ with the complex-valued function $f(m,t) = \exp(i\theta m + \frac{1}{2}\theta^2 t)$, we find

$$dX_t = X_t i\theta\, dM_t + \tfrac{1}{2} X_t (i\theta)^2\, d[M]_t + X_t \tfrac{1}{2}\theta^2 dt = X_t i\theta\, dM_t,$$

since $[M]_t = t$ by assumption. It follows that $X = X_0 + i\theta X \cdot M$ and hence X is a (complex-valued) local martingale. Because $|X_t|$ is actually bounded for every fixed t, X is a martingale. The martingale relation $\mathrm{E}(X_t |\, \mathcal{F}_s) = X_s$ can be rewritten in the form

$$\mathrm{E}\Big(e^{i\theta(M_t - M_s)} |\, \mathcal{F}_s\Big) = e^{-\frac{1}{2}\theta^2 (t-s)}, \qquad \text{a.s.}, \qquad s < t.$$

This implies that $M_t - M_s$ is independent of $\mathcal{F}_s$ and possesses the normal distribution with mean zero and variance $t - s$. (Cf. Exercise 2.7.) ∎

2.7 EXERCISE. Let X be a random variable on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{F}_0 \subset \mathcal{F}$ a sub σ-field such that $\mathrm{E}(e^{i\theta X} |\, \mathcal{F}_0)$ is equal to a constant $c(\theta)$ for every $\theta \in \mathbb{R}$. Show that X is independent of $\mathcal{F}_0$.

2.3 Cameron-Martin-Girsanov Theorem

Given a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ and a probability measure $\tilde{\mathbb{P}}$ that is absolutely continuous relatively to $\mathbb{P}$, let $d\tilde{\mathbb{P}}/d\mathbb{P}$ be a version of the Radon-Nikodym density of $\tilde{\mathbb{P}}$ relative to $\mathbb{P}$. The process L defined by

$$L_t = \mathrm{E}\Big(\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}} |\, \mathcal{F}_t\Big) \tag{2.8}$$

is a nonnegative, uniformly integrable martingale with mean $\mathrm{E}L_t = 1$. For $t \to \infty$ it converges in mean to $\mathrm{E}\big(d\tilde{\mathbb{P}}/d\mathbb{P} |\, \mathcal{F}_\infty\big)$, for $\mathcal{F}_\infty$ the σ-field generated by $\cup_t \mathcal{F}_t$. Conversely, every nonnegative, uniformly integrable martingale L with mean 1 possesses a terminal variable L_∞, and can be used to define a probability measure $\tilde{\mathbb{P}}$ by $d\tilde{\mathbb{P}} = L_\infty\, d\mathbb{P}$. This measure satisfies (2.8), i.e. has density process L relative to $\mathbb{P}$. Thus there is a one-to-one relationship between "absolutely continuous changes of measure" on $\mathcal{F}_\infty$ and certain uniformly integrable martingales.

If the restrictions of $\tilde{\mathbb{P}}$ and $\mathbb{P}$ to the σ-field $\mathcal{F}_t$ are denoted by $\mathbb{P}_t$ and $\mathbb{P}_t$, then, for every $F \in \mathcal{F}_t$, by the martingale property of L,

$$\tilde{\mathbb{P}}_t(F) = \tilde{\mathbb{P}}(F) = \mathrm{E}L_\infty 1_F = \mathrm{E}L_t 1_F = \int_F L_t \, d\mathbb{P}_t.$$

This shows that the measure $\tilde{\mathbb{P}}_t$ is absolutely continuous with respect to the measure $\mathbb{P}_t$, with density

$$\frac{d\tilde{\mathbb{P}}_t}{d\mathbb{P}_t} = L_t. \tag{2.9}$$

For this reason the martingale L is also referred to as the *density process*. Its value at t gives insight in the "change of measure" of events up till time t.

It may happen that the measure $\tilde{\mathbb{P}}_t$ is absolutely continuous relative to the measure $\mathbb{P}_t$ for every $t \geq 0$, but $\tilde{\mathbb{P}}$ is not absolutely continuous relatively to $\mathbb{P}$. To cover this situation it is useful to introduce a concept of "local absolute continuity". A measure $\tilde{\mathbb{P}}$ is *locally absolutely continuous* relatively to a measure $\mathbb{P}$ if the restricted measures satisfy $\tilde{\mathbb{P}}_t \ll \mathbb{P}_t$ for every $t \in [0, \infty)$. In this case we can define a process L through the corresponding Radon-Nikodym densities, as in the preceding display (2.9). Then $\mathrm{E}L_t 1_F = \tilde{\mathbb{P}}_t(F) = \tilde{\mathbb{P}}_s(F) = \mathrm{E}L_s 1_F$, for every $F \in \mathcal{F}_s$ and $s < t$, and hence the process L is a martingale, with mean $\mathrm{E}L_t = 1$. If this (generalized) density process were uniformly integrable, then it would have a terminal variable, and we would be back in the situation as previously. In general, the density process is not uniformly integrable: the difference between absolute continuity and local absolute continuity is precisely the uniform integrability of the density process.

The martingale L possesses a cadlag version, which we use throughout this section. In the following lemma we collect some properties. Call $\tilde{\mathbb{P}}$ and $\mathbb{P}$ *locally equivalent* if the pair of measures is locally absolutely continuous in both directions.

2.10 Lemma. *If the measure $\tilde{\mathbb{P}}$ is locally absolutely continuous relative to the measure $\mathbb{P}$, and L is a cadlag version of the corresponding density process, then:*

(i) $\tilde{\mathbb{P}}(F \cap \{T < \infty\}) = \mathrm{E}L_T 1_F 1_{T<\infty}$, *for every $F \in \mathcal{F}_T$ and every stopping time T.*

(ii) If $T_n \uparrow \infty$ $\mathbb{P}$-almost surely, then $T_n \uparrow \infty$ $\tilde{\mathbb{P}}$-almost surely, for any increasing sequence of stopping times $T_1 \leq T_2 \leq \cdots$.

(iii) $L \wedge L_- > 0$ up to $\tilde{\mathbb{P}}$-evanescence; and also up to $\mathbb{P}$-evanescence if $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are locally equivalent.

(iv) There exists a stopping time T such that $L > 0$ on $[0, T)$ and $L = 0$ on $[T, \infty)$ up to $\mathbb{P}$-evanescence.

Proof. (i). For every $n \in \mathbb{N}$ the optional stopping theorem applied to the uniformly integrable martingale L^n yields $L_{T\wedge n} = \mathrm{E}(L_n|\,\mathcal{F}_T)$, $\mathbb{P}$-almost surely. For a given $F \in \mathcal{F}_T$ the set $F \cap \{T \le n\}$ is contained in both $\mathcal{F}_T$ and $\mathcal{F}_n$. We conclude that $\mathrm{E}L_T 1_F 1_{T\le n} = \mathrm{E}L_{T\wedge n} 1_F 1_{T\le n} = \mathrm{E}L_n 1_F 1_{T\le n} = \tilde{\mathrm{E}} 1_F 1_{T\le n}$. Finally, we let $n \uparrow \infty$.

(ii). Because $T = \lim T_n$ defines a stopping time, assertion (i) yields that $\tilde{\mathbb{P}}(T < \infty) = \mathrm{E}L_T 1_{T<\infty}$. If $\mathbb{P}(T = \infty) = 1$, then the right side is 0 and hence $\tilde{\mathbb{P}}(T = \infty) = 1$.

(iii). For $n \in \mathbb{N}$ define a stopping time by $T_n = \inf\{t > 0\colon L_t < n^{-1}\}$. By right continuity $L_{T_n} \le n^{-1}$ on the event $T_n < \infty$. Consequently property (i) gives $\tilde{\mathbb{P}}(T_n < \infty) = \mathrm{E}L_{T_n} 1_{T_n<\infty} \le n^{-1}$. We conclude that $\tilde{\mathbb{P}}(\inf_t L_t = 0) \le n^{-1}$ for every n, and hence $\inf_t L_t > 0$ almost surely under $\tilde{\mathbb{P}}$. Furthermore, $T_n \uparrow \infty$ almost surely under $\tilde{\mathbb{P}}$, and then by (ii) also under $\mathbb{P}$ if $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are locally equivalent. This is equivalent to the sample paths of L being bounded away from 0 on compacta up to $\mathbb{P}$-evanescence.

(iv). The stopping times T_n defined in the proof of (iii) are strictly increasing and hence possess a limit T. By definition of T_n we have $L_t \ge n^{-1}$ on $[0, T_n)$, whence $L_t > 0$ on $[0, T)$. For any m the optional stopping theorem gives $\mathrm{E}(L_{T\wedge m}|\,\mathcal{F}_{T_n\wedge m}) = L_{T_n\wedge m} \le n^{-1}$ on the event $T_n \le m$. Because $\{T_n \le m\} \in \mathcal{F}_{T_n\wedge m}$, we can conclude that $\mathrm{E}L_{T\wedge m} 1_{T\le m} \le \mathrm{E}L_{T\wedge m} 1_{T_n\le m} \le n^{-1}$ for every m and n, and hence $L_T = 0$ on the event $T < \infty$. For any stopping time $S \ge T$ another application of the optional stopping theorem gives $\mathrm{E}(L_{S\wedge m}|\,\mathcal{F}_{T\wedge m}) = L_{T\wedge m} = 0$ on the event $T \le m$. We conclude that $\mathrm{E}L_{S\wedge m} 1_{T\le m} = 0$ for every n and hence $L_S = 0$ on the event $S < \infty$. This is true in particular for $S = \inf\{t > T\colon L_t > \varepsilon\}$, for any $\varepsilon > 0$, and hence $L = 0$ on (T, ∞). ■

2.11 EXERCISE. If L is the density process of $\tilde{\mathbb{P}}$ relative to $\mathbb{P}$ and V_0 is a strictly positive $\mathcal{F}_0$-measurable random variable such that $\mathrm{E}_{\mathbb{P}} V_0 L_0 = 1$, then $V_0 L$ is also a density process. Of which measure?

2.12 EXERCISE. If $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are locally equivalent measures on a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\})$ with density process L and $\mathcal{F}_t' \subset \mathcal{F}_t$ is a sub-filtration, then there exists a cadlag version of the process $t \mapsto \mathrm{E}_{\mathbb{P}}(L_t|\,\mathcal{F}_t')$ and this is the density process of the restriction of $\tilde{\mathbb{P}}$ to $\mathcal{F}_\infty'$ relative to the restriction of $\mathbb{P}$ to $\mathcal{F}_\infty'$. Show this.

2.13 EXERCISE. If $\mathbb{P}$, $\mathbb{Q}$ and $\mathbb{R}$ are locally equivalent measures on a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\})$, then $\mathbb{R}$ has density process KL with respect to $\mathbb{P}$, for L the density process of $\mathbb{R}$ with respect to $\mathbb{Q}$ and and K the density process of $\mathbb{Q}$ relative to $\mathbb{P}$.

If M is a local martingale on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, then it typically looses the local martingale property if we use another measure $\tilde{\mathbb{P}}$.

The Cameron-Martin-Girsanov theorem shows that M is still a semimartingale under $\tilde{\mathbb{P}}$, and gives an explicit decomposition of M in its martingale and bounded variation parts.

We start with a general lemma on the martingale property under a "change of measure". We refer to a process that is a local martingale under $\mathbb{P}$ as a $\mathbb{P}$-local martingale. For simplicity we restrict ourselves to the case that $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are locally equivalent, i.e. the restrictions $\tilde{\mathbb{P}}_t$ and $\mathbb{P}_t$ are locally absolutely continuous for every t.

2.14 Lemma. *Let $\tilde{\mathbb{P}}$ and $\mathbb{P}$ be locally equivalent probability measures on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\})$ and let L be the corresponding density process. Then a stochastic process M is a $\tilde{\mathbb{P}}$-local martingale if and only if the process LM is a $\mathbb{P}$-local martingale.*

Proof. We first prove the theorem without "local". If M is an adapted $\tilde{\mathbb{P}}$-integrable process, then, for every $s < t$ and $F \in \mathcal{F}_s$,

$$\tilde{\mathrm{E}} M_t 1_F = \mathrm{E} L_t M_t 1_F,$$
$$\tilde{\mathrm{E}} M_s 1_F = \mathrm{E} L_s M_s 1_F,$$

The two left sides are identical for every $F \in \mathcal{F}_s$ and $s < t$ if and only if M is a $\tilde{\mathbb{P}}$-martingale. Similarly, the two right sides are identical if and only if LM is a $\mathbb{P}$-martingale. We conclude that M is a $\tilde{\mathbb{P}}$-martingale if and only if LM is a $\mathbb{P}$-martingale.

If M is a $\tilde{\mathbb{P}}$-local martingale and $0 \leq T_n \uparrow \infty$ is a localizing sequence, then the preceding shows that the process LM^{T_n} is a $\mathbb{P}$-martingale, for every n. Then so is the stopped process $(LM^{T_n})^{T_n} = (LM)^{T_n}$. Because T_n is also a localizing sequence under $\mathbb{P}$, we can conclude that LM is a $\mathbb{P}$-local martingale.

Because $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are locally equivalent, we can select a version of L that is strictly positive, by Lemma 2.10(iii). Then $d\mathbb{P}_t/d\tilde{\mathbb{P}}_t = L_t^{-1}$ and we can use the argument of the preceding paragraph in the other direction to see that $M = L^{-1}(LM)$ is a $\tilde{\mathbb{P}}$-local martingale if LM is a $\mathbb{P}$-local martingale. ∎

Warning. A sequence of stopping times is defined to be a "localizing sequence" if it is increasing everywhere and has almost sure limit ∞. The latter "almost sure" depends on the underlying probability measure. Thus a localizing sequence for a measure $\mathbb{P}$ need not be localizing for a measure $\tilde{\mathbb{P}}$. In view of Lemma 2.10(ii) this problem does not arise if the measures $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are locally equivalent. In the preceding lemma the "local martingale part" can be false if $\tilde{\mathbb{P}}$ is locally absolutely continuous relative to $\mathbb{P}$, but not the other way around.

If M itself is a $\mathbb{P}$-local martingale, then generally the process LM will not be a $\mathbb{P}$-local martingale, and hence the process M will not be a $\tilde{\mathbb{P}}$-local martingale. We can correct for this by subtracting an appropriate process.

We restrict ourselves to continuous local martingales M. Then a $\mathbb{P}$-local martingale becomes a $\tilde{\mathbb{P}}$ local martingale plus a "drift" $(L_-^{-1})\cdot[L,M]$, which is of locally bounded variation.

2.15 Theorem (Girsanov). *Let $\tilde{\mathbb{P}}$ and $\mathbb{P}$ be locally equivalent probability measures on $(\Omega,\mathcal{F},\{\mathcal{F}_t\})$ and let L be the density process of $\tilde{\mathbb{P}}$ relative to $\mathbb{P}$. If M is a continuous $\mathbb{P}$-local martingale, then $M - L_-^{-1}\cdot[L,M]$ is a $\tilde{\mathbb{P}}$-local martingale.*

Proof. By Lemma 2.10(ii) the process L_- is strictly positive under both $\tilde{\mathbb{P}}$ and $\mathbb{P}$, whence the process L_-^{-1} is well defined. Because it is left-continuous, it is locally bounded, so that the integral $L_-^{-1}\cdot[L,M]$ is well defined. We claim that the two processes

$$\begin{aligned}&LM-[L,M]\\&L(L_-^{-1}\cdot[L,M])-[L,M]\end{aligned}$$

are both $\mathbb{P}$-local martingales. Then, taking the difference, we see that the process $L(M-L_-^{-1}\cdot[L,M])$ is a $\mathbb{P}$-local martingale and hence the theorem is a consequence of Lemma 2.14.

That the first process in the display is a $\mathbb{P}$-local martingale is an immediate consequence of properties of the quadratic variation. For the second we apply the integration-by-parts (or Itô's) formula to see that

$$d\big(L(L_-^{-1}\cdot[L,M])\big) = (L_-^{-1}\cdot[L,M])\,dL + L_-\,d(L_-^{-1}\cdot[L,M]).$$

No "correction term" appears at the end of the display, because the quadratic covariation between the process L and the continuous process of locally bounded variation $L_-^{-1}\cdot[L,M]$ is zero. The integral of the first term on the right is a stochastic integral (of $L_-^{-1}\cdot[L,M]$) relative to the $\mathbb{P}$-martingale L and hence is a $\mathbb{P}$-local martingale. The integral of the second term is $[L,M]$. It follows that the process $L(L_-^{-1}\cdot[L,M])-[L,M]$ is a local martingale. ∎

* **2.16** EXERCISE. In the preceding theorem suppose that M is not necessarily continuous. Show that:

(i) If L is continuous, then the theorem is true as stated.
(ii) If the predictable quadratic covariation $\langle L,M\rangle$ is well defined, then the process $M-L_-^{-1}\cdot\langle L,M\rangle$ is a $\tilde{\mathbb{P}}$-local martingale, even if L and M are cadlag, but discontinuous.

For cadlag local L_2-martingales L and M the predictable quadratic covariation process $\langle L,M\rangle$ is defined as the unique predictable process of locally bounded variation such that $LM-\langle L,M\rangle$ is a local martingale, and can be shown to be equal to $[L,M]-\sum\Delta L\Delta M$. More generally we can consider $\langle L,M\rangle$ well defined if there exists a predictable process of locally bounded variation $\langle L,M\rangle$ such that $[L,M]-\langle L,M\rangle$ is a local martingale. This is certainly the case if the process $[L,M]$ is locally integrable, so that it has a compensator by the Doob-Meyer decomposition. [Hint: (ii) can be proved following the proof of the preceding theorem, but substituting $\langle L,M\rangle$ for $[L,M]$. In the second part of the proof we do obtain a term $d\big[L, L_-^{-1}\cdot\langle L,M\rangle\big]$. This may be nonzero, but can be written as $L_-^{-1}\Delta\langle L,M\rangle\,dL$ (Cf. Jacod and Shiryaev, I.4.49) and contributes another local martingale part.

* **2.17** EXERCISE. Any semimartingale X can be written as $X = X_0 + M + A$ for a local martingale with bounded jumps M and a process of locally bounded variation A. It can be shown that $[L,M]$ is locally integrable for any local martingale L and local martingale with bounded jumps M. (Localize by the minimum of $T_n = \inf\{t>0\colon \int_0^t |d[L,M]| > n\}$ and a stopping time making the jumps of L integrable. See Jacod and Shriyaev, III.3.14.) Deduce from this and the (ii) of the preceding problem that any $\mathbb{P}$-semimartingale is a $\tilde{\mathbb{P}}$-semimartingale for any equivalent probability measure $\tilde{\mathbb{P}}$.

The quadratic covariation process $[L,M]$ in the preceding theorem was meant to be the quadratic covariation process under the orginal measure $\mathbb{P}$. Because $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are assumed locally equivalent and a quadratic covariation process can be defined as a limit of inner products of increments, it is actually also the quadratic variation under $\tilde{\mathbb{P}}$.

Because $L_-^{-1}\cdot[L,M]$ is continuous and of locally bounded variation, the process $M - L_-^{-1}\cdot[L,M]$ possesses the same quadratic variation process $[M]$ as M, where again it does not matter if we use $\mathbb{P}$ or $\tilde{\mathbb{P}}$ as the reference measure. Thus even after correcting the “drift” due to a change of measure, the quadratic variation remains the same.

The latter remark is particularly interesting if M is a $\mathbb{P}$-Brownian motion process. Then both M and $M - L_-^{-1}\cdot[L,M]$ possess quadratic variation process the identity. Because $M - L_-^{-1}\cdot[L,M]$ is a continuous local martingale under $\tilde{\mathbb{P}}$, it is a Brownian motion under $\tilde{\mathbb{P}}$ by Lévy's theorem. This proves the following corollary.

2.18 Corollary. *Let $\tilde{\mathbb{P}}$ and $\mathbb{P}$ be locally equivalent probability measures on $(\Omega,\mathcal{F},\{\mathcal{F}_t\})$ and let L be the corresponding density process. If B is a $\mathbb{P}$-Brownian motion, then $B - L_-^{-1}\cdot[L,B]$ is a $\tilde{\mathbb{P}}$-Brownian motion.*

Many density processes L arise as exponential processes. In fact, given a strictly positive, continuous martingale L, the process $X = L_-^{-1}\cdot L$ is well defined and satisfies $L_-\,dX = dL$. The exponential process is the unique solution to this equation, whence $L = L_0\mathcal{E}(X)$. Girsanov's theorem takes a particularly simple form if formulated in terms of the process X.

2.19 Corollary. *Let $\tilde{\mathbb{P}}$ and $\mathbb{P}$ be locally equivalent probability measures on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\})$ and let the corresponding density process L take the form $L = L_0\mathcal{E}(X)$ for a continuous local martingale X, 0 at 0, and a strictly positive $\mathcal{F}_0$-measurable random variable L_0. If M is a continuous $\mathbb{P}$-local martingale, then $M - [X, M]$ is a $\tilde{\mathbb{P}}$-local martingale.*

Proof. The exponential process $L = L_0\mathcal{E}(X)$ satisfies $dL = L_-\,dX$, or equivalently, $L = L_0 + L_- \cdot X$. Hence $L_-^{-1} \cdot [L, M] = L_-^{-1} \cdot [L_- \cdot X, M] = [X, M]$. The corollary follows from Theorem 2.15. ∎

A special case arises if $L = \mathcal{E}(Y \cdot B)$ for Y a predictable process and B a Brownian motion. Then

$$\frac{d\tilde{\mathbb{P}}_t}{d\mathbb{P}_t} = \mathcal{E}(Y \cdot B)_t = e^{\int_0^t Y_s\,dB_s - \frac{1}{2}\int_0^t Y_s^2\,ds} \qquad \text{a.s.}. \tag{2.20}$$

By the preceding corollaries (with $X = Y \cdot B$ and $M = B$) the process

$$t \mapsto B_t - \int_0^t Y_s\,ds$$

is a Brownian motion under $\tilde{\mathbb{P}}$. This is the original form of Girsanov's theorem. The following exercise asks to derive the vector-valued form of this theorem.

2.21 EXERCISE **(Vector-valued Girsanov)**. Let $W = (W^{(1)}, \ldots, W^{(d)})$ be a d-dimensional $\mathbb{P}$-Brownian motion process and let $\mathbb{Q}$ be a probability measure that is absolutely continuous relative to $\mathbb{P}$ with density process of the form $L = \mathcal{E}(\sum_{j=1}^d \theta^{(j)} \cdot W^{(j)})$ relative to $\mathbb{P}$, for some predictable process $(\theta^{(1)}, \ldots, \theta^{(d)})$. Show that the process $\tilde{W}$ with coordinates $\tilde{W}^{(i)} = W^{(i)} - \int_0^\cdot \theta_s^{(i)}\,ds$ is a d-dimensional $\mathbb{Q}$-Brownian motion.

It is a fair question why we would be interested in "changes of measure" of the form (2.20). We shall see some reasons when discussing stochastic differential equations or option pricing in later chapters. For now we can note that in the situation that the filtration is the completion of the filtration generated by a Brownian motion any change to an equivalent measure is of the form (2.20).

* **2.22 Lemma.** *Let $\{\mathcal{F}_t\}$ be the completion of the natural filtration of a Brownian motion process B defined on $(\Omega, \mathcal{F}, \mathbb{P})$. If $\tilde{\mathbb{P}}$ is a probability measure on $(\Omega, \mathcal{F})$ that is equivalent to $\mathbb{P}$, then there exists a predictable process Y with $\int_0^t Y_s^2\,ds < \infty$ almost surely for every $t \geq 0$ such that the restrictions $\tilde{\mathbb{P}}_t$ and $\mathbb{P}_t$ of $\tilde{\mathbb{P}}$ and $\mathbb{P}$ to $\mathcal{F}_t$ satisfy (2.20).*

Proof. The density process L is a $\mathbb{P}$-martingale relative to the filtration $\{\mathcal{F}_t\}$. Because this is a Brownian filtration, Theorem 3.2 and Example 3.3

imply that L permits a continuous version. Because L is positive, the process L^{-1} is well defined, predictable and locally bounded. Hence the stochastic integral $Z = L^{-1} \cdot L$ is a well-defined local martingale, relative to the Brownian filtration $\{\mathcal{F}_t\}$. By Example 3.3 it can be represented as $Z = Y \cdot B$ for a predictable process Y as in the statement of the lemma. The definition $Z = L^{-1} \cdot L$ implies $dL = L\,dZ$. Because $\mathcal{F}_0$ is trivial, the density at zero can be taken equal to $L_0 = 1$. This pair of equations is solved uniquely by $L = \mathcal{E}(Z)$. (Cf. Exercise 2.3.) ∎

In many applications of Girsanov's theorem the process Y is actually given first, and the purpose is to "remove a drift" of the form $\int_0^t Y_s\,ds$ from a given Brownian motion B. Then we would like to construct the new measure $\tilde{\mathbb{P}}$ starting from $\mathbb{P}$, B, and Y. From the preceding we see that this is achieved by constructing $\tilde{\mathbb{P}}$ so as to have density process $L = \mathcal{E}(Y \cdot B)$ relative to $\mathbb{P}$. This requires the exponential process $\mathcal{E}(Y \cdot B)$ to be at least a martingale. By Novikov's theorem a sufficient condition for this is that, for every $t > 0$,

$$\mathrm{E}e^{\frac{1}{2}\int_0^t Y_s^2\,ds} < \infty.$$

If $\mathcal{E}(Y \cdot B)$ is a uniformly integrable martingale, then we can define $d\tilde{\mathbb{P}} = \mathcal{E}(Y \cdot B)_\infty\,d\mathbb{P}$, and we conclude that the process $B_t - \int_0^t Y_s\,ds$ is a Brownian motion.

Under just Novikov's condition the exponential process $\mathcal{E}(Y \cdot B)$ is not necessarily uniformly integrable, but it is a martingale and hence uniformly integrable if restricted to a finite interval $[0, T]$. Consequently, the stopped exponential process $\mathcal{E}(Y \cdot B)^T = \mathcal{E}\big((Y1_{[0,T]}) \cdot B\big)$ is uniformly integrable. Then the corresponding density process is given by (2.20) with $Y1_{[0,T]}$ replacing Y. We can conclude that the process $\{B_t - \int_0^{T\wedge t} Y_s\,ds\colon t \geq 0\}$ is a Brownian motion under the measure $\tilde{\mathbb{P}}$. In particular, the process $B_t - \int_0^t Y_s\,ds$ is a Brownian motion on the restricted time interval $[0, T]$ relative to the measure $\tilde{\mathbb{P}}_T$ with density $\mathcal{E}(Y \cdot B)_T$ relative to $\mathbb{P}$.

If Novikov's condition is satisfied for every $t > 0$, then we can obtain this conclusion for every $T > 0$. The measures $\tilde{\mathbb{P}}_T$ depend on T, but are clearly related. We conclude this section by a discussion of conditions under which they can be put together to a single measure.

* 2.3.1 Local to Global

If a probability measure $\tilde{\mathbb{P}}$ is locally absolutely continuous relative to a probability measure $\mathbb{P}$, then the corresponding density process is a nonnegative $\mathbb{P}$-martingale with mean 1. We may ask if, conversely, every nonnegative martingale L with mean 1 on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ arises as the density process of a measure $\tilde{\mathbb{P}}$ relative to $\mathbb{P}$. In the introduction of this section we have seen that the answer to this question is positive if the martingale is uniformly integrable, but the answer is negative in general.

Given a martingale L and a measure $\mathbb{P}$ we can define for each $t \geq 0$ a measure $\tilde{\mathbb{P}}_t$ on the σ-field $\mathcal{F}_t$ by

$$\frac{d\tilde{\mathbb{P}}_t}{d\mathbb{P}_t} = L_t.$$

If the martingale is nonnegative with mean value 1, then this defines a probability measure for every t. The martingale property ensures that the collection of measures $\tilde{\mathbb{P}}_t$ is consistent in the sense that $\tilde{\mathbb{P}}_s$ is the restriction of $\tilde{\mathbb{P}}_t$ to $\mathcal{F}_s$, for every $s < t$. The remaining question is whether we can find a measure $\tilde{\mathbb{P}}$ on $\mathcal{F}_\infty$ for which $\tilde{\mathbb{P}}_t$ is its restriction to $\mathcal{F}_t$.

Such a "projective limit" of the system $(\tilde{\mathbb{P}}_t, \mathcal{F}_t)$ does not necessarily exist under just the condition that the process L is a martingale. A sufficient condition is that the filtration be generated by some appropriate process. Then we can essentially use Kolmogorov's consistency theorem to construct $\tilde{\mathbb{P}}$.

2.23 Theorem. *Let L be a nonnegative martingale with mean 1 on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. If $\mathcal{F}_t$ is the filtration $\sigma(Z_s: s \leq t)$ generated by some stochastic process Z on $(\Omega, \mathcal{F})$ with values in a Polish space, then there exists a probability measure $\tilde{\mathbb{P}}$ on $\mathcal{F}_\infty$ whose restriction to $\mathcal{F}_t$ possesses density L_t relative to $\mathbb{P}$ for every $t > 0$.*

Proof. Define a probability measure $\tilde{\mathbb{P}}_t$ on $\mathcal{F}_\infty$ by its density L_t relative to $\mathbb{P}$, as before. For $0 \leq t_1 < t_2 < \cdots < t_k$ let $R_{t_1,\ldots,t_k}$ be the distribution of the vector $(Z_{t_1}, \ldots, Z_{t_k})$ on the Borel σ-field $\mathcal{D}^k$ of the space $\mathbb{D}^k$ if $(\Omega, \mathcal{F}_\infty)$ is equipped with $\tilde{\mathbb{P}}_{t_k}$. This system of distributions is consistent in the sense of Kolmogorov and hence there exists a probability measure R on the space $(\mathbb{D}^{[0,\infty)}, \mathcal{D}^{[0,\infty)})$ whose marginal distributions are equal to the measures $R_{t_1,\ldots,t_k}$.

For a measurable set $B \in \mathcal{D}^{[0,\infty)}$ now define $\tilde{\mathbb{P}}\big(Z^{-1}(B)\big) = R(B)$. If this is well defined, then it is not difficult to verify that $\tilde{\mathbb{P}}$ is a probability measure on $\mathcal{F}_\infty = Z^{-1}(\mathcal{D}^{[0,\infty)})$ with the desired properties.

The definition of $\tilde{\mathbb{P}}$ is well posed if $Z^{-1}(B) = Z^{-1}(B')$ for a pair of sets $B, B' \in \mathcal{D}^{[0,\infty)}$ implies that $R(B) = R(B')$. Actually, it suffices to show that this is true for every pair of sets B, B' in the union $\mathcal{A}$ of all cylinder σ-fields in $\mathbb{D}^{[0,\infty)}$ (the collection of all measurable sets depending on only finitely many coordinates). Then $\tilde{\mathbb{P}}$ is well defined and σ-additive on $\cup_t \mathcal{F}_t = Z^{-1}(A)$, which is an algeba, and hence possesses a unique extension to the σ-field $\mathcal{F}_\infty$, by Carathéodory's theorem.

The algebra $\mathcal{A}$ consists of all sets B of the form $B = \big\{z \in \mathbb{D}^{[0,\infty)}: (z_{t_1}, \ldots, z_{t_k}) \in B_k\big\}$ for a Borel set B_k in $\mathbb{R}^k$. If $Z^{-1}(B) = Z^{-1}(B')$ for sets $B, B \in \mathcal{A}$, then there exist k, coordinates $t_1, \ldots, t_k$, and Borel sets B_k, B_k' such that $\{(Z_{t_1}, \ldots, Z_{t_k}) \in B_k\} = \{(Z_{t_1}, \ldots, Z_{t_k}) \in B_k'\}$ and hence $R_{t_1,\ldots,t_k}(B_k) = R_{t_1,\ldots,t_k}(B_k')$, by the definition of $R_{t_1,\ldots,t_k}$. ∎

The condition of the preceding lemma that the filtration be the natural filtration generated by a process Z does not permit that the filtration is complete under $\mathbb{P}$. In fact, completion may cause problems, because, in general, the measure $\tilde{\mathbb{P}}$ will not be absolutely continuous relative to $\mathbb{P}$. This is illustrated in the following simple problem.

* **2.24 Example (Brownian motion with linear drift).** Let B be a Brownian motion on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, which is assumed to satisfy the usual conditions. For a given constant $\mu > 0$ consider the process L defined by

$$L_t = e^{\mu B_t - \frac{1}{2}\mu^2 t}.$$

The process L can be seen to be a $\mathbb{P}$-martingale, either by direct calculation or by Novikov's condition, and it is nonnegative with mean 1. Therefore, for every $t \geq 0$ we can define a probability measure $\tilde{\mathbb{P}}_t$ on $\mathcal{F}_t$ by $d\tilde{\mathbb{P}}_t = L_t\, d\mathbb{P}$. Because by assumption the Brownian motion B is adapted to the given filtration, the natural filtration $\mathcal{F}_t^o$ generated by B is contained in the filtration $\mathcal{F}_t$. The measures $\tilde{\mathbb{P}}_t$ are also defined on the filtration $\mathcal{F}_t^o$. By the preceding lemma there exists a probability measure $\tilde{\mathbb{P}}$ on $(\Omega, \mathcal{F}_\infty^o)$ whose restriction to $\mathcal{F}_t^o$ is $\tilde{\mathbb{P}}_t$, for every t.

We shall now show that:

(i) There is no probability measure $\tilde{\mathbb{P}}$ on $(\Omega, \mathcal{F}_\infty)$ whose restriction to $\mathcal{F}_t$ is equal to $\tilde{\mathbb{P}}_t$.

(ii) The process $B_t - \mu t$ is a Brownian motion on $(\Omega, \mathcal{F}_\infty^o, \{\mathcal{F}_t^o\}, \tilde{\mathbb{P}})$ (and hence also on the completion of this filtered space).

It was argued previously using Girsanov's theorem that the process $\{B_t - \mu t: 0 \leq t \leq T\}$ is a Brownian motion on the "truncated" filtered space $(\Omega, \mathcal{F}_T, \{\mathcal{F}_t \cap \mathcal{F}_T\}, \tilde{\mathbb{P}}_T)$, for every $T > 0$. Because the process is adapted to the smaller filtration $\mathcal{F}_t^o$, it is also a Brownian motion on the space $(\Omega, \mathcal{F}_T^o, \{\mathcal{F}_t^o \cap \mathcal{F}_T^o\}, \tilde{\mathbb{P}}_T)$. This being true for every $T > 0$ implies (ii).

If there were a probability measure $\tilde{\mathbb{P}}$ on $(\Omega, \mathcal{F}_\infty)$ as in (i), then the process $B_t - \mu t$ would be a Brownian motion on the filtered space $(\Omega, \mathcal{F}_\infty, \{\mathcal{F}_t\}, \tilde{\mathbb{P}})$, by Girsanov's theorem. We shall show that this leads to a contradiction. For $n \in \mathbb{R}$ define the event

$$F_\nu = \Big\{\omega \in \Omega: \lim_{t \to \infty} \frac{B_t(\omega)}{t} = \nu\Big\}.$$

Then $\mathcal{F}_\nu \in \mathcal{F}_\infty^o$ and $F_\nu \cap F_{\nu'} = \emptyset$ for $\nu \neq \nu'$. Furthermore, by the ergodic theorem for Brownian motion, $\mathbb{P}(F_0) = 1$ and hence $\mathbb{P}(F_\mu) = 0$. Because $B_t - \mu t$ is a Brownian motion under $\tilde{\mathbb{P}}$, also $\tilde{\mathbb{P}}(F_\mu) = 1$ and hence $\tilde{\mathbb{P}}(F_0) = 0$. Every subset F of F_μ possesses $\mathbb{P}(F) = 0$ and hence is contained in $\mathcal{F}_0$, by the (assumed) completeness of the filtration $\{\mathcal{F}_t\}$. If $B_t - \mu t$ would be a Brownian motion on $(\Omega, \mathcal{F}_\infty, \{\mathcal{F}_t\}, \tilde{\mathbb{P}})$, then $B_t - \mu t$ would be independent (relative to $\tilde{\mathbb{P}}$) of $\mathcal{F}_0$. In particular, B_t would be independent of the event $\{B_t \in C\} \cap F_\mu$ for every Borel set C. Because $\tilde{\mathbb{P}}(F_\mu) = 1$, the variable B_t

would also be independent of the event $\{\mathcal{B}_t \in C\}$. This is only possible if B_t is degenerate, which contradicts the fact that $B_t - \mu t$ possesses a normal distribution with positive variance. We conclude that $\tilde{\mathbb{P}}$ does not exist on $\mathcal{F}_\infty$.

The problem (i) in this example is caused by the fact that the projective limit of the measures $\tilde{\mathbb{P}}_t$, which exists on the smaller σ-field $\mathcal{F}_\infty^o$, is orthogonal to the measure $\mathbb{P}$. In such a situation completion of a filtration under one of the two measures effectively adds all events that are nontrivial under the other measure to the filtration at time zero. This is clearly undesirable if we wish to study a process under both probability measures. □

3
Martingale Representation

The pricing theory for financial derivatives is based on "replicating strategies". At an abstract level such strategies are representers in "martingale representation theorems". In this chapter we discuss a general, abstract theorem and some examples.

3.1 Representations

We shall say that a local martingale M on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ has the *representing property* if every cadlag local martingale N on this filtered space can be written as a stochastic integral $N = N_0 + H \cdot M$ relative to M, for some predictable process H. Thus the infinitesimal increment dN_t of an arbitrary cadlag local martingale N is a multiple $H_t\, dM_t$ of the corresponding increment of M. Intuitively, the predictability of H means that the quantity H_t is known "just before t", and hence "all the randomness in dN_t is contained in dM_t". A constructive interpretation of the representation $N = N_0 + H \cdot M$ is to view a sample path of N as evolving through extending it at every time t by a multiple $H_t\, dM_t$ of the increment dM_t, where the multiplication factor H_t can be considered a constant and the increment dM_t is a random variable "generated" at time t.

Obviously, the representing property is a very special property. Its validity depends on both the local martingale M and the filtration. The classical example is given by the pair of a Brownian motion and its augmented natural filtration. In this chapter we deduce this result from a characterization of the representing property through the uniqueness of a "martingale measure" and Girsanov's theorem. This characterization also implies the representing property of the continuous martingale part of a solution to a stochastic differential equation, if this is weakly unique.

3.1 Definition. *A (local) martingale measure corresponding to a given (local) martingale M on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ is a probability measure $\mathbb{Q}$ on $(\Omega, \mathcal{F}_\infty)$ such that*

(i) $\mathbb{Q} \ll\gg \mathbb{P}$.
(ii) $\mathbb{Q}_0 = \mathbb{P}_0$.
(iii) M *is a* $\mathbb{Q}$*-(local) martingale.*

By assumption the original measure $\mathbb{P}$ is a martingale measure. The following theorem characterizes the representing property of M through the uniqueness of $\mathbb{P}$ as a martingale measure.

Warning. Other authors say that a local martingale N possesses the *representation property* relative to M if N can be represented as a stochastic integral $N = N_0 + H \cdot M$. We shall not use this phrase, but note that M possesses the "representing property" if all local martingales N possess the "representation property" relative to M.

Warning. Other authors (e.g. Jacod and Shiryaev) define the representing property of a process M with jumps through representation of each local martingale N as $N = N_0 + G \cdot M^c + H * (\mu^M - \nu^M)$, where M^c is the continuous, local martingale part of M, $\mu^M - \nu^M$ is its "compensated jump measure", and (G, H) are suitable predictable processes. This type of representation permits more flexibility in using the jumps than representations of the form $N = N_0 + H \cdot M$. Lévy processes do for instance possess the representing property in this extended sense, whereas most Lévy processes with jumps do not possess the representing property as considered in this chapter. Within the financial context representation through a stochastic integral $H \cdot M$ is more appropriate, as this corresponds to the gain process of a trading strategy.

3.2 Theorem. *A continuous (local) martingale M on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ possesses the representing property if and only if $\mathbb{P}$ is the unique (local) martingale measure on $\mathcal{F}_\infty$ corresponding to M and $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$.*

The proof of this theorem is given in a series of steps in Section 3.4. The theorem applies both to martingales and local martingales M. If M is a martingale, then a "martingale measure" may be understood to be a measure satisfying (i)–(ii) of the preceding definition under which M is a martingale, not a local martingale.

Warning. Some authors do not include property (i) in the definition of a martingale measure. In that case the correct characterization is that $\mathbb{P}$ is an extreme point in the set of all martingale measures. If (i) is included in the definition of a martingale measure (as is done here), then the latter description is also correct, but the "extremeness" of $\mathbb{P}$ is a trivial consequence of the fact that it is the only martingale measure.

That the representing property can be characterized by uniqueness of a martingale measure is surprising, but can be explained informally from Girsanov's theorem and the fact that a density process is a martingale. If M is a continuous $\mathbb{P}$-local martingale and $\mathbb{Q}$ an equivalent probability measure with density process L relative to $\mathbb{P}$, then the process $M - L_-^{-1} \cdot [L, M]$ is a $\mathbb{Q}$-local martingale, by Girsanov's theorem. Hence M is a $\mathbb{Q}$-local martingale, and $\mathbb{Q}$ a martingale measure, if only if $[L, M] = 0$. This expresses some sort of "orthogonality" of L and M. We shall see that the process M possesses the representing property if and only there exists no nontrivial local martingale L that is orthogonal to M in this sense. Because a density process is a local martingale, this translates into uniqueness of $\mathbb{P}$ as a martingale measure.

The theorem is true for vector-valued local martingales M, provided that the stochastic integrals $H \cdot M$ are interpreted appropriately. The correct interpretation is not entirely trivial.

For a one-dimensional, continuous local martingale the integral $H \cdot M$ is defined for any predictable process H such that $\int_0^t H_s^2 \, d[M]_s < \infty$ almost surely for every t.† Consequently, if $H = (H^{(1)}, \ldots, H^{(d)})$ and $M = (M^{(1)}, \ldots, M^{(d)})$ are vectors of predictable processes and continuous local martingales such that $\int (H_s^{(i)})^2 \, d[M^{(i)}]_s < \infty$ almost surely for every i, then every of the stochastic integrals $H^{(i)} \cdot M^{(i)}$ is well defined and we define

$$H \cdot M = \sum_{i=1}^{n} H^{(i)} \cdot M^{(i)}.$$

If the $(d \times d)$-matrix $[M]$ of quadratic variations $[M^{(i)}, M^{(j)}]$ is "uniformly nonsingular", then this set of predictable integrands is appropriate and the representing theorem is true as stated. For instance, this is the case for M equal to multivariate Brownian motion, when $[M]_t$ is t times the identity matrix. However, in general the set of processes $H \cdot M$ obtained in this way is not large enough to make the representing theorem true. In Section 3.5 we define the stochastic integral $H \cdot M$ for a larger class of predictable processes H, essentially through a "closure operation". With this extended definition Theorem 3.2 is correct for multivariate local martingales, as stated.

The classical example of a martingale with the representing property is Brownian motion. There are nice direct proofs of this fact, but in the following example we deduce it from the preceding theorem.

---

† Under the latter condition the processes H and M can be localized by the stopping times

$$T_n = \inf\{t > 0 \colon |M_t| > n, \int_0^t H_s^2 \, d[M]_s > n\}.$$

The process $H1_{[0,T_n]}$ is in the space $L_2([0,\infty) \times \Omega, \mathcal{P}, \mu_{M^{T_n}})$ for every n (where $\mathcal{P}$ is the predictable σ-field and μ_M the Doléans measure of an L_2-martingale M), and hence $H1_{[0,T_n]} \cdot M^{T_n}$ is well defined; its almost sure limit as $n \to \infty$ exists and is the stochastic integral $H \cdot M$.

3.3 Example (Brownian motion). Let M be Brownian motion on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, with $\mathcal{F}_t$ equal to the augmented natural filtration generated by M. In this setting Brownian motion possesses the representing property.

To prove this it suffices to establish uniqueness of $\mathbb{P}$ as a martingale measure. If $\mathbb{Q}$ is another martingale measure, then the quadratic variation of M under $\mathbb{Q}$ is the same as under $\mathbb{P}$, because by definition $\mathbb{Q}$ and $\mathbb{P}$ are equivalent. Because M is a $\mathbb{P}$-Brownian motion, this quadratic variation is the identity function. Hence, by Lévy's theorem, M is also a $\mathbb{Q}$-Brownian motion. We conclude that the induced law of M on $\mathbb{R}^{[0,\infty)}$ is the same under both $\mathbb{P}$ and $\mathbb{Q}$, or, equivalently, the measures $\mathbb{P}$ and $\mathbb{Q}$ are the same on the σ-field $\mathcal{F}_\infty^o := \sigma(M_t : t \geq 0)$. Then they also agree on the completion of this σ-field, which is $\mathcal{F}_\infty$, by assumption. □

The preceding theorem extends to local martingales with jumps, at least to locally bounded ones and for representing nonnegative local martingales.

3.4 Theorem. *If $\mathbb{P}$ is a unique local martingale measure for the locally bounded, cadlag, local martingale M on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, then every nonnegative local martingale N on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ can be written as $N = N_0 + H \cdot M$ for some predictable process H.*

The stochastic integral $H \cdot M$ must be interpreted appropriately, and we must allow a "maximal" set of possible integrands H to make the theorem true. Vectors of locally bounded predictable processes are of course valid integrands, but the process H in the theorem is only restricted to be "M-integrable". The "maximal" extension of the domain of the stochastic integral for processes M with jumps is even more technical than for continuous processes. See Section 3.5 for a discussion.

The preceding theorem can be derived from the following theorem, which extends the representation to supermartingales. Because a nonnegative stochastic integral relative to a local martingale is automatically a local martingale, a "true" nonnegative supermartingale N can typically not be represented as $N = N_0 + H \cdot M$. However, the following theorem shows that there is a process H such that the difference $C = N - N_0 - H \cdot M$ has nondecreasing sample paths. Thus the decreasing nature (in the mean) of the supermartingale N is captured by the (pointwise) increasing process C.

The theorem also replaces the condition that the martingale measure is unique by the assumption that the process N is a supermartingale under every local martingale measure.

3.5 Theorem. *Let M be a locally bounded, cadlag, local martingale on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. If the nonnegative, cadlag process*

N is a $\mathbb{Q}$-supermartingale for every local martingale measure $\mathbb{Q}$ for M, then there exist a predictable process H and an adapted, nondecreasing process C such that $N = N_0 + H \cdot M - C$.

Proof. For a proof of Theorem 3.5 see D.O. Kramkov, *Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets*, Probability Theory and Related Fields 105, 1996, 459–479. The definition of a martingale measure in this paper does not include (ii) of Definition 3.1. However, if N is a (super)martingale relative to every martingale measure satisfying (i)+(ii)+(iii), then N is automatically a (super)martingale relative to every measure $\mathbb{Q}$ that satisfies only (i)+(iii). Indeed, given $\mathbb{Q}$ satisfying (i)+(iii) the measure $\tilde{\mathbb{Q}}$ with density $d\mathbb{P}_0/d\mathbb{Q}_0$ relative to $\mathbb{Q}$ satisfies (i)+(ii), and also (iii) because LM is a $\mathbb{Q}$-local martingale for $L = d\mathbb{P}_0/d\mathbb{Q}_0$ (constant in time) the density process of $\tilde{\mathbb{Q}}$ relative to $\mathbb{Q}$. Because $\tilde{\mathbb{Q}}$ is a martingale measure in the sense of Definition 3.1, the process N is a $\tilde{\mathbb{Q}}$-supermartingale by assumption, which implies that it is a $\mathbb{Q}$-supermartingale.

Theorem 3.4 can be derived from Theorem 3.5 by first noting that a nonnegative local martingale is a supermartingale. Therefore, according to Theorem 3.5 the local martingale N of Theorem 3.4 can be written $N = N_0 + H \cdot M - C$ for a nondecreasing process C. The process C is the difference of two local martingales and hence a local martingale itself. (We use that the stochastic integral $H \cdot M$ is a local martingale, which is not automatic with the extended definition of stochastic integral, but true in this case as the process is nonnegative.) If it is localized by stopping times T_n, then C^{T_n} is a martingale and hence $\mathrm{E}C_t^{T_n} = \mathrm{E}C_0 = 0$. By nonnegativity $C_t^{T_n} = 0$ almost surely, whence $C = 0$. ∎

* **3.6** EXERCISE. Suppose that the process N in Theorem 3.5 is a nonnegative, cadlag process which is a $\mathbb{Q}$-supermartingale for every martingale measure $\mathbb{Q}$ for M. Investigate whether the assertion of the theorem is still true. [Can the class of martingale measures be smaller than the class of local martingale measures?]

3.2 Stability

In this section we discuss two situations in which the representing property of a given local martingale is inherited by another local martingale constructed from it.

The first situation concerns a change of measure. If M is a continuous $\mathbb{P}$-local martingale and $\tilde{\mathbb{P}}$ an equivalent probability measure with density process L relative to $\mathbb{P}$, then the process $\tilde{M} = M - L_-^{-1} \cdot [L, M]$ is a $\tilde{\mathbb{P}}$-martingale, by Girsanov's theorem. Because M and $\tilde{M}$ differ only by a process of bounded variation, and this process is chosen to retain the martingale property, we should expect that $\tilde{M}$ possesses the representing property for $\tilde{\mathbb{P}}$-local martingales if M possesses the representing property for $\mathbb{P}$-local martingales. The following lemma shows that this expectation is justified.

3.7 Lemma. *Let M be a continuous $\mathbb{P}$-local martingale on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ and let $\tilde{\mathbb{P}}$ be a probability measure that is equivalent to $\mathbb{P}$ with density process L. Then M possesses the representing property on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ if and only if the continuous $\tilde{\mathbb{P}}$-local martingale $\tilde{M} = M - L_-^{-1} \cdot [L, M]$ possesses the representing property on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \tilde{\mathbb{P}})$.*

Proof. Suppose that $\tilde{M}$ possesses the representing property on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \tilde{\mathbb{P}})$, and let N be a cadlag $\mathbb{P}$-local martingale. The process $1/L$ is the density process of $\mathbb{P}$ relative to $\tilde{\mathbb{P}}$, and hence is a $\tilde{\mathbb{P}}$-martingale. Thus it can be represented as a stochastic integral relative to $\tilde{M}$ and hence is continuous. The process $\tilde{N} = N - L_-^{-1} \cdot [L, N]$ is a $\tilde{\mathbb{P}}$-local martingale, by Girsanov's theorem, and hence there exists a predictable process ϕ such that $\tilde{N} = \tilde{N}_0 + \phi \cdot \tilde{M}$. This implies that $N - N_0 - \phi \cdot M = L_-^{-1} \cdot [L, N] - \phi \cdot L_-^{-1} \cdot [L, M]$. The left side is a cadlag $\mathbb{P}$-local martingale, 0 at 0, and the right side is a continuous process of bounded variation. It follows that both sides are identically zero, whence $N = N_0 + \phi \cdot M$.

The converse is proved similarly. ∎

Next suppose that $N = \sigma \cdot M$ for a predictable process σ and a local martingale M that possesses the representing property on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. If σ is never zero, then we can write $H \cdot M = (H/\sigma) \cdot N$ and hence every stochastic integral relative to M can be expressed as a stochastic integral relative to N. Thus N inherits the representing property from M in this case.

3.8 EXERCISE. Verify that H/σ is a good integrand for $N = \sigma \cdot M$ whenever H is a good integrand for the continuous local martingale M, where "good" in the second case means that $\int_0^t H_s^2 \, d[M]_s < \infty$ almost surely, for every $t > 0$.

This remains true if the local martingale M is vector-valued and σ is a matrix-valued predictable process. For simplicity suppose that the local martingale M is d-dimensional and that σ takes its values in the set of $d \times d$-matrices, so that N is d-dimensional as well. As intuition suggests, the correct condition is that σ_t is invertible for every t. To make this true we need the extended definition of a stochastic integral relative to a multivariate local martingale, discussed in Section 3.5. It is implicitly assumed in the following lemma that each of the stochastic integrals $\sigma^{(i,\cdot)} \cdot M$, with

$\sigma^{(i,\cdot)}$ the ith row of the matrix σ, is well defined in the sense discussed in Section 3.5.

3.9 Lemma. *Suppose that the continuous, local martingale M possesses the representing property on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. If $N = \sigma \cdot M$ for a predictable process σ taking its values in the invertible matrices, then N also possesses the representing property on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$.*

Proof. Let $[M]_t = \int_0^t C_s \, d|[M]|_s$ be a representation of $[M]$ as used in Section 3.5. Then by assumption, for every i,

$$\int_0^t \sum_k \sum_l \sigma_s^{(i,k)} \sigma_s^{(i,l)} C_s^{(k,l)} \, d|[M]|_s < \infty, \qquad \text{a.s..}$$

This is exactly the ith diagonal element of the analogous representation of $[N]$, which is given by $[N]_t = \int_0^t \sigma_s C_s \sigma_s^T \, d|[M]|_s$. It follows readily that the process $(\sigma^T)^{-1} H$ is a good integrand for N if and only if the process H is a good integrand for M, i.e.

$$\int_0^t \bigl((\sigma_s^T)^{-1} H_s\bigr)^2 \, d[N]_s < \infty, \quad \text{iff} \quad \int_0^t H_s^2 \, d[M]_s < \infty, \qquad \text{a.s..}$$

Furthermore, it can be verified that $H \cdot (\sigma \cdot M) = (\sigma^T H) \cdot M$, so that $H \cdot M = \bigl((\sigma^T)^{-1} H\bigr) \cdot N$, whenever these integrals are well defined. ∎

3.3 Stochastic Differential Equations

Consider a stochastic differential equation of the form

$$dX_t = \mu(t, X_t) \, dt + \sigma(t, X_t) \, dW_t. \tag{3.10}$$

Here μ and σ are given vector-valued and matrix-valued, measurable functions and W is a vector-valued Brownian motion process on a given filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. By definition a solution of this equation is an adapted vector-valued process X with continuous sample paths such that

$$X_t = X_0 + \int_0^t \mu(s, X_s) \, ds + \int_0^t \sigma(s, X_s) \, dW_s. \tag{3.11}$$

Here it is implicitly required that the two integrals on the right, a Lebesgue integral and a stochastic integral, respectively, are well defined. This is the case if, almost surely, for every $t > 0$,

$$\int_0^t \bigl|\mu(s, X_s)\bigr| \, ds < \infty, \qquad \int_0^t \bigl\|\sigma(s, X_s)\bigr\|^2 \, ds < \infty.$$

Ostensibly, under these conditions, the two integrals in (3.11) give the decomposition of X into its bounded variation part and its local martingale part. In this section we show that the local martingale part possesses the representing property for all $\mathcal{F}_t^X$-adapted martingales if the solution to the stochastic differential equation is "weakly unique".

As is customary we can discern "strong" and "weak solutions" of the stochastic differential equation (3.10). For a *strong solution* the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ is given together with a pair (W, ξ) of a Brownian motion W and an $\mathcal{F}_0$-measurable random variable ξ defined on it, and the solution X is a continuous, adapted process defined on the same filtered space satisfying (3.11) and $X_0 = \xi$. Often the precise definition also includes further measurability requirements on X, such as adaptation to the augmented filtration generated by (W, ξ). For a *weak solution* only the functions μ and σ are given a-priori, and the solution consists of a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ and a Brownian motion W and a process X defined on it satisfying (3.11) and with X_0 having a predescribed law.

The conceptually simplest sufficient conditions for weak uniqueness are the "Itô Lipschitz conditions", which also guarantee existence of a strong solution. Alternatively, there are sufficient conditions in terms of the generator of the diffusion equation. (See e.g. Karatzas and Shreve, Theorem 4.28.)

To the notion of a weak solution corresponds a notion of "weak uniqueness", which is also referred to as "uniqueness-in-law". The "law" of a solution X may be understood to be the set of distributions of all vectors $(X_{t_1}, \ldots, X_{t_k})$ for $t_1, \ldots, t_k$ varying over $\mathbb{R}^k$, and its law at time 0 is the law of X_0. The stochastic differential equation (3.11) allows a *weakly unique* solution if any two solutions X and $\tilde{X}$ with the same law at time 0, possibly defined relatively to different Brownian motions on different filtered probability spaces, possess the same laws. Equivalently, the solution is weakly unique if for any two solutions X and $\tilde{X}$ such that X_0 and $\tilde{X}_0$ are equal in distribution, the distributions of the vectors $(X_{t_1}, \ldots, X_{t_k})$ and $(\tilde{X}_{t_1}, \ldots, \tilde{X}_{t_k})$ on $\mathbb{R}^k$ are equal for every k and every set of time points $0 \leq t_1 \leq \cdots \leq t_k$. (We may think of the law of a solution X as the probability distribution induced by the map $X\colon \Omega \to C[0,\infty)^n$ on the Borel sets of the metric space $C[0,\infty)^n$ equipped with the topology of uniform convergence on compacta. Then "uniqueness-in-law" means that every weak solution induces the same law on the "canonical" space $C[0,\infty)^n$.)

By definition a solution X defined on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ is adapted to the given filtration $\mathcal{F}_t$, but it is of course also adapted to its natural filtration $(\mathcal{F}_t^X)^o$. Let $\mathcal{F}_t^X$ be the completion of this filtration, and assume that it is right-continuous. The process $A_t = \int_0^t \mu(s, X_s)\, ds$ is adapted to $\mathcal{F}_t^X$ by Fubini's theorem, and hence so is the process $X_t^c = \int_0^t \sigma(s, X_s)\, dW_s$, because it is the difference of X and A. Because X^c is a continuous $\mathcal{F}_t$-local martingale, it is also an $\mathcal{F}_t^X$-local martingale (Exercise 3.12) and hence X^c is also the continuous martingale part of X

relative to the filtration $\mathcal{F}_t^X$. We pose the martingale representing problem for X^c relative to the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{P})$.

3.12 EXERCISE. If X is a continuous local martingale on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ that is adapted to the smaller filtration $\mathcal{G}_t \subset \mathcal{F}_t$, then X is also a $\mathcal{G}_t$-local martingale. [Hint: verify this first without "local". Next show that a localizing sequence can be chosen $\mathcal{G}_t$-adapted.]

** **3.13** EXERCISE. Is this still true without "continuous"?

3.14 Theorem. *If the solution to the stochastic differential equation (3.10) defined by μ and σ is weakly unique, then for any solution X on a given filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ equipped with a given Brownian motion W the local martingale part X^c possesses the representing property relative to the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{P})$.*

Proof. It suffices to show that $\mathbb{P}$ is a unique martingale measure for X^c on $(\Omega, \mathcal{F}_\infty^X)$. We give two proofs, a short indirect one and a longer direct proof.

The first proof proceeds by showing that for any martingale measure $\mathbb{Q}$ the process X is a solution to the "martingale problem" on $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{Q})$, i.e. for any twice continuously differentiable function f with compact support the process

$$f(X_t) - f(X_0) - \int_0^t \Bigl(f'(X_s)\mu(s, X_s) + \tfrac{1}{2} f''(X_s)\sigma\sigma^T(s, X_s)\Bigr)\, ds$$

is a $\mathbb{Q}$-martingale relative to $\mathcal{F}_t^X$. If this is the case, then there exists an extension of the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{Q})$ with a Brownian motion $\tilde{W}$ defined on it that together with the extension $\tilde{X}$ of X provides a weak solution to the stochastic differential equation (3.11) (see e.g. Karatzas and Shreve, Proposition 4.6) and the law of $\tilde{X}$ coincides with the law of X under $\mathbb{Q}$ (Karatzas and Shreve, Corolllary 4.8). For a martingale measure $\mathbb{Q}$ for X^c on $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\})$ the measures $\mathbb{Q}$ and $\mathbb{P}$ are the same on $\mathcal{F}_0^X$ by assumption and hence the initial law of X_0 is identical under $\mathbb{P}$ and $\mathbb{Q}$. By the assumed weak uniqueness the laws of the process X under the measures $\mathbb{P}$ and $\mathbb{Q}$ must agree. Equivalently, the measures $\mathbb{P}$ and $\mathbb{Q}$ agree on the σ-field $\sigma(X_t: t \geq 0) = (\mathcal{F}_\infty^X)^o$, and hence also on the completion $\mathcal{F}_\infty^X$, because $\mathbb{P}$ and $\mathbb{Q}$ are assumed equivalent.

To prove that X solves the martingale problem on $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{Q})$, we use Itô's formula to see that

$$\begin{aligned} df(X_t) &= f'(X_t)\, dX_t + \tfrac{1}{2} f''(X_t)\, d[X]_t \\ &= f'(X_t)\sigma(t, X_t)\, dW_t + \Bigl(f'(X_t)\mu(t, X_t) + \tfrac{1}{2} f''(X_t)\sigma\sigma^T(t, X_t)\Bigr)\, dt \end{aligned}$$

The assumption that $\mathbb{Q}$ is a martingale measure entails that $X_t^c = \int_0^t \sigma(s, X_s)\, dW_s$ is a $\mathbb{Q}$-local martingale relative to $\mathcal{F}_t^X$, and hence so is $\int_0^t f'(X_s)\sigma(s, X_s)\, dW_s = \int_0^t f'(X_s)\, dX_s^c$.

Rather than referring to the general results on the martingale problem we can also give a direct proof. First consider the case that X is one-dimensional. Then σ is a real-valued function and we can define a process

$$\bar{W}_t = \int_0^t 1_{\sigma(s,X_s)>0}\, ds = \int_0^t \frac{1}{\sigma(s,X_s)}\, 1_{\sigma(s,X_s)>0}\, dX_s^c.$$

The second representation shows that $\bar{W}$ is adapted to $\mathcal{F}_t^X$. Furthermore, if $\mathbb{Q}$ is a martingale measure, then X^c is a $\mathbb{Q}$-local martingale relative to $\mathcal{F}_t^X$ and hence so is $\bar{W}$. Suppose that we can define a Brownian motion $\bar{\bar{W}}$ on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{Q})$ that is independent of $\bar{W}$, and set

$$\tilde{W}_t = \bar{W}_t + \int_0^t 1_{\sigma(s,X_s)=0}\, d\bar{\bar{W}}_s.$$

Then $\tilde{W}$ is a continuous local martingale with quadratic variation process equal to the identity (check!) and hence $\tilde{W}$ is a Brownian motion, by Lévy's theorem. Because $\sigma(t,X_t)\, dW_t$ and $\sigma(t,X_t)\, d\tilde{W}_t$ are identical, the stochastic differential equation (3.11) holds with W replaced by $\tilde{W}$ and hence the process X together with the Brownian motion $\tilde{W}$ provide a weak solution to (3.11). By the assumed weak uniqueness we then obtain that $\mathbb{P} = \mathbb{Q}$ on $\mathcal{F}_\infty^X$, as before.

It may not be possible to construct an independent Brownian motion $\bar{\bar{W}}$ on the original filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{Q})$, and then the preceding construction is not possible. However, we may always replace this filtered space by (the completion of) a space of the form $(\Omega \times \bar{\Omega}, \mathcal{F} \times \bar{\mathcal{F}}, \{\mathcal{F}_t^X\} \times \bar{\mathcal{F}}_t, \mathbb{Q} \times \bar{\mathbb{Q}})$, with a Brownian motion $\bar{\bar{W}}$ defined on $(\bar{\Omega}, \bar{\mathcal{F}}, \{\bar{\mathcal{F}}_t\})$. We can view $(X, \bar{W})$ and $\bar{\bar{W}}$ as processes defined on this product space (depending only on the first and second coordinates of $(\omega, \bar{\omega})$ respectively), and then have a weak solution in this extended setting. The proof can then be completed as before.

In the multi-dimensional case we follow a similar argument, but the process $\tilde{W}$ must be constructed with more care. The Brownian motion W is d-dimensional and the process σ takes its values in the $(n \times d)$-matrices.

There exist matrix-valued, continuous, $\mathcal{F}_t^X$-adapted process O, Λ, and U such that, for every t, (cf. Karatzas and Shreve ???)

(i) $\sigma_t := \sigma(t, X_t) = U_t O_t \sqrt{\Lambda_t} O_t^T$.
(ii) O_t is a $(d \times d)$-orthogonal matrix.
(iii) Λ_t is a $(d \times d)$ diagonal (matrix with the eigenvalues of the matrix $\sigma_t \sigma_t^T$ on the diagonal).
(iv) U_t is an $(n \times d)$ matrix with $U_t^T U_t = I$ on the range of the matrix $\sigma_t \sigma_t^T$.

Given a d-dimensional Brownian motion $\bar{\bar{W}}$, defined on an extension of the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t^X\}, \mathbb{Q})$ and independent of (X, W), define

$$\tilde{W}_t = \int_0^t O_s 1_{\Lambda_s > 0} O_s^T\, dW_s + \int_0^t O_s 1_{\Lambda_s = 0} O_s^T\, d\bar{\bar{W}}_s.$$

Here $1_{\Lambda_t>0}$ is the diagonal matrix with 1 or 0 on the diagonal if the corresponding diagonal element of Λ_t is positive or 0, and $1_{\Lambda_t=0} = I - 1_{\Lambda_t>0}$. Then

$$\begin{aligned}\sigma(t,X_t)\,d\tilde{W}_t &= U_t O_t \sqrt{\Lambda_t} O_t^T O_t \left(1_{\Lambda_t>0} O_t^T\,dW_t + 1_{\Lambda_t=0} O_t^T\,d\bar{\bar{W}}_t\right)\\ &= \sigma(t,X_t)\,dW_t.\end{aligned}$$

It follows that X satisfies the stochastic differential equation (3.11) with W replaced by $\tilde{W}$.

By a similar calculation, with $\Lambda_t^{-1/2} 1_{\Lambda_t>0}$ the diagonal matrix with diagonal entries 0 or $\lambda^{-1/2}$ if the corresponding entropy of Λ_t is λ,

$$\Lambda_t^{-1/2} 1_{\Lambda_t>0} O_t^T U_t^T\,dX_t^c = \Lambda_t^{-1/2} 1_{\Lambda_t>0} O_t^T U_t^T \sigma_t\,dW_t = 1_{\Lambda_t>0} O_t^T\,dW_t.$$

It follows that the process $\int_0^t O_s 1_{\Lambda_s>0} O_s^T\,dW_s$, and hence also the process $\tilde{W}$, is adapted to $\mathcal{F}_t^X$. Because the quadratic variation can be computed to be $[\tilde{W}]_t = tI$, it follows by Lévy's theorem that $\tilde{W}$ is a $\mathbb{Q}$-Brownian motion.

Thus X together with $\tilde{W}$ and the (possibly) extended filtered probability space is a weak solution to the stochastic differential equation (3.11). We can finish the proof as before. ∎

3.15 EXERCISE. Compare the assertion of Theorem 3.14 with the assertion that can be obtained from Lemma 3.9. Do both approaches have advantages?‡

* 3.4 Proof of Theorem 3.2

Given a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, let $\mathcal{M}^2$ be the set of all cadlag L_2-bounded martingales on this space, and let $\mathcal{M}_0^2$ be the subset of such martingales that are 0 at 0. Every L_2-bounded martingale M possesses an almost sure limit M_∞ as $t \to \infty$, which is also an L_2-limit. The space $\mathcal{M}^2$ can be alternatively described as the set of all martingales M with

$$\sup_t \mathrm{E} M_t^2 < \infty, \qquad \mathrm{E} M_\infty^2 < \infty, \qquad \mathrm{E} \sup_t M_t^2 < \infty.$$

By the maximal inequalities for cadlag martingales the three finiteness conditions in this display are equivalent. The set $\mathcal{M}^2$ is a Hilbert space relative to the inner product $(M,N) = \mathrm{E} M_\infty N_\infty$.

Let M be a given continuous, local martingale on the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. The Doléans measure of M is defined by

$$d\mu_M(t,\omega) = d[M]_t(\omega)\,d\mathbb{P}(\omega).$$

‡ How much of the special structure of (3.11) is actually needed?

Let $L_2(M)$ denote the set of predictable stochastic processes $H\colon [0,\infty)\times\Omega\to\mathbb{R}$ such that $\int H^2\,d\mu_M<\infty$. By the isometry defining the stochastic integral,

$$\int H^2\,d\mu_M=\mathrm{E}(H\cdot M)_\infty^2.$$

Thus the set $L_2(M)$ consists exactly of all predictable processes H such that the stochastic integral $H\cdot M$ is an L_2-bounded martingale, and the map $H\mapsto H\cdot M$ from $L_2(M)$ to $\mathcal{M}_0^2$ is an isometry.

The local martingale M can "represent" all L_2-bounded cadlag martingales if this isometry is onto. In that case, for every cadlag L_2-bounded martingale N there exists a process H in $L_2(M)$ such that $N-N_0=H\cdot M$. Because the map $H\mapsto H\cdot M$ is an isometry, it possesses a closed range space, given by

$$I^2(M){:=}\bigl\{N\in\mathcal{M}_0^2{:}\, N=H\cdot M,\ \text{for some } H\in L_2(M)\bigr\}.$$

By the projection theorem for Hilbert spaces, the Hilbert space $\mathcal{M}_0^2$ can be decomposed as $\mathcal{M}_0^2=I^2(M)\oplus I^2(M)^\perp$, and $I^2(M)$ is equal to $\mathcal{M}_0^2$ if and only if its orthocomplement $I^2(M)^\perp$ is zero. In other words, the local martingale M can represent all cadlag L_2-bounded martingales if and only if every cadlag L_2-bounded martingale N with $N_0=0$ and $(N,M)=0$ is identically zero. The following lemmas translate this orthogonality into a martingale property.

3.16 Lemma. *The spaces $I^2(M)$ and $I^2(M)^\perp$ are closed under stopping: if T is a stopping time and $N\in I^2(M)$, then $N^T\in I^2(M)$, and similarly for $I^2(M)^\perp$.*

Proof. For the space $I^2(M)$ the assertion is immediate from the representation $N^T=(H1_{[0,T]})\cdot M$ if $N=H\cdot M$. For the space $I^2(M)^\perp$ we deduce the assertion from the equalities

$$\begin{aligned}(N^T,H\cdot M)&=\mathrm{E}(N^T)_\infty(H\cdot M)_\infty=\mathrm{E}N_T(H\cdot M)_\infty=\mathrm{E}N_T(H\cdot M)_T\\&=\mathrm{E}N_\infty(H\cdot M)_T=\mathrm{E}N_T(H\cdot M)^T_\infty=0,\end{aligned}$$

because we already noted that $(H\cdot M)^T$ is contained in $I^2(M)$. ∎

3.17 Lemma. *For every cadlag L_2-bounded martingale N we have $N\perp I^2(M)$ if and only if NM is a local martingale. In this case, if N is bounded and M is a martingale, then NM is also a martingale.*

Proof. Suppose that $N\perp I^2(M)$. Then $N^T\perp I^2(M)$ by Lemma 3.16, for every stopping time T. There exists a localizing sequence T_n such that M^{T_n} is a bounded martingale for every fixed n. Because $M^{T_n}=1_{[0,T_n]}\cdot M$, the process M^{T_n} is contained in $I^2(M)$, and hence also $M^{T_n\wedge T}$, by Lemma 3.16, for any stopping time T. We conclude that

$$\mathrm{E}N_T(M^{T_n})_T = \mathrm{E}(N^T)_\infty(M^{T_n\wedge T})_\infty = (N^T, M^{T_n\wedge T}) = 0.$$

Because $|N_T(M^{T_n})_T| \leq \sup_t |N_t| \sup_t |M_t^{T_n}|$ is contained in L_2, as N is L_2-bounded and M is bounded, these expectations indeed exist and the variables $N_T(M^{T_n})_T$ are integrable. Because this is true for every stopping time T, the process NM^{T_n} is a uniformly integrable martingale. Consequently, the process $(NM)^{T_n} = (NM^{T_n})^{T_n}$ is a martingale, whence the process NM is a local martingale.

Conversely, if NM is a local martingale, then $[N,M] = 0$, by the uniqueness of the Doob-Meyer decomposition. Consequently, for every predictable process H such that $H\cdot M$ is well defined $[N, H\cdot M] = H\cdot[N,M] = 0$, which implies that $N(H\cdot M)$ is a local martingale. Because $N(H\cdot M) \leq (N^2 + (H\cdot M)^2)/2$, it is dominated, and hence also a uniformly integrable martingale. By the martingale property $(N, H\cdot M) = \mathrm{E}N_\infty(H\cdot M)_\infty = \mathrm{E}N_0(H\cdot M)_0 = 0$.

This concludes the proof of the first assertion. For the second assertion it suffices to show that the stopped process $(NM)^n$ is a (uniformly integrable) martingale for every $n\in\mathbb{N}$. We have already shown that NM is a local martingale. If N is bounded by the constant C, then $|(NM)^n| \leq C|M^n|$. Because M^n is a uniformly integrable martingale, it is of class D, whence the process $(NM)^n$ is also of class D, and hence $(NM)^n$ is a martingale. (See the following exercise.) ∎

3.18 EXERCISE. A process M is said to be *of class D* if the collection of random variables $\{M_T\colon T \text{ finite stopping time}\}$ is uniformly integrable. Show that:

(i) If M is a local martingale of class D, then M is a uniformly integrable martingale.
(ii) A uniformly integrable martingale is of class D.
(iii) If M is of class D and $|N|\leq|M|$, then N is of class D.

3.19 Lemma. *If $\mathbb{P}$ is a unique martingale measure for M, then $I^2(M) = \mathcal{M}_0^2$.*

Proof. If $N\in\mathcal{M}_0^2$ is bounded in absolute value by $1/2$, then we can define a measure $\mathbb{Q}$ through $d\mathbb{Q} = (1+N_\infty)\,d\mathbb{P}$. By the martingale property $\mathrm{E}_\mathbb{P}N_\infty = \mathrm{E}_\mathbb{P}N_0 = 0$, and hence $\mathbb{Q}$ is a probability measure. The density process of $\mathbb{Q}$ relative to $\mathbb{P}$ is $1+N$, which is positive and equal to 1 at zero. Hence $\mathbb{Q}$ and $\mathbb{P}$ are equivalent and $\mathbb{Q}_0 = \mathbb{P}_0$. If $N\perp I^2(M)$, then NM is a $\mathbb{P}$-(local) martingale by Lemma 3.17, and hence so is the process $(1+N)M = M+NM$. We conclude that M is a $\mathbb{Q}$-(local) martingale, and hence $\mathbb{Q}$ is a (local) martingale measure. By the uniqueness of $\mathbb{P}$, it follows that $\mathbb{Q}=\mathbb{P}$, or, equivalently, $1+N=1$. Thus $N=0$.

A minor extension of this argument shows that there exists no nonzero bounded, cadlag martingale N with $N_0 = 0$ that is orthogonal to $I^2(M)$. It suffices to multiply N by a suitably small constant and apply the preceding argument.

For a given nontrivial martingale $N \in \mathcal{M}^2$ with uniformly bounded jumps and $T = \inf\{t \geq 0\colon |N_t| > 1\}$, the stopped process N^T is a bounded martingale, which is nontrivial, as $|N_T| \geq 1$. If N is orthogonal to $I^2(M)$, then so is N^T, by Lemma 3.16. Thus no nontrivial local martingale with uniformly bounded jumps is orthogonal to $I^2(M)$.

In particular, no continuous martingale is orthogonal to $I^2(M)$. The set $\mathcal{M}_0^{2,c}$ of all L_2-bounded, continuous martingales, 0 at 0, is a closed subspace of $\mathcal{M}_0^2$, and contains $I^2(M)$. It follows that $\mathcal{M}_0^{2,c} = I^2(M)$.

The elements of the orthocomplement of $\mathcal{M}_0^{2,c}$ in $\mathcal{M}_0^2$ are by definition the L_2-bounded, "purely discontinuous martingales", 0 at 0. Every such martingale N can be orthogonally decomposed as a series $\sum_n (N_n - A_n)$ of "compensated jumps", where each N_n can be taken of the form $N_n = \Delta N_{T_n} 1_{[T_n,\infty)}$ for T_n a stopping time, $|\Delta N_{T_n}|$ bounded, and the compensator A_n a continuous process. Consequently, if there exists a nontrivial purely discontinuous martingale in $\mathcal{M}_0^2$, then there also exists such a process with bounded jumps. As in the present case there exists no nontrivial martingale with bounded jumps orthogonal to $I^2(M)$, it follows that the orthocomplement of $\mathcal{M}_0^{2,c}$ in $\mathcal{M}_0^2$ is 0, and hence $I^2(M) = \mathcal{M}_0^2$. ∎

Proof of Theorem 3.2. Suppose first that $\mathbb{P}$ is a unique (local) martingale measure. Then, by the preceding lemma, every cadlag L_2-bounded martingale on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ is continuous. We first extend this to cadlag local martingales. By localization it suffices to show that every cadlag uniformly integrable martingale is continuous. Given a uniformly integrable cadlag martingale N, let N_∞^n be the random variable N_∞ truncated to the interval $[-n, n]$, and let N^n be the cadlag version of the process $\mathrm{E}(N_\infty^n \,|\, \mathcal{F}_t)$. Because this is a bounded martingale, it is contained in $\mathcal{M}^2$, and hence it is continuous, by Lemma 3.19. By dominated convergence $N_\infty^n \to N_\infty$ in L_1 and hence $\sup_t |N_t^n - N_t| \to 0$ in probability by the L_1-maximal inequality for submartingales. It follows that N is continuous as well.

For a given continuous, local martingale N there exists a localizing sequence $0 \leq T_n \uparrow \infty$ such that N^{T_n} is a bounded martingale, for every n. By the preceding lemma there exists $H_n \in L_2(M)$ such that $N^{T_n} = N_0 + H_n \cdot M$, for every n. Because N^{T_m} and N^{T_n} agree on $[0, T_m]$ for $m \leq n$, so do the stochastic integrals $H_m \cdot M$ and $H_n \cdot M$. This implies that $\mathrm{E}\int (H_n - H_m)^2 1_{[0,T_m]}\, d[M] = 0$, or equivalently that $H_n = H_m$ almost surely on $[0, T_m]$ under the Doléans measure μ_M. This shows that H defined by $H = H_n$ on $(T_{n-1}, T_n]$ coincides up to an μ_M-null set with H_n on $(0, T_n]$, whence $H \cdot M = H_n \cdot M = N - N_0$ on $[0, T_n]$.

Conversely, suppose that M has the representing property, and let $\mathbb{Q}$ be a martingale measure. The density process L of $\mathbb{Q}$ relative to $\mathbb{P}$ is a

$\mathbb{P}$-martingale and hence can be written as as a stochastic integral $L = L_0 + H \cdot M$. Because M is a $\mathbb{Q}$-(local) martingale, the process LM is a $\mathbb{P}$-(local) martingale by Lemma 3.17, and hence so is the process $(L - L_0)M = LM - L_0 M$. This implies that $0 = [L - L_0, M] = H \cdot [M]$ almost surely. Consequently, the process H is zero almost surely under the Doléans measure of M, and hence $L = L_0 + H \cdot M = L_0$. The variable L_0 is the density of $\mathbb{Q}_0$ relative to $\mathbb{P}_0$ and hence is 1 almost surely by the assumption that $\mathbb{Q}_0 = \mathbb{P}_0$. Thus $L = 1$ and $\mathbb{Q} = \mathbb{P}$ on $\mathcal{F}_\infty$. ■

* 3.5 Multivariate Stochastic Integrals

In this section we give the proper definition of the stochastic integral $H \cdot X$ for vector-valued predictable processes $H = (H^{(1)}, \ldots, H^{(d)})$ and a vector-valued semimartingale $X = (X^{(1)}, \ldots, X^{(d)})$. If each coordinate process $H^{(i)}$ is locally bounded, then the stochastic integral $H \cdot X$ is just the sum

$$H \cdot X := \sum_{i=1}^{d} H^{(i)} \cdot X^{(i)}$$

of the integrals of the coordinates. For the purpose of the representation theorem this definition must be extended to a larger class of integrands H. This involves both dropping the local boundedness and taking care of interactions between the coordinate integrals.

We present the extension first for continuous local martingales X and next for general semimartingales.

3.5.1 Continuous Local Martingales

In this section we define the stochastic integral $H \cdot M$ of a vector-valued predictable process H and a vector-valued local martingale M. We first define the integral for suitably integrable processes H through an L_2-isometry, and next extend by localization. Recall that the "extension" is null if M is a multivariate Brownian motion, so that for most purposes this section is not needed.

Suppose $M = (M^{(1)}, \ldots, M^{(d)})$ is a vector-valued, continuous martingale defined on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. For each coordinate $M^{(i)}$ define $L_2(M^{(i)})$ as the set of predictable processes $H^{(i)}$ such that $\mathrm{E} \int_0^\infty (H_s^{(i)})^2 \, d[M]_s < \infty$ (as in Section 3.4 except that presently we do not assume that M is L_2-bounded). If H is a vector-valued predictable process, such that $H^{(i)} \in L_2(M^{(i)})$ for every i, then every of the stochastic integrals $H^{(i)} \cdot M^{(i)}$ is a well-defined L_2-bounded martingale and hence we can define another L_2-bounded martingale by

$$H\cdot M := \sum_{i=1}^d H^{(i)}\cdot M^{(i)}.$$

This defines a map $H \mapsto H\cdot M$ from the product space $L_2(M^{(1)})\times\cdots\times L_2(M^{(d)})$ to the set $\mathcal{M}^2$ of L_2-bounded martingales. A main reason that we would like to extend the definition of the stochastic integral $H\cdot M$ to a wider class of predictable processes H is that the range of this map is not necessarily closed. Even though each of the classes of processes $H^{(i)}\cdot M^{(i)}$, when $H^{(i)}$ ranges over $L_2(M^{(i)})$, is closed in $\mathcal{M}^2$, being the image of the Hilbert space $L_2(M^{(i)})$ under an isometry, their sumspace is not necessarily closed. If M is to have the representing property, then we must add the missing elements.

If $H^{(i)}\in L_2(M^{(i)})$ for every i, then the process $(H\cdot M)^2 - [H\cdot M]$ is a uniformly integrable martingale, zero-at-zero, from which we can infer that $\mathrm{E}(H\cdot M)^2_\infty = \mathrm{E}[H\cdot M]^2_\infty$. By the bilinearity of the quadratic variation and the rule $d[H^{(i)}\cdot M^{(i)}, H^{(j)}\cdot M^{(j)}] = H^{(i)}H^{(j)} d[M^{(i)}, M^{(j)}]$, this can be written in the form

$$\mathrm{E}(H\cdot M)^2_\infty = \sum_i\sum_j \mathrm{E}\int_0^\infty H_s^{(i)}H_s^{(j)}\, d[M^{(i)}, M^{(j)}]_s.$$

All individual terms of the double sum on the right side are finite, and we can interchange the order of double sum and expectation and integration. The potential for extension of the stochastic integral $H\cdot M$ is that the double sum may possess a finite integral even if not all individual terms possess a finite integral. To operationalize this we view the multivariate quadratic variation process $t\mapsto [M]_t = ([M^{(i)}, M^{(j)}]_t)$ as a matrix of distribution functions of (random) signed measures, and write it as an integral

$$[M]_t = \int_0^t C_s\, d|[M]|_s$$

of a predictable process C with values in the set of nonnegative-definite, symmetric matrices relative to a univariate continuous, adapted increasing process $|[M]|$. (This is always possible. See e.g. Jacod and Shiryaev, II.2.9. We can choose $|[M]|$ equal to the sum of the absolute variations of all components in the matrix, as suggested by our notation. The components $[M^{(i)}, M^{(j)}]$ are clearly absolutely continuous and hence have densities $C^{(i,j)}$. That C can be chosen predictable requires proof.) Then the right side of the preceding display can be written as $\mathrm{E}\int_0^\infty H_s^T C_s H_s\, d|[M]|_s$.

We now define the space $L_2(M)$ as the set of all predictable processes H with values in $\mathbb{R}^d$ for which $\mathrm{E}\int_0^\infty H_s^T C_s H_s\, d|[M]|_s < \infty$. This is a linear space, which can be equipped with the inner product, for $G, H\in L_2(M)$,

$$(G,H) = \mathrm{E}\int_0^\infty G_s^T C_s H_s\, d|[M]|_s.$$

The space $L_2(M)$ is complete for the corresponding norm and hence is a semi-Hilbert space. The product space $L_2(M^{(1)}) \times \cdots \times L_2(M^{(d)})$ forms a dense subspace of $L_2(M)$, and the map $H \mapsto H \cdot M$ is an isometry of this product space into $\mathcal{M}^2$. We define the stochastic integral $H \cdot M$ for every $H \in L_2(M)$ by the unique continuous extension of the map $H \mapsto H \cdot M$ to $L_2(M)$.

The stochastic integral can be further extended to sufficiently integrable vector-valued predictable processes H by localization. If $0 \leq T_1 \leq T_2 \leq \cdots$ is a sequence of stopping times such that $T_n \uparrow \infty$ almost surely, and $H1_{[0,T_n]} \in L_2(M)$ for every n, then the integral $(H1_{[0,T_n]}) \cdot M$ is well defined for every n. We define $H \cdot M$ to be the almost limit of the sequence $(H1_{[0,T_n]}) \cdot M$ as $n \to \infty$. It can be shown in the usual way that for $m \leq n$ the processes $(H1_{[0,T_m]}) \cdot M$ and $(H1_{[0,T_n]}) \cdot M$ are identical up to evanescence on $[0, T_m]$ and hence that this limit exists. Furthermore, it can be shown that the definition is independent of the choice of the localizing sequence.

An appropriate sequence of stopping times exists provided, for every $t > 0$,

$$\int_0^t H_s^T C_s H_s \, d|[M]|_s < \infty, \qquad \text{a.s..}$$

In that case the stopping times $T_n = \inf\{t > 0\colon \int_0^t H_s^T C_s H_s \, d|[M]|_s > n\}$ form a localizing sequence. The integrability condition is certainly satisfied if $\int_0^t (H_s^{(i)})^2 \, d[M^{(i)}]_s < \infty$ for every i, in which case the stochastic integrals $H^{(i)} \cdot M^{(i)}$ are well defined and the extended integral $H \cdot M$ coincides with $\sum_{i=1}^d H^{(i)} \cdot M^{(i)}$. We collect this fact in the following lemmas, together with some other properties of the extended stochastic integral.

For simplicity of notation set, with C and $|[M]|$ as before,

$$\int_0^t H_s^2 \, d[M]_s = \int_0^t H_s^T C_s H_s \, d|[M]|_s.$$

3.20 Lemma. *Let M be a continuous, vector-valued local martingale and H a vector-valued, predictable process with $\int H_s^2 \, d[M]_s < \infty$ almost surely.*

(i) $H \cdot M = \sum_{i=1}^d H^{(i)} \cdot M^{(i)}$ *if all the integrals on the right side are well defined.*

(ii) $H \cdot M$ *is a continuous, local martingale.*

(iii) $[H \cdot M]_t = \int_0^t H_s^2 \, d[M]_s$.

(iv) $[N, H \cdot M] = H \cdot [N, M]$ *for every local L_2-martingale N.*

(v) $G \cdot (H \cdot M) = (GH) \cdot M$ *for every predictable process G and vector-valued predictable process H for which these integrals are well defined.*

(vi) $(H1_{[0,T]}) \cdot H = (H \cdot M)^T$ *for every stopping time T.*

(vii) $\Delta(H \cdot M) = H \Delta M$.

3.21 Lemma (Dominated convergence). *Let M be a continuous, vector-valued local martingale and $H^{(n)}$ a sequence of vector-valued predictable processes that converges pointwise on $[0,\infty)\times\Omega$ to 0. If $\int_0^t (H_s^{(n)})^T C_s H_s^{(n)}\,d|[M]|_s \to 0$ in probability for every t and $\|H^{(n)}\| \le H$ for a predictable process H such that $\int_0^t H_s^2 \|C_s\|\,d|[M]|_s < \infty$ almost surely for every t and n, then $\sup_{s\le t} |H^{(n)}\cdot M_s| \to 0$ in probability for every t.*

Proofs. See Jacod and Shiryaev, Section III.4.a. Note that Jacod and Shiryaev allow a general predictable process for $|[M]|$, where we use a continuous, adapted process. Their condition 4.3 that the process $(H\cdot C\cdot H)\cdot|[M]|$ is locally integrable is equivalent to our condition of finiteness of $\int_0^t H_s^T C_s H_s\,d|[M]|_s$, as a continuous, nonnegative increasing process is automatically locally integrable (and even locally bounded).

The right side of (iv) must be read as $\int_0^t \sum_i H_s^{(i)} D_s^{(i)}\,d|[M]|_s$ for $[Y, X^{(i)}] = D^{(i)}\cdot|[M]|$. (The sum may be integrable, even if the individual terms are not.)

That the processes $(H1_{[0,T_m]})\cdot M$ and $(H1_{[0,T_n]})\cdot M$ are identical up to evanescence on $[0,T_m]$ for stopping times $T_m \le T_n$ follows from the identities $(H1_{[0,T_m]})\cdot M = (H1_{[0,T_m]}1_{[0,T_n]})\cdot M = \bigl((H1_{[0,T_n]})\cdot M\bigr)^{T_m}$, where the last equality follows from (vi).

Jacod and Shiryaev state (v) for locally bounded predictable G, but it can be extended by approximation. Because $H\cdot M$ is a continuous local martingale with quadratic variation $\int_0^t H_s^T C_s H_s\,d|[M]|_s$, by (iii) the left side $G\cdot(H\cdot M)$ of (iv) is well defined if $\int_0^t G_s^2 H_s^T C_s H_s\,d|[M]|_s$ is finite almost surely for every t. The right side is defined under the condition that $\int_0^t (G_sH_s)^T C_s (G_sH_s)\,d|[M]|_s$ is finite almost surely for every t, which is equivalent. Equality (v) can be extended from locally bounded predictable processes G, H to the general case by approximation, using the second lemma.

For the proof of the second lemma let T_m be a localizing sequence such that $\mathrm{E}\int_0^\infty H_s^2 1_{[0,T_m]}\|C_s\|\,d|[M]|_s < \infty$ for every m. Because $\|H^{(n)}\| \le H$ it follows that $\mathrm{E}\int_0^\infty (H_s^{(n)}1_{[0,T_m]})^T C_s H_s^{(n)} 1_{[0,T_m]}\,d|[M]|_s < \infty$ for all m, whence $(H^{(n)}1_{[0,T_m]})\cdot M$ is a well-defined L_2-martingale. By Doob's inequality

$$\mathrm{E}\sup_{s\le t}\bigl|(H^{(n)}1_{[0,T_m]})\cdot M_s\bigr|^2 \le 4\mathrm{E}\bigl|(H^{(n)}1_{[0,T_m]})\cdot M_t\bigr|^2 = 4\mathrm{E}\int_0^{T_m} (H_s^{(n)})^T C_s H_s^{(n)}\,d|[M]|_s.$$

The integrand in the last double integral tends to zero pointwise, and is bounded by the integrable function $H^2\|C\|1_{[0,T_m]}$. By the dominated convergence theorem the double integral tends to zero. This shows that

$\sup_{s\le t}|(H^{(n)}\cdot M)_s^{T_m}=\sup_{s\le t}|(H^{(n)}1_{[0,T_m]}\cdot M)_s|$ tends to zero in probability, for every m. If $|Y^{(n)}_{t\wedge T_m}|=|(Y^{(n)})_t^{T_m}|$ tends to zero for every m, then $|Y_t^{(n)}|$ tends to zero. ■

3.22 Example (Brownian motion). The quadratic variation matrix $[B]$ of a multivariate Brownian motion is the diagonal matrix with the identity function on the diagonal. It follows that in the preceding discussion we can set C_t equal to the identity matrix and $|[M]|$ equal to the identity. The condition for existence of the integral $H\cdot B$ becomes finiteness of the integrals $\int_0^t H_s^T H_s\,ds$, which is the same as finiteness of the integrals corresponding to each of the components. In this case the "extension" of the multivariate integral discussed in this section is unnecessary and does not yield anything new. All stochastic integrals $H\cdot B$ are of the form $\sum_{i=1}^d H^{(i)}\cdot B^{(i)}$. □

3.23 Example. Let $M=\sigma\cdot B$ for B a e-dimensional Brownian motion and σ a $(d\times e)$-matrix-valued predictable process with $\int_0^t\|\sigma_s\|^2\,ds<\infty$. The integral $\sigma\cdot B$ is understood as $M^{(i)}=\sum_{j=1}^e\sigma^{(i,j)}\cdot B^{(j)}$, where the stochastic integrals on the right are well defined in the ordinary sense by the integrability condition on σ. Then $[M]_t=\int_0^t\sigma_s\sigma_s^T\,ds$ and hence we can take $C=\sigma\sigma^T$ and $|[M]|$ equal to the identity. The condition for existence of $H\cdot M$ reduces to finiteness of the process $\int_0^t\|\sigma_s^T H_s\|^2\,ds$. If the process σ is uniformly bounded away from infinity and singularity, then this reduces to finiteness of $\int_0^t\|H_s\|^2\,ds$.

The condition $\int_0^t\|\sigma_s^T H_s\|^2\,ds<\infty$ is the natural one if we think of $H\cdot M$ as $(\sigma^T H)\cdot B$ and apply Example 3.22. It may be weaker than the condition that $\int\|H_s\|^2\,d[M]_s<\infty$, as shown in the exercise below. □

3.24 EXERCISE. In Example 3.23 let $\sigma=\begin{pmatrix}1&0\\K&1-K\end{pmatrix}$ for a predictable process K.

(i) Show that we can take $C=\begin{pmatrix}1&K\\K&K^2+(1-K)^2\end{pmatrix}$ and $|[M]|$ equal to the identity.
(ii) Show that $H=(-K,1)/(1-K)$ is contained in $L_2(M)$ for any choice of K, but $H^{(1)}\notin L_2(M^{(1)})$ for some K.

3.5.2 Semimartingales

In this section we define the integral $H\cdot X$ for a vector-valued predictable process H and a cadlag semimartingale X. If X is a continuous local martingale, then the definition agrees with the definition in the preceding section.♭

♭ Proofs and further discussion of the results in this section can be found in the papers: C.S.

By definition there exists a decomposition $X = X_0 + M + A$ of X into a local martingale M and a process of locally bounded variation A. The predictable process $H = (H^{(1)}, \ldots, H^{(d)})$ is called *X-integrable* if there exists such a decomposition such that:

(i) $H \cdot M$ exists as a stochastic integral;

(ii) $H \cdot A$ exists as a Lebesgue-Stieltjes integral.

The integral $H \cdot X$ is then defined as $H \cdot X = H \cdot M + H \cdot A$. Here it is understood that we use cadlag versions so that the integral $H \cdot X$ is a cadlag semimartingale.

Warning. This definition can be shown to be well posed: the sum $H \cdot M + H \cdot A$ does indeed not depend on the decomposition $X = X_0 + M + A$. However, the decomposition itself is allowed to depend on the process H, and different H may indeed need different decompositions. It can be shown that a decomposition such that A contains all jumps of X with $|\Delta X| > 1$ or $|H\,\Delta X| > 1$ can be used without loss of generality. In particular, if X is continuous, then we can always use $X = M + A$ with M continuous.$^\sharp$

Warning. If X is a local martingale, then we cannot necessarily use the decomposition $X = X_0 + (X - X_0) + 0$. Given H there may be a decomposition $X = X_0 + M + A$ such that (i)-(ii) hold, giving an integral $H \cdot X$, which is the sum of a stochastic and a Lebesgue-Stieltjes integral, whereas $H \cdot X$ may not exist as a stochastic integral. Similarly, if X is of locally bounded variation, then we cannot necessarily use $X = X_0 + 0 + (X - X_0)$.

Warning. If X is a local martingale, then $H \cdot X$ is not necessarily a local martingale. This is because the Lebesgue-Stieltjes integral $H \cdot A$ for a local martingale of locally bounded variation A (as in (ii)) is not necessarily a local martingale. The Lebesgue-Stieltjes integral is defined pathwise, whereas the (local) martingale property requires some integrability. The stochastic integral $H \cdot M$ in (i) is always a local martingale. If X is a local martingale and $H \cdot X$ is bounded below by a constant, then $H \cdot X$ is a local martingale.

Warning. If $\mathcal{G}_t$ is a filtration with $\mathcal{F}_t \subset \mathcal{G}_t$, X is a semimartingale relative to $\mathcal{F}_t$ (and hence $\mathcal{G}_t$) and H is predictable relative to $\mathcal{F}_t$, then it may happen that $H \cdot X$ exists relative to $\mathcal{F}_t$, but not relative to the filtration $\mathcal{G}_t$. (Conversely, if it exists for $\mathcal{G}_t$, then also for $\mathcal{F}_t$.)

Chou, P.A. Meyer, C. Stricker: *Sur les intégrales stochastiques de processus prévisibles non bornés*, Lecture Notes in Mathematics 784, 1980, 128–139; and J. Jacod: *Intégrales stochastiques par rapport à une semimartingale vectorielle et changements de filtration*, Lecture Notes in Mathematics 784, 1980, 161–172.

$^\sharp$ If H is X-integrable, then both X and $H \cdot X$ are cadlag semimartingales, and hence each sample path of H or $H \cdot X$ has at most finitely many jumps of absolute value bigger than 1; the jumps of $H \cdot X$ are $H\,\Delta X$. It follows that the process of "big jumps"

$$\bar{X}_t = \Sigma_{s \le t} \Delta X_s 1_{|\Delta X_s| > 1 \text{ or } |H_s\, \Delta X_s| > 1}$$

is well defined. Then $X - \bar{X}$ is a semimartingale, that can be decomposed as $X - \bar{X} = X_0 + M + \tilde{A}$, so that $X = X_0 + M + A$ with $A = \tilde{A} + \bar{X}$.

The "existence" of the integrals in (i) and (ii) means the following:

(i) A stochastic integral $H \cdot M$ of a predictable process H relative to a local martingale M exists if and only if the process $t \mapsto \sqrt{\int_0^t H_s^2\, d[M]_s}$ is locally integrable.

(ii) A Lebesgue-Stieltjes integral $H \cdot A$ of a predictable process relative to a process of locally bounded variation exists if $\int_0^t |H_s|\, d|A|_s < \infty$ almost surely for every t.

If M is continuous or locally bounded, then its jump process is locally bounded, and hence the process $t \mapsto \int_0^t H_s^2\, d[M]_s$ is locally bounded and is certainly locally integrable to any order, for any H for which it is finite.

An alternative (equivalent) definition of the stochastic integral is through a limit of integrals of locally bounded processes. A predictable process H is said to be X-*integrable* if the sequence of processes $(H1_{|H|\leq n}) \cdot X$ converges as $n \to \infty$ to a limit Y, which is denoted $H \cdot X$, in the sense that, for every t,

$$\sup_{|G|\leq 1} \mathrm{E}\big|(G \cdot (H1_{|H|\leq n}) \cdot X)_t - (G \cdot Y)_t\big| \wedge 1 \to 0.$$

Here the supremum is taken over all predictable processes G with values in the interval $[-1, 1]$. Because we can choose $G = 1_{[0,t]}$, this implies that $\big((H1_{|H|\leq n}) \cdot X\big)_t \to (H \cdot X)_t$ in probability, for every t.

Thus the stochastic integral is a limit relative to the collection of semimetrics

$$d_t(X, Y) = \sup_{|G|\leq 1} \mathrm{E}\big|(G \cdot X)_t - (G \cdot Y)_t\big| \wedge 1, \qquad t > 0.$$

The topology generated by these metrics on the class of semimartingales is called the *semimartingale topology*. This topology is metrizable (restrict t to the natural numbers to reduce to a countable collection), and the class of semimartingales can be shown to form a complete metric space.

From the definition of $H \cdot X$ through the semimartingale topology it is clear that the stochastic integral is invariant under an equivalent change of measure: if $\mathbb{Q}$ is equivalent to $\mathbb{P}$ then the class of X-integrable processes H on $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{Q})$ and $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{P})$ are the same, and so is the stochastic integral $H \cdot X$.

3.25 Lemma. *Let G, H be vector-valued predictable processes and X and Y be cadlag semimartingales.*

(i) If H is both X- and Y-integrable, then H is $X + Y$-integrable and $H \cdot (X + Y) = H \cdot X + H \cdot Y$.

(ii) If both G and H are X-integrable, then so is $(G+H)$ and $(G+H) \cdot X = G \cdot X + H \cdot X$.

(iii) If H is X-integrable, then G is $H \cdot X$-integrable if and only if GH is X-integrable, and in that case $G \cdot (H \cdot X) = (GH) \cdot X$.

(iv) If H is X-integrable, then $\int_0^t |H_s|\,|d[X,Y]|_s$ is finite almost surely and $[H\cdot X, Y] = H\cdot [X,Y]$.
(v) If H is X-integrable, then $\Delta(H\cdot X) = H\,\Delta X$.

3.26 Lemma (Dominated convergence). *Let X be a semimartingale and H_n a sequence of predictable processes H_n that converges pointwise to a process H. If $|H_n| \leq K$ for an X-integrable process K, then $H_n\cdot X$ tends to $H\cdot X$ in the semimartingale topology.*

The notation in the preceding suggests one-dimensional processes H and X rather than vector-valued processes. The results can be interpreted in a vector-valued sense after making the following notational conventions. The vector-valued semimartingale X can be decomposed coordinatewise as $X = X_0 + M + A$ for a vector of local martingales M and a vector of processes of locally bounded varation A. The matrix-valued quadratic variation $[M]$ of M can be written as $[M] = C\cdot |[M]|$ for a process C with values in the symmetric, nonnegative-definite matrices and a nondecreasing, real-valued, adapted process $|[M]|$. The vector-valued process A can be written $A = D\cdot |A|$ for C a vector-valued optional process and $|A|$ a nondecreasing, adapted process. We then say that $H\cdot M$ and $H\cdot A$ exist if:
(i) The process $(H^TCH\cdot |[M]|)^{1/2}$ is locally integrable.
(ii) The Lebesgue-Stieltjes integral $|H^TD|\cdot |A|$ is finite.
The stochastic integral $H\cdot M$ is by definition the unique local martingale such that $[H\cdot M, N] = (H^TK)\cdot |[M]|$ for any local martingale N such that $[M^{(i)}, N] = K^{(i)}\cdot |[M]|$. The integral $H\cdot A$ is by definition the Lebesgue-Stieltjes integral $H^TD\cdot A$.

4
Finite Economies

In this chapter we consider an "economy" consisting of a vector $A = (A^{(1)}, \ldots, A^{(n)})$ of n "asset price processes". Throughout the chapter we assume that these processes are semimartingales defined on a given filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ that satisfies the usual conditions.

The σ-field $\mathcal{F}_t$ models our relevant knowledge at time t. In many applications this filtration is taken to be the augmented natural filtration of the process A, denoted by $\mathcal{F}_t^A$, meaning that we know the past evolution of all the asset prices (and no more). It is often assumed that the σ-field $\mathcal{F}_0$ is trivial (up to null sets).

For the general theory it is not necessary to make further assumptions on the asset price processes or the underlying filtered space. One typical more concrete specification would be that the assets satisfy a stochastic differential equation of the type, for given measurable functions μ and σ,

$$dA_t = \mu(t, A_t)\, dt + \sigma(t, A_t)\, dW_t,$$

where W is a multi-dimensional Brownian motion. Under reasonable conditions (e.g. the Itô conditions) the process A will then be adapted to the filtration generated by W. We can then take the filtration $\mathcal{F}_t$ equal to the natural filtration generated by A, the natural filtration generated by W, or possibly a still bigger filtration. The choice of filtration is only essential for Theorem 4.31, the other results in this chapter being at a more abstract level and being true for a general filtration. A single filtration $\mathcal{F}_t$ is fixed throughout the chapter.

Throughout the chapter we work with a finite time horizon, meaning that all processes need to be defined on a finite interval $[0, T]$ only. Properties of processes should be interpreted to refer to the time interval $[0, T]$ only. Alternatively, we may think of all processes being stopped at time T.

4.1 Strategies and Numeraires

At any point in time we invest in the assets A. The number of assets of every type as a process over time is a "trading strategy" or simply "strategy".

4.1 Definition.
(i) *A strategy is a predictable process ϕ with values in $\mathbb{R}^n$ such that the stochastic integral $\phi \cdot A$ is well defined.*
(ii) *A strategy ϕ is self-financing if $\phi A = \phi_0 A_0 + \phi \cdot A$.*

Products between vectors of the type ϕA are to be understood as inner products; for instance

$$\phi_t A_t = \sum_{i=1}^{n} \phi_t^{(i)} A_t^{(i)}.$$

The dot-notation $\cdot$ is reserved for the stochastic integral, where $\phi \cdot A$ is also to be understood as a linear combination $\sum_{i=1}^n \phi^{(i)} \cdot A^{(i)}$ of stochastic integrals if each of the stochastic integrals $\phi^{(i)} \cdot A^{(i)}$ is well defined.

The stochastic integral $\phi \cdot A$ of a locally bounded, predictable process ϕ is well defined relative to any semimartingale A. Thus locally bounded, predictable processes are always strategies. For some purposes it is necessary to allow a larger set of strategies, which may depend on A. In particular, for results on completeness involving the representing theorem for martingales, the stochastic integral $\phi \cdot A$ must be understood in the extended sense discussed in Section 3.5, which allows strategies that are not locally bounded. We say that a predictable process ϕ is *A-integrable* if the stochastic integral $\phi \cdot A$ is well defined.

Because the theory employs also other probability measures besides the "true world measure" $\mathbb{P}$, it is important to note that stochastic integrals (and hence the set of strategies) do not depend on the underlying measure, as long we use equivalent measures only. That this is true follows from the fact that the stochastic integral can always be written as a limit of integrals of simple integrands, which are Riemann sums and hence independent of any measure.

We interpret the strategy ϕ_t as the numbers of units of assets kept in an "investment portfolio" at time t. Thus the value of the portfolio at time t is given by

$$V_t = \phi_t A_t.$$

A strategy is not restricted to be nonnegative. Owning a negative amount of an asset is referred to as "taking a short position" in that asset. This is possible in real markets, up to some limitations. For instance, you can borrow money from the bank, as long as the bank is confident that you will be able to pay the interest and/or return the money eventually.

The predictability of a strategy can be interpreted as meaning that, for each t, the content of the portfolio at the time t is determined based on knowledge of the development of the asset prices before t only. (This intuitive interpretation should not be taken too seriously. For instance, not much would change if we would allow more general adapted processes in the case that the process A is continuous.) The self-financing property of the strategy can be more concisely written in differential notation as:

$$d(\phi_t A_t) = \phi_t dA_t.$$

The self-financing property ensures that the reshuffling of the contents of the portfolio over time is carried out without "money import". The relation in the preceding display requires that "a change in the value ϕA of the portfolio is solely due to changes dA_t in the values of the underlying assets". Thus we reconstitute the portfolio "just before time t" using the capital V_{t-} of the portfolio at that time. Next the value may change due to changes in value of the underlying assets. The resulting *gain process* $\phi \cdot A$ has "increments" $\phi_t \, dA_t$ and gives the cumulative increase or decrease of the portfolio value.

4.2 EXERCISE. For given stopping times $T_0 = 0 < T_1 < \cdots < T_k = T$ and $\mathcal{F}_{T_i}$-measurable random variables ψ_i consider the process $\phi = \psi_{-1} 1_{\{0\}} + \sum_i \psi_i 1_{(T_i, T_{i+1}]}$. Show that ϕ is a strategy, determine its value process, and show that ϕ is self-financing if and only if $\psi_{i-1} A_{T_i} = \psi_i A_{T_i}$ for every $i \geq 0$. Interpret this intuitively!

4.3 EXERCISE. Let ϕ and ψ be self-financing strategies and S a stopping time such that $\phi_S A_S = \psi_S A_S$. Show that the strategy $\phi 1_{[0,S)} + \psi 1_{[S,T]}$ is self-financing.

4.4 Definition. *A numeraire is a semimartingale of the form $N = \alpha_0 A_0 + \alpha \cdot A$ for some self-financing strategy α and such that both N and N_- are strictly positive. A numeraire is* special *if the strategy α can be chosen locally bounded.*

4.5 EXERCISE. Show that $N = A^{(1)}$ is a special numeraire provided that it is strictly positive.

Numeraires turn out to play an essential role in financial analysis. They will be used to write down the "fair" prices of options. Furthermore, "completeness" of an economy will be characterized by the existence of a numeraire of locally bounded variation.

For now we may just think of numeraires as special units to measure our wealth. Rather than in absolute units such as euros or guilders, we can express asset prices and our portfolio relative to the value of the numeraire. If the asset prices are A_t in absolute units, then they are $A_t^N := A_t/N_t$ if quoted relative to the numeraire N. The following lemma states the intuitively obvious fact that the self-financing property of a strategy is retained if the value process is quoted in different units.

4.6 Lemma (Unit invariance). *For any strategy ϕ we have $d(\phi_t A_t) = \phi_t\, dA_t$ if and only if $d(\phi_t A_t^N) = \phi_t\, dA_t^N$ for every numeraire N.*

Proof. By two applications of the partial integration formula,

$$(4.7)\qquad \begin{aligned} d\Big(\phi A \frac{1}{N}\Big) &= \frac{1}{N_-} d(\phi A) + (\phi A)_-\, d\Big(\frac{1}{N}\Big) + d\Big[\phi A, \frac{1}{N}\Big],\\ \phi\, d\Big(A \frac{1}{N}\Big) &= \phi\Big(\frac{1}{N_-}\, dA + A_-\, d\Big(\frac{1}{N}\Big) + d\Big[A, \frac{1}{N}\Big]\Big). \end{aligned}$$

The self-financing property $d(\phi A) = \phi\, dA$ implies that $d[\phi A, M] = \phi\, d[A, M]$ for every semimartingale M. Therefore, if ϕ is self-financing, then the difference of the right sides of the display is equal to

$$\big((\phi A)_- - \phi A_-\big)\, d\Big(\frac{1}{N}\Big) = -\big(\Delta(\phi A) - \phi \Delta A\big)\, d\Big(\frac{1}{N}\Big).$$

This is zero, as the self-financing property also implies that the processes $\phi \cdot A$ and ϕA have the same jumps.

The preceding argument uses that N is a strictly positive semimartingale, but no other property of a numeraire. We can change back to absolute units by repeating the argument with A^N and $1/N$ instead of A and N.

We complete the proof by showing that the preceding manipulations are indeed justified. A strategy ϕ is by definition an A-integrable predictable process. We shall show that a self-financing strategy is also automatically A^N-integrable. In view of the second line of (4.7) it suffices to show that ϕ is $(1/N_-) \cdot A$-integrable, $A_- \cdot (1/N)$ integrable and $[A, 1/N]$-integrable. The first follows because the process $1/N_-$ is locally bounded, and the third because A-integrability implies $[A, X]$-integrability for any semimartingale X. (See Lemma 3.25(iv).) If ϕ is self-financing, then $\phi A_- = (\phi A)_-$ by the preceding display. This shows that the process ϕA_- is left-continuous, whence locally bounded and $1/N$-integrable. Consequently ϕA_- is $A_- \cdot (1/N)$-integrable. (See Lemma 3.25(iii).) ∎

We can quote the value of our portfolio in arbitrary units. Using a numeraire, a special type of unit, has the great advantage that we need worry less about the self-financing property: for any strategy there exists a self-financing strategy with the same gain process if the gain is measured in the numeraire.

4.8 Lemma. *For any $\mathcal{F}_0$-measurable variable V_0, any numeraire N and any A^N-integrable predictable process ϕ there exists a self-financing strategy ψ such that $\phi_0 A_0^N = V_0$ and $\phi \cdot A^N = \psi \cdot A^N$.*

Proof. Suppose that $N = \alpha_0 A_0 + \alpha \cdot A = \alpha A$ for some self-financing strategy α. Then $1 = N/N = \alpha A^N$ and hence $0 = d(\alpha A^N) = \alpha\, dA^N$, by the self-financing property of α and unit invariance. Set

$$\psi = \phi + \kappa\alpha, \qquad \kappa = V_0 + \phi \cdot A^N - \phi A^N.$$

Because $\phi \cdot A^N - \phi A^N = (\phi \cdot A^N)_- - \phi(A^N)_-$, and ϕ is predictable, the process κ is predictable.

Because $(\kappa\alpha)\cdot A^N = \kappa\cdot(\alpha\cdot A^N) = \kappa\cdot 0 = 0$, it follows that $\psi\cdot A^N = \phi\cdot A^N$ and hence $\psi\, dA^N = \phi\, dA^N$. Furthermore $\psi A^N = \phi A^N + \kappa\alpha A^N = \phi A^N + \kappa = V_0 + \phi \cdot A^N$, by the definition of κ. In particular $(\psi A^N)_0 = V_0$ and $d(\psi A^N) = \phi\, dA^N$. Combining the preceding we see that $\psi\, dA^N = d(\psi A^N)$ and hence ψ is self-financing relative to the numeraire N. By unit invariance ψ is self-financing relative to any unit. ∎

4.9 Definition. *A numeraire pair* $(\mathbb{N}, N)$ *consists of a probability measure* $\mathbb{N}$ *on* $(\Omega, \mathcal{F})$ *that is equivalent to* $\mathbb{P}$ *and a numeraire* N *such that* A^N *is an* $\mathbb{N}$*-local martingale.*

Warning. Hunt and Kennedy require that A^N is an $\mathbb{N}$-martingale. We shall call the numeraire pair a *martingale numeraire pair* in that case.

The measure $\mathbb{N}$ in a numeraire pair is also referred to as a *(local) martingale measure*. It is a different type of martingale measure as considered in Chapter 3. Presently the processes A or its “rebased” version $A^N = A/N$ are typically not martingales under the initial measure $\mathbb{P}$. The change to the martingale measure $\mathbb{N}$ ensures that the process A^N is a local martingale.

Not every economy admits a numeraire pair. In Lemma 4.16 we shall see that existence of a numeraire pair precludes the possibility of riskless gains (arbitrage). Conversely, an appropriate form of absence of no arbitrage implies existence of a numeraire pair. (This is called the “fundamental theorem of asset pricing”.)

In this chapter existence of a numeraire pair is always assumed. Then there are automatically many numeraire pairs. We shall see that uniqueness of the martingale measure $\mathbb{N}$ going with a given numeraire N is equivalent to “completeness of the economy”.

4.10 Example (Black-Scholes). The classical Black-Scholes economy consists of two assets, which for simplicity of notation we shall write as $A = (R_t, S_t)$. The process S_t corresponds to a risky asset, such as a stock, and is assumed to satisfy the differential equation, for a given Brownian motion W,

$$dS_t = \mu S_t\, dt + \sigma S_t\, dW_t.$$

The parameter σ is assumed to be positive. The asset R_t is risk-free and satisfies the equations

$$dR_t = rR_t\, dt, \qquad R_0 = 1.$$

The last equation can be solved to give $R_t = e^{rt}$. Similarly, the stochastic differential equation can be solved explicitly to give

$$S_t = S_0 e^{(\mu - \frac{1}{2}\sigma^2)t + \sigma W_t}.$$

We shall show that the process $N_t = e^{rt}$ is a numeraire, and shall find a corresponding martingale measure.

That N is positive is clear. Furthermore, it is equal to one of the two processes in the economy and hence is the value process of a self-financing strategy. (Cf. Exercise 4.5). Thus N is a numeraire.

By the partial integration formula

$$d\big(e^{-rt}S_t\big) = e^{-rt}\,dS_t + S_t(-re^{-rt})\,dt + 0 = e^{-rt}\sigma S_t\,d\tilde W_t, \tag{4.11}$$

for the process $\tilde W$ defined by $\tilde W_t = W_t - (r-\mu)t/\sigma$. Thus to ensure that the process $A^N = (1, S^N)$ is an $\mathbb{N}$-local martingale, it suffices to determine the measure $\mathbb{N}$ such that the process $\tilde W$ is an $\mathbb{N}$-local martingale.

By Novikov's condition, or direct calculation, the exponential process $\mathcal{E}(\theta\cdot W)$ for $\theta_t = (r-\mu)/\sigma$ is a $\mathbb{P}$-martingale, with mean 1. Therefore, we can define a probability measure $\mathbb{N}$ by $d\mathbb{N} = \mathcal{E}(\theta\cdot W)_T\,d\mathbb{P}$. By Girsanov's theorem the process $\tilde W$ is then an $\mathbb{N}$-Brownian motion (for the time parameter restricted to $[0,T]$), and hence an $\mathbb{N}$-martingale.

Thus we have shown that the process $A^N = (1, S^N)$ is an $\mathbb{N}$-local martingale. From the explicit expression $S^N_t = S_0\exp(\sigma\tilde W_t - \frac12\sigma^2 t)$ it follows that it is actually also an $\mathbb{N}$-martingale.

In the preceding we have not made the filtration that we work with explicit. An "initial" filtration $\mathcal{G}_t$ is implicit in the assumption that W is a Brownian motion. The present numeraire N is the value process of a deterministic strategy, which is predictable relative to any filtration. In view of the invertible relationship between A and W the augmented natural filtration $\mathcal{F}^A_t$ generated by A is equal to the augmented filtration $\mathcal{F}^W_t$ of the driving Brownian motion W, which may be smaller than $\mathcal{G}_t$. The process A^N is an $\mathbb{N}$-martingale relative to the bigger filtration $\mathcal{G}_t$, where by definition of a numeraire pair it suffices that it is a martingale relative to the filtration $\mathcal{F}_t$ we work with.

Thus the preceding shows the existence of a numeraire pair relative to any filtration $\mathcal{F}_t$ that is sandwiched between $\mathcal{F}^A_t$ and the initial filtration $\mathcal{G}_t$. The standard choice for $\mathcal{F}_t$ is $\mathcal{F}^A_t$, for which the Black-Scholes economy is "complete", as we shall see. □

4.12 EXERCISE. Extend the preceding example to the situation that the stock price process S satisfies a stochastic differential equation of the form $dS_t = \mu(t, S_t)S_t\,dt + \sigma(t, S_t)S_t\,dW_t$. Find regularity conditions (e.g. boundedness) on the functions μ and σ that ensure the existence of a numeraire pair.

* **4.13 Example (Lévy processes)**. A *Lévy process* X is a cadlag stochastic process with stationary and independent increments with initial value $X_0 = 0$. We interpret the independence of the increments relative to a general filtration $\mathcal{F}_t$. Thus for each $s < t$ the increment $X_t - X_s$ is independent

of $\mathcal{F}_s$ and possesses the same distribution as X_{t-s}. Consider an economy consisting of the asset processes $R_t = e^{rt}$ and $S_t = e^{\mu t+\sigma X_t}$ for some constants r, μ, σ.

Because Brownian motion is a Lévy process, the Black-Scholes economy is a special example. Brownian motion is the only Lévy process with continuous sample paths. More general Lévy processes have been introduced to introduce jumps in the asset processes. We shall exhibit a martingale measure $\mathbb{R}$ with the numeraire R under the condition that there exists a solution v to the equation $\psi(v+\sigma)/\psi(v) = e^{(r-\mu)v}$, for $\psi(u) = \mathrm{E}e^{uX_1}$ the Laplace transform of X_1 (assumed to be finite in an interval).†

The key observation is that the process $t \mapsto e^{vX_t}/\psi(v)^t$ is a $\mathbb{P}$-martingale for any v such that $\psi(v) < \infty$. Indeed, the stationarity and independence of the increments implies that $\mathrm{E}e^{uX_t} = \psi(u)^t$ and hence, for $s < t$,

$$\mathrm{E}\Big(\frac{e^{vX_t}}{\psi(v)^t}\,|\,\mathcal{F}_s\Big) = \mathrm{E}\Big(\frac{e^{v(X_t-X_s)}}{\psi(v)^{t-s}}\,|\,\mathcal{F}_s\Big)\frac{e^{vX_s}}{\psi(v)^s} = \frac{e^{vX_s}}{\psi(v)^s},$$

where we use the independence of $X_t - X_s$ from $\mathcal{F}_s$, and the fact that $X_t - X_s$ is distributed as X_{t-s}.

For any v the $\mathbb{P}$-martingale $L_t := e^{vX_t}/\psi(v)^t$ is positive with mean 1 and hence defines a density process of a measure $\mathbb{N}$ relative to $\mathbb{P}$. By Lemma 2.14 the discounted process S/R is an $\mathbb{N}$-martingale if LS/R is a $\mathbb{P}$-martingale. Because

$$L_t\frac{S_t}{R_t} = e^{(v+\sigma)X_t}\frac{e^{(\mu-r)t}}{\psi(v)^t},$$

this process is a $\mathbb{P}$-martingale if $e^{(\mu-r)t}\psi(v)^{-t} = \psi(v+\sigma)^{-t}$.

We shall see later that for most Lévy processes there exist infinitely many other martingale measures with the numeraire R. ◻

4.14 EXERCISE. In the preceding example show that the process X is also a Lévy process under the measure $\mathbb{N}$.

4.2 Arbitrage and Pricing

It was seen in Example 1.7 that without some further restriction on the set of strategies it will be possible to make certain profits. Thus we allow only "admissible" strategies.

† The condition of existence of a solution v to the equation $\psi(v+\sigma)/\psi(v) = e^{(r-\mu)v}$ is not automatic.

4.15 Definition. *A strategy is admissible if for any numeraire pair* $(\mathbb{N}, N)$ *the process* $\phi \cdot A^N$ *is an* $\mathbb{N}$*-martingale.*

If there exists no numeraire pair, then the preceding definition does not make sense. We are not interested in this situation, but for consistency we define the collection of admissible strategies to be empty if no numeraire pair exists.

By the definition of a numeraire pair the process A^N is an $\mathbb{N}$-local martingale, and hence the process $\phi \cdot A^N$ is typically an $\mathbb{N}$-local martingale for any self-financing strategy ϕ.‡ The special feature of an admissible strategy is that it is an $\mathbb{N}$-martingale. Thus admissibility adds integrability properties to the gain processes $\phi \cdot A^N$. It prevents them from becoming too extreme.

Warning. Definition 4.15 follows Hunt and Kennedy. Most authors define a strategy to be admissible if its value process is lower bounded by 0. Lemma 4.16 remains valid under the latter definition of admissibility. In fact, it is valid as soon as the process $\phi \cdot A^N$ is a supermartingale. Any local martingale that is lower bounded by a martingale, in particular a nonnegative local martingale, is a supermartingale. One might think of the preceding definition as giving a two-sided sense of admissibility, giving control for both buyer and seller.

The following lemmas show that admissible strategies never yield arbitrage. Furthermore, if they lead to the same value at time T, then they have identical value processes throughout $[0, T]$.

4.16 Lemma (No arbitrage). *If* ϕ *is an admissible, self-financing strategy, then it cannot happen that:*
(i) $\phi_0 A_0 < 0$, *but* $\phi_T A_T \geq 0$ *almost surely, or:*
(ii) $\phi_0 A_0 = 0$, *but* $\phi_T A_T \geq 0$ *almost surely and* $\mathbb{P}(\phi_T A_T > 0) > 0$.

4.17 Lemma (Unique value). *If* ϕ *and* ψ *are admissible, self-financing strategies with* $\phi_T A_T = \psi_T A_T$, *then* $\phi_t A_t = \psi_t A_t$ *for every* $t \in [0, T]$.

Proofs. By convention the existence of an admissible strategy implies the existence of a numeraire pair $(\mathbb{N}, N)$. The self-financing property of a strategy ϕ is, by unit invariance, equivalent to the identity $\phi A^N = \phi_0 A_0^N + \phi \cdot A^N$. If ϕ is admissible, then this is an $\mathbb{N}$-martingale and taking expectations left and right under the martingale measure yields

$$\phi_0 \frac{A_0}{N_0} = \mathrm{E}_{\mathbb{N}}(\phi_0 A_0^N) = \mathrm{E}_{\mathbb{N}}(\phi_T A_T^N) = \mathrm{E}_{\mathbb{N}}\Big(\phi_T \frac{A_T}{N_T}\Big).$$

If the strategy ϕ would satisfy (i), then the left and right sides of this identity would be negative and nonnegative, respectively, which is impossible.

‡ The local martingale property is automatic if A^N is continuous, but may fail if A^N or $\phi \cdot A^N$ possess too big jumps. See Section 3.5.2.

Similarly, if ϕ would be as in (ii), then the two sides would be zero and strictly positive, respectively. This concludes the proof of the first lemma.

Under the conditions of the second lemma the process $(\phi-\psi)\cdot A^N$ is an $\mathbb{N}$-martingale. The self-financing property yields, in view of unit invariance, that $(\phi-\psi)A^N = (\phi_0-\psi_0)A_0^N + (\phi-\psi)\cdot A^N$, and hence the process $(\phi-\psi)A^N$ is an $\mathbb{N}$-martingale. By assumption it vanishes at time T, whence it is identically zero throughout $[0,T]$. ∎

The first lemma is interpreted as saying that "the economy is arbitrage-free": no admissible, self-financing strategy leads to sure profit. With the present definitions this is true without conditions. This is somewhat at odds with the literature, but is due to the convention to define the set of admissible strategies to be empty if there exists no numeraire pair.

The second lemma shows that the value of an admissible, self-financing portfolio during the interval $[0,T]$ is uniquely determined by its terminal value. This property is the justification for the no-arbitrage pricing principle. Consider an $\mathcal{F}_T$-measurable random variable X, interpreted as the value of a derivative contract at expiration time T, for which there exists a *replicating strategy*: an admissible, self-financing strategy ϕ such that

$$X = \phi_T A_T.$$

Then the *no-arbitrage principle* leads us to define the "just price" for the claim at time t to be the value $\phi_t A_t$ of the replicating portfolio at time t. The preceding lemma shows that this definition is independent of the replicating portfolio as long as this is required to be self-financing and admissible.

We can express this value using a numeraire pair $(\mathbb{N}, N)$ with $\mathrm{E}_{\mathbb{N}}|X/N_T| < \infty$. Let ϕ be a replicating strategy, so that $X = \phi_T A_T$. By unit invariance and the self-financing property of ϕ,

$$\phi_t A_t^N = \phi_0 A_0^N + (\phi \cdot A^N)_t.$$

Because ϕ is admissible, the right side, and hence the left side is an $\mathbb{N}$-martingale. Thus its value at $t \in [0,T]$ is determined by its final value at T: by the martingale property

$$\phi_t A_t = N_t \phi_t A_t^N = N_t \mathrm{E}_{\mathbb{N}}\Big(\phi_T A_T^N \,|\, \mathcal{F}_t\Big).$$

The left side is the value V_t of the claim at time t, and the variable inside the conditional expectation on the right is equal to $\phi_T A_T^N = X/N_T$. Thus we obtain the pricing formula for the value V_t at t of a derivative with claim X at time T:

$$V_t = N_t \mathrm{E}_{\mathbb{N}}\Big(\frac{X}{N_T} \,|\, \mathcal{F}_t\Big). \tag{4.18}$$

For $t = 0$ the conditional expectation reduces to an ordinary expectation, in the case that $\mathcal{F}_0$ is trivial.[♭]

We have derived the pricing formula using a numeraire pair $(\mathbb{N}, N)$. At first sight the value given by the formula appears to depend on the numeraire pair, but this is not true, because the formula gives merely a representation of the value process of a replicating strategy, which is unique by Lemma 4.17. Thus we may use any numeraire pair such that $\mathrm{E}_{\mathbb{N}}|X/N_T| < \infty$.

On the other hand, beware of thinking of (4.18) as "the price" of the claim X. Even though this is not visible in the formula, the formula is conditional on the existence of a replicating strategy. There are examples of economies with multiple numeraire pairs for which formula (4.18) evaluates to multiple values. Of course, in these examples there cannot exist a replicating strategy for the claim X, in which case we have not said what we mean by "fair price" in the first place. We discuss this further in Section 4.4.

4.3 Completeness

The economy is said to be complete if there exists a numeraire pair that allows to price all contingent claims in this way.

4.19 Definition. *The economy is complete if there exists a numeraire pair* $(\mathbb{N}, N)$ *such that for every* $\mathcal{F}_T$*-measurable random variable* X *with* $\mathrm{E}_{\mathbb{N}}|X/N_T| < \infty$ *there exists an admissible, self-financing strategy* ϕ *with* $X = \phi_T A_T$.

This definition suggests that completeness requires a special numeraire pair, but this is misleading. Given completeness the set of claims for which $\mathrm{E}_{\mathbb{N}}|X/N_T| < \infty$ is independent of the numeraire pair $(\mathbb{N}, N)$. Hence the choice of numeraire pair is irrelevant, and we may choose a convenient one. In fact, under completeness the measure $F \mapsto \mathrm{E}_{\mathbb{N}}(1_F/N_T)$ is the same on the "final" σ-field $\mathcal{F}_T$ for every numeraire pair $(\mathbb{N}, N)$ with given initial measure $\mathbb{N}_0$ (the restriction of $\mathbb{N}$ to $\mathcal{F}_0$) and initial numeraire value N_0. Essentially, this follows because $\mathrm{E}_{\mathbb{N}}(1_F/N_T)$ is the fair price at time 0 of the claim 1_F, which is given by the value at 0 of a replicating portfolio and hence is independent of the numeraire pair. A more formal argument is given in the following lemma.

4.20 Lemma. *If the economy is complete and* $(\mathbb{M}, M)$ *and* $(\mathbb{N}, N)$ *are arbitrary numeraire pairs with* $\mathbb{M}_0 = \mathbb{N}_0$ *and* $M_0 = N_0$*, then the process* M/N

[♭] This argument does not seem to need that N is a numeraire. It suffices that it is a unit and that the processes $\phi \cdot A^N$ are $\mathbb{N}$-martingales.

is the density process of $\mathbb{M}$ *relative to* $\mathbb{N}$. *In particular,* $M_T^{-1}\, d\mathbb{M} = N_T^{-1}\, d\mathbb{N}$ *on* $\mathcal{F}_T$.

Proof. Suppose that the economy is complete and let $(\mathbb{N}, N)$ be a numeraire pair as in the definition of completeness. For a given event $F \in \mathcal{F}_T$ the claim $X = N_T 1_F$ is $\mathcal{F}_T$-measurable and satisfies $\mathrm{E}_{\mathbb{N}}|X/N_T| < \infty$. Therefore, by completeness there exists an admissible, self-financing strategy ϕ such that $N_T 1_F = \phi_T A_T$. The self-financing property and unit invariance show that the process Z defined by $Z_t = \phi_t A_t^N$ can also be represented as $Z = \phi_0 A_0^N + \phi \cdot A^N$ and hence is an $\mathbb{N}$-martingale, by the assumed admissibility of ϕ. Clearly $Z_T = 1_F$.

If $(\mathbb{M}, M)$ is another numeraire pair with $M_0 = N_0$, then, again by self-financing and unit invariance, the process Y defined by $Y_t = \phi_t A_t^M$ can be written as $Y = \phi_0 A_0^M + \phi \cdot A^M$ and hence is an $\mathbb{M}$-martingale, by the admissibility of ϕ. Clearly $Y = (N/M)Z$ and $Y_0 = Z_0$.

Being martingales, the processes Z and Y have constant means, and hence

$$\begin{aligned}
\mathrm{E}_{\mathbb{N}} 1_F &= \mathrm{E}_{\mathbb{N}} Z_T = \mathrm{E}_{\mathbb{N}} Z_0 = \mathrm{E}_{\mathbb{M}} Z_0 \\
&= \mathrm{E}_{\mathbb{M}} Y_0 = \mathrm{E}_{\mathbb{M}} Y_T = \mathrm{E}_{\mathbb{M}} Z_T \frac{N_T}{M_T} = \mathrm{E}_{\mathbb{M}} 1_F \frac{N_T}{M_T},
\end{aligned}$$

where we use that $\mathbb{M}_0 = \mathbb{N}_0$ in the last equality of the first line. This being true for an arbitrary event $F \in \mathcal{F}_T$ shows that $\mathbb{N}$ possesses density N_T/M_T relative to $\mathbb{M}$ on $\mathcal{F}_T$, or, equivalently, that $\mathbb{M}$ possesses density M_T/N_T relative to $\mathbb{N}$, or that $M_T^{-1}\, d\mathbb{M} = N_T^{-1}\, d\mathbb{N}$. This is true for an arbitrary numeraire pair $(\mathbb{M}, M)$ and the special numeraire pair $(\mathbb{N}, N)$. It follows that the measure $M_T^{-1}\, d\mathbb{M}$ is the same for every numeraire pair $(\mathbb{M}, M)$.

Because $d\mathbb{M}/d\mathbb{N} = M_T/N_T$, the density process of $\mathbb{M}$ relative to $\mathbb{N}$ is the process $\mathrm{E}_{\mathbb{N}}(M_T/N_T|\, \mathcal{F}_t)$. This process coincides with the process M/N if (and only if) the process M/N is an $\mathbb{N}$-martingale.

Because M is a numeraire, there exists a self-financing strategy α with $M = \alpha A$. By unit invariance and self-financing $M/N = \alpha A^N = \alpha_0 A_0^N + \alpha \cdot A^N$. Because $(\mathbb{N}, N)$ is a numeraire pair, the process A^N is an $\mathbb{N}$-local martingale and hence $\alpha \cdot A^N$ and M/N are $\mathbb{N}$-local martingales. (In the case that A^N has jumps and α is not locally bounded, we use that M/N is nonnegative to ensure the local martingale property.) A positive, local martingale is a supermartingale, by Fatou's lemma, and (hence) is a martingale if its mean is constant. The mean of M/N satisfies $\mathrm{E}_{\mathbb{N}}(M_T/N_T) = \mathrm{E}_{\mathbb{M}} 1 = 1 = \mathrm{E}_{\mathbb{N}}(M_0/N_0)$, by assumption. Thus M/N is an $\mathbb{N}$-martingale. ∎

4.21 EXERCISE. Suppose that we drop the conditions $\mathbb{M}_0 = \mathbb{N}_0$ and $M_0 = N_0$. Show that $M/N(N_0/M_0)(d\mathbb{M}_0/d\mathbb{N}_0)$ is the density process of $\mathbb{M}$ relative to $\mathbb{N}$.

The last assertion of the lemma implies that in a complete economy $\mathrm{E}_{\mathbb{M}}(X/M_T) = \mathrm{E}_{\mathbb{N}}(X/N_T)$ for every contingent claim X and any pair of numeraire pairs. Thus for pricing contingent claims the choice of numeraire pair in formula (4.18) is irrelevant, although one choice may lead to easier computations than another.

Because a density process is a martingale, the first assertion of the lemma shows that in a complete economy the quotient M/N of two numeraires that are part of numeraire pairs $(\mathbb{M}, M)$ and $(\mathbb{N}, N)$ is an $\mathbb{N}$-martingale. Conversely, if M is a numeraire and M/N is an $\mathbb{N}$-martingale, then there exists a martingale measure $\mathbb{M}$ corresponding to M, and we can construct it as in the preceding lemma.

4.22 Lemma. *If $(\mathbb{N}, N)$ is a numeraire pair and M a numeraire with $M_0 = N_0$ such that M/N is an $\mathbb{N}$-martingale, then $(\mathbb{M}, M)$ is a numeraire pair for $\mathbb{M}$ the probability measure on $\mathcal{F}_T$ with $d\mathbb{M} = M_T/N_T\, d\mathbb{N}$.*

Proof. By the martingale property $\mathrm{E}_{\mathbb{N}}(M_T/N_T) = \mathrm{E}_{\mathbb{N}}(M_0/N_0) = 1$, whence $\mathbb{M}$ is a probability measure. Because $(\mathbb{N}, N)$ is a numeraire pair, the process $A^N = (M/N)A^M$ is an $\mathbb{N}$-local martingale. By the definition of $\mathbb{M}$ and the assumption that M/N is an $\mathbb{N}$-martingale, the process M/N is the density process of $\mathbb{M}$ relative to $\mathbb{N}$. Therefore, A^M is an $\mathbb{M}$-local martingale, by Lemma 2.14. ∎

4.23 Example (Positive asset as numeraire). If it is strictly positive, then a fixed component $M = A^{(i)}$ of the asset process A is a numeraire. If there exists a numeraire pair $(\mathbb{N}, N)$ with $\mathbb{N}$ a martingale measure rather than a local martingale measure for A^N, then A^N and hence its component $M/N = (A^{(i)})^N$ is an $\mathbb{N}$-martingale. Lemma 4.22 then guarantees the existence of a measure $\mathbb{M}$ such that $(\mathbb{M}, M)$ is a numeraire pair. Thus if there exists a numeraire pair with martingale measure, then we can always use a numeraire pair with as numeraire a strictly positive component of the asset price process, if there is one. □

** **4.24** EXERCISE. In an incomplete economy the quotient M/N of two numeraires need not be an $\mathbb{N}$-martingale even if both numeraires are part of numeraire pairs $(\mathbb{M}, M)$ and $(\mathbb{N}, N)$??

The definition of completeness requires that every claim can be replicated by an admissible strategy. The definition of an admissible strategy ϕ requires that the process $\phi \cdot A^N$ be an $\mathbb{N}$-martingale for *every* numeraire pair $(\mathbb{N}, N)$. This is rather inconvenient, as there are many numeraire pairs. Actually, the requirement of admissibiliy can be relaxed: it suffices that $\phi \cdot A^N$ is an $\mathbb{N}$-martingale for just a single numeraire pair $(\mathbb{N}, N)$. An economy is complete as soon as every sufficiently integrable claim can be replicated by a self-financing strategy ϕ such that $\phi \cdot A^N$ is an $\mathbb{N}$-martingale for some given numeraire pair $(\mathbb{N}, N)$. Also, in a complete economy any strategy ϕ for which there exists a numeraire pair $(\mathbb{N}, N)$ such that $\phi \cdot A^N$ is an $\mathbb{N}$-martingale is automatically admissible.

4.25 Lemma. *The economy is complete if (and only if) there exists a numeraire pair* $(\mathbb{N}, N)$ *such that for every nonnegative* $\mathcal{F}_T$*-measurable random variable* X *with* $\mathrm{E}_{\mathbb{N}}(X/N_T) < \infty$ *there exists a self-financing strategy* ϕ *with* $X = \phi_T A_T$ *and such that* $\phi \cdot A^N$ *is an* $\mathbb{N}$*-martingale.*

4.26 Lemma. *If the economy is complete and* ϕ *is a self-financing strategy such that* $\phi \cdot A^N$ *is an* $\mathbb{N}$*-martingale for some numeraire pair* $(\mathbb{N}, N)$*, then* $\phi \cdot A^M$ *is an* $\mathbb{M}$*-martingale for every numeraire pair* $(\mathbb{M}, M)$*.*

Proofs. Suppose that we can prove that for any other numeraire pair $(\mathbb{M}, M)$ the process M/N is a multiple of the density process of $\mathbb{M}$ relative to $\mathbb{N}$. Then by Lemma 2.14 for any self-financing strategy ϕ the process ϕA^M is an $\mathbb{M}$-martingale if and only if the process $(M/N)\phi A^M = \phi A^N$ is an $\mathbb{N}$-martingale. Equivalently by unit invariance, the process $\phi \cdot A^M$ is an $\mathbb{M}$-martingale if and only if the process $\phi \cdot A^N$ is an $\mathbb{N}$-martingale. The assertion of the second lemma then follows. Furthermore, it is then clear that under the conditions of the first lemma any nonnegative claim can be replicated by an admissible strategy. The completeness of the economy follows then by splitting a general claim X with $\mathrm{E}_{\mathbb{N}}(|X|/N_T) < \infty$ into its positive and negative parts X^+ and X^-.

The proof that M/N is a multiple of the density process of $\mathbb{M}$ with respect to $\mathbb{N}$ is in part a repetition of the proof of Lemma 4.20, with the difference that we cannot assume that $\phi \cdot A^M$ is automatically an $\mathbb{M}$-martingale. This difficulty is overcome by a stopping argument. For simplicity of notation assume that $M_0 = N_0$, which can always be arranged by scaling.

Fix an event $F \in \mathcal{F}_T$ and define a stopping time by $T_n = \inf\{t > 0\colon N_t/M_t > n\}$. Because the claim $X = N_T 1_F 1_{T_n > T}$ is $\mathcal{F}_T$-measurable and satisfies $0 \le X \le N_T$, under the conditions of the lemmas there exists a self-financing strategy ϕ such that $\phi \cdot A^N$ is an $\mathbb{N}$-martingale and $N_T 1_F 1_{T_n > T} = \phi_T A_T$. The self-financing property and unit invariance show that the process Z defined by $Z_t = \phi_t A_t^N$ can also be represented as $Z = \phi_0 A_0^N + \phi \cdot A^N$ and hence is an $\mathbb{N}$-martingale, by assumption. Clearly $Z_T = 1_F 1_{T_n > T}$.

If $(\mathbb{M}, M)$ is another numeraire pair, then, again by self-financing and unit invariance, the process $Y = (N/M)Z = \phi A^M$ can be written as $Y = \phi_0 A_0^M + \phi \cdot A^M$ and hence the stopped process Y^{T_n} can be represented as a stochastic integral as well. Because $Y = (N/M)Z$ and $Z_t = \mathrm{E}_{\mathbb{N}}(Z_T|\,\mathcal{F}_t)$ is bounded in absolute value by 1, this process Y^{T_n} is bounded by n on the interval $(0, T_n)$, by the definition of T_n. Furthermore, by optional stopping $Z_{T \wedge T_n} = \mathrm{E}_{\mathbb{N}}(Z_T|\,\mathcal{F}_{T_n}) = 1_{T_n > T}\mathrm{E}_{\mathbb{N}}(1_F|\,\mathcal{F}_{T_n})$, because T_n is $\mathcal{F}_{T_n}$-measurable, which implies that $Z_{T \wedge T_n} = Z_{T_n} = 0$ if $T_n \le T$, whence $Y_{T_n} = 0$ on this event. We conclude that that Y^{T_n} is bounded by n on $(0, T]$ and hence is an $\mathbb{M}$-martingale, for every fixed n.♯

♯ Hunt and Kennedy claim that the process Y is bounded by n on $[0, T]$. If that is true the stopping is not necessary and the proof can be simplified.

Let $V_0 = d\mathbb{N}_0/d\mathbb{M}_0$. Since V_0 is $\mathcal{F}_0$-measurable, the process $V_0 Y^{T_n}$ is an $\mathbb{M}$-martingale, just like the process Y^{T_n}. Being martingales, the processes Z and $V_0 Y^{T_n}$ have constant means, and hence

$$\begin{aligned} \mathrm{E}_{\mathbb{N}} 1_F 1_{T_n > T} &= \mathrm{E}_{\mathbb{N}} Z_T = \mathrm{E}_{\mathbb{N}} Z_0 \\ &= \mathrm{E}_{\mathbb{M}} V_0 Y_0^{T_n} = \mathrm{E}_{\mathbb{M}} V_0 Y_T^{T_n} = \mathrm{E}_{\mathbb{M}} 1_F 1_{T_n > T} \frac{V_0 N_T}{M_T}. \end{aligned}$$

Letting $n \to \infty$ and applying the monotone convergence theorem yields $\mathrm{E}_{\mathbb{N}} 1_F = \mathrm{E}_{\mathbb{M}} 1_F (V_0 N_T / M_T)$. Repeating this for an arbitrary event $F \in \mathcal{F}_T$, we see that $d\mathbb{N}/d\mathbb{M} = V_0 N_T/M_T$. Hence the density process of $\mathbb{N}$ relative to $\mathbb{M}$ is $\mathrm{E}_{\mathbb{M}}\bigl(V_0 N_T/M_T \,|\, \mathcal{F}_t\bigr)$.

Because N is a numeraire, $N = \alpha A$ for some self-financing strategy α and hence $N/M = \alpha A^M = \alpha_0 A_0^M + \alpha \cdot A^M$ by unit invariance. Because N/M is nonnegative it follows that it is an $\mathbb{M}$-super martingale, and so is the process $V_0 N/M$. But $\mathrm{E}_{\mathbb{M}} V_0 N_T/M_T = 1 = \mathrm{E}_{\mathbb{M}} V_0 = \mathrm{E}_{\mathbb{M}} V_0 N_0/M_0$ and hence the process $V_0 N/M$ is an $\mathbb{M}$-martingale. Combining this with the last conclusion of the preceding paragraph shows that the process $V_0 N/M$ is the density process of $\mathbb{N}$ relative to $\mathbb{M}$. ∎

4.27 EXERCISE. Let $(\mathbb{N}, N)$ be a numeraire pair and suppose that ϕ is a self-financing strategy whose discounted value process ϕA^N is bounded below by a constant and which replicates a claim X in that $\phi_T A_T = X$. Show that $\phi_0 A_0 \geq N_0 \mathrm{E}_{\mathbb{N}}(X/N_T)$. Thus the starting value of any replicating self-financing strategy with lower bounded value process (not necessarily admissible) is at least the price given by (4.18). Show this.

In a complete economy numeraire pairs $(\mathbb{N}, N)$ and $(\mathbb{M}, M)$ are connected through the change of measure relation $M_T^{-1}\, d\mathbb{M} = N_T^{-1}\, d\mathbb{N}$, by Lemma 4.20. This immediately implies that there can be at most one martingale measure for every given numeraire. If the discounted asset processes are locally bounded (in particular if the asset processes are continuous), then this uniqueness property characterizes completeness.

4.28 Theorem (Completeness). *Assume that the economy permits a numeraire pair $(\mathbb{N}, N)$ such that the process A^N is locally bounded. Then the following statements are equivalent:*
(i) The economy is complete.
(ii) If $(\mathbb{N}', N)$ is also a numeraire pair and $\mathbb{N}_0' = \mathbb{N}_0$, then $\mathbb{N}' = \mathbb{N}$ on $\mathcal{F}_T$.
(iii) For every nonnegative $\mathbb{N}$-local martingale M there exists an A^N-integrable predictable process ϕ such that $M = M_0 + \phi \cdot A^N$.
(iv) For every nonnegative $\mathbb{N}$-local martingale M there exists a self-financing strategy ϕ such that $M = \phi_0 A_0^N + \phi \cdot A^N$.

Proof. (i) ⇒ (ii). By Lemma 4.20 $\mathbb{N}'$ possesses density process $N/N = 1$ relative to $\mathbb{N}$.

(ii) ⇒ (iii). Assumption (ii) entails that there exists a unique probability measure on $\mathcal{F}_T$ with given initial measure such that the process A^N is an $\mathbb{N}$-local martingale. Therefore, by Theorem 3.2 (if A^N is continuous) or Theorem 3.4 (for general locally bounded A^N) the process A^N possesses the martingale representing property (iii).

(iii) ⇒ (iv). This follows from Lemma 4.8.

(iv) ⇒ (i). Given a nonnegative claim X with $\mathrm{E}_{\mathbb{N}}(X/N_T) < \infty$, we apply (iv) to the martingale M defined by $M_t = \mathrm{E}_{\mathbb{N}}\bigl(X/N_T \mid \mathcal{F}_t\bigr)$. Then we obtain a self-financing strategy ϕ such that $M = \phi_0 A_0^N + \phi \cdot A^N$, which implies that the process $\phi \cdot A^N$ is a martingale. By unit invariance and self-financing we have that $M = \phi A^N$. Evaluating this identity at time T, we see that $X = M_T N_T = \phi_T A_T$. The economy is complete by Lemma 4.25. ∎

A possible disadvantage of the characterization of completeness provided by the preceding theorem is that its conditions (ii) or (iii) may be difficult to verify. Uniqueness of the "martingale measure" $\mathbb{N}$ going with a numeraire N means uniqueness of the probability measure $\mathbb{N}$ making the process A^N into an $\mathbb{N}$-martingale. Because it is often the process A that is directly described in a model for the economy, not the rebased process A^N, standard theorems may not directly apply.

In particular, consider the situation that the asset processes are modelled through a stochastic differential equation of the form

$$dA_t = \mu(t, A_t)\, dt + \sigma(t, A_t)\, dW_t. \tag{4.29}$$

Then the continuous martingale part of A possesses the representing property relative to $(\Omega, \mathcal{F}, \{\mathcal{F}_t^A\}, \mathbb{P})$ under the reasonable condition that the solution to the equation is weakly unique, by Theorem 3.14. However, (iii) of the preceding theorem requires that we establish that the process A^N possesses the representing property relative to the filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t^A\}, \mathbb{N})$.

If the numeraire is a function of the assets A, then we can approach this by first deriving a diffusion equation for the process A^N. Because A^N is an $\mathbb{N}$-martingale, this diffusion equation should have a zero drift term if written relative to an $\mathbb{N}$-Brownian motion. We can next deduce the desired result with the help of Theorem 3.14. Similarly, the disadvantage may disappear if the numeraire N takes some other concrete form.

4.30 Example (Black-Scholes). The Black-Scholes economy of Example 4.10 permits a numeraire pair $(\mathbb{R}, R)$ with the numeraire $R_t = e^{rt}$, and the rebased stock price process satisfies

$$d\bigl(e^{-rt} S_t\bigr) = e^{-rt} \sigma S_t\, d\tilde{W}_t,$$

where $\tilde{W}_t = W_t + (\mu - r)t/\sigma$ is an $\mathbb{R}$-Brownian motion. The process $e^{-rt}\sigma S_t$ is strictly positive, and Brownian motion possesses the representing property relative to its augmented filtration. It follows from this that (iii) of the preceding theorem holds for $A = (R, S)$ and the filtration equal to the augmented filtration of Brownian motion. Thus the Black-Scholes economy is complete in this setting.

In Example 4.10 it was seen that $(\mathbb{R}, R)$ is a numeraire pair for the Black-Scholes economy for any given filtration for which the driven Brownian motion is indeed a Brownian motion. However, for a bigger filtration the Black-Scholes economy is not complete. By enlarging the filtration we add claims, which by definition are $\mathcal{F}_T$-measurable random variables, or equivalently local martingales M as in (iii) of Theorem 4.28. These may not be replicable or representable in terms of the Brownian motion, as they are not described in terms of this Brownian motion.

In practice one might think of an option on Shell stocks that cannot be replicated using only Philips stocks, even if the stocks satisfy the Black-Scholes model. So if the filtration is generated by both stocks and we use as asset process only the Shell stocks, then the economy will not be complete. A solution in this case could be to consider the Shell and Philips stocks jointly. However, one can easily imagine other examples of filtrations containing relevant information available in the market beyond the information in the stocks itself. □

The representing property of the continuous martingale part A^c (relative to $(\Omega, \mathcal{F}, \{\mathcal{F}_t^A\}, \mathbb{P})$) of the asset price processes is similar to the representing property of the rebased process A^N as in (iii) of the preceding lemma, but it does not imply it. Because the process A^c possesses the representing property in many situations, it is useful to investigate which additional conditions are necessary to ensure the representing property of A^N and hence completeness. The following theorem shows that, given the existence of a numeraire pair, it suffices to add the existence of a numeraire of bounded variation.[†]

4.31 Theorem (Completeness). *If the continuous martingale part $A^{c,\mathbb{P}}$ of a continuous economy A possesses the representing property relative to $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ and the economy permits a numeraire pair $(\mathbb{N}, N)$, then the following statements are equivalent:*
(i) The economy is complete.
(ii) There exists a numeraire of bounded variation.
(iii) There exists a strategy ϕ with $\phi \cdot A^{c,\mathbb{N}} = 0$ and $\phi A = N$.

Proof. Unless indicated differently, we interpret all statements and notation relative to the martingale measure $\mathbb{N}$. In particular, we abbreviate $A^c = A^{c,\mathbb{N}}$. Being the value process of a strategy, any numeraire is continuous.

[†] We write $X^{c,\mathbb{N}}$ for the local martingale part of a continuous semi-martingale X relative to a measure $\mathbb{N}$. This is the unique continuous local martingale such that $X - X^{c,\mathbb{N}}$ is an adapted continuous process of locally bunded variation.

(i) ⇒ (iii). Because N is a numeraire, there exist a self-financing strategy α with $N = \alpha A = \alpha_0 A_0 + \alpha \cdot A$. This implies that the local martingale part of N is given by $N^c = \alpha \cdot A^c$, for A^c the local martingale part of A. By completeness and Theorem 4.28(iv) there exists a self-financing strategy ψ such that $N^c = \psi \cdot A^N = \psi A^N$. Now set

$$\phi = \alpha - \frac{\psi}{N} + \Bigl(\frac{\psi A}{N^2}\Bigr)\alpha.$$

The process ψA is predictable, as it is adapted and continuous, and hence ϕ is predictable.[‡] Because $\alpha A = N$,

$$\begin{aligned}\phi A &= \alpha A - \frac{\psi A}{N} + \frac{\psi A}{N}\frac{\alpha A}{N} = \alpha A = N,\\ \phi\, dA &= \alpha\, dA - \frac{\psi}{N}\, dA + \frac{\psi A}{N^2}\alpha\, dA = dN - \frac{d(NN^c)}{N} + \frac{N^c}{N}\, dN,\end{aligned}$$

where for the last equality we use that $\psi\, dA = d(\psi A) = d(N\psi A^N) = d(NN^c)$. Using the partial integration formula on $d(NN^c)$ and the fact that $[N, N^c] = [N]$, we can rewrite the right side as

$$dN - \frac{N\, dN^c}{N} - \frac{N^c\, dN}{N} - \frac{d[N]}{N} + \frac{N^c}{N}\, dN = d(N - N^c) - \frac{d[N]}{N}.$$

We conclude that $\phi \cdot A = N - N^c - N^{-1} \cdot [N] - N_0$ is a continuous process of bounded variation, and hence the continuous martingale part $\phi \cdot A^c$ of the process $\phi \cdot A$ is zero.

(i)+(iii) ⇒ (ii). The process $M = \exp\bigl((\phi/N) \cdot A\bigr)$ is well defined, strictly positive, continuous, and of locally bounded variation (as $\phi \cdot A^c = 0$).[♭] Because $N = \phi A$, we have $(\phi/N)A = 1$ and hence $M = M(\phi/N)A$. By the definition of M and Itô's formula, where the second order term does not appear because M is of locally bounded variation, $dM = M(\phi/N)\, dA$. Together these equalities show that $M = \beta A$ for the self-financing strategy $\beta = M(\phi/N)$. Hence M is a numeraire of bounded variation.

(ii) ⇒ (i). In view of Girsanov's theorem, the continuous martingale part of A relative to the measure $\mathbb{N}$ is given by $A^{c,\mathbb{N}} = A^{c,\mathbb{P}} - L_-^{-1} \cdot [L, A^{c,\mathbb{P}}]$, if L is the density process of $\mathbb{N}$ relative to $\mathbb{P}$. By assumption the process $A^{c,\mathbb{P}}$ possesses the representing property for $\mathbb{P}$-local martingales. Therefore, by Lemma 3.7 the $\mathbb{N}$-local martingale $A^c = A^{c,\mathbb{N}}$ possesses the representing property for $\mathbb{N}$-local martingales. If we can show that $A^c = \delta \cdot A^N$ for some predictable process δ, then A^N inherits the representing property of A^c, and completeness follows by Theorem 4.28(iii).

[‡] It is A-integrable, as α is A-integrable, ψ is A-integrable (being self-financing and A^N-integrable), and $1/N$ and $\psi A/N$ are locally bounded.

[♭] Hunt and Kennedy claim that $(\phi/N) \cdot A$ is the bounded variation part of $\log N$. Is that true??

Finally, we represent A^c as a stochastic integral relative to A^N, using a bounded variation numeraire M with $M = \beta A = \beta_0 A_0 + \beta \cdot A$ for some self-financing strategy β. By unit invariance $d(M/N) = \beta\, dA^N$. By the partial integration formula,

$$dM = d\Big(\frac{M}{N} N\Big) = \frac{M}{N}\, dN + N\, d\Big(\frac{M}{N}\Big) + \Big[\frac{M}{N}, N\Big].$$

Because M is continuous and of bounded variation, its continuous martingale part is zero. Similarly, the continuous martingale part of the last term on the right $[M/N, N]$ is zero. Comparing the continuous martingale parts of the two sides of the display, we deduce that

$$\frac{M}{N}\, dN^c = -N\, d\Big(\frac{M}{N}\Big)^c = -N\beta\, d(A^N)^c = -N\beta\, dA^N,$$

since A^N is a continuous $\mathbb{N}$-local martingale, by assumption.

Again by the partial integration formula, $dA = d(NA^N) = N\, dA^N + A^N\, dN + [N, A^N]$. Hence $dA^c = N\, dA^N + A^N\, dN^c$. Combination with the preceding display gives the desired representation of dA^c as a multiple of dA^N. ∎

Warning. Existence of a numeraire pair is an assumption of the theorem, which is not implied by the existence of a numeraire of bounded variation. Even under the conditions of the theorem, the numeraire of finite variation whose existence is guaranteed in (iii) is not necessarily part of a numeraire pair.

The preceding theorem applies in particular to economies described through a stochastic differential equation of the type (4.29). If the filtration is chosen equal to the augmented natural filtration $\mathcal{F}_t^A$ of the asset prices (assumed right-continuous), then the condition that A^c possesses the representing property is verified if the solution to (4.29) is weakly unique, by Theorem 3.14. This is true for “nice” functions μ and σ.

4.32 Example (Black-Scholes). The Black-Scholes economy given in Example 4.10 is described through a stochastic differential equation that satisfies the Itô conditions. Hence the solution to the equation is weakly unique and possesses the representing property for the filtration generated by the assets. It permits a numeraire pair relative to this filtration.

The process $R_t = e^{rt}$ is a numeraire of bounded variation. Thus the Black-Scholes economy is complete. □

4.4 Incompleteness

Real markets are often thought to be not complete. Any replicable claim in an incomplete market can still be priced by the no-arbitrage principle of the preceding sections, and possesses fair price (4.18). However, by definition, in an incomplete market some claims are not replicable and a different approach is necessary. The main message of this section is that the no-arbitrage principle allows an interval of possible prices, and each of these prices can be written as in formula (4.18) for some martingale measure $\mathbb{N}$. By Theorem 4.28 an incomplete market allows multiple martingale measures $\mathbb{N}$ with a given numeraire N. Every of these measures defines a possible "fair" price.

As before let A be a vector-valued semimartingale and let A^N be its value expressed in terms of a given numeraire N. Let $\mathcal{N}$ be the set of all probability measures $\mathbb{N}$ such $(\mathbb{N}, N)$ is a numeraire pair. Given a claim X, let

$$\overline{V}_t = \operatorname*{ess\,sup}_{\mathbb{N}\in\mathcal{N}} \mathrm{E}_{\mathbb{N}}\Big(\frac{X}{N_T}\,|\,\mathcal{F}_t\Big). \tag{4.33}$$

Under integrability conditions (nonnegativeness of the claim X suffices) there exist versions of these essential suprema that define a cadlag process $\overline{V}$ that is an $\mathbb{N}$-supermartingale for every measure $\mathbb{N} \in \mathcal{N}$. (See Lemma 4.36.) The process $\overline{V}$ in all of the following will be understood to be this version, and it will silently be understood that it is finite.

4.34 EXERCISE **(Essential supremum)**. For any set $\{X_\alpha\colon \alpha \in A\}$ of random variables, there exists a random variable X (possibly with value ∞) such that
(i) $X \geq X_\alpha$ almost surely for every α.
(ii) if $Y \geq X_\alpha$ almost surely for every α, then $Y \geq X$ almost surely.
This random variable X, which is unique up to null sets, is called the *essential supremum* of the collection $\{X_\alpha\colon \alpha \in A\}$ and is denoted $\operatorname{ess\,sup}_\alpha X_\alpha$. [Hint: For any countable $B \subset A$ set $X_B = \sup_{\alpha\in B} X_\alpha$. Let B_n be a sequence such that $\mathrm{E}\,\mathrm{arctg}\, X_{B_n} \to \sup_B \mathrm{E}\,\mathrm{arctg}\, X_B$ and set $X = \sup_n X_{B_n}$.]

We shall show that $N_0\overline{V}_0$ is the minimum initial investment needed to superreplicate the claim with certainty: it is the minimal capital needed to start a self-financing strategy that is certain to yield (at least) the value of the claim X at time T. In other words, $N_0\overline{V}_0$ is the minimal price that is risk-free for the seller of the option: after selling the option the seller could buy the portfolio and implement a self-financing strategy that is certain to yield a capital equivalent to the claim at expiry time.

Conversely, the amount $N_0\underline{V}_0$, for $\underline{V}$ the corresponding essential infimum, is the biggest price that is risk-free for the buyer. Any price not in the interval $[N_0\underline{V}_0, N_0\overline{V}_0]$ allows risk-free profit (for buyer or seller), whereas any price inside the interval is arbitrage-free and requires risk taking on the part of buyer and/or seller.

The class $\mathcal{N}$ of all local martingale measures is convex, and the map $\mathbb{N} \mapsto \mathrm{E}_{\mathbb{N}}(X/N_T)$ is linear. It follows that the range of this map is convex, and hence an interval in the real line. Thus, if $\mathcal{F}_0$ is trival, then any price in the interval $(N_0\underline{V}_0, N_0\overline{V}_0)$ can be written in the form $\mathrm{E}_{\mathbb{N}}(X/N_T)$ for some $\mathbb{N} \in \mathcal{N}$.

4.35 EXERCISE. Suppose that A is a local martingale on the filtered space $(\mathcal{X}, \mathcal{F}, \{\mathcal{F}_t\})$, both when equipped with $\mathbb{P}$ and when equipped with $\mathbb{Q}$. Then A is also a local martingale under the measure $\lambda\mathbb{P} + (1-\lambda)\mathbb{Q}$ for any $\lambda \in [0,1]$. [Hint: consider the density process of the convex combination and use Lemma 2.14.]

The final value of $\overline{V}$ satisfies $\overline{V}_T = X/N_T$ almost surely. The process $\overline{V}$ is the smallest supermartingale with this property.

4.36 Lemma. *If X is a nonnegative $\mathcal{F}_T$-measurable random variable such that $\mathrm{E}_{\mathbb{N}}(X/N_T) < \infty$ for all $\mathbb{N} \in \mathcal{N}$, then*
(i) the process $\overline{V}$ is an $\mathbb{N}$-supermartingale for every $\mathbb{N} \in \mathcal{N}$ and permits a cadlag version.
(ii) any process V with $V_T \geq X/N_T$ which is an $\mathbb{N}$-supermartingale for every $\mathbb{N} \in \mathcal{N}$ satisfies $V_t \geq \overline{V}_t$ almost surely, for every t.

4.37 Corollary. *Suppose $X \geq 0$. Any self-financing strategy ϕ with $(\phi A)_T \geq X$ and which is admissible or has nonnegative value process ϕA satisfies $\phi_t A_t \geq N_t \overline{V}_t$ almost surely.*

Proofs. For (ii) of the lemma we note that an $\mathbb{N}$-supermartingale V satisfies $V_t \geq \mathrm{E}_{\mathbb{N}}(V_T|\,\mathcal{F}_t)$ almost surely, which is bounded below by $\mathrm{E}_{\mathbb{N}}(X/N_T|\,\mathcal{F}_t)$ if $V_T \geq X/N_T$. This being true for any $\mathbb{N} \in \mathcal{N}$ gives that $V_t \geq \overline{V}_t$ almost surely by the definition of the essential supremum.

The value process in terms of a numeraire $V = \phi A^N = (\phi A^N)_0 + \phi \cdot A^N$ of a self-financing strategy ϕ as in the corollary is either an $\mathbb{N}$-martingale (if ϕ is admissible) or at least an $\mathbb{N}$-supermartingale (if ϕA is nonnegative). It satisfies $V_T \geq X/N_T$ if $(\phi A)_T \geq X$. Therefore $(\phi A)_t = N_t(\phi A^N)_t \geq N_t \overline{V}_t$ almost surely by (ii).

For the proof of (i) of the lemma fix an arbitrary measure $\mathbb{Q}$ in $\mathcal{N}$ for reference, and write expectations $\mathrm{E}_{\mathbb{Q}}$ relative to $\mathbb{Q}$ as E. For simplicity of notation write X for X/N_T and A for A^N.

Fix some $t \in [0,T]$. If $\mathbb{N}$ has density process L relative to $\mathbb{Q}$, then

$$\mathrm{E}_{\mathbb{N}}(X|\,\mathcal{F}_t) = \frac{\mathrm{E}(L_T X|\,\mathcal{F}_t)}{L_t} = \mathrm{E}(K_T X|\,\mathcal{F}_t), \tag{4.38}$$

for K the process $K = 1_{[0,t)} + (L/L_t)1_{[t,T]}$. The process K is a positive martingale with mean 1 and hence defines a density process of a probability measure $\tilde{\mathbb{N}}$ relative to $\mathbb{Q}$. This measure $\tilde{\mathbb{N}}$ is contained in $\mathcal{N}$ (i.e. A is an $\tilde{\mathbb{N}}$-local martingale), because the process $KA = 1_{[0,t)}A + (LA/L_t)1_{[t,T]}$ is a $\mathbb{Q}$-local martingale, both the processes A and LA being $\mathbb{Q}$-local martingales. It follows that the essential supremum in the definition of $\overline{V}_t$ can be restricted to local martingale measures $\mathbb{N}$ with density process identically equal to 1 on $[0,t]$. Write $\mathcal{K}_t$ for the corresponding collection of density processes.

Because an essential supremum can always be written as a supremum over a countable subset, it follows that $\overline{V}_t = \sup_n \mathrm{E}(K_T^{(n)}X|\,\mathcal{F}_t)$ for a sequence of density processes $K^{(n)} \in \mathcal{K}_t$. In the present case we can choose the countable subset in such a way that the sequence $\mathrm{E}(K_T^{(n)}X|\,\mathcal{F}_t)$ is nondecreasing. Indeed, given density processes $K^{(1)}$ and $K^{(2)}$ in $\mathcal{K}_t$, we can define

$$K = K^{(1)}1_C + K^{(2)}1_{C^c}, \qquad C = \bigl\{\mathrm{E}(K_T^{(1)}X|\,\mathcal{F}_t) > \mathrm{E}(K_T^{(2)}X|\,\mathcal{F}_t)\bigr\}. \tag{4.39}$$

The process K is a $\mathbb{Q}$-martingale, as $K^{(1)}, K^{(2)}$ are $\mathbb{Q}$-martingales which coincide on $[0,t]$ and $C \in \mathcal{F}_t$. Furthermore, the process KA is a $\mathbb{Q}$-local martingale by the same reasoning. The process K is positive, has mean one and is identically one on $[0,t]$. Therefore, K belongs to $\mathcal{K}_t$. By construction

$$\begin{aligned}\mathrm{E}(K_TX|\,\mathcal{F}_t) &= 1_C\mathrm{E}(K_T^{(1)}X|\,\mathcal{F}_t) + 1_{C^c}\mathrm{E}(K_T^{(2)}X|\,\mathcal{F}_t)\\ &= \mathrm{E}(K_T^{(1)}X|\,\mathcal{F}_t) \vee \mathrm{E}(K_T^{(2)}X|\,\mathcal{F}_t).\end{aligned}$$

Thus if $\mathrm{E}(K_T^{(n)}X|\,\mathcal{F}_t)$ is not increasing for the original sequence $K^{(n)}$, we can transform it in an increasing sequence by taking successive linear combinations of the type (4.39).

It follows that there exists a sequence $K^{(n)}$ in $\mathcal{K}_t$ with $0 \leq \mathrm{E}(K_T^{(n)}X|\,\mathcal{F}_t) \uparrow \overline{V}_t$. By the monotone convergence theorem, for $s < t$,

$$\mathrm{E}(K_T^{(n)}X|\,\mathcal{F}_s) = \mathrm{E}\bigl(\mathrm{E}(K_T^{(n)}X|\,\mathcal{F}_t)|\,\mathcal{F}_s\bigr) \uparrow \mathrm{E}(\overline{V}_t|\,\mathcal{F}_s), \qquad \text{a.s.}.$$

Because each $K^{(n)}$ is a density process of a measure in $\mathcal{N}$ and $K_s^{(n)} = 1$, the left side is $\mathrm{E}_{\mathbb{N}}(X|\,\mathcal{F}_s)$, for some $\mathbb{N} \in \mathcal{N}$, for every n. (Cf. equation (4.38) with s in place of t.) Therefore it is bounded above by $\overline{V}_s$, whence $\overline{V}_s \geq \mathrm{E}(\overline{V}_t|\,\mathcal{F}_s)$ almost surely. This concludes the proof that the process $\overline{V}$ is a supermartingale.

By general theory the supermartingale $\overline{V}$ permits a cadlag version if its mean $t \mapsto \mathrm{E}\overline{V}_t$ is right-continuous. The monotone convergence theorem applied again to the increasing sequence $0 \leq \mathrm{E}(K_T^{(n)}X|\,\mathcal{F}_t) \uparrow \overline{V}_t$, but this time unconditionally, gives that $\mathrm{E}\overline{V}_t = \sup_n \mathrm{E}(K_T^{(n)}X)$. Because $\overline{V}_t \geq \mathrm{E}(K_TX|\,\mathcal{F}_t)$ almost surely for every $K \in \mathcal{K}_t$, it follows that the supremum does not increase by allowing all $K \in \mathcal{K}_t$, i.e. $\mathrm{E}\overline{V}_t = \sup_{K\in\mathcal{K}_t}\mathrm{E}(K_TX)$. For $K \in \mathcal{K}_t$ and $s > t$, the process $B_sK = 1_{[0,s)} + (K/K_s)1_{[s,T]}$ is contained in $\mathcal{K}_s$. (Previously K was obtained from L by the operator B_t, and hence the claim follows by preceding arguments.) We have that

$(B_sK)_T = K_T/K_s \to K_T/K_t = K_T$ as $s \downarrow t$, by right-continuity of K. By Fatou's lemma

$$\mathrm{E}\overline{V}_t = \sup_{K\in\mathcal{K}_t} \mathrm{E}(K_TX) \le \sup_{K\in\mathcal{K}_t} \liminf_{s\downarrow t} \mathrm{E}((B_sK)_TX) \le \liminf_{s\downarrow t} \mathrm{E}\overline{V}_s.$$

The supermartingale property of $\overline{V}$ immediately yields the reverse inequality, whence $t \mapsto \mathrm{E}\overline{V}_t$ is right-continuous. ∎

4.40 EXERCISE. If M and N are (local) martingales, then
(i) $M1_{[0,t)} + N(M_t/N_t)1_{[t,\infty)}$ is a (local) martingale.
(ii) if M and N agree on $[0,t]$ for some t and $C \in \mathcal{F}_t$, then $M1_C + N1_{C^c}$ is a (local) martingale.

The lemma shows that any admissible strategy ϕ that yields with certainty at least the value of the claim X at expiry time will cost at least $N_0\overline{V}_0$ to start at time 0. The next result shows that there exists a strategy at this price that superreplicates the claim.

4.41 Theorem. *Assume that $X \ge 0$ and that the process A^N is locally bounded. Then there exists a self-financing strategy ϕ such that $\overline{V} = \phi A^N - C$ for a cadlag, adapted, nondecreasing process C with $C_0 = 0$.*

Proof. Because the process $\overline{V}$ is an $\mathbb{N}$-super martingale for every $\mathbb{N} \in \mathcal{N}$, Theorem 3.5 shows that there exist an A^N-integrable predictable process ψ and a cadlag, adapted, nondecreasing process C such that $\overline{V} = \overline{V}_0 + \psi \cdot A^N - C$. By Lemma 4.8, there exist a self-financing strategy ϕ with $(\phi A^N)_0 = V_0$ and $\phi \cdot A^N = \psi \cdot A^N$. This satisfies $\overline{V} = \phi A^N - C$ and hence fulfills the requirements. ∎

The strategy ϕ in the lemma is not necessarily admissible in that the process $\phi \cdot A^N$ is an $\mathbb{N}$-martingale. However, the strategy is "one-sided" admissible in that its value process $\phi A^N = \overline{V} + C \ge \overline{V}$ is nonnegative if $X \ge 0$.

We interpret the process C in Theorem 4.41 as a cumulative *consumption process*: at each time instant t the current value of the portfolio is used to form a new portfolio and to extract an amount dC_t, yielding the portfolio value (in terms of the numeraire) at time t

$$\overline{V}_t = (\phi A^N)_0 + (\phi \cdot A^N)_t - C_t.$$

The extracted amount dC_t is nonnegative, because the process C is nondecreasing. The final value of the portfolio is

$$(\phi A)_T = N_T(\phi A^N)_T = N_T(\overline{V}_T + C_T) \ge N_T\overline{V}_T = X.$$

Thus notwithstanding the possible consumption during the term of the contract, the portfolio ends up superreplicating the claim. The initial cost

$\phi_0 A_0 = N_0 \overline{V}_0$ of the portfolio should be an acceptable price to the seller of the claim X. An initial cost higher than $N_0 \overline{V}_0$ would provide the seller an opportunity for riskless gain (arbitrage), as would a price of exactly $N_0 \overline{V}_0$ if $C_T > 0$.

Of course, the price $N_0 \overline{V}_0$ may not be to the liking of the buyer of the claim. Set

$$\underline{V}_t = \operatorname*{ess\,inf}_{\mathbb{N} \in \mathcal{N}} \mathrm{E}_{\mathbb{N}}\Bigl(\frac{X}{N_T} \,|\, \mathcal{F}_t\Bigr).$$

By similar arguments we can argue that $N_0 \underline{V}_0$ is the highest price for the buyer to be sure that he can hedge away the complete risk of buying the claim X. If the prices $N_0 \underline{V}_0$ and $N_0 \overline{V}_0$ do not agree, seller and/or buyer must take some risk, and the price of the claim will be some number between the two extremes.

To argue the riskless price $N_0 \overline{V}_0$ for the seller we have assumed that the claim X is nonnegative. The analogous argument from the point of view of the buyer is the seller's argument applied to minus the claim. Application of the preceding would require the assumption that the claim $-X$ is nonpositive, which is incompatible. It is therefore necessary to relax the assumptions. By adding and subtracting it can be seen that Theorem 4.41 is true for any claim X that is lower-bounded by a claim that can be replicated by an admissible strategy. This shows that both the buyer's and the seller's argument are valid for any claim X for which there exist replicable claims $\underline{X}$ and $\overline{X}$ with $\underline{X} \le X \le \overline{X}$.

4.42 Corollary. *Suppose that $X \ge \underline{X}$ for a claim $\underline{X}$ for which there exists a self-financing strategy H with $\underline{X} = (HA)_T$ such that $H \cdot A^N$ is an $\mathbb{N}$-martingale for every $\mathbb{N} \in \mathcal{N}$. Then the assertion of Theorem 4.41 remains true.*

Proof. We apply Theorem 4.41 to the nonnegative claim $X - \underline{X}$. The supermartingale attached to this claim through formula (4.33) is $V(X - \underline{X}) = \overline{V}(X) - HA^N$ for $\overline{V}(X)$ the supermartingale attached to X, exactly as given in (4.33). The corollary follows by adding and subtracting, with the desired strategy ϕ equal to $H + \phi(X - \underline{X})$ for $\phi(X - \underline{X})$ the strategy provided for the claim $X - \underline{X}$ by Theorem 4.41. ∎

4.43 Example (Black-Scholes with stochastic volatility). Consider an economy consisting of two asset processes, a deterministic fixed income process $R_t = e^{rt}$ and a process S satisfying the coupled SDEs

$$\begin{aligned} dS_t &= S_t(\mu\, dt + \sigma_t\, dW_t^{(1)}), \\ d\sigma_t &= \sigma_t\, dW_t^{(2)}, \end{aligned}$$

for a two-dimensional Brownian motion $W = (W^{(1)}, W^{(2)})$. This economy differs from the Black-Scholes model in that the volatility σ_t is a random process rather than a constant. It is an example of a *stochastic volatility* model. Economies with stochastic volatility are not necessarily incomplete, but this example is. We shall show that the economy (R, S) permits a numeraire pair, but is not complete relative to the augmented filtration generated by S. Note that the volatility process σ is used to describe the economy, but is not part of the economy (R, S) itself.

Two applications of Itô's formula show that the processes S and σ can be written in the form

$$\begin{aligned} S_t &= S_0 \exp\Bigl(\mu t - \tfrac{1}{2}\int_0^t \sigma_s^2\, ds + \int_0^t \sigma_s\, dW_s^{(1)}\Bigr),\\ \sigma_t &= \sigma_0 \exp(W_t^{(2)} - \tfrac{1}{2}t). \end{aligned}$$

In particular these processes are strictly positive (provided the initial values are positive).

Let $\mathcal{F}_t$ be the augmented natural filtration of the driving Brownian motion W. The preceding display shows that (S, σ) is adapted to this filtration (which is also an implicit assumption in saying that these processes solve the SDE). The second equation of the display shows directly that $W^{(2)}$ is adapted to the filtration generated by σ. The positivity of S and σ allows to write $W^{(1)}$ as a stochastic integral of (S, σ). Thus W is adapted to the filtration generated by (S, σ), whence the filtration $\mathcal{F}_t$ is the augmented natural filtration of (S, σ). From the fact that the process $Y = S^{-1} \cdot S$ has quadratic variation $[Y]_t = \int_0^t \sigma_s^2\, ds$ it follows that the integral of the square volatility process is adapted to the filtration generated by S, which implies that σ itself is also adapted. We conclude that the augmented filtrations generated by W, by S, or by (S, σ) all coincide with the single filtration $\mathcal{F}_t$.

We choose $R_t = e^{rt}$ as a numeraire. The rebased process S/R can be written in the form

$$\frac{S_t}{R_t} = S_0 \exp\Bigl(\int_0^t \sigma_s\, d\tilde W_s^{(1)} - \tfrac{1}{2}\int_0^t \sigma_s^2\, ds\Bigr),$$

for $\tilde W_t^{(1)} = W_t^{(1)} - \int_0^t (r-\mu)/\sigma_s\, ds$. If $\tilde W^{(1)}$ is an $\mathbb{R}$-Brownian motion, then this process is exactly the exponential process $S_0 \mathcal{E}(\sigma \cdot W^{(1)})$ and hence S/R is an $\mathbb{R}$-martingale. It follows that $(\mathbb{R}, R)$ is a numeraire pair for the economy (R, S) for any equivalent measure $\mathbb{R}$ that makes $\tilde W^{(1)}$ an $\mathbb{R}$-Brownian motion.

We can construct such martingale measures with Girsanov's theorem. For any given suitably integrable predictable process $Y = (Y^{(1)}, Y^{(2)})$ the process $\mathcal{E}(Y \cdot W)$ is a strictly positive martingale and hence defines a density process of an equivalent measure $\mathbb{N}$ relative to $\mathbb{P}$. (We write $Y \cdot W = Y^{(1)} \cdot W^{(1)} + Y^{(2)} \cdot W^{(2)}$, as usual.) By the vector-valued version of Girsanov's theorem the process $\tilde W$ given by $\tilde W_t = W_t - \int_0^t Y_s\, ds$ is a two-dimensional $\mathbb{N}$-Brownian motion. It follows that any process Y such that $Y^{(1)} = (r-\mu)/\sigma$ generates a numeraire measure $\mathbb{R}$. This means that the first coordinate of Y is fully determined, but the second component is free (up to integrability). Thus the numeraire measure is not unique and hence the economy is incomplete, by Theorem 4.28.

Sufficient "integrability" of the process Y can be verified by application of Novikov's criterion and/or direct arguments. Indeed, the process $\mathcal{E}(Y \cdot W)$ is a local martingale as soon as the stochastic integral $Y \cdot W$ is well defined. It is a martingale as soon as its mean is constant. In the case that $Y^{(2)} = 0$, the latter can be seen directly by conditioning on $W^{(2)}$ (and hence on σ), because, for every t, by the independence of $W^{(1)}$ from $W^{(2)}$ (and hence σ),

$$\mathrm{E}\Bigl(e^{\int_0^t (r-\mu)/\sigma_s\, dW_s^{(1)} - \frac{1}{2}\int_0^t (r-\mu)^2/\sigma_s^2\, ds} \,|\, \sigma\Bigr) = 1.$$

This calculation can also be applied to more general $Y^{(2)}$ that depend on $W^{(2)}$ only, in particular any constant process $Y^{(2)}$. Because $[W^{(1)}, W^{(2)}] = 0$, the process $L = \mathcal{E}(Y \cdot W)$ factorizes as $L = L^{(1)} L^{(2)}$ for $L^{(i)} = \mathcal{E}(Y^{(i)} \cdot W^{(i)})$, and we can apply Novikov's criterion to see that $L^{(2)}$ is a martingale, and next the conditioning argument to see that L is a martingale.

In contrast, the economy (R, S, σ) in which the volatility process itself is tradable is complete. (Compare Theorem 5.7.) Thus a claim X, a measurable function of the process $(S_t\colon 0 \le t \le T)$, may be replicable using the assets (R, S, σ), but not using only (R, S). This is a little surprising, because the volatility process σ is itself a measurable function of the process S, so that it is observed if (R, S) are observed. The explanation is that even though an integral $\phi \cdot \sigma$ is a function of S, it cannot necessarily be written as an integral $\psi \cdot S$ with respect to S.

It is instructive to compute the price interval $[\underline{V}_0, \overline{V}_0]$ for a European option $X = (S_T - K)^+$ in the economy (R, S). By definition the boundaries of the price interval are infimum and supremum relative to all numeraire measures, and hence if $\mathbb{N}_y$ is the numeraire measure constructed previously with $Y^{(2)} = y$ constant, then

$$\begin{aligned} \underline{V}_0 &\le \inf_y \mathrm{E}_{\mathbb{N}_y} e^{-rT} X, \\ \overline{V}_0 &\ge \inf_y \mathrm{E}_{\mathbb{N}_y} e^{-rT} X. \end{aligned}$$

The expectations in the right sides can be written as

$$\begin{aligned} e^{-rT} \mathrm{E}_{\mathbb{N}_y} &\Bigl(e^{rT} S_0 e^{\int_0^T \sigma_s\, d\tilde W_s^{(1)} - \frac{1}{2}\int_0^T \sigma_s^2\, ds} - K\Bigr)^+ \\ &= e^{-rT} \mathrm{E}_{\mathbb{N}_y} \Bigl(e^{rT} S_0 e^{\tau Z - \frac{1}{2}\tau^2} - K\Bigr)^+, \end{aligned}$$

where Z is a standard normal variable independent of $\tau^2 = \int_0^T \sigma_s^2\, ds$, which is distributed according to the SDE

$$d\sigma_t = \sigma_t\, dW_t^{(2)} = \sigma_t (y\, dt + d\tilde W_t^{(2)}),$$

for $\tilde{W}^{(2)}$ a Brownian motion under $\mathbb{N}_y$. The constant y, which is the drift of the diffusion σ_t under $\mathbb{N}_y$, ranges over $\mathbb{R}$. For $y \uparrow \infty$ the sample paths of the process σ increase to infinity and hence $\tau \uparrow \infty$, whereas for $y \downarrow -\infty$ the process tends to zero and $\tau \downarrow 0$. For a given value of τ the right side of the display evaluates to the standard Black-Scholes formula, given in Chapter 1, with volatility σ satisfying $\sigma\sqrt{T} = \sqrt{\tau}$. This formula tends to $(S_0 - Ke^{-rT}) \vee 0$ and S_0 as $\tau \downarrow 0$ or $\tau \uparrow \infty$, respectively. It follows that $\underline{V}_0 \leq (S_0 - Ke^{-rT}) \vee 0 \leq S_0 \leq \overline{V}_0$.

We can show that the first and last inequality is actually an equality by exhibiting two explicit strategies. By Corollary 4.37 and Theorem 4.41 the variable $\overline{V}_0$ is the smallest amount needed to superreplicate the claim $(S_t - K)^+$, and similarly $-\underline{V}_0$ is the smallest amount needed to superreplicate minus the claim. Consider the following two hedging schemes:

(i) At time $t = 0$ buy one stock. Do nothing until time $t = T$.

(ii) At time $t = 0$ sell one stock and place an amount Ke^{-rT} in the savings account. Do nothing until time $t = T$.

Strategy (i) costs S_0 to start, and has value S_T at expiry time. Since $S_T \geq (S_T - K)^+$ the portfolio overreplicates the claim and hence $\overline{V}_0 \leq S_0$, by Corollary 4.37. Portfolio (ii) costs $-S_0 + Ke^{-rT}$ to start and is worth $-S_T + K$ at expiry time. Since the latter is more than $-(S_T - K)^+$ it follows that $-\underline{V}_0 \leq -S_0 + Ke^{-rT}$, by Corollary 4.37. Together with the trivial bound $\underline{V}_0 \geq 0$ this shows that $\underline{V}_0 \geq (S_0 - Ke^{-rT})^+$.

We conclude that the boundaries of the interval $[\underline{V}_0, \overline{V}_0]$ correspond to strategies that do no require any trading during $(0, T]$. This seems a bit disappointing. Apparently in the present stochastic volatility model the no-arbitrage approach excludes only the most unrealistic prices and gives no further guidance about a "fair price". □

** **4.44** EXERCISE. Extend the calculation of the price interval $[\underline{V}_0, \overline{V}_0]$ in the preceding example to more general claims of the type $X = f(S_T)$ for f a convex function.

** **4.45** EXERCISE. In the preceding example let $(W^{(1)}, W^{(2)})$ be *correlated Brownian motion*: given two independent Brownian motions $B^{(1)}$ and $B^{(2)}$ and a constant $\rho \in (-1, 1)$ let $W^{(1)} = B^{(1)}$ and let $W^{(2)} = \rho B^{(1)} + (1 - \rho^2)^{1/2} B^{(2)}$. Investigate whether this makes a difference for the conclusions.

Thus the fair price of a claim X can be written as $\mathrm{E}_{\mathbb{N}}(X/N_T)$ for some local martingale measure $\mathbb{N} \in \mathcal{N}$ (if $\mathcal{F}_0$ is trivial). There are at least two approaches for selecting the martingale measure that determines the price. The first is to say that "the market determines the martingale measure". The prices of commonly traded options can be observed on the option market, and the martingale measure can be inferred by calibrating the observed prices to the prices given by (4.18). The second approach digs deeper in the theory of economic behaviour and is based on utility arguments.

* 4.5 Utility-based Pricing

Fix some nondecreasing function $u: \mathbb{R} \mapsto \mathbb{R}$, which we shall refer to as a *utility function.* The *expected utility* (at time 0) of receiving an amount Y at time T is defined as the expected value $\mathrm{E}_{\mathbb{P}} u(Y/N_T)$ computed under the real world probability measure $\mathbb{P}$. The future payment Y is discounted by a given numeraire N and Y is assumed to be an $\mathcal{F}_T$-measurable random variable. Given an initial wealth x at time 0 we aim at maximizing expected utility, where our possible strategies are to invest in the assets and/or to buy or sell an option with a given claim X.

If we do not trade in the option, but invest the total initial wealth x in the assets, then the maximal expected utility is given by

$$U(x) = \sup_{\phi\in\Phi:\phi_0 A_0 = x} \mathrm{E}_{\mathbb{P}} u\bigl(\phi_T A_T^N\bigr) = \sup_{\phi\in\Phi_x} \mathrm{E}_{\mathbb{P}} u\bigl(x/N_0 + \phi\cdot A_T^N\bigr).$$

The first supremum is computed over all strategies with initial investment x belonging to the collection Φ of self-financing strategies with nonnegative value process ϕA. The second supremum may be computed over the same set of strategies, but in view of Lemma 4.8 it remains the same if computed over the collection Φ_x of all A^N-integrable, predictable processes ϕ with $x/N_0 + \phi\cdot A^N \geq 0$. The limitation to strategies with a nonnegative value process is meant to ensure that the expected utility does not reduce to the trivial value $u(\infty)$.[♯]

Alternatively, we could buy one option and invest the remainder of our initial capital in the assets. If we buy the option at price y and start with an initial capital x, then the maximal expected utility is

$$U_b(x;y,X) = \sup_{\phi\in\Phi:\phi_0 A_0 = x-y} = \mathrm{E}_{\mathbb{P}} u\Bigl(\phi_T\cdot A_T^N + \frac{X}{N_T}\Bigr).$$

From the buyer's point of view an acceptable price y of the option at time 0 should satisfy $U_b(x;y,X)\geq U(x)$. The maximal value y for which the buyer is interested to buy the option is $p_b(x;X) = \sup\{y: U_b(x;y,X)\geq U(x)\}$ and is called the *buyer's price.*

The seller of the option has to deliver the amount X at time T, but collects the price y of the option at time 0. Hence from the seller's point of view the expected utility of an initial capital x is

$$U_s(x;y,X) = \sup_{\phi\in\Phi:\phi_0 A_0 = x+y} \mathrm{E}_{\mathbb{P}} u\Bigl(\phi_T A_T^N - \frac{X}{N_T}\Bigr).$$

[♯] A qualitative restriction such as "admissibility" in the sense of Definition 4.15 may not achieve this. Suppose that ϕ is a self-financing strategy with $\phi_0 A_0 = x$ such that $\phi\cdot A^N$ is an $\mathbb{N}$-martingale. Then for any $\lambda\in\mathbb{R}$ the process $\lambda\phi$ is A^N-integrable, whence by Lemma 4.8 there exists a self-financing strategy ψ with $\psi_0 A_0 = x$ and $\psi\cdot A^N = (\lambda\phi)\cdot A^N$. In particular $\psi\cdot A^N$ is an $\mathbb{N}$-martingale. It appears $\mathrm{E}_{\mathbb{P}} u(\psi_T A_T^N) = \mathrm{E}_{\mathbb{P}} u(x/N_0 + \lambda\phi\cdot A_T^N)$ can be maximized to an extreme value by choosing $\lambda\to\pm\infty$, at least if u is unbounded.

For the seller a price y is acceptable only if $U_s(x;y,X) \geq U(x)$. The *seller's price* of the option is the acceptable minimal price, given by $p_s(x;X) = \inf\{y\colon U_s(x;y,X) \geq U(x)\}$. In general buyer's and seller's prices are not the same (??). Of course, there is also no reason why the utility functions of buyer and seller would agree, although they are the same in our notation.

It follows easily from the definitions that $p_b(x;X) = -p_s(x;-X)$. This is what one would intuitively expect: buying is the same as selling the negative at the negated price. In general, the prices depend on the initial capital x, as indicated in the notation, but also on the real world probability measure $\mathbb{P}$. The prices are consistent with the no-arbitrage prices found before.

4.46 Theorem. *Assume that A^N is locally bounded and that $X \geq 0$. Then for any strictly increasing utility function u such that U is strictly monotone both the buyer's price $p_b(x;X)$ and the seller's price $p_s(x;X)$ are bounded above by $N_0\overline{V}_0$.*

Proof. By Theorem 4.41 there exists a self-financing strategy ψ with $\overline{V} = \psi A^N - C$ for an increasing, adapted process C with $C = 0$. The value process ψA of this strategy is nonnegative and satisfies $\psi_0 A_0 = N_0 \overline{V}_0$ and $\psi_T A_T^N \geq X/N_T$.

Given any $\phi \in \Phi$, the process $\phi + \psi$ is contained in Φ as well and $(\phi+\psi)_T A_T^N - X/N_T \geq \phi_T A_T^N$.

If $\phi_0 A_0 = x$, then $(\phi+\psi)_0 A_0 = x + N_0\overline{V}_0$. Thus for every strategy ϕ in the definition of U the strategy $\phi+\psi$ is a strategy in the supremum defining $U_s(x;N_0\overline{V}_0;X)$ that has at least the same (expected) utility for the seller. We conclude that $U_s(x;N_0\overline{V}_0;X) \geq U(x)$, so that $p_s(x;X) \leq N_0\overline{V}_0$.

If $\phi_0 A_0 = x - y$ then $(\phi+\psi)_0 A_0 = x - y + N_0\overline{V}_0$ and $(\phi+\psi)_T A_T^N \geq \phi_T A_T^N + X/N_T$. We conclude that $U(x - y + N_0\overline{V}_0) \geq U_b(x;y;X)$, so that $U(x) > U_b(x;y;X)$ for $y > N_0\overline{V}_0$ by the assumed strict monotonicity of U. This implies that $p_b(x;X) \leq y$, whence $p_b(x;X) \leq \overline{N}_0 V_0$. ∎

We can prove by an analogous argument that the prices $p_b(x;X)$ and $p_s(x;X)$ are lower bounded by $N_0\underline{V}_0$. In fact, this follows from the upper bound given in the preceding lemma applied to the claim $-X$ and the identities $p_b(x;X) = -p_s(x;-X)$, $p_s(x;X) = -p_b(x;-X)$, together with the fact that $-\underline{V}_0$ relates to $-X$ as $\overline{V}_0$ relates to X. Of course, we cannot assume that both X and $-X$ are nonnegative, and hence the preceding lemma would need to be extended to more general claims. This seems to require consideration of strategies with value processes that can be negative (but are bounded below by a suitable process).

4.6 Early Payments

The pricing formula (4.18) gives the value of a contract that consists of a single payment at an "expiry time" T. Many contracts include payments made during the term of the contract. We can extend the pricing formula to such contracts by replacing a payment of Y at time $S < T$ by an equivalent amount paid at time T. Given a numeraire N the "equivalent" payment is YN_T/N_S, where the factor N_T/N_S is typically larger than 1 and accounts for the time value of money. We also arrive at this amount if we think of the payment Y as being invested in Y/N_S units of the numeraire at time S, so that it grows to the value $(Y/N_S)N_T$ at time T.

If this reasoning is correct, then the just price at time t for a claim consisting of a payment Y at time S is given by (4.18) applied to $X = YN_T/N_S$, i.e. $N_t\mathrm{E}_{\mathbb{N}}\bigl(Y/N_S|\,\mathcal{F}_t\bigr)$. Repeating this argument, we see that the value at time t of a contract with payments $Y_1, \ldots, Y_n$ at times $T_1 < \cdots T_n < T$ is given by, for $t < T_1$,

$$N_t\sum_{i=1}^{n}\mathrm{E}_{\mathbb{N}}\Bigl(\frac{Y_i}{N_{T_i}}|\,\mathcal{F}_t\Bigr). \tag{4.47}$$

Thus each of the payments is discounted with the discount factor that is current at the time of payment.

This reasoning is believable, but perhaps not as convincing as the arguments based on replication used to "prove" the pricing formula (4.18). In Chapter 7 we find the same formula using replication arguments, within a more general framework allowing a continuous flow of payments over an interval.

* 4.7 Pricing Kernels

The pricing formula (4.18) employs a change of measure to a martingale measure to express the value process of a claim. Because the martingale measure is equivalent to the original or "true" measure $\mathbb{P}$, an expectation relative to it can also be viewed as a weighted expectation under the original measure. This is formalized by the concept of a pricing kernel.

In this section a "numeraire pair" is understood to be a "martingale numeraire pair" and the assets are assumed continuous.

4.48 Definition. *A pricing kernel Z is a strictly positive, cadlag semimartingale such that ZA is a $\mathbb{P}$-martingale.*

If $(\mathbb{N}, N)$ is a numeraire pair, then $Z = L/N$, for L the density process of $\mathbb{N}$ relative to $\mathbb{P}$, is a pricing kernel. This follows, because the process

A^N is an $\mathbb{N}$-martingale if and only if the process $LA^N = (L/N)A$ is a $\mathbb{P}$-martingale, by Lemma 2.14.[†]

In terms of this kernel the pricing formula (4.18) can be rewritten as

$$V_t = N_t \frac{\mathrm{E}_{\mathbb{P}}\big(L_T(X/N_T)|\,\mathcal{F}_t\big)}{\mathrm{E}_{\mathbb{P}}(L_T|\,\mathcal{F}_t)} = \frac{1}{Z_t}\mathrm{E}_{\mathbb{P}}\big(Z_T X|\,\mathcal{F}_t\big).$$

The first equality is a consequence of the formula for re-expressing a conditional expectation using a different underlying measure.

Other properties of numeraire pairs can also be translated in pricing kernels. In particular, an economy is complete if there exists a unique pricing kernel.[‡]

4.49 Theorem. *Suppose that one of the asset price processes is strictly positive. Then:*

(i) *There exists a pricing kernel if and only if there exists a numeraire pair.*

(ii) *A strategy ϕ is admissible if and only if $\phi \cdot (ZA)$ is a $\mathbb{P}$-martingale for every pricing kernel Z.*

(iii) *The economy is complete if and only if the pricing kernel is unique on $[0,T]$ up to a multiplicative constant (and evanescence).*

Proof. (i). We have already argued that $Z = L/N$ is a pricing kernel if $(\mathbb{N}, N)$ is a numeraire pair and L is the density process of $\mathbb{N}$ relative to $\mathbb{P}$. Conversely, a strictly positive component N of the asset price process is always a numeraire. For a pricing kernel Z the process $L = NZ$, a coordinate of the process AZ, is a $\mathbb{P}$-martingale. If necessary we can multiply N by a constant to ensure that its (constant) mean is equal to 1. Then we can define a probability measure $\mathbb{N}$ by $d\mathbb{N} = (NZ)_T\, d\mathbb{P}$, and L is the density process of $\mathbb{N}$ relative to $\mathbb{P}$. Because $LA^N = (NZ)A^N = ZA$ is a $\mathbb{P}$-martingale, it follows that A^N is an $\mathbb{N}$-martingale, whence $(\mathbb{N}, N)$ is a numeraire pair.

(ii). For a given numeraire pair $(\mathbb{N}, N)$ there exists a pricing kernel Z with $NZ = L$ equal to the density process L of $\mathbb{N}$ relative to $\mathbb{P}$. Conversely, for every pricing kernel Z there exists a constant c and a numeraire pair $(\mathbb{N}, N)$ such that again $N(cZ) = L$. We can always scale the pricing kernel such that $c = 1$.

[†] This process Z is continuous if, for instance, the asset price is a weakly unique solution of the SDE. Then A^c possesses the representation property, so that the martingale L is a stochastic integral relative to A^c. Continuity appears not to be guaranteed without some sort of condition on A, even though Theorem 7.48 of Hunt and Kennedy appears to make this claim.

[‡] In the following theorem we assume that one of the coordinates of the asset price process is strictly positive, so that it can serve as a numeraire. Theorem 7.48 of Hunt and Kennedy makes the same claims without this assumption. The numeraire is constructed by switching from one nonzero asset to another. Should they not at least assume that at any time point there exists a nonzero asset?

The inverse $1/Z$ of a pricing kernel Z is a unit. For any self-financing strategy ϕ unit invariance gives that $\phi \cdot A^N$ is an $\mathbb{N}$-martingale if and only if ϕA^N is an $\mathbb{N}$-martingale if and only if $L\phi A^N = \phi A^{1/Z}$ is a $\mathbb{P}$-martingale if and only if $\phi \cdot A^{1/Z} = \phi \cdot (ZA)$ is a $\mathbb{P}$-martingale.

(iii). Completeness is equivalent to uniqueness of the martingale measure $\mathbb{N}$ corresponding to a numeraire N. If Z_1 and Z_2 are pricing kernels, then the construction of (i) gives two numeraire pairs $(\mathbb{N}_1, N)$ and $(\mathbb{N}_2, N)$ with $NZ_i = L_i$, for L_i the density process of $\mathbb{N}_i$ relative to $\mathbb{P}$. If $\mathbb{N}_1 = \mathbb{N}_2$, then it follows that $L_1 = L_2$ and hence $Z_1 = Z_2$. Conversely, if $(\mathbb{N}_1, N)$ and $(\mathbb{N}_2, N)$ are numeraire pairs, then $Z_i = L_i/N$ are two pricing kernels ($i = 1, 2$). If the pricing kernel is unique, then $L_1 = L_2$ and hence $\mathbb{N}_1 = \mathbb{N}_2$. ∎

Under the condition of the preceding theorem there exists only one pricing kernel in a complete economy and this has relation $Z = cL/N$, for a constant c, to every given numeraire pair $(\mathbb{N}, N)$, for L the density process of $\mathbb{N}$ relative to $\mathbb{P}$.

5
Extended Black-Scholes Models

In this chapter we consider the "standard finance model", an economy consisting of one "risk-free" asset and finitely many other assets. This model is an extension of the Black-Scholes model of Chapter 1 to more than two assets and with greater flexibility in the parameters. Existence of a numeraire pair and completeness of the economy is shown to be equivalent to existence of the "market price of risk" process.

The assets are denoted by (R, S), where R is a special "risk-free" asset and $S = (S^{(1)}, \ldots, S^{(n)})$ is vector-valued. The stochastic process (R, S) is assumed to satisfy the system of stochastic differential equations, for given scalar-, n-vector- and $(n \times d)$-matrix-valued predictable processes r, $\mu = (\mu^{(i)})$ and $\sigma = (\sigma^{(i,j)})$,

(5.1)
$$\begin{aligned} dR_t &= R_t r_t \, dt, \\ dS_t^{(i)} &= S_t^{(i)} \mu_t^{(i)} \, dt + S_t^{(i)} \sum_{j=1}^d \sigma_t^{(i,j)} \, dW_t^{(j)}, \qquad i = 1, \ldots, n. \end{aligned}$$

We shall abbreviate the second equation to $dS_t = S_t \, (\mu_t \, dt + \sigma_t \, dW_t)$, where it is understood that the product of S_t with the vector inside brackets on the right is taken coordinatewise. The significance of factoring the drift and diffusion coefficients as the product of the asset and another process is that this allows to write the solution to the equations in the form

$$\begin{aligned} R_t &= R_0 e^{\int_0^t r_s \, ds}, \\ S_t &= S_0 e^{\int_0^t \mu_s \, ds + \int_0^t \sigma_s \, dW_s - \frac{1}{2} \int_0^t \sigma_s^2 \, ds}. \end{aligned}$$

$\flat$ This follows by the uniqueness of the exponential process as the solution

$\flat$ The exponential function is applied coordinatewise and $\int_0^t \sigma_s^2 \, ds$ is understood to be a vector-valued process with ith coordinate equal to the sum $\Sigma_{j=1}^d \int_0^t (\sigma_s^{(i,j)})^2 \, ds$.

of the Doléans differential equation. As a consequence, the asset processes are strictly positive.

We assume that the process $W = (W^{(1)}, \ldots, W^{(d)})$ is a d-dimensional Brownian motion on a given filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, and interpret the predictability of the processes r, μ and σ relative to the filtration $\mathcal{F}_t$. The process (R, S) is also assumed adapted to the filtration $\mathcal{F}_t$, but it may itself generate a smaller filtration. Thus we can study arbitrage and completeness relative to three filtrations: the "original" filtration $\mathcal{F}_t$, and the augmented natural filtrations generated by (R, S), and W. Unless mentioned otherwise conditions and assertions are understood to be relative to the original filtration $\mathcal{F}_t$.

The "risk-free" asset R may well be a stochastic process, and it may depend on W through the function r: "risk-free" does not mean the same as "deterministic". The difference between the two types of assets is better described by the fact that the risk-free assets are of bounded variation, whereas the risky assets are directly driven by a diffusion term. Nevertheless, it is common to speak of the "risk-free" asset. For brevity we shall also refer to the asset S as the "stocks", even though the model applies equally well to other financial processes, including bonds.

The risk-free asset R is a numeraire for the economy consisting of the asset price process $A = (R, S)$, but without further conditions there does not exist a corresponding martingale measure. We shall show that the existence of a numeraire pair requires the existence of a predictable, vector-valued process $\theta = (\theta^{(1)}, \ldots, \theta^{(d)})$ such that, for Lebesgue almost every t,

$$\sigma_t \theta_t = r_t 1 - \mu_t, \qquad \text{a.s.}. \tag{5.2}$$

(Here $r_t 1$ is the n-dimensional process $(r_t, \ldots, r_t)$.) Furthermore, the existence of such a process θ, called the *market price of risk*, together with some integrability conditions is sufficient for the existence of a martingale measure accompanying the numeraire R, and renders the economy complete relative to the natural filtration generated by (R, S).

5.1 Arbitrage

The existence of the "market price of risk" process requires that the vector $r_t 1 - \mu_t$ is contained in the range space of the $(n \times d)$-matrix σ_t, for almost every t. This is immediate if the rank of σ_t is equal to the number n of stocks in the economy, as the range of σ_t is all of $\mathbb{R}^n$ in that case. If the rank of σ_t is smaller than the number of stocks, then existence of the market price of risk process requires a relationship between the three parameters σ, r and μ. This situation is certain to arise if the number of components of the driving Brownian motion is smaller than the number of risky assets, i.e. $d < n$. Hence we can interpret the condition of existence of a process θ as in the preceding display as implying that the "random inputs" $W^{(i)}$ to the market should be at least as numerous as the (independent) risky assets".

Writing a portfolio for the asset $A = (R, S)$ in the form (ψ, ϕ), we can express the value process $V_t = \psi_t R_t + \phi_t S_t$ of a self-financing strategy (ψ, ϕ) as

$$V_t - V_0 = (\psi \cdot R)_t + (\phi \cdot S)_t = \int_0^t \psi_s R_s r_s \, ds + \int_0^t \phi_s \, dS_s$$
$$= \int_0^t V_s r_s \, ds + \int_0^t \phi_s \, (dS_s - S_s r_s \, ds).$$

By the partial integration formula and the fact that $dR_t = r_t R_t \, dt$,

$$(5.3) \qquad d\Big(\frac{V}{R}\Big)_t = -V_t \frac{1}{R_t^2} \, dR_t + \frac{1}{R_t} \, dV_t = \frac{1}{R_t} \phi_t \, (dS_t - S_t r_t \, dt),$$

by the preceding display. Hence the discounted value process takes the form, in view of (5.1),

$$(5.4) \qquad \frac{V_t}{R_t} = \frac{V_0}{R_0} + \int_0^t \frac{S_s \phi_s}{R_s} \big(\sigma_s \, dW_s - (r_s 1 - \mu_s) \, ds\big).$$

[♯] This formula does not make explicit reference to the amount ψ invested in the numeraire, which has been eliminated. A "partial strategy" ϕ defines a discounted gain process $V_t/R_t - V_0/R_0$ through the right side of (5.3), and given ϕ we can define a process ψ from the equation $V_t = \psi_t R_t + \phi_t S_t$. By retracing the calculations the resulting strategy (ψ, ϕ) can be seen to be self-financing and to possess value process V_t.

Nonexistence of a market price of risk process implies that the vector $r_t 1 - \mu_t$ is not contained in the range of σ_t, for a positive set of times t. Then there exists a vector ϕ_t such that $(S\phi)_t$ is orthogonal to this range (i.e. $(S\phi)_t \sigma_t = 0$) and such that the inner product $(S\phi)_t (r_t 1 - \mu_t)$ is strictly negative. We can arrange it so that the latter inner product is never positive and hence, by the preceding display, the corresponding discounted gain process will be zero at time 0 and strictly positive at time T. This suggests that nonexistence of a market price of risk process creates a potential for arbitrage. Because we can scale ϕ appropriately, this is not caused by a lack of integrability, i.e. inadmissibility. The only possibility is that there does not exist a numeraire pair. Lemma 5.5 below makes this reasoning rigorous.

[♯] In the multivariate case the process $S\phi$ in the display is understood to be $(1 \times n)$-vector-valued with coordinates $S^{(i)} \phi^{(i)}$; it is multiplied as a vector versus the $(n \times d)$-matrix σ_t giving a $(1 \times d)$ vector which is next multiplied with the $(d \times 1)$ vector W; it is also multiplied with the $(n \times 1)$ vector $r1 - \mu$.

On the other hand, if the market price of risk process θ exists, then the gains process in the preceding display can be written as a stochastic integral relative to the process $\sigma \cdot \tilde{W}$, for

$$\tilde{W}_t = W_t - \int_0^t \theta_s \, ds.$$

By Girsanov's theorem, the process $\tilde{W}$ will be a Brownian motion after an appropriate change of measure, and hence the discounted gains process will be a local martingale. The new measure will be the martingale measure corresponding to the numeraire R.

5.5 Lemma. *If there exists a numeraire pair, then there exists a predictable process θ with values in $\mathbb{R}^d$ such that $\sigma_t\theta_t = r_t 1 - \mu_t$ for Lebesgue almost all t, almost surely.*

Proof. Let ϕ_t' be the orthogonal projection of the vector $r_t 1 - \mu_t$ onto the orthocomplement of the range of σ_t. Define a process ϕ by setting $\phi_t = -(\phi_t'/S_t)\alpha_t$ for a given scalar, positive process α_t, where the quotient is interpreted coordinatewise. Then $(S\phi)_t\sigma_t = 0$ and $-(S\phi)_t(r_t 1 - \mu_t) = \|\phi_t'\|^2\alpha_t > 0$ for every t such that $r_t 1 - \mu_t$ is not contained in the range of σ_t. It can be shown that $\phi_t' = f(r_t 1 - \mu_t, \sigma_t)$ for a measurable map $f\colon \mathbb{R}^n \times \mathbb{R}^{nd} \to \mathbb{R}^n$ and hence the process ϕ is predictable. (Cf. Karatzas and Shreve I.6.9, p26.) The strategy ϕ can be made suitably integrable, by choice of α.

The process ϕ defines a discounted gains process V through (5.4), given by $V_t/R_t = \int_0^t (1/R_s)\|\phi_s'\|^2\alpha_s \, ds \geq 0$ with strict inequality for $t = T$ unless $r_t 1 - \mu_t$ is contained in the range of σ_t for almost all t. The corresponding strategy (ψ, ϕ) is self-financing and possesses value process $V_t = \psi_t R_t + \phi_t S_t$. If there existed a numeraire paire $(\mathbb{N}, N)$, then by self-financing and unit invariance the process V^N would be a local martingale and hence a supermartingale, because it is bounded below by 0. Because $V_0 = 0$ and $V_T \geq 0$ we must have $V_t = 0$ with probability 1 and hence $\int_0^T (1/R_s)\|\phi_s'\|^2\alpha_s \, ds = 0$ almost surely. This implies that $\phi_t' = 0$, whence $r_t 1 - \mu_t$ is in the range of σ_t almost surely, for almost all t. ∎

5.6 Lemma. *If there exists a predictable process θ with values in $\mathbb{R}^d$ such that $\sigma_t\theta_t = r_t 1 - \mu_t$ for Lebesgue almost all t and* $\mathrm{E}\exp\frac{1}{2}\int_0^T \|\theta_s\|^2 \, ds < \infty$, *then there exists a local martingale measure $\mathbb{R}$ corresponding to R. If moreover* $\mathrm{E}\exp\frac{1}{2}\int_0^T \|\sigma_s + \theta_s\|^2 \, ds < \infty$, *then there exists a martingale measure $\mathbb{R}$ as well.*

Proof. Because Novikov's condition is satisfied by assumption, the exponential process $\mathcal{E}(\theta \cdot W)$ is a martingale, and hence the measure $\mathbb{R}$ defined by $d\mathbb{R} = \mathcal{E}(\theta \cdot W)_T \, d\mathbb{P}$ is a probability measure. By Girsanov's theorem the process $\tilde W$ defined by $\tilde W_t = W_t - \int_0^t \theta_s\,ds$ is a Brownian motion process under $\mathbb{R}$. (See Exercise 2.21 for the vector-valued case.)

The solutions R and S to (5.1) can be written in exponential form, and consequently the discounted stock process can be represented as

$$\frac{S_t}{R_t} = e^{\int_0^t (\mu_s - r_s 1)\,ds + \int_0^t \sigma_s\,dW_s - \frac{1}{2}[\sigma\cdot W]_t}.$$

In view of the property (5.2) of the market price of risk θ, the right side is equal to,

$$e^{-\int_0^t \sigma_s\theta_s\,ds + \int_0^t \sigma_s\,dW_s - \frac{1}{2}[\sigma\cdot W]_t} = e^{(\sigma\cdot\tilde W)_t - \frac{1}{2}[\sigma\cdot W]_t} = \mathcal{E}(\sigma\cdot\tilde W)_t.$$

It follows that the discounted stock process S/R is an $\mathbb{R}$-local martingale, and hence $(\mathbb{R}, R)$ is a numeraire pair.

Because S/R is nonnegative it is also an $\mathbb{R}$-supermartingale and hence is an $\mathbb{R}$-martingale if its mean is constant. The mean process can be written as

$$\mathrm{E}_{\mathbb{R}}\mathcal{E}(\sigma\cdot\tilde W)_t = \mathrm{E}_{\mathbb{P}}\mathcal{E}(\sigma\cdot\tilde W)_t\mathcal{E}(\theta\cdot W)_t = \mathrm{E}_{\mathbb{P}}\mathcal{E}\big((\sigma+\theta)\cdot W\big)_t.$$

The assumptions imply that Novikov's condition is satisfied for the exponential process in the right side, and hence the process $\mathcal{E}\big((\sigma+\theta)\cdot W\big)$ is a $\mathbb{P}$-martingale, whence its mean function is constant. We conclude that S/R is an $\mathbb{R}$-martingale. ∎

5.2 Completeness

In this section we assume the existence of the market price of risk process and the integrability conditions of Lemma 5.6, so that there exists a local martingale measure to the numeraire R. We consider the completeness of the economy both relative to the augmented filtration generated by the asset processes, and relative to the augmented filtration generated by the driving Brownian motion. We assume that the first filtration is right-continuous (as is the second filtration automatically).

5.7 Theorem. *Suppose that the conditions of Lemma 5.6 hold. If the stochastic differential equation (5.1) is given by processes $r_t = r(t, R_t, S_t)$, $\mu_t = \mu(t, R_t, S_t)$ and $\sigma_t = \sigma(t, R_t, S_t)$, for measurable functions r, μ and σ, and possesses a weakly unique solution, then there exists a numeraire pair and the economy is complete relative to the augmented natural filtration generated by the process (R, S).*

Proof. By Lemma 5.6 there exists a numeraire pair. Because the process R is a numeraire of bounded variation, the completeness of the economy follows from Theorem 4.31 upon noting that the martingale part $(0, S\,\sigma\cdot W)$ of the asset process possesses the representing property by Theorem 3.14.

An alternative proof that does not use Theorem 4.31 is to note first that the process $\sigma\cdot W$ possesses the representing property for $\mathbb{P}$-local martingales, by Theorem 3.14. As shown in the proof of Lemma 5.6 the discounted stock process can be written as $S/R = \mathcal{E}(\sigma\cdot\tilde W)$, so that $d(S/R) = (S/R)\sigma\, d\tilde W$. Here $\sigma\cdot\tilde W$ inherits the representing property for $\mathbb{R}$-local martingales from the representing property of $\sigma\cdot W$, by Lemma 3.7, and hence S/R has the representing properrty, by Lemma 3.9. ■

5.8 EXERCISE. Show that a complete "extended Black-Scholes" economy can have at most one risk-free asset: if the ith row of σ in (5.1) is identically zero, then $S^{(i)} = S_0^{(i)}R$.

5.9 EXERCISE. Consider the economy consisting of a risk-free asset $R_t = e^{rt}$ and two additional assets $S^{(1)}$ and $S^{(2)}$ such that $dS_t^{(i)} = S_t^{(i)}(\mu_t^{(i)}\,dt + \sigma_t^{(i)}\,dW_t)$ for a one-dimensional Brownian motion W. For what parameter values is this economy complete?

5.10 EXERCISE. Suppose that in Theorem 5.7 we replace the condition that the stochastic differential equation (5.1) possesses a weakly unique solution by the assumption that the equation $dS_t = S_t(r_t 1\,dt + \sigma_t\,d\tilde W_t)$ possesses a weakly unique solution, for $\tilde W_t = W_t - \int_0^t \theta_s\,ds$. Is the theorem still valid?

The natural filtration generated by the asset processes, as used in the preceding theorem, is probably the most natural filtration for use in connection with completeness. An alternative, used by many authors, is the augmented filtration $\mathcal{F}_t^W$ generated by the driving Brownian motion. Because this may be bigger than the natural filtration of the asset price processes, completeness relative to $\mathcal{F}_t^W$ may be more restrictive. The following theorem requires the existence of a market price of risk process, and also that the number of risky assets is not smaller than the number of driving Brownian motions (i.e. $n \ge d$). In contrast (by Theorem 5.7) relative to the filtration generated by the assets the market is typically complete as soon as the market price of risk process exists, which roughly requires that the number of risky assets is no larger than the number of driving Brownian motions (i.e. $n \le d$). Together this forces $n = d$.

The difference between the two set-ups is seen easily by adding an additional "driving" Brownian motion to the model, but letting it not drive anything, by setting the corresponding column of σ equal to zero. It will be impossible to replicate claims that are a function of this extra Brownian motion using a portfolio consisting only of the asset price processes. If this were possible, then the extra Brownian could be written as a stochastic integral relative to the other Brownian motions, which is not possible.

5.11 Theorem. *Suppose that the conditions of Lemma 5.6 hold. If the number of stocks is equal to the dimension of the Brownian motion W and the process σ takes its values in the invertible matrices, then there exists a numeraire pair and the economy is complete relative to the augmented natural filtration $\mathcal{F}_t^W$ generated by W.*

Proof. By Lemma 5.6 there exists a numeraire pair. Because the process R is a numeraire of bounded variation, the completeness of the economy follows from Theorem 4.31 upon noting that the martingale part $(0, S\,\sigma\cdot W)$ of the asset process possesses the representing property by Lemma 3.9. ∎

5.3 Partial Differential Equations

Under the conditions of Theorem 5.7, the process $\tilde W$ defined by $\tilde W_t = W_t - \int_0^t \theta_s\,ds$ is a Brownian motion under the martingale measure $\mathbb{R}$ corresponding to the numeraire R. Because option prices can be written as expectations under $\mathbb{R}$, it is useful to rewrite the system of stochastic differential equations (5.1) in terms of the process $\tilde W$. If we also assume that the processess r and σ take the forms $r_t = \bar r(t, R_t, S_t)$ and $\sigma_t = \sigma(t, R_t, S_t)$, then the equations take the form

$$
\begin{aligned}
dR_t &= R_t\,\bar r(t, R_t, S_t)\,dt,\\
dS_t &= S_t\,\bar r(t, R_t, S_t)1\,dt + S_t\,\sigma(t, R_t, S_t)\,d\tilde W_t.
\end{aligned}
\tag{5.12}
$$

As usual we assume that (R, S) is adapted to the augmented natural filtration $\mathcal{F}_t^{\tilde W}$ of $\tilde W$. Then, under regularity conditions on r and σ, the process (R, S) will be Markovian relative to this filtration. If we assume in addition that σ is invertible, then $\tilde W$ can be expressed in (R, S) by inverting the second equation, and hence the filtrations $\mathcal{F}_t$ and $\mathcal{F}_t^{\tilde W}$ generated by (R, S) and $\tilde W$ are the same. The process (R, S) is then Markovian relative to its own filtration $\mathcal{F}_t$. In that case a conditional expectation of the type $\mathrm{E}_{\mathbb{R}}(X|\,\mathcal{F}_t)$ of a random variable X that is a measurable function of $(R_s, S_s)_{s\geq t}$ can be written as $F(t, R_t, S_t)$ for a measurable function F.

This observation can be used to characterize the value processes of certain options through a partial differential equation. The value process of a claim that is a function $X = g(S_T)$ of the final value S_T of the stocks takes the form

$$V_t = R_t \mathrm{E}_{\mathbb{R}}\Bigl(\frac{g(S_T)}{R_T}|\,\mathcal{F}_t\Bigr) = \mathrm{E}_{\mathbb{R}}\Bigl(e^{-\int_t^T \bar r(s,R_s,S_s)\,ds} g(S_T)|\,\mathcal{F}_t\Bigr).$$

If the process (R, S) is Markovian as in the preceding paragraph, then we can write $V_t = F(t, R_t, S_t)$ for a measurable function F. We assume that this function possesses continuous partial derivatives up to the second order. For simplicity of notation we also assume that S is one-dimensional. Then, by Itô's formula,

$$dV_t = F_t\,dt + F_r\,dR_t + F_s\,dS_t + \tfrac{1}{2}F_{ss}\,d[S]_t.$$

Here F_t, F_r, F_s are the first order partial derivatives of F relative to its three arguments, F_{ss} is the second order partial derivative relative to its third argument, and for brevity we have left off the argument (t, R_t, S_t) of these functions. (Beware of the different meaning of the subscript t in expressions such as F_t, R_t or S_t!) A second application of Itô's formula and substitution of the diffusion equation for (R, S) yields

$$d\Bigl(\frac{V_t}{R_t}\Bigr) = \frac{1}{R_t}\Bigl(-F\bar{r} + F_t + F_r R_t \bar{r} + F_s S_t r + \tfrac{1}{2}F_{ss}S_t^2\sigma^2\Bigr)\,dt + \frac{1}{R_t}F_s \mathcal{S}_t \sigma\,d\tilde{W}_t.$$

The process V_t/R_t is the discounted value process of an admissible, self-financing strategy (replicating the claim $X = g(S_T)$) and hence is an $\mathbb{R}$-local martingale. Because the process $\tilde{W}$ is a Brownian motion, this can only be true if the drift term on the right side of the preceding display is zero, i.e.

$$\begin{aligned} &-(F\bar{r})(t,r,s) + F_t(t,r,s) + (\bar{r}F_r)(t,r,s)r \\ &\qquad + (\bar{r}F_s)(t,r,s)s + \tfrac{1}{2}(\sigma^2 F_{ss})(t,r,s)s^2 = 0. \end{aligned}$$

This is the *Black-Scholes partial differential equation.* It was obtained by Black and Scholes in 1973 within the context of the simple Black-Scholes model of Chapter 1, strangely enough by a reasoning that appears not quite correct from today's standpoint.

The preceding derivation is closely related to derivation of the *Feynman-Kac formula*, which, however, goes in the opposite direction. This formula expresses the value of a given function that satisfies a certain partial differential equation as the expectation of a certain function of a Brownian motion. (Cf. e.g. Karatzas and Shreve, Chapter 4.4.)

The Black-Scholes partial differential equation is useful for the numerical computation of option prices. Even though the equation is rarely explicitly solvable, a variety of numerical methods permit to approximate the solution F. The equation depends only on the functions $\bar{r}$ and σ defining the stochastic differential equation (5.12). Hence it is the same for every option with a claim of the type $X = g(S_T)$, the form of the claim only coming in to determine the boundary condition. Because $X = g(S_T) = F(T, R_T, S_T)$, this takes the form

$$F(T, r, s) = g(s).$$

For instance, for a European call option on the stock S, this becomes $F(T, r, s) = (s - K)^+$.

The preceding approach is possible only if the value process V is a smooth function $V_t = F(t, R_t, S_t)$ of time and the underlying assets. The regularity of the value process depends on the functions r and σ, determining the evolution of the assets. In the simple Black-Scholes model where these functions are constants the smoothness can be verified. In that case, the conditional distribution of the variable

$$S_T = S_t e^{\sigma(\tilde{W}_T - \tilde{W}_t) - \frac{1}{2}\sigma^2(T-t)}$$

given $\mathcal{F}_t$ is log normal, and we can write the value process $V_t = R_t \mathrm{E}_{\mathbb{R}}(X/R_T |\, \mathcal{F}_t)$ for the claim $X = g(S_T)$ in the form $V_t = F(t, S_t)$, for the function, with Z a standard normal variable,

$$F(t,s) = e^{-r(T-t)} \mathrm{E} g\bigl(s e^{\sigma\sqrt{T-t}\, Z - \frac{1}{2}\sigma^2(T-t)}\bigr).$$

5.13 EXERCISE. Verify that the function F is infinitely smooth in both t and s.

6
American Options

An *American option*, as opposed to a European option, is a contract derived from a *claim process* X, which we take to be a cadlag semimartingale. The contract entitles the holder to a cash payment X_τ at a time τ chosen by the holder of the contract, restricted to be a stopping time with values in an $[0, T]$, which is fixed in the contract. A standard example is an American option on a stock S, which has claim process $X = (S - K)^+$.

Naturally, the holder of the contract tries to maximize the payment X_τ by a clever choice of the stopping time τ. The value of the contract turns out to be the solution to an optimal stopping problem. However, as for European options, the problem must be formulated relative to a martingale measure, not the measure of the true world. This surprising fact is again the result of a no-arbitrage argument. We shall first express the value of an American option in the value of a "replicating portfolio" through a no-arbitrage argument, and next relate this value to an optimal stopping problem. We work in the finite economy model described in Chapter 4.

6.1 Replicating Strategies

Suppose that there exists a self-financing strategy ϕ with $\phi A \geq X$ throughout $[0, T]$ and $\phi_\tau A_\tau = X_\tau$ for some stopping time τ with values in $[0, T]$. Then the fair price of the American contract X at time 0 is $\phi_0 A_0$. We corroborate this by an economic argument.

Because $\phi A \geq X$ and we can sell our portfolio at any given stopping time, just as we are allowed to cash X at any stopping time, it is obvious that we should prefer the portfolio over the contract X. This shows that the American contract X is not worth more than $\phi_0 A_0$ at time 0.

Conversely, if the value of the option were strictly less than $\phi_0 A_0$, then we could create a positive cash flow at time 0 by buying the American contract X and selling the portfolio ϕ_0. Until the time τ we could reshuffle the portfolio $-\phi_0$ according to the strategy $-\phi$, which can be achieved free of money input. At time τ we could cash the American claim X_τ, and buy the portfolio ϕ_τ. Because $\phi_\tau A_\tau = X_\tau$ by assumption, this would leave us with no money and no obligations, except for the money cashed at time 0. Assuming that it is impossible to create certain profit, we conclude that the initial assumption is wrong and hence the just price of the American claim at time 0 is $\phi_0 A_0$.

We can repeat this argument to find the value at other times $t \in [0,T]$. If there exists a self-financing portfolio ϕ such that $\phi A \geq X$ on $[t,T]$ and $\phi_\tau A_\tau = X_\tau$ for some stopping time τ taking its values in $[t,T]$, then the just price of the American claim X is equal to $\phi_t A_t$. Here the replicating portfolio ϕ is permitted to (and will typically) depend on t.

The preceding no-arbitrage argument is not a proof of a mathematical theorem, but it appears to be convincing. One serious attack on the validity of this reasoning is that the strategies ϕ that it is based on may not be unique. This challenge to the argument is solved if we insist that the replicating portfolio be admissible, in view of the following lemma, which shows that the value derived from a replicating portfolio is unique.

6.1 Lemma (Unique value). *If ϕ and ψ are self-financing admissible strategies with $\phi A \geq X$ and $\psi A \geq X$ on $[t,T]$ and such that $\phi_\sigma A_\sigma = X_\sigma$ and $\psi_\tau A_\tau = X_\tau$ for stopping times σ and τ that take their values in $[t,T]$, then $\phi_t A_t = \psi_t A_t$.*

Proof. The existence of admissible strategies implies the existence of a numeraire pair $(\mathbb{N}, N)$. By the definition of admissibility, self-financing and unit invariance, the processes ϕA^N and ψA^N are $\mathbb{N}$-martingales. By the optional stopping theorem applied to the identity $(\psi A^N)_\tau = X_\tau^N$, we see that $(\psi A^N)_{\sigma\wedge\tau} = \mathrm{E}_{\mathbb{N}}\big(X_\tau^N \,|\, \mathcal{F}_\sigma\big)$. By assumption the right side is bounded above by $\mathrm{E}_{\mathbb{N}}\big((\phi A^N)_\tau \,|\, \mathcal{F}_\sigma\big) = (\phi A^N)_{\sigma\wedge\tau}$, by the optional stopping theorem. Thus $(\psi A^N)_{\sigma\wedge\tau} \leq (\phi A^N)_{\sigma\wedge\tau}$, and by symmetry also equality. This shows that the $\mathbb{N}$-martingales ψA^N and ϕA^N are equal at $\sigma \wedge \tau$. By optional stopping they are the same on the interval $[0, \sigma\wedge\tau]$, which includes the point t. ∎

We shall call a *replicating strategy from t* for the claim process X an admissible, self-financing strategy ϕ such that $\phi A \geq X$ on $[t,T]$ and $\phi_\tau A_\tau = X_\tau$ for some stopping time τ taking its values in $[t,T]$.

Our next aim is to give a more concrete expression of the value of the American claim through the solution of an optimal stopping problem. We first recall some generalities on optimal stopping. Let $\mathcal{T}$ be the set of all stopping times with values in $[0,T]$.

6.2 Optimal Stopping

Suppose that X is a nonnegative, cadlag adapted process, indexed by $[0, T]$, on a given filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ such that[†]

$$\mathrm{E} \sup_{0 \le t \le T} X_t < \infty. \tag{6.2}$$

The optimal stopping problem is concerned with finding the stopping time τ that maximizes $\mathrm{E}X_\tau$ over all stopping times with values in $[0, T]$, and finding the value of the maximum.

If the time set were a finite set $0 < t_1 < t_2 < \cdots < t_n < T$, rather than a continuum, then the solution of this problem could be found easily through *backwards programming.* This consists of recursively optimizing over all stopping times with values in $\{t_n\}, \{t_{n-1}, t_n\}, \ldots$. The first step is to decide for any possible realization of X_{t_n} whether it is better to stop at t_n (and collect X_{t_n}) or to continue (and collect X_T). Given the best strategy from time t_n onwards, encoded in a stopping time τ_n with values in $\{t_n, T\}$, the second step is to decide at time t_{n-1} whether to stop (and collect $X_{t_{n-1}}$) or to continue, in which case we would follow the optimal strategy τ_n. This scheme continues backwards in time to time t_1. Section 6.4 describes the method in more detail. The essential object to formalize the algorithm is the "Snell envelope", whose value at k gives the optimal (conditional) expectation when the stopping times are restricted to $\{t_k, \ldots, t_n\}$.

The continuous time solution is similar, but technically much more complicated. The *Snell envelope* is the smallest cadlag supermartingale Z with $Z \ge X$, where "smallest" can be understood in that $Z \le Z'$ up to evanescence for every other supermartingale Z' with $Z' \ge X$. There exists a version of this supermartingale such that, for every t,

$$Z_t = \operatorname{ess\,sup}\bigl\{\mathrm{E}(X_\tau \mid \mathcal{F}_t) : \tau \ge t, \tau \in \mathcal{T}\bigr\}, \tag{6.3}$$

where the supremum is taken over all stopping times τ with values in $[t, T]$. (Equation (6.3) is even true for t a stopping time.) We can view the variable in the display as the optimal expected payoff X_τ that can be achieved by stopping the payoff process in $[t, T]$, seen from the time t point of view. In particular, if $\mathcal{F}_0$ is the trivial σ-field, then Z_0 is the maximal value of $\mathrm{E}X_\tau$ over all stopping times τ.

If X is continuous, then there exists an optimal stopping time, given by the first time that the Snell envelope Z and the process X coincide. More generally, we have $Z_t = \mathrm{E}(X_{\tau_t} \mid \mathcal{F}_t)$ for τ_t the stopping time

$$\tau_t = \inf\{s \ge t : Z_s = X_s\}.$$

Thus τ_t is the optimal stopping time restricted to $[t, T]$ evaluated from the perspective at time t, conditionally on the past. By (right) continuity

[†] Is it necessary to assume that $\mathcal{F}_0$ is trivial?

$Z_{\tau_t} = X_{\tau_t}$ and $Z_s > X_s$ on $[t, \tau_t)$. Thus an optimal time to stop in $[t, T]$ is the first time in this interval that the Snell envelope and the process X coincide.

The Snell envelope can be shown to be of class D and hence possesses a Doob-Meyer decomposition $Z = M - \Lambda$ for a uniformly integrable martingale M and a nondecreasing, cadlag, adapted process Λ with $\Lambda_0 = 0$ and $\mathrm{E}\Lambda_T < \infty$. It will be important to know that $\Lambda_{\tau_t} = \Lambda_t$. We collect some relevant properties in the following proposition.

6.4 Proposition. *Let X be a continuous adapted, nonnegative process with $\mathrm{E}\sup_{0\le t\le T} X_t < \infty$. Then there exists a cadlag supermartingale Z of class D satisfying (6.3) and $Z_t = \mathrm{E}(X_{\tau_t}|\,\mathcal{F}_t)$ for $\tau_t = \inf\{s \ge t\colon Z_s = X_s\}$. The process Z can be decomposed as $Z = M - \Lambda$ for a uniformly integrable martingale M and a continuous, nondecreasing process Λ satisfying $\Lambda_{\tau_t} = \Lambda_t$. For each t the process Z restricted to $[t, T]$ is the smallest supermartingale with $Z \ge X$.*

Proof. See e.g. Karatzas and Shreve, Appendix D. ∎

6.3 Pricing and Completeness

We apply the optimal stopping theory to the rebased process X^N for a given numeraire pair $(\mathbb{N}, N)$, and X the claim process of an American option. We assume that the claim process X is nonnegative. The following theorem shows that the fair price process of the option is given by the Snell envelope of the process X^N on the filtered probability space $(\Omega, \mathcal{F}, \{F_t\}, \mathbb{N})$. At least this is true if the "fair price" can be defined through a replicating strategy, as previously. (Otherwise, we do not have a notion of "fair price".) If the market is complete (in the sense of Chapter 4), then this is true for every claim such that $\mathrm{E}_{\mathbb{N}}\sup_{0\le t\le T}(X_t/N_t) < \infty$.

6.5 Theorem. *If ϕ is a replicating strategy for t for the continuous, nonnegative claim process X, then, for any numeraire pair $(\mathbb{N}, N)$ such that $\mathrm{E}_{\mathbb{N}}\sup_t(X_t/N_t) < \infty$,*

$$\phi_t A_t = N_t \operatorname{ess\,sup}\Big\{\mathrm{E}_{\mathbb{N}}\Big(\frac{X_\tau}{N_\tau}\,|\,\mathcal{F}_t\Big)\colon \tau \ge t, \tau \in \mathcal{T}\Big\}.$$

6.6 Theorem (Completeness). *If the market is complete with numeraire pair $(\mathbb{N}, N)$, then for every continuous, nonnegative claim process X such that $\mathrm{E}_{\mathbb{N}}\sup_t(X_t/N_t) < \infty$, there exists for every $t \in [0, T]$ a self-financing, admissible strategy ϕ such that $\phi A \ge X$ on $[t, T]$ and $\phi_\tau A_\tau = X_\tau$ for some stopping time τ with values in $[t, T]$.*

Proofs. Let Z be the Snell envelope of the process $X^N = X/N$ on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{N})$, and let $Z = M - \Lambda$ be its Doob-Meyer decomposition.

Let ϕ be a replicating strategy as in the first theorem. Because ϕA^N is an $\mathbb{N}$-martingale with $\phi A^N \geq X^N$ on $[t, T]$, and the restriction of Z to $[t, T]$ is the minimal $\mathbb{N}$-super martingale with $Z \geq X^N$ on $[t, T]$, it follows that $\phi A^N \geq Z$ on this interval. On the other hand, because ϕ is a replicating strategy, there exists a stopping time τ with values in $[t, T]$, such that $\phi_\tau A_\tau^N = X_\tau^N \leq Z_\tau$, whence by the optimal stopping theorem, $\phi_t A_t^N = \mathrm{E}_{\mathbb{N}}\bigl(\phi_\tau A_\tau^N |\, \mathcal{F}_t\bigr) \leq \mathrm{E}_{\mathbb{N}}(Z_\tau |\, \mathcal{F}_t) \leq Z_t$, by the supermartingale property of Z and the optional stopping theorem. By combining we see that $\phi_t A_t^N = Z_t$, which is the assertion of the first theorem.

By completeness there exists a replicating strategy for the claim $N_T(M_T - \Lambda_t)$, i.e. an admissible, self-financing strategy ϕ with $N_T(M_T - \Lambda_t) = \phi_T A_T$, or, equivalently, $M_T - \Lambda_t = \phi_T A_T^N$. The left- and right sides of this equation are the values at T of $\mathbb{N}$-martingales on $[t, T]$. Taking conditional expectations, we conclude that $M - \Lambda_t = \phi A^N$ on the interval $[t, T]$. This shows that $\phi A^N = Z + \Lambda - \Lambda_t \geq Z \geq X^N$ on $[t, T]$, because Λ is nondecreasing, whence $\phi A \geq X$ on $[t, T]$. Furthermore, for $\tau_t = \inf\{s \geq t\colon Z_s = X_s^N\}$ we have that $(\phi A^N)_{\tau_t} = Z_{\tau_t} + \Lambda_{\tau_t} - \Lambda_t = Z_{\tau_t} = X_{\tau_t}^N$, by Proposition 6.4. This shows that $(\phi A)_{\tau_t} = X_{\tau_t}$. ∎

The American option with claim process X is worth more than the European option with claim X_T, as the latter can be viewed as the American contract with the restriction that early stopping is not allowed. The difference in value is

$$N_t\biggl(\operatorname{ess\,sup}\Bigl\{\mathrm{E}_{\mathbb{N}}\Bigl(\frac{X_\tau}{N_\tau} |\, \mathcal{F}_t\Bigr)\colon \tau \geq t, \tau \in \mathcal{T}\Bigr\} - \mathrm{E}_{\mathbb{N}}\Bigl(\frac{X_T}{N_T} |\, \mathcal{F}_t\Bigr)\biggr).$$

This is strictly positive in general. A case of interest where the values are the same is when the process X/N is an $\mathbb{N}$-submartingale. Then, by optional stopping, $\mathrm{E}_{\mathbb{N}}\bigl(X_\tau/N_\tau |\, \mathcal{F}_t\bigr) \leq \mathrm{E}_{\mathbb{N}}\bigl(X_T/N_T |\, \mathcal{F}_t\bigr)$ almost surely for every stopping time $\tau \geq t$, and hence early stopping yields no advantage. (On the other hand, it need not be detrimental, because an optimal stopping time need not be unique, and hence early stopping may be preferable for other reasons.)

6.7 Example (American call option). The claim process of an *American call option* on an asset S is given by $X = (S - K)^+$, for K a constant fixed in the contract, referred to as the *strike price*. The claim value $X_T = (S_T - K)^+$ at expiry time is the value of a European call option on S.

For American call options stopping is typically not helpful and the value of an American call option is the same as the value of a European call option. This is the case if there exists a numeraire pair $(\mathbb{N}, N)$ with $\mathbb{N}$ a martingale measure and such that $1/N$ is an $\mathbb{N}$-supermartingale. In particular, this is true if N is nondecreasing, as is the case for instance for the standard numeraire in the Black-Scholes model.

To see this it suffices to show that the rebased claim process X/N is an $\mathbb{N}$-submartingale. By Jensen's inequality, for $s < t$,

$$\mathrm{E}_{\mathbb{N}}\Bigl(\frac{(S_t - K)^+}{N_t}\,|\,\mathcal{F}_s\Bigr) \geq \Bigl(\mathrm{E}_{\mathbb{N}}\Bigl(\frac{S_t - K}{N_t}\,|\,\mathcal{F}_s\Bigr)\Bigr)^+ = \Bigl(\frac{S_s}{N_s} - K\mathrm{E}_{\mathbb{N}}\Bigl(\frac{1}{N_t}\,|\,\mathcal{F}_s\Bigr)\Bigr)^+,$$

because S/N is an $\mathbb{N}$-martingale. If $1/N$ is a supermartingale, then this can be further bounded below by $(S_s - K)^+/N_s$. $\square$

6.8 EXERCISE. Suppose that the economy is complete. Show that if $(\mathbb{N}, N)$ is a numeraire pair such that $1/N$ is an $\mathbb{N}$-supermartingale, then $1/M$ is a $\mathbb{M}$-supermartingale for every numeraire pair $(\mathbb{M}, M)$.

6.9 EXERCISE. Express the difference in value between the American and European options as $N_t\bigl(\mathrm{E}(\Lambda_T|\,\mathcal{F}_t) - \Lambda_t\bigr)$ for Λ the nondecreasing process in the Doob-Meyer decomposition $Z = M - \Lambda$ of the Snell envelope Z of X^N on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{N})$.

If the prices of European and American options do not agree, then the computation of the value of an American option may not be easy, and explicit formulas are rarely (or never) available. Typically the optimal stopping problem can be rewritten as a variational problem, but this then needs to be solved numerically. An alternative is to discretize time and compute the Snell envelope of te discretized claim process, possibly by stochastic simulation.

6.4 Optimal Stopping in Discrete Time

For an intuitive understanding (and possibly also for numerical implementation) it is helpful to consider the optimal stopping problem in discrete time. Given a time set $0 < t_1 < t_2 < \cdots < t_N = T$ and integrable random variables $X_1, \ldots, X_N$ adapted to a filtration $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots \subset \mathcal{F}_N$, we wish to compute the supremum $\sup_{\tau\in\mathcal{T}} \mathrm{E}X_\tau$ of $\mathrm{E}X_\tau$ over the set $\mathcal{T}$ of all stopping times τ with values in $\{t_1, \ldots, t_N\}$, and also the suprema $\sup_{\tau\in\mathcal{T}:\tau\geq k} \mathrm{E}(X_\tau|\,\mathcal{F}_k)$ for stopping after t_k.

The optimal stopping time is computed backwards in time, starting at time t_N. If we do not stop at times $t_1 < \cdots < t_{N-1}$, then we must stop at time t_N, resulting in a payment X_N. We encode this by defining a random variable and stopping time by

$$Z_N = X_N, \qquad \tau_N = N.$$

Next consider an optimal strategy if we stop either at time $N-1$ or at time N. If we stop at time $N-1$, then we receive payment X_{N-1}, whereas if we do not stop we receive payment Z_N, whose expected value from time $N-1$ perspective is $\mathrm{E}(Z_N|\,\mathcal{F}_{N-1})$. To maximize the expected payment from the time $N-1$ perspective we thus decide to stop if $X_{N-1}\geq \mathrm{E}(Z_N|\,\mathcal{F}_{N-1})$ and to continue otherwise. This gives expected payment from time $N-1$ perspective equal to $X_{N-1}\vee\mathrm{E}(Z_N|\,\mathcal{F}_{N-1})$. We encode this optimal strategy in the random variable and stopping time given by

$$Z_{N-1}=X_{N-1}\vee \mathrm{E}(Z_N|\,\mathcal{F}_{N-1}),\qquad \tau_{N-1}=\begin{cases} N-1 & \text{if } Z_{N-1}=X_{N-1}\\ \tau_N & \text{if } Z_{N-1}>X_{N-1}\end{cases}.$$

Next we proceed to time $N-2$. If we do not stop at this time, then the best strategy is to play the optimal strategy τ_{N-1} from time $N-1$ onwards, which has expected pay-off $\mathrm{E}(Z_{N-1}|\,\mathcal{F}_{N-2})$ from the time $N-2$ perspective. Thus we decide to stop if $X_{N-2}\geq \mathrm{E}(Z_{N-1}|\,\mathcal{F}_{N-2})$, giving the optimal expected payment $X_{N-2}\vee\mathrm{E}(Z_{N-1}|\,\mathcal{F}_{N-2})$ from the time $N-2$ perspective. We repeat this argument down to time 1.

We record the strategy at time k in the random variable and stopping time given by

$$Z_k=X_k\vee \mathrm{E}(Z_{k+1}|\,\mathcal{F}_k),\qquad \tau_k=\begin{cases} k & \text{if } Z_k=X_k\\ \tau_{k+1} & \text{if } Z_k>X_k\end{cases}.$$

The gain of stopping at time k rather than continuing to time $k+1$ is equal to $X_k-\mathrm{E}(Z_{k+1}|\,\mathcal{F}_k)$ if this variable is positive, and there is not gain otherwise. We record this in the jump of a stochastic process Λ, as

$$\Delta\Lambda_{k+1}=\bigl(X_k-\mathrm{E}(Z_{k+1}|\,\mathcal{F}_k)\bigr)^+.$$

The nondecreasing process Λ is then equal to $\Lambda_k=\sum_{j=2}^k\Delta\Lambda_j$.

Thus we have defined a discrete time stochastic process $Z=(Z_1,\ldots,Z_N)$, a sequence of stopping times τ_k, and a nondecreasing process Λ. The variable Z_k gives the value of the optimal stopping problem from the perspective at time t_k, and the time τ_k is the optimal stopping time for stopping after time t_k. These stopping times have a nice interpretation as the first time after t_k that the processes Z and X coincide.

6.10 Theorem.

(i) Z *is the smallest supermartingale with* $Z\geq X$.
(ii) $Z_k=\mathrm{E}(X_{\tau_k}|\,\mathcal{F}_k)=\sup_{\tau\in\mathcal{T}:\tau\geq t_k k}\mathrm{E}(X_\tau|\,\mathcal{F}_k)$ *for every* k.
(iii) $\tau_k=\min\{t_j: j\geq k, X_j=Z_j\}$ *for every* k.
(iv) $Z=M-\Lambda$ *for a martingale* M.
(v) $\Lambda_k=\Lambda_{\tau_k}$ *for every* k.

Proof. Item (i) is immediate from the definition of Z.

Item (ii) is proved by backward induction on k. The statement is clearly true for $k = N$. For $k < N$ the definitions give

$$\begin{aligned} \mathrm{E}(X_{\tau_k}|\,\mathcal{F}_k) &= \mathrm{E}\bigl(X_k 1_{X_k \geq \mathrm{E}(Z_{k+1}|\mathcal{F}_k)} + X_{\tau_{k+1}} 1_{X_k < \mathrm{E}(Z_{k+1}|\mathcal{F}_k)}|\,\mathcal{F}_k\bigr) \\ &= X_k 1_{X_k \geq \mathrm{E}(Z_{k+1}|\mathcal{F}_k)} + \mathrm{E}(X_{\tau_{k+1}}|\,\mathcal{F}_k) 1_{X_k < \mathrm{E}(Z_{k+1}|\mathcal{F}_k)}. \end{aligned}$$

Using the induction assumption $X_{\tau_{k+1}} = \mathrm{E}(Z_{k+1}|\,\mathcal{F}_{k+1})$, we can reduce the right side to

$$X_k 1_{X_k \geq \mathrm{E}(Z_{k+1}|\mathcal{F}_k)} + \mathrm{E}(Z_{k+1}|\,\mathcal{F}_k) 1_{X_k < \mathrm{E}(Z_{k+1}|\mathcal{F}_k)} = Z_k.$$

This concludes the proof of the first equality in (ii). The second equality in (ii) follows from the fact that $Z_k \geq \mathrm{E}(Z_\tau|\,\mathcal{F}_k)$ for any stopping time $\tau \geq t_k$, because Z is a supermartingale, where $\mathrm{E}(Z_\tau|\,\mathcal{F}_k)$ is bounded below by $\mathrm{E}(X_\tau|\,\mathcal{F}_k)$, because $Z \geq X$. The first equality in (ii) shows that these inequalities are equalities for $\tau = \tau_k$.

Item (iii) is clearly true for $k = N$. For $k < N$ it is also proved by backward induction on k. From the definitions we see that $\tau_k = k$ if and only if $X_k = Z_k$. Therefore $\tau_k = k$ implies that $\tau_k = \min\{t_j : j \geq k, X_j = Z_j\}$. Furthermore, $\tau_k > k$ implies that $X_k \neq Z_k$ and $\tau_k = \tau_{k+1}$. By the induction hypothesis $\tau_{k+1} = \min\{t_j : j \geq k+1, X_j = Z_j\}$, which is the same as $\min\{t_j : j \geq k, X_j = Z_j\}$ if $X_k \neq Z_k$.

The process M in (iv) clearly must be defined as $Z + \Lambda$, and hence it suffices to show that $\mathrm{E}(\Delta Z_k + \Delta\Lambda_k|\,\mathcal{F}_{k-1}) = 0$ for every k. But by the definitions $\Delta Z_k + \Delta\Lambda_k$ is equal to $Z_k - Z_{k-1} + 0 = Z_k - \mathrm{E}(Z_k|\,\mathcal{F}_{k-1}) + 0$ if $X_{k-1} < \mathrm{E}(Z_k|\,\mathcal{F}_{k-1})$, and it is equal to $Z_k - Z_{k-1} + X_{k-1} - \mathrm{E}(Z_k|\,\mathcal{F}_{k-1}) = Z_k - \mathrm{E}(Z_k|\,\mathcal{F}_{k-1})$ in the other case. Hence it is equal to $Z_k - \mathrm{E}(Z_k|\,\mathcal{F}_{k-1})$ in both cases and clearly a martingale increment.

The definition of $\Delta\Lambda_k$ shows that $\Delta\Lambda_k = 0$ if $\tau_{k-1} = \tau_k$. Clearly $\tau_j = \tau_k$ for any j with $k < j \leq \tau_k$. Therefore $0 = \sum_{j=k+1}^{\tau_k} \Delta\Lambda_j = \Lambda_{\tau_k} - \Lambda_k$, proving (v). ∎

The supermartingale property of Z expresses the fact that stopping later yields lower expected gain. The decomposition $Z = M - \Lambda$ makes the (expected) decrease of the supermartingale Z visible through the decreasing sample paths of the process $-\Lambda$. Property (v) could be rephrased as saying that Z follows the martingale M on any interval $(k, \tau_k]$ (which may well be empty of course): from time k onwards we stop only at τ_k; the gains $\Delta\Lambda_j$ of stopping earlier are zero.

The process Λ is predictable, and the decomposition $Z = M - \Lambda$ is the *Doob decomposition* of Z, the discrete time equivalent of the Doob-Meyer decomposition.

7
Payment Processes

, In Chapter 4 we found the fair price of a contract that yields a single pay-off X at an "expiry time" T. Several financial instruments yield payments at multiple times during an interval $[0, T]$. In this chapter we extend the pricing formula to general payment processes.‡

We shall obtain this extension in a more general framework, allowing the asset processes A to be general cadlag semimartingales, possibly discontinuous. Throughout we assume that $A = (A^{(1)}, \ldots, A^{(n)})$ is a vector of semimartingales defined on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ that satisfies the usual conditions. As before, a *strategy* is a predictable process ϕ such that the stochastic integral $\phi \cdot A$ is well defined, and a numeraire is a strictly positive semimartingale N such that $N = \alpha_0 A_0 + \alpha \cdot A = \alpha A$ for some (self-financing) strategy α. In the present chapter a numeraire is a cadlag process that may have jumps.

A *payment process* is defined to be a predictable semimartingale X. The value X_t at time t model is interpreted as the cumulative payments on a contract up to and including time t, and could be viewed as the sum (or integral) of a series of payments of sizes dX_t to the holder of the contract. (If X is not of bounded variation, then the latter intuitive interpretation should not be taken too seriously.) In practice a contract often consists of an agreement of a sequence of payments at finitely may predetermined times, whose values depend on the history of the asset process up to the time of payment. Thus the most interesting paymenst processes are of bounded variation, with finitely many jumps.

A *replicating strategy* for X is a predictable process ϕ with

$$\begin{aligned} \phi_0 A_0 + \phi \cdot A &= \phi A + X, \\ \phi_T A_T &= 0. \end{aligned}$$

‡ This chapter has not been changed in the November 2005 version. It may not be fully consistent with the preceding material.

As before, the strategy ϕ is interpreted as the contents of an investment portfolio. At time t the value of this portfolio changes by an amount $\phi_t\, dA_t$ due to the movement of the asset process. This change in value is used both to make a payment of size dX_t and to finance a change $d(\phi A)_t$ in the value of the portfolio. We are allowed to reshuffle the contents of the portfolio, as long as we fulfill these financing requirements. At expiry time T the portfolio is liquidated (i.e. $\phi_T A_T = 0$) and a final payment of ΔX_T is made to the holder of the contract.

The requirement that a payment process be predictable can be interpreted as saying that the size of a payment is known just before the time that the payment is made. If the asset process A is continuous, then the filtration $\mathcal{F}_t^A$ is left-continuous and hence this does not mean much as "just before time t" is the same as "at time t". In any case it will be seen below that any payment process that can be replicated is necessarily predictable.

7.1 Example. For given deterministic times $0 < T_1 < T_2 < \cdots < T_n = T$ and $\mathcal{F}^A_{T_i-}$-measurable random variables Y_i, the process $X = \sum_{i=1}^n Y_i 1_{[T_i,T]}$ is a payment process. It corresponds to the series of payments $Y_1, Y_2, \ldots, Y_n$ made at the times $T_1, T_2, \ldots, T_n$.

To verify the claim it suffices to show that the process X is predictable, or equivalently that a process of the type $Y 1_{[S,T]}$ is predictable for every $Y \in \mathcal{F}^A_{S-}$. For Y an indicator function $Y = 1_F$ of a set $F \in \mathcal{F}^A_s$ for some $s < S$, this follows because $1_F 1_{[S,T]}$ is the limit as $n \to \infty$ of the left-continuous processes $1_F 1_{(S-n^{-1}<T]}$, which are adapted and hence predictable for sufficiently large n. The set of Y such that $Y 1_{[S,T]}$ is predictable is a vector space that is closed under taking monotone limits. By the preceding it contains all indicators 1_F of sets in the algebra $\cup_{s<S}\mathcal{F}^A_s$. Thus the claim follows by a monotone class theorem.

More generally, the process $X = \sum_{i=1}^n Y_i 1_{[T_i,T]}$ is a payment process for predictable stopping times $0 < T_1 < T_2 < \cdots < T_n$ and random variables $Y_i \in \mathcal{F}^A_{T_i-}$. (Cf. Jacod and Shiryaev 2.12b).) □

In the derivation of the pricing formula (4.18) "unit invariance" plays a pervasive role. The idea is to express the asset price A relative to given numeraire, giving A^N. In the present situation we need to reexpress the value of both the asset and the payment process. Interpreting dX_t as the payment at time t, we see that the payment is dX_t/N_t if measured in the unit N. This suggests that the cumulative payment at time t is given by $\int_0^t (1/N_s)\, dX_s$. If N is continuous and X is of bounded variation, this is also the term that appears in the following lemma. For more general numeraires and payment processes the "rebased payment" process can better be defined as

$$\int_0^t \frac{1}{N_{s-}}\, dX_s + \Big\langle \Big(\frac{1}{N}\Big)^c, X^c \Big\rangle_t.$$

With this interpretation the "self-financing" property of a replicating strategy remains true, as shown in the following lemma.

7.2 Lemma (Unit invariance). *For any strictly positive semimartingale N and replicating strategy ϕ for X,*

$$\phi_0 A_0^N + \phi \cdot A^N = \phi A^N + \frac{1}{N_-} \cdot X + \Big\langle \Big(\frac{1}{N}\Big)^c, X^c \Big\rangle.$$

Proof. The defining relation $\phi_0 A_0 + \phi \cdot A = \phi A + X$ of a replicating strategy immediately implies that

$$\phi\, dA = d(\phi A) + dX,$$
$$\phi\, d\Big[A, \frac{1}{N}\Big] = d\Big[\phi A + X, \frac{1}{N}\Big].$$

The first equation is also true with "Δ" instead of "d". From the resulting equation we deduce that

$$\phi A_- = \phi A - \phi \Delta A = (\phi A)_- - \Delta X.$$

Therefore, by two applications of the partial integration formula,

$$\begin{aligned}
\phi\, d\Big(A\frac{1}{N}\Big) &= \phi\Big(\frac{1}{N_-}\, dA + A_-\, d\Big(\frac{1}{N}\Big) + d\Big[A, \frac{1}{N}\Big]\Big)\\
&= \frac{1}{N_-}\big(d(\phi A) + dX\big) + \big((\phi A)_- - \Delta X\big) d\Big(\frac{1}{N}\Big) + d\Big[\phi A + X, \frac{1}{N}\Big],\\
d\Big(\phi A\frac{1}{N}\Big) &= \frac{1}{N_-}\, d(\phi A) + (\phi A)_-\, d\Big(\frac{1}{N}\Big) + d\Big[\phi A, \frac{1}{N}\Big].
\end{aligned}$$

The difference of the two right hand sides can be written as

$$\frac{dX}{N_-} - \Delta X\, d\Big(\frac{1}{N}\Big) + d\Big[X, \frac{1}{N}\Big].$$

In view of the predictability of X the second and third terms together are equal to $d\langle X^c, (1/N)^c\rangle$. (Cf. Jacod & Shiryaev, ??). Integrating this gives the result. ■

Because a replicating strategy produces exactly the same payment process as the contract X, the choice between the replicating portfolio and the contract should be immaterial. Thus the initial value $\phi_0 A_0$ of the portfolio is the "just" price at time 0 for the contract with payment process X.

If we acquire the contract at an intermediate time $t \in (0, T)$, then we do not receive the payments that have been made under the contract before t. Given the amount $\phi_t A_t$, we could buy the portfolio ϕ_t at time t and next replicate the payments after t with certainty. Therefore, the value of the claim X at time t is given by $\phi_t A_t$.

This value can be expressed in the claim using a numeraire pair. If $(\mathbb{N}, N)$ is a numeraire pair and ϕ is an admissible strategy, then $\phi \cdot A^N$ is an $\mathbb{N}$-martingale. Consequently, if ϕ replicates the payment process X, then the process

$$Z{:}= \phi A^N + \frac{1}{N_-} \cdot X + \Big\langle \Big(\frac{1}{N}\Big)^c, X^c \Big\rangle = \phi_0 A_0^N + \phi \cdot A^N$$

is an $\mathbb{N}$-martingale. Using the fact that the final value of the process is given by $Z_T = (N_-^{-1} \cdot X)_T + \langle X^c, (1/N)^c \rangle_T$, we can write the martingale relationship $\mathrm{E}(Z_T - Z_t |\, \mathcal{F}_t) = 0$ in the form

$$\begin{aligned} \phi_t A_t &= N_t(\phi_t A_t^N) \\ &= N_t \mathrm{E}_{\mathbb{N}} \Big(\int_{(t,T]} \frac{1}{N_{s-}}\, dX_s + \int_{(s,T]} d\Big\langle X^c, \Big(\frac{1}{N}\Big)^c \Big\rangle |\, \mathcal{F}_t \Big). \end{aligned} \tag{7.3}$$

Thus the just price at time t involves the conditional expectation under the martingale measure of the "discounted payments" that are still to be received. This is similar to the formula (4.18) for a single payment at expiry time. If either the numeraire N or the payment process X is of bounded variation, then the value reduces to

$$\phi_t A_t = N_t \mathrm{E}_{\mathbb{N}} \Big(\int_{(t,T]} \frac{1}{N_{s-}}\, dX_s |\, \mathcal{F}_t \Big).$$

7.4 EXERCISE. Show that if ϕ satisfies the relationship in the preceding lemma, then $\phi_0 A_0 + \phi \cdot A = \phi A + X$. [Hint: if $Z = N_-^{-1} \cdot X + \langle X^c, (1/N)^c \rangle$, then $X = N_- \cdot Z + \langle Z^c, N^c \rangle$.]

7.5 EXERCISE. Specialize the pricing formula to the case that X consists of a single payment Y at expiry time T. Compare to formula (4.18).

7.6 Example (Single payment). A single payment of $Y \in \mathcal{F}_{S-}$ at time $S \in [0, T]$ is modelled through the payment process $X = Y 1_{[S,T]}$. The price process takes the form

$$V_t = N_t \mathrm{E}_{\mathbb{N}} \Big(\frac{Y}{N_{S-}} 1_{t<S} |\, \mathcal{F}_t \Big) = 1_{t<S} N_t \mathrm{E}_{\mathbb{N}} \Big(\frac{Y}{N_{S-}} |\, \mathcal{F}_t \Big).$$

This is the same as the price obtained by applying formula (4.18) to a claim consisting of a payment of $Y N_{T-}/N_{S-}$ at time T. We can think of N_{T-}/N_{S-} as a "premium factor" for late payment. Alternatively, we can think of $Y N_{T-}/N_{S-}$ as the capital obtained at time T by investing the payment of Y at time S in Y/N_{S-} units of the numeraire, and letting this "grow" to the value $(Y/N_{S-})N_{T-}$ at time T. □

A final question to be answered is under what conditions all payment processes are replicable. This turns out to be the case exactly if the market is complete in the sense introduced in Chapter 4 for payments only at expiry time.

7.7 Theorem. *If the economy is complete, then every predictable payment process X such that $\mathrm{E}_{\mathbb{N}}|(N_-^{-1}\cdot X)_T+\langle X^c,(1/N)^c\rangle_T|<\infty$ can be replicated by an admissible strategy.*

Proof. Let Z be the process $Z = N_-^{-1}\cdot X + \langle X^c,(1/N)^c\rangle$, which is predictable by the predictability of X. By Theorem 4.28, completeness of the economy is equivalent to the existence of a numeraire pair $(\mathbb{N}, N)$ for which the process A^N possesses the representing property. Thus the martingale M defined by $M_t = \mathrm{E}_{\mathbb{N}}(Z_T|\,\mathcal{F}_t)$ can be represented as $M = \phi_0 A_0^N + \phi\cdot A^N$ for some predictable process ϕ. It suffices to show that we can choose ϕ such that $\phi_T A_T = 0$ and $\phi_0 A_0 + \phi\cdot A = \phi A + X$. By unit invariance the second equation is equivalent to $M = \phi_0 A_0^N + \phi\cdot A^N = \phi A^N + Z$, i.e. $M_t - Z_t = \phi_t A_t^N$.

Because N is a numeraire, there exists a self-financing strategy α such that $N = \alpha A = \alpha_0 A_0 + \alpha\cdot A$. By self-financing and unit invariance $1 = N/N = \alpha A^N = \alpha_0 A_0 + \alpha\cdot A$ and hence $\alpha\cdot A^N = 0$. Define a process ψ by

$$\psi_t = \phi_t + \Bigl(M_t - Z_t - \phi_t A_t^N\Bigr)\alpha_t.$$

Then $\psi_t A_t^N = M_t - Z_t$, as desired, $\psi_T A_T = 0$, and $\psi\cdot A^N = \phi\cdot A^N$.

It suffices to show that ψ is predictable. Because ϕ, Z and α are predictable, this is certainly the case if the process $M - \phi A^N = \phi\cdot A^N - \phi A^N$ is predictable. This process is certainly adapted. If ϕ is left-continuous, then $\Delta(\phi\cdot A^N - \phi A^N) = \phi\,\Delta A^N - \phi\Delta A^N = 0$, and hence the process $\phi\cdot A^N - \phi A^N$ is continuous. Thus the process $\phi\cdot A^N - \phi A^N$ is predictable for every left-continuous strategy ϕ. The set of all strategies for which this process is predictable is a vector space, and is closed under monotone uniform limits. Thus it is predictable for all strategies, by a monotone class theorem. ∎

8
Infinite Economies

Even though real economies contain only finitely many assets, it is of interest to consider also economies with infinitely many assets. For instance, popular models for the bond market, considered in Chapter 9, assume the availability of bonds of an arbitrary maturity date $T > 0$, and hence in principle uncountably many assets. In this chapter we extend the pricing theory of Chapter 4 to economies with infinitely many assets. The extension is only modest in that the definitions are chosen such that the set-up essentially reduces to that of finite economies.[♭]

We write the asset processes as $A = (A^i : i \in \mathcal{I})$, where $\mathcal{I}$ is an arbitrary index set, and each A^i is a continuous semimartingale on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. We refer to the family of processess A as the "economy $\mathcal{E}$". For a subset $I \subset \mathcal{I}$ we let $\mathcal{E}^I$ be the "sub-economy" consisting of the family of asset processes $A^I = (A_i : i \in I)$, defined on the same filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. For a finite subset $I \subset \mathcal{I}$ the economy $\mathcal{E}^I$ is a finite economy of the type considered in Chapter 4. We are interested in pricing derivatives with a finite time horizon $T > 0$, and hence the asset processes are important on the time interval $[0, T]$ only.

A "strategy" in the economy $\mathcal{E}$ should be a certain family $(\phi^i : i \in \mathcal{I})$ of predictable processes, with the interpretation that at time t we keep ϕ^i_t assets of type A^i in our portfolio. We greatly simplify the set-up by allowing each agent in the economy to trade in only finitely many assets, and define the collection of all strategies in $\mathcal{E}$ as the union of all strategies available in the subeconomies $\mathcal{E}^I$ with I ranging over the finite subsets of $\mathcal{I}$. Thus each agent may trade in finite, but arbitrarily many assets throughout $[0, T]$. For a family of predictable processes $(\phi^i : i \in \mathcal{I})$ and a subset $I \subset \mathcal{I}$ set $\phi^I = (\phi^i : i \in I)$.

[♭] This chapter has not been changed in the November 2005 version. It may not be fully consistent with the preceding material.

8.1 Definition. *A strategy ϕ in $\mathcal{E}$ is a family $(\phi^i: i \in \mathcal{I})$ of predictable processes such that for some finite subset $I \subset \mathcal{I}$:*
(i) $\phi^i = 0$ *for every* $i \notin I$.
(ii) $\phi^I = (\phi^i: i \in I)$ *is a strategy in* $\mathcal{E}^{\mathcal{I}}$.
The strategy ϕ is said to be self-financing if ϕ^I is self-financing in $\mathcal{E}^I$. The value process of ϕ is the value process of ϕ^I in $\mathcal{E}^I$.

8.2 EXERCISE. Verify that the preceding definitions of self-financing and value process are well posed, in that they do not depend on the choice of finite subset I. Also verify that if ϕ^I is a strategy in $\mathcal{E}^I$ and we set $\phi^i = 0$ for $i \notin I$, then ϕ^J is a strategy in $\mathcal{E}^J$ for every finite set $J \supset I$. [Hint: $\phi^I A^I = \phi^J A^J$ and $\phi^I \cdot A^I = \phi^J \cdot A^J$ whenever I and J are subsets of $\mathcal{I}$ that differ only by $i \in \mathcal{I}$ such that $\phi^i = 0$.]

By the preceding definition (and exercise) the value process of a self-financing strategy ϕ is the process $\phi^I A^I = \phi_0^I A_0^I + \phi^I \cdot A^I$ for every finite subset $I \subset \mathcal{I}$ such that $\phi^i = 0$ for every $i \notin I$. We shall write the value process also as $\phi A = \phi_0 A_0 + \phi \cdot A$.

We define a *numeraire* exactly as in Chapter 4 as a strictly positive semimartingale that is the value process of some self-financing strategy. Then, in view of the preceding definition, every numeraire in $\mathcal{E}$ is also a numeraire in some finite subeconomy $\mathcal{E}^I$.

8.3 Definition.
(i) A numeraire is a strictly positive semimartingale that is the value process of some self-financing strategy.
(ii) A numeraire pair is a pair $(\mathbb{N}, N)$ consisting of a numeraire N and a probability measure $\mathbb{N}$ on $(\Omega, \mathcal{F})$ that is equivalent to $\mathbb{P}$ and such that the process $t \mapsto A_t^i/N_t$ is an $\mathbb{N}$-martingale for every $i \in \mathcal{I}$.
(iii) A pricing process Z is a strictly positive, cadlag semimartingale such that ZA^i is a $\mathbb{P}$-martingale for every $i \in \mathcal{I}$.

8.4 Definition. *A strategy ϕ is admissible if for every numeraire pair $(\mathbb{N}, N)$ the process $\phi \cdot (A/N)$ is an $\mathbb{N}$-martingale.*

By notational convention the "stochastic integral" $\phi \cdot (A/N)$ in the definition of admissibility is the integral $\phi^I \cdot (A^I/N)$ for an (arbitrary) finite subset $I \subset \mathcal{I}$ such that $\phi^i = 0$ for $i \notin I$. This might suggest that admissibility might be the same as admissibility in finite subeconomies $\mathcal{E}^I$. This appears to be false, because the definition only requires the process $\phi \cdot (A/N)$ to be an $\mathbb{N}$-martingale for numeraire pairs $(\mathbb{N}, N)$ in $\mathcal{E}$, and not for every numeraire pair in the finite subexperiment $\mathcal{E}^I$. Not every numeraire pair in a finite subexperiment need be "extendible" to a numeraire pair in the infinite experiment $\mathcal{E}$. This subtlety makes the following theorems, which are otherwise straigthforward extensions of results for finite economies, worth the effort.

** **8.5** EXERCISE. Investigate this point.

8.6 Theorem (No arbitrage). *If there exists a numeraire pair* $(\mathbb{N}, N)$ *in* $\mathcal{E}$ *and* ϕ *and* ψ *are self-financing, admissible strategies with* $\phi_T A_T = \psi_T A_T$, *then* $\phi A = \psi A$ *on* $[0, T]$.

Proof. There exists a finite subset $I \subset \mathcal{I}$ such that $\phi^i = \psi^i = 0$ for every $i \notin I$. By assumption ϕ and ψ are self-financing in $\mathcal{E}$, and this implies that ϕ^I and ψ^I are self-financing in $\mathcal{E}^I$. By admissibility $\phi(A/N) = \phi^I(A^I/N) = \phi_0^I(A_0^I/N) + \phi^I \cdot (A^I/N)$ is an $\mathbb{N}$-martingale, and hence it is completely determined by its final value $\phi_T(A_T/N_T)$, and similarly for ψ. Because the final values are the same, we have $\phi(A/N) = \psi(A/N)$, and hence $\phi A = \psi A$. ∎

As for finite economies the preceding theorem justifies to define the "fair price" of a European option with claim X at T to be the value of a self-financing, admissible strategy ϕ such that $X = \phi_T A_T$, if there exists such a strategy. This value can be expressed as an expectation under a numeraire pair $(\mathbb{N}, N)$ or pricing process Z, as

$$\phi_t A_t = N_t \mathrm{E}_{\mathbb{N}}\Big(\frac{X}{N_T} |\, \mathcal{F}_t\Big) = \frac{1}{Z_t} \mathrm{E}_{\mathbb{P}}\big(X Z_T |\, \mathcal{F}_t\big).$$

These formulas are identical to the formulas for finite economies, and are immediate from the martingale properties of the processes $\phi(A/N)$ under $\mathbb{N}$, or ZA under $\mathbb{P}$. We silently assume that the variables X/N_T and XZ_T are integrable under $\mathbb{N}$ and $\mathbb{P}$, respectively.

We shall show that this formula is available for every sufficiently integrable claim $X \in \mathcal{F}_T$ as soon as there exists a finite sub-economy that is complete. The following definitions and results all refer to a fixed time horizon $T > 0$.

8.7 Definition. *The economy* $\mathcal{E}$ *is complete if there exists a numeraire pair* $(\mathbb{N}, N)$ *such that for* $\mathcal{F}_T$*-measurable random variable* X *with* $\mathrm{E}_{\mathbb{N}}|X/N_T| < \infty$ *there exists a self-financing, admissible strategy* ϕ *with* $X = \phi_T A_T$.

8.8 Theorem (Completeness). *If there exists a numeraire pair* $(\mathbb{N}, N)$ *in* $\mathcal{E}$ *and there exists a finite subset* $I \subset \mathcal{I}$ *such that* $\mathcal{E}^I$ *is complete and contains a strictly positive asset process, then* $\mathcal{E}$ *is complete.*

Proof. We first show that the completeness of $\mathcal{E}^I$ implies the completeness of $\mathcal{E}^J$ for every finite set $J \subset \mathcal{I}$ with $J \supset I$ such that N is a numeraire in $\mathcal{E}^J$. Indeed, under these conditions the pair $(\mathbb{N}, N)$ is a numeraire pair in $\mathcal{E}^J$ and $\mathcal{E}^J$ contains a strictly positive asset process. If Z_1 and Z_2 are pricing processes in $\mathcal{E}^J$, then they are also pricing processes in $\mathcal{E}^I$. By completeness of $\mathcal{E}^I$ and the fact that $\mathcal{E}^I$ contains a strictly positive asset process, it follows that $Z_1 = Z_2$. Thus the pricing process in $\mathcal{E}^J$ is unique, whence the economy $\mathcal{E}^J$ is complete.

By definition a numeraire is a value process of some strategy, and hence there exists a finite subset $J \subset \mathcal{I}$ such that N is a numeraire in $\mathcal{E}^J$. Without loss of generality we can choose $J \supset I$. Then $(\mathbb{N}, N)$ is a numeraire pair in the experiment $\mathcal{E}^J$ and by the preceding $\mathcal{E}^J$ is complete. Therefore, for every X with $\mathrm{E}_\mathbb{N}\left|X/N_T\right| < \infty$ there exists a self-financing, admissible strategy ϕ^J in $\mathcal{E}^J$ such that $X = \phi_T^J A_T^J$. If we set $\phi^i = 0$ for $i \notin J$, then ϕ is a self-financing strategy in $\mathcal{E}$.

The strategy ϕ is also admissible in $\mathcal{E}$, but this requires proof. The admissibility of ϕ^J in $\mathcal{E}^J$ guarantees that $\phi^J \cdot (A^J/N^J)$ is an $\mathbb{N}^J$-martingale for every numeraire pair $(\mathbb{N}^J, N^J)$ in $\mathcal{E}^J$, but it must be verified that $\phi \cdot (A/M)$ is an $\mathbb{M}$-martingale for every numeraire pair $(\mathbb{M}, M)$ in $\mathcal{E}$. For every such pair $(\mathbb{M}, M)$ there exists a finite subset $H \subset \mathcal{I}$ with $H \supset J$ such that M is a numeraire in $\mathcal{E}^H$. By the first paragraph of the proof it follows that $\mathcal{E}^H$ is complete. By the admissibility of ϕ^J in $\mathcal{E}^J$, the process $\phi^H \cdot (A^H/N) = \phi^J \cdot (A^J/N)$ is an $\mathbb{N}$-martingale. Thus strategy ϕ^H is a self-financing strategy in $\mathcal{E}^H$ such that $\phi^H \cdot (A^H/N)$ is an $\mathbb{N}$-martingale. By Lemma 4.25 the process $\phi^H \cdot (A^H/M) = \phi \cdot (A/M)$ is an $\mathbb{M}$-martingale. ∎

The preceding theorem is limited, as it is only usable if the infinite economy is not richer than some finite subeconomy. However, it is good enough for most applications.

We finish this chapter with some results on pricing processes.

8.9 Lemma. *Suppose that there exists a strictly positive asset process in $\mathcal{E}$. Then there exists a pricing process in $\mathcal{E}$ if and only if there exists a numeraire pair in $\mathcal{E}$.*

Proof. If Z is a pricing process and N is a strictly positive asset process in $\mathcal{E}$, then $L = ZN/\mathrm{E}_\mathbb{P}(NZ)_T$ is a nonnegative $\mathbb{P}$-martingale with mean 1. Hence $d\mathbb{N} = L_T\, d\mathbb{P}$ defines a probability measure with density process L relative to $\mathbb{P}$. Because $L(A^i/N) = ZA^i/\mathrm{E}_\mathbb{P}(NZ)_T$ is a $\mathbb{P}$-martingale, the process A^i/N is an $\mathbb{N}$-martingale, for every $i \in \mathcal{I}$.

Conversely, if $(\mathbb{N}, N)$ is a numeraire pair in $\mathcal{E}$ and L is the density process of $\mathbb{N}$ relative to $\mathbb{P}$, then $Z = L/N$ is a pricing process in $\mathcal{E}$. ∎

8.10 Lemma. *If there exists a numeraire pair in $\mathcal{E}$ and there exists a finite subset $I \subset \mathcal{I}$ such that $\mathcal{E}^I$ is complete and contains a strictly positive asset process, then $\mathcal{E}$ possesses a unique pricing process.*

Proof. There exists a pricing process in $\mathcal{E}$ by the preceding lemma. If Z_1 and Z_2 are both pricing processes in $\mathcal{E}$, then they are also pricing processes in $\mathcal{E}^I$. Because $\mathcal{E}^I$ is complete and contains a positive asset process, its pricing process is unique, and hence $Z_1 = Z_2$. ∎

** **8.11** EXERCISE. Investigate whether the existence of a unique pricing process in $\mathcal{E}$ implies the completeness of $\mathcal{E}$.

9
Term Structures

A *zero coupon bond*, also known as a *pure discount bond*, is a contract that guarantees a payment of 1 unit at a given time T in the future. The bond is said to "mature" at the *maturity time* T.

In a real market we can typically buy bonds of many different maturities at any given date. Even though in any finite interval only finitely many bonds may be "active", the "maturity periods" of the different bonds overlap, and the totality of these contracts is most naturally modelled on an infinite time horizon $[0, \infty)$.

We denote by $D_{t,T}$ the value at t of a zero coupon bond maturing at time T. Obviously, the bond will be worthless after time T and hence we may either set $D_{t,T}$ equal to zero for $t > 0$, or think of $t \mapsto D_{t,T}$ as a process defined on $[0, T]$ only. Because $D_{t,T}$ models the value at time $t < T$ of one unit to be received at time T, it is also called a *discount rate*, expressing the "time value of money".

A mathematical model for the discount rates $D_{t,T}$ is called a *term structure model*. We always assume that each process $(D_{t,T}: 0 \leq t \leq T)$ is a nonnegative, cadlag semimartingale defined on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$, with $D_{T,T} = 1$. The requirement that the process $D_{t,T}$ be adapted models the fact that the discount bonds are for sale at time t, whence their prices are known at that time. However, the true meaning depends crucially on the choice of filtration, which we have not fixed.

In this chapter we consider a number of standard models for specifying a complete term structure. In practice one is often interested in special derivatives of the discount rates, and a full model for all discount rates may be unnecessary or even undesired. Ad-hoc models for the entities that are relevant for the particular derivative may suffice, as long as it is clear that these are compatible with some term structure, and consistent with other ad-hoc models. We shall see examples of such partial models in later chapters.

In most models the discount rates $D_{t,T}$ depend smoothly on the maturity time T, but are of unbounded variation as a function of the time parameter t. This models the fact that at every given time the interest attainable for the period $[t,T]$ should not vary much with T, but will be highly sensitive to changes on the market, caused for instance by the arrival new information.

9.1 Short and Forward Rates

Discount rates have close ties to "interest rates" in a wide sense. If at every moment in time it would be possible to obtain an arbitrage-free, continuously compounded, fixed interest of rate $r > 0$ on an investment, then the only reasonable term structure model would be $D_{t,T} = e^{-r(T-t)}$. This is the value of one unit discounted for the fact that the bond has cash value (equal to one) at time T and not at time t. This example shows that $-\log D_{t,T}$ is a natural transformation of the discount rates. The process

$$Y_{t,T} = \frac{-\log D_{t,T}}{T-t}$$

is called the *yield*, and can be viewed as a fixed interest rate over the period $[t,T]$, contracted at time t. Buying a bond with yield $Y_{t,T}$ at time t is equivalent to putting money in a savings account for the period $[t,T]$ against the fixed (continuously compounded) rate $Y_{t,T}$. In particular, if $D_{t,T} = e^{-r(T-t)}$, then the yield is constant and equal to r. The yield $T \mapsto Y_{t,T}$ as a function of T is known as the *yield curve* at time t. Typically, the yield curve is increasing, more distant maturities giving higher returns, but this is not necessarily the case, "inverted yield curves" having been observed also.

For a general term structure model, the *short rate* is defined as the limit, if it exists,

$$r_t = -\lim_{h \downarrow 0} \frac{1}{h} \log D_{t,t+h}. \tag{9.1}$$

Equivalently, the short rate is the limit of $Y_{t,T}$ as $T \downarrow t$, so that the short rate is an "infinitesimal yield" at time t. The short rate has the interpretation of an interest rate (or "yield") on a deposit during the period $[t, t+dt]$, i.e. a deposit that is "withdrawn immediately". In practice the short rate is often identified with the yield on bonds with a short term, such one month or a week.

A *short rate model* takes the short rate as its point of departure. It consists of a model for the evolution of the short rate process r, often a stochastic differential equation, and a description allowing to recover the processes $D_{t,T}$ from the short rate. Such a description typically consists of the assumption that, for some probability measure $\mathbb{R}$ on the probability space $(\Omega, \mathcal{F})$ on which r is defined,

$$D_{t,T} = \mathrm{E}_{\mathbb{R}}\Bigl(e^{-\int_t^T r_s\,ds} \,|\, \mathcal{F}_t\Bigr). \tag{9.2}$$

This formula together with a specification of the process r could be taken as a definition of a term structure, but then carries little intuition.

The formula can be interpreted through the pricing formula (4.18) if we assume that the market contains, besides the discount bonds, an asset R that evolves according to the differential equation (with r_t the short rate)

$$dR_t = r_t R_t\,dt, \qquad R_0 = 1.$$

Then the process R, which takes the form $R_t = \exp\bigl(\int_0^t r_s\,ds\bigr)$, is a numeraire. If the market (assumed to consist of at least the assets R and the discount bonds) is complete, then there exists a corresponding (local) martingale measure $\mathbb{R}$ making $(\mathbb{R}, R)$ into a numeraire pair, by Example 4.23, and by (4.18) the price at t of a given claim X at time T is equal to

$$R_t\,\mathrm{E}_{\mathbb{R}}\Bigl(\frac{X}{R_T} \,|\, \mathcal{F}_t\Bigr) = \mathrm{E}_{\mathbb{R}}\Bigl(X e^{-\int_t^T r_s\,ds} \,|\, \mathcal{F}_t\Bigr).$$

Because a discount bond $D_{t,T}$ corresponds to a payment of $X = 1$ at time T, we arrive at the short rate model (9.2). Based on this interpretation the measure $\mathbb{R}$ is referred to as the *risk neutral measure*.

In view of the definition (9.1) of the short rate r the asset R satisfying $dR_t = r_t R_t\,dt$ could be interpreted as the result of investing at each time t the current amount R_t in a zero coupon bond $D_{t,t+dt}$ at the "current interest rate" r_t. The existence in the real world of an asset consisting of rolling over bonds with infinitesimal contract periods is questionable, but it is a standard model assumption. It is similar to the assumption of existence of a "risk-free" asset in the extended Black-Scholes model, except that presently the rate r_t is derived from the discount rates by (9.1).

Actually the preceding argument for the short rate model (9.2) is valid as soon as there exists a numeraire pair $(\mathbb{R}, R)$ corresponding to the process $R_t = \exp(\int_0^t r_s\,ds)$ and the market is complete. Thus the process R must be "tradable" in that it is the value process of a portfolio, but it does not need to be one of the basic assets itself. The argument can be generalized to arrive at conclusion (9.2) under the assumption of existence of a numeraire with differentiable sample paths (and market completeness). If $(\mathbb{N}, N)$ is an arbitrary numeraire pair in a complete economy containing the discount rate processes, then the pricing formula (4.18) implies that

$$D_{t,T} = N_t \mathrm{E}_{\mathbb{N}}\Bigl(\frac{1}{N_T} \,|\, \mathcal{F}_t\Bigr).$$

If the numeraire N possesses absolutely continuous sample paths, then it can be shown that the short rate exists and that $N = N_0 R$, and hence we again arrive at a short rate model.

9.3 Lemma. *Suppose that $D_{t,T} = \zeta_t^{-1}\mathrm{E}_{\mathbb{R}}(\zeta_T|\,\mathcal{F}_t)$ for a strictly positive semimartingale ζ defined on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{R})$, with absolutely continuous sample paths with derivative process $t \mapsto \dot\zeta_t$ that is continuous in $\mathbb{R}$-mean. Then the short rate (9.1) exists as a limit in probability, and (9.2) is valid with $R = \zeta_0/\zeta$.*

Proof. By assumption we can write $\zeta_T = \zeta_t + \int_t^T \dot\zeta_s\,ds$ for a derivative process $\dot\zeta$ such that $s \mapsto \mathrm{E}_{\mathbb{R}}|\dot\zeta_s - \dot\zeta_t|$ is continuous, and hence $\mathrm{E}_{\mathbb{R}}\int_0^T |\dot\zeta_s|\,ds < \infty$, for every $T > 0$. Thus we can rewrite the discount rate processes as

$$D_{t,T} = \frac{1}{\zeta_t}\mathrm{E}_{\mathbb{R}}\Big(\zeta_t + \int_t^T \dot\zeta_s\,ds|\,\mathcal{F}_t\Big) = 1 + \frac{1}{\zeta_t}\int_t^T \mathrm{E}_{\mathbb{R}}(\dot\zeta_s|\,\mathcal{F}_t)\,ds,$$

by Fubini's theorem. This shows that the sample paths of the process $T \mapsto D_{t,T}$ are absolutely continuous on $[t,\infty)$ with derivative $s \mapsto \zeta_t^{-1}\mathrm{E}_{\mathbb{R}}(\dot\zeta_s|\,\mathcal{F}_t)$. Because $\dot\zeta_t = \mathrm{E}_{\mathbb{R}}(\dot\zeta_t|\,\mathcal{F}_t)$, we have

$$\mathrm{E}_{\mathbb{R}}\Big|\frac{1}{h}\int_t^{t+h} \mathrm{E}_{\mathbb{R}}(\dot\zeta_s|\,\mathcal{F}_t)\,ds - \dot\zeta_t\Big| \le \frac{1}{h}\int_t^{t+h} \mathrm{E}_{\mathbb{R}}|\dot\zeta_s - \dot\zeta_t|\,ds \to 0.$$

Combining the two preceding displays, we conclude that $h^{-1}(D_{t,t+h} - 1) \to \dot\zeta_t/\zeta_t$ in $\mathbb{R}$-probability as $h \downarrow 0$, and hence $h^{-1}\log D_{t,t+h} \to \dot\zeta_t/\zeta_t$ in probability, by a first order Taylor expansion of the logarithm (the "delta-method"). Thus the short rate as in (9.1) exists as a limit in probability and is given by $r_t = -\dot\zeta_t/\zeta_t$.

The absolute continuity and positivity of ζ imply the absolute continuity of $t \mapsto -\log\zeta_t$, with derivative $t \mapsto -\dot\zeta_t/\zeta_t$, by the chain rule. By the preceding paragraph this is equal to $r_t = f_{t,t}$. Hence $\zeta_t = \zeta_0\exp\big(-\int_0^t r_s\,ds\big)$ and (9.2) is verified. ∎

Warning. We might use the equation $D_{t,T} = N_t\mathrm{E}_{\mathbb{N}}(N_T^{-1}|\,\mathcal{F}_t)$ to define a term structure, rather than deduce the equation from the pricing formula. Then we first construct a strictly positive semimartingale N on a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{N})$ and next define an economy to consist of (at least) the processes given in the display, and possibly the process N itself. If we define discount rates through $D_{t,T} = N_t\mathrm{E}_{\mathbb{N}}(N_T^{-1}|\,\mathcal{F}_t)$ for a given strictly positive semimartingale N on a given filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{N})$, then $t \mapsto D_{t,T}/N_t$ is an $\mathbb{N}$-martingale, for every T. However, it is not necessarily true that N is a numeraire in the economy consisting of all discount rate processes $t \mapsto D_{t,T}$, for $T > 0$. For example, let $N = 1/\mathcal{E}(W)$ for W a Brownian motion on an arbitrary filtered probability space. Then $D_{t,T} =$

$\mathcal{E}(W)_t^{-1}\mathrm{E}_{\mathbb{N}}\bigl(\mathcal{E}(W)_T|\,\mathcal{F}_t\bigr) = 1$, for every T. This does not allow to self-finance any nontrivial value process. This example also shows that in this type of model the initial filtration $\mathcal{F}_t$ can be strictly bigger than the filtration generated by the discount rates.

Under the short rate model (9.2) the "discounted discount rate" $D_{t,R}/R_t$ takes the form $D_{t,T}/R_t = \mathrm{E}_{\mathbb{R}}(R_T^{-1}|\,\mathcal{F}_t)$ and hence is an $\mathbb{R}$-martingale. We conclude that the pair $(\mathbb{R}, R)$ is a numeraire pair for the economy consisting of all discount rates provided that R is a numeraire. In that case the term structure model defined through (9.2) automatically fulfills the important requirement of no-arbitrage. Whether the economy consisting of all discount rates (and possibly the process R) is complete depends on the further specification of the term structure, for which relation (9.2) provides only the skeleton. Only if both the process r and an underlying filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{R})$ are specified, we can appeal to (9.2) to recover the discount rates.

A *short rate model* in the strict sense consists of (9.2) and a specification of the short rate process r as a one-dimensional diffusion process, with the filtration $\mathcal{F}_t$ taken equal to the natural filtration of the driving Brownian motion. We discuss these models further in Section 9.2. An advantage of this type of model is that it leads to relatively simple pricing formulas for many options based on the discount rates. Because r becomes a Markov process, the option prices for claims X that are a function of the evolution of r past some fixed time T, which are conditional expectations of the type $\mathrm{E}_{\mathbb{R}}(X/R_T|\,\mathcal{F}_t)$, will be functions of (t, r_t) only, for $t \leq T$. This permits the use of partial differential equations, including numerical methods, to compute these prices in concrete cases.

The simplicity of a short rate model in the strict sense is also its weakness. A short rate model models the complete bond market through the single stochastic process r. The Markovian nature of this process could be interpreted as implying that at every time t the state of the bond market is described by the single number r_t. This seems unrealistic.

More general term structure models can be build on the *forward rate*, which is the partial derivative, if it exists,

$$f_{t,T} = -\frac{\partial}{\partial T} \log D_{t,T}. \tag{9.4}$$

This is interpreted as an interest rate that can be contracted at time t on investments into a savings account at time T. The short rate can be written $f_{t,t}$ and is the interest that can be contracted at time t on investments that are deposited. immediately.

To motivate this inpretation of the forward process, consider an owner of an S-bond at a time $t < S$. The S-bond guarantees a payment of 1 at time S, but suppose that the owner needs the money only at a time $T > S$. One strategy would be to exchange the S-bond at time t for $D_{t,S}/D_{t,T}$ units of T-bonds, which would give the guaranteed payment of $D_{t,S}/D_{t,T}$

at time T. An alternative strategy would be to keep the S-bond until its maturity at time S, and invest the payment of one unit received at that time into a risk-free account during the period $[S,T]$. The *forward rate for* $[S,T]$ *contracted at* t is by definition the (imaginary) fixed interest rate $R = R_{t,S,T}$ such that this investment would give exactly the value $D_{t,S}/D_{t,T}$ at time T. In other words, the forward rate $R_{t,S,T}$ is the number R solving the equation $e^{R(T-S)} = D_{t,S}/D_{t,T}$, or

$$R_{t,S,T} = -\frac{\log D_{t,T} - \log D_{t,S}}{T-S}.$$

Thus the forward rate is an (imaginary) constant interest rate over the interval $[S,T]$, fixed in a contract entered at a time $t < S$ to the left of the interval.

In the special case that $t = S$, the forward rate is also called the "spot rate", and is exactly the yield process $Y_{t,T} = R_{t,t,T}$ encountered previously. More importantly, the *instantaneous forward rate* $f_{t,S}$ is the limit as $T \downarrow S$ of the process $R_{t,S,T}$. This is also referred to as just the "forward rate", and is the process defined in (9.4).

If the function $T \mapsto -\log D_{t,T}$ is continuously differentiable, then the discount rates can be recovered from the forward rates through

$$D_{t,T} = e^{-\int_t^T f_{t,S}\, dS}. \tag{9.5}$$

In that case, the short rate (9.1) also exists, and is given by $r_t = f_{t,t} := \lim_{T\downarrow t} f_{t,T}$.

The forward rate $f_{t,T}$ has the intuitive interpretation as the "interest rate" over a period $[T, T+dT]$ in the future obtainable at time $t < T$. Unlike is the case for the short rate, it depends on a time horizon T, and hence at every given time t there exist many forward rates. Even though the notion of an instantaneous rate remains a theoretical construct, the existence of multiple interest rates at each given time appears to be realistic. Thus the forward rates permit an attractive, different starting point for defining a model for the term structure.

In principle the forward rates permit to describe the bond economy at time t through the infinite-dimensional "state vector" consisting of the set $(f_{t,T}: T > t)$ of all forward rates set at time t. In practice, it may be realistic to reduce the effective dimension of this state vector. One standard approach is to model the forward rates as diffusions relative to a given finite-dimensional Brownian motion.

We may still assume that the process $R_t = \exp\left(\int_0^t r_s\, ds\right)$ based on the short rate is a numeraire for the bond economy. If the economy is complete, then, as before, this forces the equality (9.2), where $\mathbb{R}$ is the martingale measure corresponding to R. Then the discount rates $D_{t,T}$ are related to the forward rates both through (9.2) and (9.5), where in the former equation $r_t = f_{t,t}$. This implies restrictions on the specification of a model for the forwards. For instance, we shall see that if the forwards are modelled through diffusion equations

$$df_{t,T} = \mu_{t,T}\, dt + \sigma_{t,T}\, dW_t,$$

where W is a multivariate Brownian motion under $\mathbb{R}$, then necessarily $\mu_{t,T} = -\sigma_{t,T} \int_t^t \sigma_{t,S}\, dS$. Thus if using this type of diffusion model we can only freely specify the diffusion function.

9.2 Short Rate Models

In a wide sense every term structure model defined through the equation (9.2) with r the short rate process as in (9.1) is a "short rate model". In a more narrow sense a short rate model is a model which also assumes that the short rate process r is defined through a diffusion model on a "standard Brownian space". This is the type of short rate model we discuss in this section.

We assume given a one-dimensional Brownian motion W defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{R})$, with the filtration $\mathcal{F}_t$ equal to the augmented filtration generated by W. The short rate process r is assumed to satisfy a diffusion equation of the form

$$dr_t = \mu(t, r_t)\, dt + \sigma(t, r_t)\, dW_t.$$

Under appropriate conditions on the functions μ and σ, a solution r to this equation will indeed exist and be adapted to the filtration $\mathcal{F}_t$. If $\sigma(t, r_t)$ is strictly positive, then the equation can be inverted to express W into r, and hence the filtrations generated by r and W are the same. Under (possibly) additional conditions the process r will be a strong Markov process. We assume that all this is the case, and also that the variable $\exp\bigl(-\int_0^t r_s\, ds\bigr)$ is $\mathbb{R}$-integrable for every t (which is trivially true if the process r is nonnegative).

Given the process r we define discount rate processes through (9.2), i.e. with the process R defined by $R_t = \exp\bigl(\int_0^t r_s\, ds\bigr)$,

$$D_{t,T} = R_t\, \mathrm{E}_{\mathbb{R}}\Bigl(\frac{1}{R_T} \,|\, \mathcal{F}_t\Bigr).$$

We assume that the process R is a numeraire. It then follows that $(\mathbb{R}, R)$ is a numeraire pair for the economy consisting of all discount bonds.

Unlike is the case for the Black-Scholes model, which is the accepted model for pricing options on stocks, no single short rate model has gained universal acclaim. Some particular models are:

$$
\begin{aligned}
dr_t &= (\theta - \alpha r_t)\,dt + \sigma\,dW_t, && \text{Vasiček,}\\
dr_t &= (\theta - \alpha r_t)\,dt + \sigma\sqrt{r_t}\,dW_t, && \text{Cox-Ingersoll-Ross,}\\
dr_t &= \alpha r_t\,dt + \sigma r_t\,dW_t, && \text{Dothan,}\\
dr_t &= (\theta - \alpha\sqrt{r_t})\,dt + \sigma\sqrt{r_t}\,dW_t, && \text{Longstaff,}\\
d\log r_t &= (\theta - \alpha\log r_t)\,dt + \sigma\,dW_t.
\end{aligned}
$$

Here θ, α, σ are given positive constants, which describe the evolution of the short rate process under the "risk-neutral measure" $\mathbb{R}$. Because these three parameters are the only degrees of freedom in the resulting formulas for all discount rate processes and their derivatives, it would be surprising if any of the five models gave a good fit to the real world. For instance, the initial bond prices $D_{0,T}$ are known from the market at time 0, but are also computable from the short rate as $\mathrm{E}_{\mathbb{R}}(1/R_T)$, and hence can be expressed in the parameters θ, α, σ of the short rate model. Similarly, the prices of many derivatives are fixed by the model, but also observable in the market. Because in practice it is impossible to choose the three parameters so that the theoretical bond and derivative prices are consistent with the observed prices, it has been proposed to replace the fixed parameters by functions of the time variable. The versions of the Vasiček and Cox-Ingersoll-Ross model in which θ, α and σ are functions of t are both known as the *Hull-White* model. The time-dependent version of the Dothan model is the *Black-Derman-Toy* model. Finally, the *Ho-Lee* model is given by

$$dr_t = \theta_t\,dt + \sigma\,dW_t.$$

This model also appears as a forward rate model, as will be seen in the next section. As the Vasiček-Hull-White models, it does not preclude the short rate from becoming negative, which is perhaps an undesirable feature.

In these models the time-dependent drift function can typically be chosen so that the resulting theoretical *initial yield curve* $T \mapsto Y_{0,T} = -T^{-1}\log D_{0,T}$ can be fitted exactly to a given yield curve observed in the market (at time 0). This is called *calibration*. The remaining parameters of the model can next be calibrated from observed prices on derivatives, or fitted by statistical analysis of the history of the discount rates (less common in practice). In practice the initial yield curve $T \mapsto Y_{0,T}$ is only observed for a finite number of maturities T, and the model is calibrated to an interpolated initial yield curve.

The question of completeness of the short rate models is usually easy to resolve, because a sub-economy consisting of the numeraire R and a single (active) discount rate $D_{t,T}$ forms a standard Black-Scholes model, as discussed in Chapter 5, and is typically complete for the Brownian filtration $\mathcal{F}_t$. Because the process $D_{t,T}/R_t = \mathrm{E}_{\mathbb{R}}(R_T^{-1} \,|\, \mathcal{F}_t)$ is a martingale relative to the Brownian filtration $\mathcal{F}_t$, it can be represented as a stochastic integral, by the representation theorem for Brownian motion. If $d(D_{t,T}/R_t) = H_t\,dW_t$ for a predictable process H, then, by Itô's formula,

$$dD_{t,T} = d\Bigl(\frac{D_{t,T}}{R_t}R_t\Bigr) = D_{t,T}r_t\,dt + R_tH_t\,dW_t.$$

By Theorem 5.11 the economy consisting of the processes $t \mapsto (R_t, D_{t,T})$ is complete relative to the filtration $\mathcal{F}_t$ provided that the process H is nonzero.

** **9.6** EXERCISE. Investigate whether this is automatically the case if σ is strictly positive.

Because by our assumptions the short rate process is Markovian and generates the filtration $\mathcal{F}_t$, the conditional law of $(r_s\colon s \geq t)$ given $\mathcal{F}_t$ depends on r_t only. It follows that the value process

$$V_t = \mathrm{E}_{\mathbb{R}}\bigl(Xe^{-\int_t^T r_s\,ds}|\,\mathcal{F}_t\bigr)$$

of a contract with claim X that is a function of r_T can be written as a function $V_t = F(t, r_t)$, for $t \leq T$. If F is smooth, then we can use Itô's formula to compute

$$d\Bigl(\frac{V_t}{R_t}\Bigr) = \frac{1}{R_t}\bigl(F_t + F_r\mu + \tfrac{1}{2}F_{rr}\sigma^2 - Fr\bigr)\,dt + \frac{1}{R_t}F_r\sigma\,dW_t.$$

Here W is an $\mathbb{R}$-Brownian motion and, if the claim is replicable, the discounted value process V/R is an $\mathbb{R}$-martingale. This is possible only if the drift term of the preceding diffusion equation vanishes, whence we obtain the *term structure equation*

$$F_t(t,r) + \mu(t,r)F_r(t,r) + \tfrac{1}{2}\sigma^2(t,r)F_{rr}(t,r) - rF(t,r) = 0.$$

The corresponding boundary condition is $F(T,r) = g(r)$ if $X = g(R_T)$.

In particular, we can write the discount rates in the form $D_{t,T} = F(t, r_t)$ for a function F on $[0,T] \times \mathbb{R}$ satisfying the term structure equation and satisfying the boundary condition $F(T,r) = 1$.

9.7 Example (Affine structure). A short rate model is said to possess *affine structure* if the drift and diffusion functions take the forms

$$\mu(t,r) = \alpha(t)r + \beta(t),$$
$$\sigma^2(t,r) = \gamma(t)r + \delta(t).$$

Then the coefficients in the term structure equation depend also linearly on r and a solution ought to take the form $F(t,r) = e^{A(t)-B(t)r}$. Inserting this and the equations for μ and σ^2 into the term structure equation yields an equation of the form $C(t) + D(t)r = 0$ for certain functions C and D. This equation is satisfied identically in (t,r) if and only if $C = D = 0$, which takes the concrete form

$$\begin{aligned} B_t(t) + \alpha(t)B(t) - \tfrac{1}{2}\gamma(t)B^2(t) &= -1,\\ A_t(t) - \beta(t)B(t) + \tfrac{1}{2}\delta(t)B^2(t) &= 0. \end{aligned}$$

The first equation is a Riccati differential equation for B, while given B the second equation is an ordinary first order differential equation for A. The boundary condition $F(T,r) = g(r)$ for every r translates into equations $e^{A(T)} = g(0)$ and $e^{B(T)} = g(0)/g(1)$.

In particular, the discount rates take the form $D_{t,T} = F(t, r_t; T)$ for F given by $F(t,r) = e^{A(t,T)-B(t,T)r}$ for certain processes $A(t,T)$ and $B(t,T)$, which as functions of t, for fixed T, must satisfy the differential equations under the boundary condition $A(T,T) = B(T,T) = 0$. The yields are a logarithmic transformation of the discount rates and satisfy the attractive "linear" equation $(T-t)Y_{t,T} = -A(t,T) + B(t,T)r_t$. □

9.8 Example (Vasiček-Hull-White). For the Vasiček-Hull-White model the equations derived in the Example 9.7 are especially simple and lead to explicit formulas. In this model we have $\mu(t,r) = \bigl(\theta(t) - \alpha r\bigr)$ for a deterministic function θ and $\sigma(t,r) = \sigma$ independent of (t,r). The equation for B reduces to a first order linear differential equation equation and hence can be solved explicitly, after which the function A can be found by one integration.

In particular, we can write the discount rates in the explicit form $D_{t,T} = e^{A(t,T)-B(t,T)r_t}$ for the functions

$$\begin{aligned} A(t,T) &= \tfrac{1}{2}\sigma^2 \int_t^T B^2(s,T)\,ds - \int_t^T \theta(s)B(s,T)\,ds,\\ B(t,T) &= \alpha^{-1}(1 - e^{-\alpha(T-t)}). \end{aligned}$$

In the special case of the Vasiček model, the function θ is assumed to be constant, and the integral defining $A(t,T)$ can be evaluated analytically.

Rather than using the approach of Example 9.7 we can also derive the distribution of the short rate process and next employ equation (9.2). It can be verified that the solution to the Vasiček-Hull-White diffusion equation is given by

$$r_t = e^{-\alpha t} r_0 + e^{-\alpha t}\int_0^t \theta(s)e^{\alpha s}\,ds + \sigma e^{-\alpha t}\int_0^t e^{\alpha s}\,dW_s.$$

This is the sum of a deterministic function and the *Ornstein-Uhlenbeck process* $y_t = \sigma e^{-\alpha t}\int_0^t e^{\alpha s}\,dW_s$. The latter process satisfies the reduced diffusion equation $dy_t = -\alpha y_t\,dt + \sigma\,dW_t$ and is well studied in probability theory.

Because it is the integral of a deterministic function relative to Brownian, the Ornstein-Uhlenbeck process is a zero-mean Gaussian process. Its covariance function can be computed to be

$$\mathrm{E} y_s y_t = \sigma^2 e^{-\alpha(s+t)} \int_0^{s\wedge t} e^{2\alpha u}\,du = \frac{\sigma^2}{2\alpha} e^{-\alpha(s+t)} e^{2\alpha s\wedge t} = \frac{\sigma^2}{2\alpha} e^{-\alpha|s-t|}.$$

It follows that the short rate process is also a Gaussian process, both unconditionally and conditionally given its past. Consequently, by the affine structure, the discount rates $D_{t,T}$ are log normally distributed, whereas the yields are normal. We can also verify this from (9.2) using the fact that the integral $\int_t^T r_s\,ds$ is again conditionally normally distributed given $\mathcal{F}_t$. □

9.9 EXERCISE. Carry out the calculations as indicated in the last sentence of the preceding example, i.e. derive the conditional mean and variance of the variable $\int_t^T r_s\,ds$ given r_t and employ (9.2) to verify the formula for the discount rates.

9.10 EXERCISE. Is the Vasiček-Hull-White model complete?

9.11 EXERCISE. Show that it is possible to determine a parameter θ (a function $\theta\colon [0,\infty) \to \mathbb{R}$) in the Hull-White model such that the corresponding initial yield curve $T \mapsto Y_{0,T}$ is exactly equal to a given function. [The significance is that the model can be exactly calibrated to the observed bond rates on the market at time 0.] Is this still possible if θ is restricted to be constant?

9.12 Example (Cox-Ingersoll-Ross). Bessel process. □

The preceding short rate models are called "single-factor" models, because they are driven by a single Brownian motion. A *multi-factor short rate model* describes the short rate process r as a measurable function $r = g(s)$ of a multi-dimensional diffusion process s. An example is the *two-factor Hull-White model*, in which the short rate r is the first coordinate of the two-dimensional diffusion process (r, s) satisfying, for a two-dimensional Brownian motion (V, W),

$$\begin{aligned} dr_t &= (\theta_t + s_t - \alpha r_t)\,dt + \sigma\,dV_t,\\ ds_t &= -\beta s_t\,dt + \tau\,dV_t + \rho\,dW_t. \end{aligned}$$

The Markov structure of a multi-factor short rate model permits a characterization of the discount rates through a partial differential equation, much in the same way as for a single-factor model.

9.13 EXERCISE. Derive this equation for the two-dimensional Hull-White model.

9.3 Forward Rate Models

The best known term structures that take their point of departure in the forward rates are the *Heath-Jarrow-Morton models.* In these models the forward rate processes $t \mapsto f_{t,T}$ are assumed to satisfy stochastic differential equations of the type

$$df_{t,T} = \mu_{t,T}\,dt + \sigma_{t,T}\,dW_t. \tag{9.14}$$

Here $t \mapsto \mu_{t,T}$ and $t \mapsto \sigma_{t,T}$ are stochastic processes that may depend on the horizon T, but W is a single, multivariate Brownian motion that is common to all forward processes.♯ The differentials $df_{t,T}$ are understood to be relative to the argument t, with T being fixed. The corresponding initial conditions, one for each value of $T > 0$, consist of the specification of a complete curve $T \mapsto f_{0,T}$, known as the *initial yield curve.*

We shall assume that there exist versions of the solutions $t \mapsto f_{t,T}$ to (9.14) such that the processes $(t,T) \mapsto f_{t,T}$ are jointly continuous. Then the functions $T \mapsto f_{t,T}$ are integrable, and we can define the discount rate processes $D_{t,T}$ by (9.5). Furthermore, the short rate, defined as the limit (9.1), exists, and is given by $r_t = f_{t,t}$.

We also assume that the process R defined by $R_t = \exp\bigl(\int_0^t r_s\,ds\bigr)$ is a numeraire, with corresponding martingale measure $\mathbb{R}$, so that the discount rates also satisfy equation (9.2). The following theorem shows that this necessitates a relation between the drift and diffusion parameters in (9.14). To make this as transparent as possible, we shall interprete (9.14) as a diffusion equation relative to the risk-neutral measure $\mathbb{R}$ as in (9.2), i.e. the driving process W is an $\mathbb{R}$-Brownian motion. In that case the drift parameters $\mu_{t,T}$ are completely determined by the diffusion parameters $\sigma_{t,T}$.

9.15 Theorem. *Let W be a multivariate Brownian motion defined on a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{R})$, and suppose that the process $t \mapsto f_{t,T}$ is a solution to (9.14), for every $T > 0$, for given continuous processes $(t,T) \mapsto \mu_{t,T}$ and $(t,T) \mapsto \sigma_{t,T}$. Then the process $t \mapsto D_{t,T}/R_t$, with $D_{t,T}$ defined by (9.5) and $R_t = \exp(\int_0^t f_{s,s}\,ds)$, is an $\mathbb{R}$-local martingale if and only if,*

♯ If W is d-dimenionsal, then $t \mapsto \sigma_{t,T}$ is an $\mathbb{R}^d$-valued process, for every fixed $T > 0$, and $\sigma_{t,T}\,dW_t$ is understood as an inner product.

with $\Sigma_{t,T} = \int_t^T \sigma_{t,S}\, dS$,†

$$\mu_{t,T} = \sigma_{t,T}\Sigma_{t,T}, \qquad \text{a.e. } t.$$

In that case, with $r_t = f_{t,t}$,

$$dD_{t,T} = D_{t,T}\bigl(r_t\, dt - \Sigma_{t,T}\, dW_t\bigr).$$

Proof. By definition $r_u = f_{u,u} = f_{0,u} + \int_0^u df_{s,u}$, where the differential is relative to s, for fixed u. Hence

$$\int_0^t r_u\, du = \int_0^t f_{0,u}\, du + \int_0^t \int_0^u df_{s,u}\, du. \tag{9.16}$$

Similarly, we can write $-\log D_{t,T} = \int_t^T f_{t,u}\, du$ in the form

$$-\log D_{t,T} = \int_t^T f_{0,u}\, du + \int_t^T \int_0^t df_{s,u}\, du. \tag{9.17}$$

The sum of the two double integrals on the far right sides of (9.16) and (9.17) gives a double integral over the area $A = \{(s,u)\colon 0 \le s \le t,\ s \le u \le T\}$. We can substitute the diffusion equation (9.14) for the forward rates $f_{t,T}$ and use the stochastic Fubini theorem to write this double integral as

$$\int\!\!\int_A df_{s,u}\, du = \int_0^t \int_s^T (\mu_{s,u}\, du\, ds + \sigma_{s,u}\, du\, dW_s).$$

The sum of the first terms on the right sides of (9.16) and (9.17) is equal to the integral $\int_0^T f_{0,u}\, du$, which is constant in t. Thus adding the equations (9.16) and (9.17) for $\int_0^t r_u\, du$ and $-\log D_{t,T}$ and next taking the differential relative to t, we find

$$r_t\, dt - d\log D_{t,T} = \Bigl(\int_t^T \mu_{t,u}\, du\Bigr)\, dt + \Bigl(\int_t^T \sigma_{t,u}\, du\Bigr)\, dW_t.$$

Abbreviating the right side to $M_{t,T}\, dt + \Sigma_{t,T}\, dW_t$, and using Itô's rule we find

$$\begin{aligned} d\Bigl(\frac{D_{t,T}}{R_t}\Bigr) &= d\exp\bigl(\log D_{t,T} - \int_0^t r_s\, ds\bigr)\\ &= \frac{D_{t,T}}{R_t}\Bigl(d\log D_{t,T} + \tfrac{1}{2} d[\log D_{t,T}] - r_t\, dt\Bigr)\\ &= \frac{D_{t,T}}{R_t}\Bigl(\bigl(-M_{t,T} + \tfrac{1}{2}\|\Sigma_{t,T}\|^2\bigr)\, dt - \Sigma_{t,T}\, dW_t\Bigr). \end{aligned}$$

By assumption W is an $\mathbb{R}$-Brownian motion. Therefore, the process $t \mapsto D_{t,T}/R_t$ is an $\mathbb{R}$-local martingale if and only the drift term on the right side

† The product $\sigma_{t,T}\Sigma_{t,T}$ is understood to be the inner product of two d-dimensional stochastic processes.

of the preceding display is zero, i.e. if and only if $M_{t,T} = \frac{1}{2}\|\Sigma_{t,T}\|^2$. This is equivalent to the equality $\mu_{t,T} = \sigma_{t,T}\Sigma_{t,T}$ for almost every t.

The last assertion of the theorem follows by another application of Itô's rule. ∎

In the preceding theorem the process W in the diffusion equation (9.14) is a Brownian motion under the measure $\mathbb{R}$. If $\mathbb{R}$ is the martingale measure corresponding to the short rate numeraire R, then the processes $t \mapsto D_{t,T}/R_t$ must be $\mathbb{R}$-local martingales. The theorem shows that in this case the drift functions $\mu_{t,T}$ are completely determined by the diffusion functions $\sigma_{t,T}$.

This is different if the diffusion equation is understood relative to another underlying measure, e.g. the "historical measure". If the probability measure $\mathbb{P}$ possesses density process L relative to $\mathbb{R}$, then the process $\tilde{W} = W - C$ for $C_t = \int_0^t L_{s-}^{-1}\, d[L,W]_s$ is a $\mathbb{P}$-Brownian motion, by Girsanov's theorem. By rewriting the diffusion equation in terms of $\tilde{W}$, we see that the drift term under $\mathbb{P}$ takes the form

$$\mu_{t,T}\, dt = \sigma_{t,T}\bigl(\Sigma_{t,T}\, dt + dC_t\bigr).$$

Because the additional term dC_t is independent of T, the drift functions are severely restricted, also if the diffusion equation (9.14) is understood relative to a general underlying measure.

If the processes $\mu_{t,T}$ and $\sigma_{t,T}$ are smooth in the variable T, then the short rate process $r_t = f_{t,t}$ also satisfies a stochastic differential equation.

Warning. In the following we denote the partial derivative of a process $h_{t,T}$ relative to T by $\dot{h}_{t,T}$.

9.18 Lemma. *Let W be a multivariate Brownian motion defined on a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{R})$, and suppose that the process $t \mapsto f_{t,T}$ is a solution to (9.14), for every $T > 0$, for given processes $(t,T) \mapsto \mu_{t,T}$ and $(t,T) \mapsto \sigma_{t,T}$ whose sample paths are partially differentiable relative to T with continuous partial derivative processes $(t,T) \mapsto \dot{\mu}_{t,T}$ and $(t,T) \mapsto \dot{\sigma}_{t,T}$. Then*

$$dr_t = df_{0,t} + \Bigl(\mu_{t,t} + \int_0^t \dot{\mu}_{s,t}\, ds + \int_0^t \dot{\sigma}_{s,t}\, dW_s\Bigr)\, dt + \sigma_{t,t}\, dW_t.$$

Proof. By definition $r_t = f_{t,t} = f_{0,t} + \int_0^t df_{s,t}$, where the differential is relative to s, for fixed t. Inserting the diffusion equation (9.14) for the forward rates, and writing $\mu_{s,t} = \mu_{s,s} + \int_s^t \dot{\mu}_{s,u}\, du$ and similarly for $\sigma_{s,t}$, we find

$$\begin{aligned} r_t &= f_{0,t} + \int_0^t (\mu_{s,t}\,ds + \sigma_{s,t}\,dW_s) \\ &= f_{0,t} + \int_0^t (\mu_{s,s}\,ds + \sigma_{s,s}\,dW_s) + \int_0^t \int_s^t (\dot\mu_{s,u}\,du\,ds + \dot\sigma_{s,u}\,du\,dW_s) \\ &= f_{0,t} + \int_0^t (\mu_{s,s}\,ds + \sigma_{s,s}\,dW_s) + \int_0^t \int_0^u (\dot\mu_{s,u}\,ds\,du + \dot\sigma_{s,u}\,dW_s\,du), \end{aligned}$$

by the stochastic Fubini theorem. Taking differentials gives the result. ∎

9.19 EXERCISE. Rewrite the assertion of the lemma as $dr_t = (\mu_{t,t} + \dot f_{t,t})\,dt + \sigma_{t,t}\,dW_t$ for $\dot f_{t,T}$ the partial derivative of $T \mapsto f_{t,T}$. [Hint: differentiate the identity $f_{t,T} = f_{0,T} + \int_0^t (\mu_{s,T}\,ds + \sigma_{s,T}\,dW_s)$ with respect to T.]

The preceding lemma should not be mistaken to imply that a Heath-Jarrow-Morton forward rate model is a multi-factor short rate model in disguise. The lemma merely asserts that the short rate process satisfies a stochastic differential equation, but this need not be of the standard diffusion type. Neither the drift function, nor the "diffusion function" need be expressible in (t, r_t), as required for a single-factor short rate model, or even in the value of an underlying vector-valued Markov process.

On the other hand, in simple examples the equation for r may well reduce to a diffusion equation.

9.20 Example (Hoo-Lee). Consider the Heath-Jarrow-Morton model (9.14) driven by a one-dimensional Brownian motion and with diffusion coefficients $\sigma_{t,T} = \sigma$ equal to a constant, for every T. Then $\Sigma_{t,T} = \sigma(T - t)$ and hence the diffusion equations for $f_{t,T}$ and $D_{t,T}$ take the forms

$$\begin{aligned} df_{t,T} &= \sigma^2 (T-t)\,dt + \sigma\,dW_t, \\ dD_{t,T} &= D_{t,T}\big(r_t\,dt - \sigma(T-t)\,dW_t\big). \end{aligned}$$

The forward curves $t \mapsto f_{t,T}$ can be viewed as random perturbations of the parabola $t \mapsto \sigma^2(Tt - \frac{1}{2}t^2)$, where the random deviations from this fixed curve are the same for every T. In economic terms "the only possible movements of the yield curve are parallel shifts", or "all rates along the yield curve fluctuate in the same way".

Given an initial forward curve $T \mapsto f_{0,T}$, the short rate $r_t = f_{t,t}$ can be computed as

$$r_t = f_{0,t} + \tfrac{1}{2}\sigma^2 t^2 + \sigma W_t.$$

If the initial yield curve is differentiable, then we can write this in differential form and obtain a short rate model with nonrandom, time-dependent drift function and diffusion term σW_t. This is known as the *Hoo-Lee model.* □

9.21 Example. The choice $\sigma_{t,T}=\sigma e^{-\gamma(T-t)}$ in the Heath-Jarrow-Morton model yields a model with an exponentially increasing influence of the diffusion term. This leads to

$$\begin{aligned}\Sigma_{t,T}&=-\frac{\sigma}{\gamma}\bigl(e^{-\gamma(T-t)}-1\bigr),\\ df_{t,T}&=-\frac{\sigma^2}{\gamma}e^{-\gamma(T-t)}\bigl(e^{-\gamma(T-t)}-1\bigr)\,dt+\sigma e^{-\gamma(T-t)}\,dW_t,\\ dr_t&=(\theta_t-\gamma r_t)\,dt+\sigma\,dW_t,\end{aligned}$$

for $\theta_t=\gamma m_t+m_t'$ and $m_t=f_{0,t}-\frac{1}{2}\sigma^2(1-e^{-\gamma t})^2/\gamma^2$. [Or +???] This is again a short rate model with time-dependent drift parameters. □

9.22 EXERCISE. Verify the calculations of Example 9.21 and investigate the completeness of this model relative to the filtration generated by W.

If the process R is a numeraire and $\mathbb{R}$ a corresponding martingale measure, then the processes $t\mapsto D_{t,T}/R_t$ must be $\mathbb{R}$-martingales. By Theorem 9.15 this can only be true for a Heath-Jarrow-Morton model (9.14) if $\mu_{t,T}=\sigma_{t,T}\Sigma_{t,T}$. This condition and some integrability is also sufficient for $(\mathbb{R},R)$ to be a numeraire pair.

9.23 Theorem. *Suppose that the conditions of Theorem 9.15 hold and that $R_t=\exp(\int_0^t r_s\,ds)$ is a numeraire. If $\mu_{t,T}=\sigma_{t,T}\Sigma_{t,T}$ and $\mathrm{E}_{\mathbb{R}}\exp\bigl(\int_0^T\frac{1}{2}\|\Sigma_{s,T}\|^2\,ds\bigr)<\infty$, then $(\mathbb{R},R)$ is a martingale numeraire pair for the economy consisting of all discount rate processes $t\mapsto D_{t,T}$.*

Proof. By Theorem 9.15 the discounted discount rate processes $t\mapsto D_{t,T}/R_t$ are $\mathbb{R}$-local martingales. It suffices to show that they are also $\mathbb{R}$-martingales.

By the last assertion of Theorem 9.15 the discount rate processes satisfy a stochastic differential equation. This equation and an application of Itô's formula give that

$$d\Bigl(\frac{D_{t,T}}{R_t}\Bigr)=-\frac{D_{t,T}}{R_t}\Sigma_{t,T}\,dW_t.$$

(Cf. the proof of Theorem 9.15.) This shows that the processes $D_{t,T}/R_t$ satisfy the stochastic differential equation $dX_t=X_t\,dM_t$ for the local martingale $M=-\Sigma_{t,T}\cdot W$. This is the Doléans equation and hence the processes $t\mapsto D_{t,T}/R_t$ can be shown to be a martingale by verification of Novikov's condition for M. This is the condition as in the theorem. ∎

The last assertion of Theorem 9.15 shows that the discount rate processes $t\mapsto D_{t,T}$ satisfy a stochastic differential equation. Therefore, the economy formed by the process R together with finitely many discount rate processes $t\mapsto D_{t,T_i}$, for $i=1,\ldots,n$, is an extended Black-Scholes model of the type discussed in Chapter 5. In the parameterization of Theorem 9.15 the diffusion model for the asset processes $t \mapsto D_{t,T_i}$ is given under the martingale measure $\mathbb{R}$ and possesses drift r instead of m as in Chapter 5. If we take this as the initial measure, then a "market price of risk" process as in Theorem 5.7 can be taken equal to zero, and hence exists. It follows that the economy is complete relative to the augmented natural filtration generated by the asset processes $t \mapsto D_{t,T_i}$ provided that these are weakly uniquely defined by the stochastic differential equation. If the initial filtration $\mathcal{F}_t$ has the property that for every $T > 0$ there exist finitely many times $T_1 < T_2 < \cdots < T_n$ such that $(\mathcal{F}_t)_{0 \leq t \leq T}$ is generated by the processes $(R_s: 0 \leq s \leq t)$ and $(D_{t,T_i}: 0 \leq t \leq T_i, i = 1, \ldots, n)$, then the Heath-Jarrow-Morton economy is also complete relative to the filtration $\mathcal{F}_t$.

Alternatively, we may apply Theorem 5.11 to address the completeness relative to the augmented natural filtration $\mathcal{F}_t^W$ of the driving Brownian motion W. If W is d-dimensional, then a sufficient condition for completeness is that for each $T > 0$ there exist times $T_1 < T_2 < \cdots < T_d$ such that the $(d \times d)$-matrices with ith row Σ_{t,T_i} are invertible, for every $t \in [0, T]$.

10
Vanilla
Interest Rate Contracts

In this chapter we consider a number of standard contracts, called "vanilla" or "over the counter", because they are commonly traded, as opposed to "exotic" contracts, which are tailored to special demands. We are interested in a description of these contracts, and particularly in their valuation. Throughout the chapter it is silently understood that the bond economy permits numeraire pairs, and if convenient also that it is complete, or embedded into a complete economy. The completeness assumption implies that to every numeraire, for instance a discount rate, exists a corresponding martingale measure. This free choice of numeraire is convenient for obtaining the valuation formulas.

All contracts considered in this chapter are derivatives of the *zero coupon bonds*, described in Chapter 9. A zero coupon bond is a contract that guarantees a cash-flow of 1 unit at a time T in the future. To make the dependence on the maturity explicit, we shall also refer to such a bond as a T*-bond.* A graphical display of this contract is given in Figure 10.1, the upward arrow indicating a payment to the owner of the contract at time T. The value of a zero coupon bond at a time $t < T$ is by definition the discount rate $D_{t,T}$.

Figure 10.1. Zero coupon bond or "T-bond". The value at time t is $D_{t,T}$.

The value of a general derivative with payments $X_1, \ldots, X_n$ at times $T_1 < \cdots < T_n$ that are functions of the discount rates up to these times is equal to

$$N_t \sum_{i=1}^{n} \mathrm{E}_{\mathbb{N}}\Big(\frac{X_i}{N_{T_i}} \,|\, \mathcal{F}_t\Big).$$

Here we are free to choose a numeraire pair $(\mathbb{N}, N)$. Rather than the obvious pair $(\mathbb{R}, R)$ corresponding to the short rate, we may choose a numeraire pair for which the evaluation of this expression is straightforward. Often one even chooses both a special numeraire pair and a special model for the distribution of the relevant discount rate that makes the calculations easy. If everything fails we can always use the numeraire R and determine the value of the derivative numerically by stochastic simulation under the corresponding martingale measure $\mathbb{R}$.

10.1 Deposits

A *deposit* is a contract that guarantees a fixed interest rate L on a given capital over a prespecified period. A graphical display is given in Figure 10.2. The downward arrow indicates a payment of a single unit by the buyer of the contract at time t, for which he in return receives a payment of $1 + \alpha L$ units at the expiry time T of the contract. The number α, referred to as an *accrual factor* or *daycount fraction* indicates the duration of the deposit. In practice this could be the number of days divided by 360, but for us it will do to think of it as an absolute number. The number L is the rate of return, and is also specified in the contract, usually as a percentage.

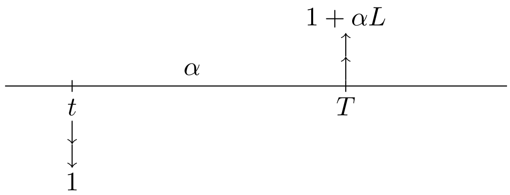

Figure 10.2. Deposit. The value at time t is zero if $L = L_t[t, T]$. The parameter α measures the length of the time interval $[t, T]$.

The value of the deposit contract at time t is positive if the return rate L is high and negative in the opposite case. The number L such that the value at time t is zero is called *LIBOR* (from "London Inter Bank Office Rate") and is denoted by $L_t[t, T]$.

Thus a deposit with return rate L set equal to the LIBOR guarantees a payment of $1 + \alpha L_t[t, T]$ at time T after making the initial investment of 1 unit at time t, and no other cashflows. Because buying $(1 + \alpha L_t[t, T])$ zero coupon bonds with maturity T returns the same payment, the no-arbitrage principle forces the costs of the two contracts to be the same. The bonds can be acquired at cost $D_{t,T}$ per bond at time t. Hence

$$
\begin{aligned}
1 &= (1 + \alpha L_t[t, T]) D_{t,T}, \\
L_t[t, T] &= \frac{1 - D_{t,T}}{\alpha D_{t,T}}.
\end{aligned}
\tag{10.1}
$$

It follows that we can think of the LIBOR as a derivative of the discount rates.

A different interpretation of the LIBOR results from noting that $\alpha L_t[t,T] = 1/D_{t,T} - 1$ is the profit made by investing 1 unit at time t in discount bonds with maturity T.

The LIBOR relates to the yield through the formula $\exp\bigl((T-t)Y_{t,T}\bigr) = 1 + \alpha L_t[t,T]$. Thus a third way of interpreting the LIBOR is to view it as a fixed interest rate for the interval $[t,T]$, which, unlike the yield, is not continuously compounded, but applied once to multiply the capital.

10.2 Forward Rate Agreements

The *forward rate agreement* (FRA) is graphically displayed in Figure 10.3. It incorporates two payments at the expiry time T, one receiving and one buying. The buying payment αK is at a fixed rate of return K, whereas the receiving payment is proportional to the LIBOR $L_S[S,T]$ for the period $[S,T]$. Because this LIBOR will only be "set" at some time S in the future, from the current time t perspective this payment is a random variable. In financial jargon it is referred to as a *floating payment*.

The purpose of the forward rate agreement is to exchange the unknown rate of return $L_S[S,T]$ for a rate K that is written in the contract at time t. Depending on K this may not be without cost. The value of K such that the FRA has zero value at time t is called the *forward LIBOR* and is denoted by $L_t[S,T]$.

The FRA enables one to obtain a certain return of $L_t[S,T]$ on a sum of money that we know will come in our possession during a time interval $[S,T]$ in the future. If the sum of money is 1, the following strategy could be used.

time t sell an FRA at the forward LIBOR rate $K = L_t[S,T]$ at no cost.
time S receive 1 unit; deposit 1 unit at rate $L_S[S,T]$ until T.
time T pay $\alpha\bigl(L_S[S,T] - L_t[S,T]\bigr)$ on FRA,
cash deposit giving $1 + \alpha L_S[S,T]$,
total value sums up to $1 + \alpha L_t[S,T]$.

Following this scheme, we are certain to receive a return rate of $L_t[S,T]$ on the money received at the future time S, a rate that is fixed at the current time t. Thus an FRA is an instrument to "swap" a random future rate for a fixed rate.

For K unequal to the forward LIBOR the FRA possesses a nonzero value process. We can determine the value of an FRA at time t by the general theory of Chapter 4, under the assumption that the economy is complete. The FRA contract guarantees a single payment of $\alpha\bigl(L_S[S,T] - K\bigr)$ at time T. Therefore, given a martingale numeraire pair $(\mathbb{N}, N)$, the

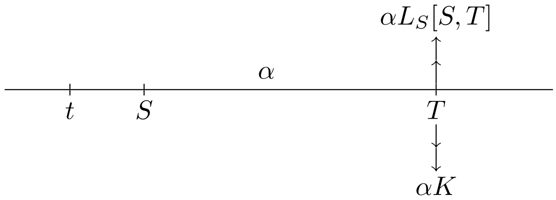

Figure 10.3. Forward rate agreement (FRA). The value of the contract at time t is zero if $K = L_t[S,T]$. The parameter α measures the length of the time interval $[S,T]$.

value of the contract at time t is given by

$$\begin{aligned} V_t &= N_t\,\mathrm{E}_{\mathbb{N}}\Big(\frac{\alpha\big(L_S[S,T]-K\big)}{N_T}\,|\,\mathcal{F}_t\Big)\\ &= N_t\,\mathrm{E}_{\mathbb{N}}\Big(\alpha\big(L_S[S,T]-K\big)\mathrm{E}_{\mathbb{N}}\Big(\frac{1}{N_T}\,|\,\mathcal{F}_S\Big)\,|\,\mathcal{F}_t\Big)\\ &= N_t\,\mathrm{E}_{\mathbb{N}}\Big(\alpha\big(L_S[S,T]-K\big)\frac{D_{S,T}}{N_S}\Big)\,|\,\mathcal{F}_t\Big). \end{aligned}$$

Here we have used the fact that, by the definition of a martingale numeraire pair $(\mathbb{N}, N)$, the process $t \mapsto D_{t,T}/N_t$ is an $\mathbb{N}$-martingale, for every fixed T, and $1/N_T = D_{T,T}/N_T$. We now express $L_S[S,T]$ into the discount rates through (10.1) and rewrite the right side as

$$\begin{aligned} V_t &= N_t\mathrm{E}_{\mathbb{N}}\Big(\frac{1 - D_{S,T} - K\alpha D_{S,T}}{N_S}\,|\,\mathcal{F}_t\Big)\\ &= D_{t,S} - (1+\alpha K)D_{t,T}, \end{aligned}$$

again by the martingale property of the discounted discount rates.

We have carried out the preceding calculation without ever specifying the numeraire pair or indicating a term structure model. This may be explained from the fact that there exists a very simple hedging strategy for an FRA, which also does not require specification of a model. For buying an FRA this takes the following steps.

Given an amount of $D_{t,S} - (1+\alpha K)D_{t,T}$ at time t:

time t	buy an S-bond at cost $D_{t,S}$,
	sell $(1+\alpha K)$ T-bonds at cost $-(1+\alpha K)D_{t,T}$.
time S	cash S-bond; deposit 1 unit at rate $L_S[S,T]$ until T.
time T	pay off $(1+\alpha K)$ T-bonds at cost $(1+\alpha K)$,
	cash deposit giving $1+\alpha L_S[S,T]$,
	total value sums up to $\alpha\big(L_S[S,T]-K\big)$.

By definition the forward LIBOR $L_t[S,T]$ is the value of K setting the value of the FRA at time t equal to zero. Solving the equation $0 = D_{t,S} - (1+\alpha K)D_{t,T}$ for K allows to express the LIBOR in the discount rates, as

$$L_t[S,T] = \frac{D_{t,S} - D_{t,T}}{\alpha D_{t,T}}. \tag{10.2}$$

We can substitute this equation back into the expression for the value of an FRA to obtain that the value at time of an FRA is given by $V_t = \alpha D_{t,T}\big(L_t[S,T] - K\big)$.

10.3 Swaps

A *swap* or *interest rate swap* can be considered an FRA with payments that are spread over multiple time points in the future. As can be seen in the graphical display given in Figure 10.4, the swap consists of a floating leg and a fixed leg. For simplicity we assume that there is a single set of payment dates $T_1 < \cdots < T_n$, although in practice the payment dates for the two legs may not match perfectly. The rate of return K in the fixed leg is the same for every payment and is fixed in the contract. The return rates on the payments in the floating leg may differ and are equal to the LIBORs setting at the beginning of the periods. Thus a swap allows to exchange random LIBORs that are set in the future for a fixed return rate written in the contract.

A swap whose value at the current time t is zero is called a *par swap*, and the corresponding value of K is the *swap rate*, denoted by $y_t[T_0, T_1, \ldots, T_n]$. If $t < T_0$, then this is a "forward swap rate", whereas for $t = T_0$ the swap rate is "spot". The spot swap rate is commonly quoted on market screens and hence is an important indicator for the bond market.

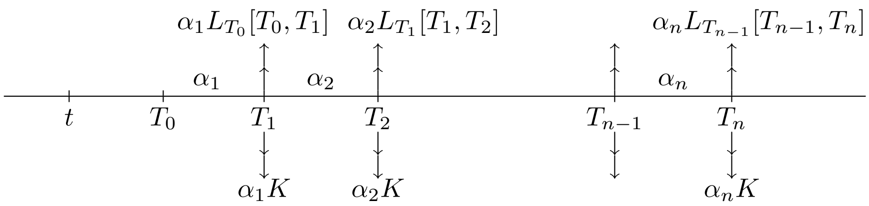

Figure 10.4. An Interest rate swap. The value of the swap at time t is zero for $K = y_t[T_0, T_1, \ldots, T_n]$. The parameter α_i measures the length of the time interval $[T_{i-1}, T_i]$.

Because a swap is a repetition of FRAs, its valuation is similar to that of an FRA. The payment at time T_i is equal to $\alpha_i\big(L_{T_{i-1}}[T_{i-1}, T_i] - K\big)$ and hence by (4.47), given a numeraire pair $(\mathbb{N}, N)$, the value of the swap at time t is equal to

$$V_t = N_t \mathrm{E}_{\mathbb{N}}\Big(\sum_{i=1}^{n} \frac{\alpha_i\big(L_{T_{i-1}}[T_{i-1}, T_i] - K\big)}{N_{T_i}} \,|\, \mathcal{F}_t\Big).$$

Each of the terms of the sum on the right side can be evaluated exactly as for an FRA. This yields

$$\begin{aligned} V_t &= \sum_{i=1}^{n} \bigl(D_{t,T_{i-1}} - (1+\alpha_i K) D_{t,T_i}\bigr) \\ &= D_{t,T_0} - D_{t,T_n} - K P_t[T_0, T_1, \ldots, T_n], \end{aligned}$$

for $P_t[T_0, T_1, \ldots, T_n]$ the "present value of unity" defined by

$$P_t[T_0, T_1, \ldots, T_n] = \sum_{i=1}^{n} \alpha_i D_{t,T_i}. \tag{10.3}$$

The present value of unity is the value of the fixed leg of the swap if the return rate K is unity. It is 10 000 times the *present value of a basis point* (PVBP), a "basis point" being 0.01 %, which is a more usual quantity quoted on the market.

By definition the value of a swap is zero if K is equal to the swap rate $y_t[T_0, T_1, \ldots, T_n]$. This readily gives the formula

$$y_t[T_0, T_1, \ldots, T_n] = \frac{D_{t,T_0} - D_{t,T_n}}{P_t[T_0, T_1, \ldots, T_n]}. \tag{10.4}$$

We can substitute this back into the value formula for a swap to see that this can also be written as $V_t = P_t[T_0, T_1, \ldots, T_n]\bigl(y_t[T_0, T_1, \ldots, T_n] - K\bigr)$.

As for an FRA there is a simple hedging strategy for a swap, which could have been used for the valuation.

Given the amount $D_{t,T_0} - D_{t,T_n} - K P_t[T_0, T_1, \ldots, T_n]$ at time t:

time t	buy an T_0-bond at cost D_{t,T_0}, sell T_n-bond at cost $-D_{t,T_n}$, sell $\alpha_i K$ T_i-bonds for $i = 1, \ldots, n$ at total cost $-K P_t[T_0, \ldots, T_n]$.
time T_0	cash T_0-bond, deposit 1 unit at rate $L_{T_0}[T_0, T_1]$ until T_1.
time T_1	deposit returns $1 + \alpha_1 L_{T_0}[T_0, T_1]$, pay out $\alpha_1 L_{T_0}[T_0, T_1]$, pay off $\alpha_1 K$ T_1-bonds at cost $\alpha_1 K$, deposit 1 unit at rate $L_{T_1}[T_1, T_2]$ until T_2.
time T_i	deposit returns $1 + \alpha_i L_{T_{i-1}}[T_{i-1}, T_i]$, pay out $\alpha_i L_{T_{i-1}}[T_{i-1}, T_i]$, pay off $\alpha_i K$ T_i-bonds at cost $\alpha_i K$, deposit 1 unit at rate $L_{T_i}[T_i, T_{i+1}]$ until T_{i+1}.
time T_n	deposit returns $1 + \alpha_n L_{T_{n-1}}[T_{n-1}, T_n]$, pay out $\alpha_n L_{T_{n-1}}[T_{n-1}, T_n]$, pay off $\alpha_n K$ T_n-bonds at cost $\alpha_n K$, pay off T_n-bond at cost 1.

10.4 Caps and Floors

Caplets and *Floorlets* are European call options and put options on the spot LIBOR rates. Their payment schemes are shown in Figures 10.5 and 10.6. *Caps* and *Floors* are repetitive caplets and floorlets, much as a swap is a repetition of FRAs. The payment scheme of a cap is shown in Figure 10.7. These contracts allow to profit from a potential rise in the value of the LIBOR, without full exposure to a possible decrease in the interest rate. Thus they are less conservative than FRAs or swaps, which are meant to take away all uncertainty by "swapping" future interest rates for a fixed rate.

Because the pay-offs of a caplet and floorlet are functions of the variable $\alpha\big(L_S[S,T]-K\big)$, which is the value at time T of a FRA, we can also view these contracts as options on an FRA. As the value of an FRA can be expressed in the discount rates, in the end caps and floors are derivatives of the discount rates.

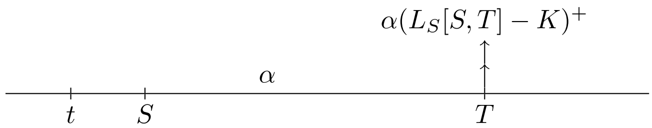

Figure 10.5. Caplet.

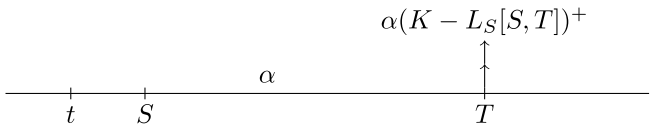

Figure 10.6. Floorlet.

Unlike is the case for FRAs and swaps, there is no universal hedging strategy for caplets and floorlets. The replicating strategy necessarily depends on the model used for the underlying discount rates. We might valuate these contracts using a model for the full term structure of the economy. However, because the payoffs on caplets and floorlets are functions of the LIBOR, their prices are expectations of the discounted LIBOR under a martingale measure, and hence an easier way to proceed is to model the LIBOR and the numeraire directly under a martingale measure. Here we are free to choose a convenient numeraire.

If we use the discount rate $N_t = D_{t,T}$ as a numeraire, then in view of (10.2) the forward LIBOR process $t \mapsto L_t[S,T]$ is a linear combination of two assets divided by the numeraire. By the definition of a martingale numeraire pair the forward LIBOR must be a martingale under a martingale measure $\mathbb{N}$ corresponding to the numeraire $N_t = D_{t,T}$. Because $N_T = D_{T,T} = 1$ the price of a derivative X based on the LIBOR takes the form

$V_t = D_{t,T}\mathrm{E}_{\mathbb{N}}(X|\,\mathcal{F}_t)$. We can evaluate this expectation as soon as we specify the distribution of the LIBOR under the martingale measure.

The model must incorporate the martingale property of the LIBOR. The simplest possible model that also ensures positivity of the LIBOR $L_t = L_t[S,T]$ is given by the differential equation $dL_t = L_t\sigma_t\,dW_t$, where W is an $\mathbb{N}$-Brownian motion relative to the filtration $\mathcal{F}_t$, and σ_t is a deterministic function. This equation is solved by the Doléans exponential $L_t = \mathcal{E}(\sigma\cdot W)$ and hence

$$L_S = L_t e^{\int_t^S \sigma_s\,dW_s - \frac{1}{2}\int_t^S \sigma_s^2\,ds}, \qquad t < S.$$

Because W possesses independent increments, the conditional distribution of L_S given $\mathcal{F}_t$ is the same as the distribution of $L_t\exp(\tau_t Z - \frac{1}{2}\tau_t^2)$, for $L_t = L_t[S,T]$ considered constant, Z a standard normal variable, and $\tau_t^2 = \int_t^S \sigma_s^2\,ds$. The value of a caplet can now be calculated as

$$\begin{aligned}
V_t &= N_t\mathrm{E}_{\mathbb{N}}\Bigl(\frac{\alpha\bigl(L_S[S,T]-K\bigr)^+}{N_T}\,|\,\mathcal{F}_t\Bigr)\\
&= D_{t,T}\mathrm{E}\alpha\Bigl(L_t e^{\tau_t Z - \frac{1}{2}\tau_t^2} - K\Bigr)^+\\
&= \alpha D_{t,T}\Bigl(L_t\Phi\Bigl(\frac{\log(L_t/K)}{\tau_t} + \tfrac{1}{2}\tau_t\Bigr) - K\Phi\Bigl(\frac{\log(L_t/K)}{\tau_t} - \tfrac{1}{2}\tau_t\Bigr).
\end{aligned}$$

This is known as *Black's formula.*

An alternative approach to valuing a caplet is to employ one of the term structure models discussed in Chapter 9. If we use the numeraire R with corresponding martingale measure $\mathbb{R}$, then we must evaluate

$$\mathrm{E}_{\mathbb{R}}\Bigl(e^{-\int_t^T r_s\,ds}\Bigl(\frac{1}{D_{S,T}} - 1 - \alpha K\Bigr)^+|\,\mathcal{F}_t\Bigr).$$

In most cases it is not possible to compute the expectation analytically. An exception is the Vasiček-Hull-White model, for which the preceding display reduces to

$$\mathrm{E}_{\mathbb{R}}\Bigl(e^{-\int_t^T r_s\,ds}\Bigl(e^{-A(t,T)+B(t,T)r_S} - 1 - \alpha K\Bigr)^+|\,\mathcal{F}_t\Bigr).$$

This can be evaluated using the fact that the random vector $(\int_t^T r_s\,ds, r_S)$ is bivariate-normally distributed given $\mathcal{F}_t$. The resulting expression depends on the parameters (θ,α,σ) of the Vasiček-Hull-White model, and will generally not agree with Black's formula.[True ???]

We can use the same approach to calculate the value of a floorlet. It is more interesting to derive this from the *put-call parity*, which is based on the identity $x^+ - (-x)^+ = x$ applied with x equal to $\alpha\bigl(L_S[S,T]-K\bigr)$. If we write the values processes of caplets and floorlets by V^{cap} and V^{floor}, then the identity and the pricing formula (4.18) yield

$$\begin{aligned} V_t^{cap} - V_t^{floor} &= N_t \mathrm{E}_{\mathbb{N}}\Big(\frac{\alpha\big(L_S[S,T]-K\big)}{N_T} \,|\, \mathcal{F}_t\Big) \\ &= N_t \alpha \mathrm{E}_{\mathbb{N}}\Big(\big(L_S[S,T]-K\big)\frac{D_{S,T}}{N_S} \,|\, \mathcal{F}_t\Big). \end{aligned}$$

This is true for every martingale numeraire pair $(\mathbb{N}, N)$. If the numeraire is again chosen equal to the discount rate $D_{t,T}$, then $D_{S,T}/N_S = 1$, and the LIBOR $t \mapsto L_t[S,T]$ is an $\mathbb{N}$-martingale. Then the right side reduces to $D_{t,T}\alpha\big(L_t[S,T]-K\big)$.

In this argument we have not used any model for the term structure, or the LIBOR. The put-call parity is true independently of the distribution of the assets.

Because a cap pays a series of caplets, the value of a cap is the sum of the values of a series of caplets. We omit the details.

10.5 EXERCISE. What is the value of a caplet at a time t with $S \leq t < T$?

10.6 EXERCISE. What is the value of a cap at time $t \in [T_i, T_{i+1})$?

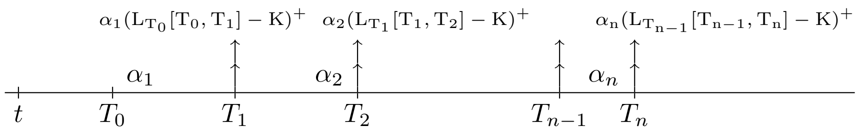

Figure 10.7. Cap.

10.5 Vanilla Swaptions

A *swaption* is a European option on a swap. In the context of swaptions the "call" and "put" forms of options are referred to as "payer's" and "buyer's" swaptions. Given fixed times $T_0 < T_1 < \cdots < T_n$, the claim of a *payer's swaption* consists of a payment at time T_0 of the amount

$$P_{T_0}[T_0,T_1,\ldots,T_n]\big(y_{T_0}[T_0,T_1,\ldots,T_n]-K\big)^+.$$

This is exactly the positive part of the value at T_0 of a swap at time T_0. The payment scheme of a swaption is displayed in Figure 10.8.

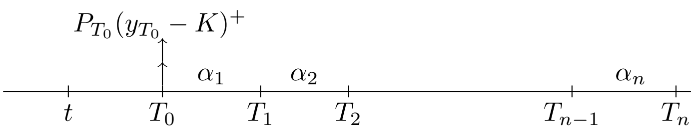

Figure 10.8. Swaption.

The swap rate $y_t = y_t[T_0, \ldots, T_n]$ is the quotient of $D_{t,T_0} - D_{t,T_n}$ and the PVBP $P_t = P_t[T_0, \ldots, T_n]$. This and the form of the claim makes the PVBP into a convenient numeraire for valuating the swaption. Let $\mathbb{N}$ be a martingale measure corresponding to P_t. Then $t \mapsto D_{t,T}/P_t$ is an $\mathbb{N}$-martingale for every discount rate $D_{t,T}$, and hence so is the swap rate y_t. The value of the swaption is equal to (choose $N = P$)

$$V_t = N_t \mathrm{E}_{\mathbb{N}}\Big(\frac{P_{T_0}(y_{T_0} - K)^+}{N_{T_0}} |\, \mathcal{F}_t\Big) = P_t \mathrm{E}_{\mathbb{N}}\big((y_{T_0} - K)^+ |\, \mathcal{F}_t\big).$$

To turn this into a concrete formula it suffices to model the distribution of y_{T_0} under the martingale measure $\mathbb{N}$. Because the swap rate is positive and is an $\mathbb{N}$-martingale, a simple model is the geometric Brownian motion, i.e.

$$y_t = e^{\int_0^t \sigma_s \, dW_s - \frac{1}{2}\sigma_s^2 \, ds},$$

for a deterministic function σ and a process W that is an $\mathbb{N}$-Brownian motion relative to the filtration $\mathcal{F}_t$. The corresponding value of the swaption is given by

$$V_t = P_t[T_0, \ldots, T_n] \mathrm{E}\Big(y_t[T_0, \ldots, T_n] e^{\tau_t Z - \frac{1}{2}\tau_t^2} - K\Big)^+.$$

Here $\tau_t^2 = \int_0^t \sigma_s^2 \, ds$ and the expectation is to be taken relative to the standard normal variable Z, keeping $y_t[T_0, \ldots, T_n]$ fixed. This can be further evaluated by similar calculations as used for caplets, leading to an expression of the same form as Black's formula.

10.7 EXERCISE. The claim of a *receiver's swaption* takes the form $P_{T_0}[T_0, T_1, \ldots, T_n]\big(y_{T_0}[T_0, T_1, \ldots, T_n] - K\big)^-$. Show that the put-call parity (or payer-receiver parity) takes the form $V_t^{pay} - V_t^{rec} = P_t(y_t - K)$, for $t < T_0$.

10.6 Digital Options

A *digital option* is an "all or nothing option", giving a fixed return if a certain event happens and no return in the opposite case. The *digital caplet* pays one unit at time T if the LIBOR $L_S[S, T]$ set at $S < T$ is above a certain prespecified level and nothing otherwise. (Cf. Figure 10.9.) For a *digital floorlet* these possibilities are exchanged.

A convenient numeraire is the discount rate $t \mapsto D_{t,T}$. Under a martingale measure corresponding to this numeraire the LIBOR $t \mapsto L_t[S, T]$ is a martingale, and hence can be reasonably modelled as a geometric Brownian motion. The approach is identical to the one taken for caplets, and, with the same notation as before, leads to the value of digital caplet given by

$$V_t = D_{t,T}\Phi\Big(\frac{\log L_t[S,T]/K - \frac{1}{2}\tau_t^2}{\tau_t}\Big).$$

10.8 EXERCISE. Derive a put-call parity between digital caplets and floorlets.

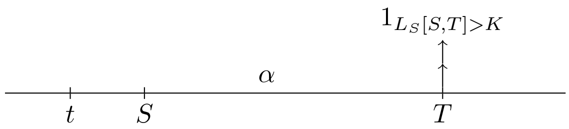

Figure 10.9. Digital caplet.

The *digital swaption* pays one unit at some time T if the swap rate $y_t[T_0, \dots, T_n]$ is bigger (digital payer's swaption) or smaller (digital receiver's swaption) than some constant, and nothing otherwise. The payment time T may or may not be one of the times T_i. The payment scheme of a digital payer's swaption is displayed in Figure 10.10. It appears that this derivative does not permit pricing with an equally simple model as before.

Figure 10.10. Digital receiver's swaption.

10.7 Forwards

A *forward* can be viewed as an FRA where the payment is a general random variable X rather than based on the LIBOR. The contract is graphically displayed in Figure 10.11.

The *forward price* is the amount of the fixed payment that makes the value of the contract 0 at time t. This can be computed to be, for an arbitrary numeraire pair $(\mathbb{N}, N)$,

$$F_t = \frac{\mathrm{E}_{\mathbb{N}}(X/N_T \mid \mathcal{F}_t)}{\mathrm{E}_{\mathbb{N}}(1/N_T \mid \mathcal{F}_t)}. \tag{10.9}$$

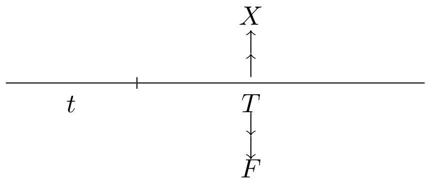

Figure 10.11. Forward for $X \in \mathcal{F}_T$. The value of the contract at time t is zero if $F = F_t$ is equal to the "forward price".

If we choose for numeraire the discount rate $N_t = D_{t,T}$, then $N_T = 1$ and the forward prices reduces to the conditional expectation $F_t = \mathrm{E}_{\mathbb{F}}(X|\,\mathcal{F}_t)$, for $\mathbb{F}$ the martingale measure corresponding to the discount rate. This is known as the *forward measure.*

10.10 EXERCISE. Verify the formula for the forward price. What is the forward price for $X = L_S[S,T]$?

11
Futures

Given a contract that includes a payment at a time in the future, there is always a risk that the other party will not fulfill their obligation. "Futures" are contracts that are designed to minimize the risk of *default* of the parties to the contract. This is achieved by allowing the parties to enter the contract at zero cost, at any time t, and next to require frequent payments in order to keep the current value of the contract near zero. Thus both parties can close out the contract at any time and at zero cost, eliminating the default risk.

In practice the default risk is further minimized by the creation of "margins" at a banking institution. Each party to a futures contract is under the obligation to deposit funds up to a margin at the bank at the beginning of the contract. The bank uses these funds to transfer money from one party to the other, depending on the evolution of the futures prices defined in the contract. If a party's deposit falls below the minimal margin, then the party is under the obligation to reinvest the deposit up to a central margin. On the other hand, if the deposit exceeds a certain upper margin, the party is permitted to withdraw the excess value, and in practice will do so. As long as the (random) payments that are agreed in the contract become not too extreme, they can always be carried out by the funds kept in deposit, and the parties cannot default. If a party fails to reinvest the deposit, then the futures contract is closed out.

Money sitting in a deposit without being used for portfolio management is "dead", and could certainly be more profitable if retracted and used otherwise. Thus this element of a futures contract is in contradiction to the no-arbitrage assumption that we make in the pricing theory. For pricing futures, we shall ignore this element of the futures contract, even though it is important for its practical usefulness.

There are many types of futures contracts. For each futures contract there is a *settlement date* T and a *settlement amount*, which is an $\mathcal{F}_T$-

measurable random variable X. Furthermore, there is a *futures process*, denoted $\{\Phi_t : t \in [0,T]\}$, which specifies the (cumulative) payments received (or paid) by a party to the contract. We describe the nature of these payments below, both in a discrete time and in a continuous time framework. The futures process is determined by the requirement that the value of the futures contract is zero at every time t.

The settlement amount is based on the price of an underlying asset or interest rate and determines the type of futures. One example are the *euro-dollar futures*, which are based on the LIBOR. The settlement amount of euro-dollar futures is $X = 100\bigl(1 - L_S[S,T]\bigr)$. Alternatively, futures may be based on stocks or commodities.

The settlement amount is the total payment received by a party holding the futures contract throughout the contract period $[0,T]$. However, the futures contract is defined so that it can be entered and settled at any time at no cost, and in practice the contract is rarely kept throughout the period.

11.1 Discrete Time

In the discrete time model the contract requires payments only at given times $T_1 < T_2 < \cdots < T_n$. (Set $T_0 = 0$.) If the contract is entered at time $t \in [T_{i-1}, T_i)$, no payment is required at that time, nor are payments due for the past times $T_j < t$. A first payment is required at time T_i, and is in the amount $\Phi_{T_i} - \Phi_t$. Next payments in the amounts $\Phi_{T_j} - \Phi_{T_{j-1}}$ are required at each of the times T_j for $j > i$. We can summarize this payment scheme by saying, if the contract is acquired at time t, then the payment at time T_j is equal to

$$\Phi_{T_j \vee t} - \Phi_{T_{j-1} \vee t}, \qquad j = 1, \ldots, n.$$

If the contract is entered at time 0, then the total payment over the interval $[0, T_i]$ is equal to $\Phi_{T_i} - \Phi_0$. Thus the futures process can be interpreted as the cumulative payments over time intervals.

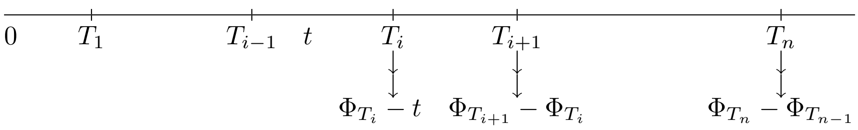

Figure 11.1. Payments required on futures contract acquired at time t.

As usual we assume the existence of a martingale numeraire pair $(\mathbb{N}, N)$ and market completeness. Then the pricing formula (4.47) shows that the value of the futures contract (to the other party in the contract) at time t

is equal to

$$V_t = N_t \mathrm{E}_{\mathbb{N}}\Big(\sum_{j=1}^n \frac{\Phi_{T_j\vee t} - \Phi_{T_{j-1}\vee t}}{N_{T_j}}\,|\, \mathcal{F}_t\Big). \tag{11.1}$$

11.2 Definition. *The process Φ is said to be a futures price process in discrete time for the settlement X if Φ is a cadlag semimartingale with $\Phi_T = X$ and $V_t = 0$ for every $t \in [0, T]$.*

Under some integrability conditions the futures price process exists and is uniquely determined. To see this we can rewrite the equation $V_t = 0$ for $t \in [T_{i-1}, T_i)$ in the form

$$\Phi_t = \frac{\mathrm{E}_{\mathbb{N}}(\Phi_{T_i}/N_{T_i}\,|\, \mathcal{F}_t) + \mathrm{E}_{\mathbb{N}}\big(\sum_{j=i+1}^n (\Phi_{T_j} - \Phi_{T_{j-1}})/N_{T_j}\,|\, \mathcal{F}_t\big)}{\mathrm{E}_{\mathbb{N}}(1/N_{T_i}\,|\, \mathcal{F}_t)}.$$

This can be solved recursively, starting with $\Phi_T = X$, next using the formula for the intervals $[T_{n-1}, T)$, $[T_{n-2}, T_{n-1})$, etc.

To obtain a concrete representation of the futures process, we must model the joint distribution of the settlement amount X and the numeraire process N under the martingale measure $\mathbb{N}$. Furthermore, the recursions may be difficult to implement in practice.

11.2 Continuous Time

In the continuous time model the payments are made continuously over time. To motivate the definition we can take a limit along a sequence of discrete time futures models with time points $0 = T_0^n < T_1^n < \cdots < T_n^n = T$. If $\max_i |T_j^n - T_{j-1}^n| \to 0$, then

$$\sum_{j=1}^n \frac{\Phi_{T_j\vee t} - \Phi_{T_{j-1}\vee t}}{N_{T_j}} \xrightarrow{\mathrm{P}} G_T - G_t,$$

for the process G defined by‡

$$G_t = \int_0^t \frac{1}{N_s}\, d\Phi_s + \Big[\frac{1}{N}, \Phi\Big]_t. \tag{11.3}$$

This suggests to replace the process V_t in (11.1) in the continuous time model by the process

$$V_t = N_t \mathrm{E}_{\mathbb{N}}\big(G_T - G_t\,|\, \mathcal{F}_t\big), \tag{11.4}$$

‡ We use that $\Sigma_i H_{s_i}(X_{s_i} - X_{s_{i-1}})$ converges to $\int_0^t H\, dX + [H, X]_t$ in probability if the mesh widths of the partition $0 = s_0 < \cdots < s_n = t$ tends to zero. The "correction term" $[H, X]$ arises because the sums use the final value of the process H in the partitioning interval.

and to define a process Φ to be a *futures process* for the claim X if Φ is a cadlag semimartingale with $V = 0$, exactly as in the preceding definition. This is equivalent to the statement $G_t = \mathrm{E}_{\mathbb{N}}(G_T |\, \mathcal{F}_t)$ for every t, and hence that the process G is an $\mathbb{N}$-martingale. In the following we assume that the numeraire N is continuous.

11.5 Definition. *The process Φ is said to be a futures price process in continuous time for the claim X if Φ is a cadlag semimartingale with $\Phi_T = X$ and such that the process G in (11.3) is an $\mathbb{N}$-martingale on $[0, T]$ for some numeraire pair $(\mathbb{N}, N)$.*

A different way to arrive at this definition is to consider $\Phi_t - \Phi_0$ to be the cumulative payments on a futures contract over the time period $[0, t]$. Then, as seen in Chapter 7, for a continuous numeraire N the cumulative discounted payment over the interval $[0, t]$ is equal to G_t, and the value at time t of the contract is equal to V_t, by (7.3). The more general approach of Chapter 7 also shows how to extend the definition to numeraires with jumps.

If the numeraire N is continuous and of bounded variation, then the term $[1/N, \Phi]$ in the definition of G vanishes. In that case we can invert the relation $G = N^{-1} \cdot \Phi$ to obtain that $\Phi = \Phi_0 + N \cdot G$. If Φ is a futures process, then G is an $\mathbb{N}$-martingale, and hence Φ is an $\mathbb{N}$-local martingale. If Φ is also an $\mathbb{N}$-martingale, then we obtain the pricing formula, since $\Phi_T = X$,

$$\Phi_t = \mathrm{E}_{\mathbb{N}}(X |\, \mathcal{F}_t).$$

This observation helps to find sufficient conditions for the existence of a futures process.

11.6 Lemma. *If $\mathrm{E}_{\mathbb{N}} X^2 < \infty$ and N is a continuous numeraire of bounded variation that is bounded away from zero, then $\Phi_t = \mathrm{E}_{\mathbb{N}}(X |\, \mathcal{F}_t)$ is a futures process with $\Phi_T = X$.*

Proof. If $\mathrm{E}_{\mathbb{N}} X^2 < \infty$, then $\Phi_t := \mathrm{E}_{\mathbb{N}}(X |\, \mathcal{F}_t)$ is an $\mathbb{N}$-martingale bounded in L_2. If N^{-1} is bounded, then $G = N^{-1} \cdot \Phi$ is also an $\mathbb{N}$-martingale bounded in L_2, and hence Φ is a futures process. ∎

** **11.7** EXERCISE. Investigate uniqueness of the futures process, for instance under the conditions of the lemma.

11.8 EXERCISE. Calculate the futures price process for the settlement equal to the price S_T of a stock in the Black-Scholes model of Chapter 1.

The futures prices are often compared to the values F_t of a forward on the same underlying asset. Let N be a numeraire of bounded variation. In view of formula (10.9),

$$\Phi_t - F_t = \mathrm{E}_{\mathbb{N}}(X|\,\mathcal{F}_t) - \frac{\mathrm{E}_{\mathbb{N}}(X/N_T|\,\mathcal{F}_t)}{\mathrm{E}_{\mathbb{N}}(1/N_T|\,\mathcal{F}_t)} = -\frac{\mathrm{cov}_{\mathbb{N}}(X, 1/N_T|\,\mathcal{F}_t)}{\mathrm{E}_{\mathbb{N}}(1/N_T|\,\mathcal{F}_t)}.$$

This is called the *futures correction.* If the settlement X and the numeraire at expiry time are positively dependent, then the futures price is higher than the forward price.

The forward is also entered at zero cost at time t, and includes a total payment of $X - F_t$ over the contract period $[t, T]$. The difference is that this payment is made in one installment at the expiry time, whereas the payments $\Phi_T - \Phi_t = X - \Phi_t$ made under the futures contract are made continuously during the contract period.

11.9 Example (Euro-dollar contract). The Euro-dollar contract has claim $X = 100\bigl(1 - L_S[S,T]\bigr)$. The forward LIBOR $L_t[S,T]$ and the short rate r_t are typically positively dependent, and hence so are the forward LIBOR and the process $N_t = \exp\bigl(\int_0^T r_s\,ds\bigr)$. This implies that the claim X and the variable $1/N_T = \exp\bigl(-\int_0^T r_s\,ds\bigr)$ are also typically positively correlated. Assuming that N is a numeraire, the futures correction shows that $\Phi_t < F_t$. □

11.10 Example (Euro-dollar, V-H-W). We can calculate the futures process and the futures correction analytically in the Vasiček-Hull-White model. In this model the short rate process is an Ornstein-Uhlenbeck process with a deterministic drift and hence is a Gaussian process, under the martingale measure $\mathbb{R}$ corresponding to the numeraire $R_t = \exp\bigl(\int_0^t r_s\,ds\bigr)$.

The discount rates can be expressed in the short rate as $D_{t,T} = A_{t,T}e^{-B_{t,T}r_t}$ for the constants $A(t,T)$ and $B(t,T)$ given in Example 9.8, and hence by (10.1) the LIBOR satisfies $\alpha L_S[S,T] = \exp\bigl(-A(S,T) + B(S,T)r_S)\bigr) - 1$. We conclude that the futures price process for the settlement amount $X = L_S[S,T]$ is given by

$$\Phi_t = \frac{1}{\alpha}e^{-A(S,T)}\mathrm{E}_{\mathbb{R}}\Bigl(e^{B(S,T)r_S}|\,\mathcal{F}_t\Bigr) - \frac{1}{\alpha}.$$

Given $\mathcal{F}_t$ the random variable r_S possesses a log normal distribution, and hence the expectation is easy to compute.

The futures correction takes the form

$$\Phi_t - F_t = -\frac{1}{\alpha}e^{-A(S,T)}\frac{\mathrm{cov}_{\mathbb{R}}\bigl(e^{B(S,T)r_S}, e^{-\int_0^T r_s\,ds}\bigr)}{D_{t,T}/N_t}.$$

This is more work to compute, but just as straightforward, the vector $\bigl(r_S, \int_0^T r_s\,ds\bigr)$ possessing a bivariate normal distribution under $\mathbb{R}$. We omit the details. □

11.11 EXERCISE. Calculate the futures correction for the settlement equal to the price S_T of a stock in the Black-Scholes model of Chapter 1.

12
Swap Rate Models

The swap rate $y_t[T_0, T_1, \ldots, T_n]$, given by (10.4), is by its definition an indicator for the "interest rate" on the bond market in a given time interval. A *swap rate model* is a term structure that models the discount rates , or functions thereof, in terms of the swap rate. This is particularly attractive if we are interested in linear combinations of discount rates with maturities in the time interval spanned by the time points $T_0 < T_1 < \cdots < T_n$ only.

Throughout the chapter we fix the time points $0 < T_0 < T_1 < \cdots < T_n$ and abbreviate the swap rate and present value of a base point (PVBP) corresponding to these time points to $y_t = y_t[T_0, T_1, \ldots, T_n]$ and $P_t = P_t[T_0, T_1, \ldots, T_n]$. This gives corresponding stochastic processes y and P. The processes P and yP are linear combinations of discount rate processes and hence are numeraires in the bond market (provided that y is positive). We silently assume that the bond market is complete, so that there exist corresponding martingale measures $\mathbb{S}$ and $\mathbb{Y}$, giving the numeraire pairs $(P, \mathbb{S})$ and $(yP, \mathbb{Y})$.

12.1 Linear Swap Rate Model

The swap rate y is a difference of bond prices divided by the numeraire PVBP, and hence the swap rate y is a martingale under the corresponding martingale measure $\mathbb{S}$. For any maturity time T the process $t \mapsto D_{t,T}/P_t$ is an $\mathbb{S}$-martingale also. The *linear swap rate model* postulates that these two martingales are affinely proportional in that, for a constant A and a deterministic function $T \mapsto B_T$,

$$\frac{D_{t,T}}{P_t} = A + B_T\, y_t. \tag{12.1}$$

This expresses the discount rate $D_{t,T}$ into the discount rates D_{t,T_i}. The model leads to particularly simple calculations for derivatives that are functions of the quotients $D_{t,T}/P_t$, because their law (under the martingale measure) will be determined as soon as we know the distribution of the swap rate.

Suitable constants A and B_T for the linear swap rate model (12.1) can be derived from the implied relations

$$\begin{aligned} 1 &= \sum_i \alpha_i \frac{D_{t,T_i}}{P_t} = \sum_i \alpha_i \bigl(A + B_{T_i} y_t\bigr), \\ D_{0,T} &= \bigl(A + B_{T_i}\, y_0\bigr) P_0. \end{aligned}$$

If we want the first relation to hold for every value of the swap rate y_t, then we are led to set $\sum_i \alpha_i \, A = 1$. The initial bond prices $D_{0,T}$ and the initial swap rate P_0 are known from the market at time 0, and hence, given A, we can solve the constant B_T from the second equation.

12.2 Exponential Swap Rate Model

The *exponential swap rate model* postulates that all discount rates are determined by a single univariate process z through a relationship of the form

$$D_{t,T} = e^{-C_T z_t},$$

for a deterministic function $T \mapsto C_T$. In view of the definition of the swap rate (10.4) this implies the relationship

$$y_t = \frac{e^{-C_{T_0} z_t} - e^{-C_{T_n} z_t}}{\sum_i \alpha_i e^{-C_{T_i} z_t}}.$$

If the constants C_T are fixed this expresses z_t into the swap rate y_t. Substituting this inverse relationship in the exponential form postulated for $D_{t,T}$ we obtain a model for the discount rates in terms of the swap rates.

The constants C_T can be derived from the known values of the bond prices and the PVBP at time 0. The martingale property of the processes $t \mapsto D_{t,T}/P_t$ under the martingale measure $\mathbb{S}$ corresponding to the numeraire $P_t[T_0, \ldots, T_n]$ gives the relations

$$\frac{D_{0,T}}{P_0} = \mathrm{E}_{\mathbb{S}} \frac{D_{t,T}}{P_t} = \mathrm{E}_{\mathbb{S}} \frac{e^{-C_T z_t}}{\sum_i \alpha_i e^{-C_{T_i} z_t}}.$$

We may use these relations for a suitable choice of t together with a model for z_t (or indirectly the swap rate y_t) under the martingale measure $\mathcal{S}_r$ and solve for the constants C_T.

12.3 Calibration

The value at time 0 of a swaption with strike price K as considered in Section 10.5 is given by

$$V_0(K) = P_0 \, \mathrm{E}_{\mathbb{S}}\big(y_{T_0} - K\big)^+.$$

Because a vanilla swaption is a commonly traded instrument, this value is known from the market at time 0 for a large number of strike prices K. A formal differentiation of $V_0(K)$ relative to K yields

$$\frac{\partial}{\partial K} V_0(K) = -P_0 \, \mathrm{E}_{\mathbb{S}} 1_{y_{T_0} > K}.$$

The expectation on the right side is the survival function of the swap rate y_{T_0}. In principle, the left side and the PVBP P_0 are known from the market at time 0. This enables us to infer the distribution of the swap rate under the martingale measure $\mathbb{S}$ from observed market prices.

An alternative is to model the swap rate as a geometric Brownian motion, as in Section 10.5. The parameters of this model can then be inferrred from the prices of a few swaptions.

12.4 Convexity Corrections

Ad-hoc models used for option pricing, as used in Chapter 10, carry the danger that different, but related derivatives are not priced consistently. This could create the possibility of arbitrage within the model, which is certainly unrealistic, as the prices are based on the assumption of no arbitrage. It is thus important to price related products using a single model set-up. In this section we give an example.

Consider pricing two ($k = 1, 2$) European options that guarantee a single payment at time T_0 of sizes, for given times $0 < T_0 < T_1^{(k)} < \cdots < T_n^{(k)}$ and a given $\mathcal{F}_{T_0}$-measurable variable Z,

$$X^{(k)} = \Big(\sum_{i=1}^n c_i^{(k)} D_{T_0, T_i^{(k)}}\Big) Z, \qquad k = 1, 2.$$

The two claims possess the same form and include the same variable Z, but the maturities $T_i^{(k)}$ of the discount rates on which they are based and the corresponding weights $c_i^{(k)}$ may be different.

One way to price these claims individually would be to choose as a numeraire the process

$$Q_t^{(k)} = \sum_{i=1}^n c_i^{(k)} D_{t, T_i^{(k)}},$$

and to model the distribution of the variable Z under the corresponding martingale measure. The option price process would then be equal to $Q_t^{(k)}\mathrm{E}_{\mathbb{N}^{(k)}}(Z|\,\mathcal{F}_t)$, where $\mathbb{N}^{(k)}$ is the martingale measure. To price both options in this manner requires to model the variable Z under two different martingale measures, and it could be hard to make this modelling consistent between the two options.

A different approach is to adopt the linear swap rate model of Section 12.1, choosing suitable times $0 < T_0 < T_1 < \cdots < T_n$. In terms of the linear swap rate model we can write the claims as

$$X^{(k)} = \Bigl(A\sum_{i=1}^n c_i^{(k)}\Bigr)\,P_{T_0}\,Z + \Bigl(\sum_{i=1}^n c_i^{(k)}B_{T_i^{(k)}}\Bigr)\,y_{T_0}P_{T_0}\,Z.$$

Here $y_t = y_t[T_0, T_1, \ldots, T_n]$ and $P_t = P_t[T_0, T_1, \ldots, T_n]$, for times $T_0 < T_1 < \cdots < T_n$ that are chosen to represent the times $T_i^{(k)}$, but need not coincide with the latter times. The claims are sums of two terms, the first taking the form of a constant multiple of the product of the PVBP P_{T_0} and the variable Z, the second having the same form except that $y_{T_0}P_{T_0}$ takes the place of the PVBP. If the claim were given by only one of the two terms, then a convenient pricing strategy would be to choose the process P, or the process yP, as a numeraire, and to model the distribution of Z under the corresponding martingale measure. In the present situation we can use both numeraires, on the two corresponding terms of the option. A minor extension of the pricing formula (4.18) permits to express the value process of the claim $X^{(k)}$ as

$$V_t^{(k)} = P_t\Bigl(A\sum_{i=1}^n c_i^{(k)}\Bigr)\mathrm{E}_{\mathbb{S}}(Z|\,\mathcal{F}_t) + y_tP_t\Bigl(\sum_{i=1}^n c_i^{(k)}B_{T_i^{(k)}}\Bigr)\mathrm{E}_{\mathbb{Y}}(Z|\,\mathcal{F}_t).$$

Here $\mathbb{S}$ and $\mathbb{Y}$ are the martingale measures corresponding to the two numeraires P and yP. Expressing the process $Q_t^{(k)}$ in the two numeraire processes according to the linear swap rate model, we can write this in the form

$$V_t^{(k)} = Q_t^{(k)}\Bigl(w_t^{(k)}\mathrm{E}_{\mathbb{S}}(Z|\,\mathcal{F}_t) + (1-w_t^{(k)})\mathrm{E}_{\mathbb{Y}}(Z|\,\mathcal{F}_t)\Bigr),$$

for the "weights" given by

$$w_t^{(k)} = \frac{P_t/\sum_i \alpha_i}{Q_t^{(k)}/\sum_i c_i^{(k)}}.$$

Both numerator and denominator of the weights are a weighted average of discount rates, and the weights are unity if the times T_i and $T_i^{(k)}$ and the raw weights α_i and $c_i^{(k)}$ agree.

We can compare the formula for the value process $V^{(k)}$ with the formula obtained if the two claims were priced individually, using the numeraires $Q^{(k)}$. As we noted the latter procedure leads to a value process of the form $\tilde{V}^{(k)} = Q_t{}^{(k)}\mathrm{E}_{\mathbb{N}^{(k)}}(Z|\,\mathcal{F}_t)$, where $N^{(k)}$ is the martingale measure corresponding to the numeraire $Q^{(k)}$. Without a specification of the conditional expectations of the variable Z under the measures $N^{(k)}$, $\mathbb{S}$, and $\mathbb{Y}$, a direct comparison of the different pricing formulas is impossible. However, if we would employ simple ad-hoc models for the distributions of Z under these measures, then quite possibly we would use the same model under the measures $N^{(k)}$ and $\mathbb{S}$, as the numeraires $Q^{(k)}$ and P are similar in form: both numeraires are a linear combination of discount rates with maturities spread over a time interval following T_0. If we would use the same model, then the conditional expectations $\mathrm{E}_{N^{(k)}}(Z|\,\mathcal{F}_t)$ and $\mathrm{E}_{\mathbb{S}}(Z|\,\mathcal{F}_t)$ would be identical, and the difference $V^{(k)} - \tilde{V}^{(k)}$ between the two price processes would be equal to

$$Q_t^{(k)}(1 - w_t^{(k)})\big(\mathrm{E}_{\mathbb{Y}}(Z|\,\mathcal{F}_t) - \mathrm{E}_{\mathbb{S}}(Z|\,\mathcal{F}_t)\big).$$

This is known as a *convexity correction.*

To turn the expression for $V^{(k)}$ into a concrete formula, it suffices to specify the distribution of the variable Z under the measures $\mathbb{S}$ and $\mathbb{Y}$. Because these are the martingale measures corresponding to the numeraires P and yP these distributions are related, under the assumption of completeness, by their $\mathcal{F}_{T_0}$-density

$$\frac{d\mathbb{Y}}{d\mathbb{S}} = \frac{y_{T_0}P_{T_0}/(y_0P_0)}{P_{T_0}/P_0} = \frac{y_{T_0}}{y_0}.$$

The density process of $\mathbb{Y}$ relative to $\mathbb{S}$ is the process y/y_0. If we specify the distribution of Z under $\mathbb{S}$, then its distribution under $\mathbb{Y}$ is fixed.

A simple possibility is to model the swap rate as a geometric Brownian motion under $\mathbb{S}$, i.e. $dy_t = \sigma_t y_t\,dW_t$ for a deterministic function σ and an $\mathbb{S}$-Brownian motion W relative to the filtration $\mathcal{F}_t$. By Girsanov's theorem the process $\tilde{W}$ given by $\tilde{W}_t = W_t - (y/y_0)^{-1}\cdot[y/y_0, W]_t$ is an $\mathbb{Y}$-Brownian motion. Because $y = (\sigma y)\cdot W$, this process takes the form $\tilde{W}_t = W_t - \int_0^t \sigma_s\,dW_s$. It follows that

$$y_t = y_0 e^{\int_0^t \sigma_s\,dW_s - \frac{1}{2}\int_0^t \sigma_s^2\,ds} = y_0 e^{\int_0^t \sigma_s\,d\tilde{W}_s + \frac{1}{2}\int_0^t \sigma_s^2\,ds}.$$

If the function the function σ is strictly positive, then the augmented filtrations generated by the swap rate y and the driving Brownian motion W are the same. If we also assume that this filtration coincides with the given filtration $\mathcal{F}_t$, then the preceding model for the swap rate is enough to determine the conditional expectations $\mathrm{E}_{\mathbb{S}}(Z|\,\mathcal{F}_t)$ and $\mathrm{E}_{\mathbb{Y}}(Z|\,\mathcal{F}_t)$. Furthermore, if s is sufficiently regular, then the swap rate will be a Markov process, and these conditional expectations will be measurable functions of the swap rate y_t at time t.

* **12.2** EXERCISE. If $\mathrm{E}_{\mathbb{S}}(Z|\,\mathcal{F}_t) = F_t(y_t)$ for a measurable function F_t, show that $\mathrm{E}_{\mathbb{Y}}(Z|\,\mathcal{F}_t) = F_t\big(y_t \exp(\int_t^{T_0}\sigma_s^2\,ds)\big)$. [True?]