Introduction to Stochastic Calculus

These notes provide an introduction to stochastic calculus, the branch of mathematics that is most identified with financial engineering and mathematical finance. We will ignore most of the "technical" details and take an "engineering" approach to the subject. We will cover more material than is strictly necessary for this course. Any material that is not required, however, should be of value for other courses such as Term Structure Models.

We make the following assumptions throughout.

- There is a probability triple $(\Omega, \mathcal{F}, P)$ where

    - P is the "true" or *physical* probability measure
    - Ω is the universe of possible outcomes. We use $\omega \in \Omega$ to represent a generic outcome, typically a sample path(s) of a stochastic process(es).
    - the set[1] $\mathcal{F}$ represents the set of possible *events* where an event is a subset of Ω.

- There is also a *filtration*, $\{\mathcal{F}_t\}_{t \geq 0}$, that models the evolution of information through time. So for example, if it is known by time t whether or not an event, E, has occurred, then we have $E \in \mathcal{F}_t$. If we are working with a finite horizon, $[0, T]$, then we can take $\mathcal{F} = \mathcal{F}_T$.

- We also say that a stochastic process, X_t, is $\mathcal{F}_t$-adapted if the value of X_t is known at time t when the information represented by $\mathcal{F}_t$ is known. All the processes we consider will be $\mathcal{F}_t$-adapted so we will not bother to state this in the sequel.

- In the continuous-time models that we will study, it will be understood that the filtration $\{\mathcal{F}_t\}_{t \geq 0}$ will be the filtration *generated* by the stochastic processes (usually a Brownian motion, W_t) that are specified in the model description.

- Since these notes regularly refer to *numeraires* and *equivalent martingale measures* (EMMs), readers should be familiar with these concepts in advance. My *Martingale Pricing Theory* lecture note, for example, covers these topics in a discrete-time, discrete-space framework. In Section 12 of these notes we will discuss martingale pricing theory in the continuous-time setting and state without proof[2] the two Fundamental Theorems of Asset pricing.

1   Martingales and Brownian Motion

Definition 1 *A stochastic process*, $\{W_t \ : \ 0 \leq t \leq \infty\}$, *is a* standard Brownian motion *if*

1. $W_0 = 0$

2. *It has continuous sample paths*

3. *It has independent, normally-distributed increments.*

---

[1] Technically, $\mathcal{F}$ is a σ-algebra.

[2] These Theorems, however, are stated and proven in the discrete-time, discrete-space setting of the *Martingale Pricing Theory* lecture notes.

Definition 2 *An n-dimensional process, $W_t = (W_t^{(1)}, \ldots, W_t^{(n)})$, is a standard n-dimensional Brownian motion if each $W_t^{(i)}$ is a standard Brownian motion and the $W_t^{(i)}$'s are independent of each other.*

Definition 3 *A stochastic process, $\{X_t : 0 \leq t \leq \infty\}$, is a martingale with respect to the filtration, $\mathcal{F}_t$, and probability measure, P, if*

1. $\mathrm{E}^P[|X_t|] < \infty$ *for all $t \geq 0$*

2. $\mathrm{E}^P[X_{t+s}|\mathcal{F}_t] = X_t$ *for all $t, s \geq 0$.*

Example 1 (Brownian martingales)

Let W_t be a Brownian motion. Then W_t , $W_t^2 - t$ and $\exp\left(\theta W_t - \theta^2 t/2\right)$ are all martingales.

The latter martingale is an example of an *exponential martingale*. Exponential martingales are of particular significance since they are positive and may be used to define new probability measures. ∎

Exercise 1 (Conditional expectations as martingales) *Let Z be a random variable and set $X_t := \mathrm{E}[Z|\mathcal{F}_t]$. Show that X_t is a martingale.*

---

2 Quadratic Variation

Consider a partition of the time interval, $[0, T]$ given by

$$0 = t_0 < t_1 < t_2 < \ldots < t_n = T.$$

Let X_t be a Brownian motion and consider the sum of squared changes

$$Q_n(T) := \sum_{i=1}^{n} [\Delta X_{t_i}]^2 \tag{1}$$

where $\Delta X_{t_i} := X_{t_i} - X_{t_{i-1}}$.

Definition 4 (Quadratic Variation*) The quadratic variation of a stochastic process, X_t, is equal to the limit of $Q_n(T)$ as $\Delta t := \max_i(t_i - t_{i-1}) \to 0$.*

Theorem 1 *The quadratic variation of a Brownian motion is equal to T with probability 1.*

The functions with which you are normally familiar, e.g. continuous differentiable functions, have quadratic variation equal to zero. Note that any continuous stochastic process or function[3] that has non-zero quadratic variation must have infinite **total variation** where the total variation of a process, X_t, on $[0, T]$ is defined as

$$\text{Total Variation} := \lim_{\Delta t \to 0} \sum_{i=1}^{n} |X_{t_k} - X_{t_{k-1}}|.$$

This follows by observing that

$$\sum_{i=1}^{n} (X_{t_k} - X_{t_{k-1}})^2 \leq \left(\sum_{i=1}^{n} |X_{t_k} - X_{t_{k-1}}|\right) \max_{1 \leq k \leq n} |X_{t_k} - X_{t_{k-1}}|. \tag{2}$$

If we now let $n \to \infty$ in (2) then the continuity of X_t implies the impossibility of the process having finite total variation and non-zero quadratic variation. Theorem 1 therefore implies that the total variation of a Brownian motion is infinite. We have the following important result which will prove very useful when we price options when there are multiple underlying Brownian motions, as is the case with quanto options for example.

---

[3]A sample path of a stochastic process can be viewed as a function.

Theorem 2 *(Levy's Theorem) A continuous martingale is a Brownian motion if and only if its quadratic variation over each interval $[0, t]$ is equal to t.*

Another interesting result is the following:

Theorem 3 *Any non-constant continuous martingale must have infinite total variation.*

We know from our discrete-time models that any arbitrage-free model must have an equivalent martingale measure. The same is true in continuous-time models. Theorem 3 then implies that if we work in continuous-time with continuous (deflated) price processes, then these processes must have infinite total variation.

---

3 Stochastic Integrals

We now discuss the concept of a stochastic integral, ignoring the various technical conditions that are required to make our definitions rigorous. In this section, we write $X_t(\omega)$ instead of the usual X_t to emphasize that the quantities in question are stochastic.

Definition 5 *A* stopping time *of the filtration $\mathcal{F}_t$ is a random time, τ, such that the event $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t > 0$.*

In non-mathematical terms, we see that a stopping time is a random time whose value is part of the information accumulated by that time.

Definition 6 *We say a process, $h_t(\omega)$, is* elementary *if it is piece-wise constant so that there exists a sequence of stopping times $0 = t_0 < t_1 < \ldots < t_n = T$ and a set of $\mathcal{F}_{t_i}$-measurable[4] functions, $e_i(\omega)$, such that*

$$h_t(\omega) \;=\; \sum_i e_i(\omega) I_{[t_i, t_{i+1})}(t)$$

where $I_{[t_i, t_{i+1})}(t) = 1$ if $t \in [t_i, t_{i+1})$ and 0 otherwise.

Definition 7 *The stochastic integral of an elementary function, $h_t(\omega)$, with respect to a Brownian motion, W_t, is defined as*

$$\int_0^T h_t(\omega) \; dW_t(\omega) \;:=\; \sum_{i=0}^{n-1} e_i(\omega) \left(W_{t_{i+1}}(\omega) - W_{t_i}(\omega) \right). \tag{3}$$

Note that our definition of an elementary function assumes that the function, $h_t(\omega)$, is evaluated at the **left-hand point** of the interval in which t falls. This is a key component in the definition of the stochastic integral: without it the results below would no longer hold. Moreover, defining the stochastic integral in this way makes the resulting theory suitable for financial applications. In particular, if we interpret $h_t(\omega)$ as a trading strategy and the stochastic integral as the gains or losses from this trading strategy, then evaluating $h_t(\omega)$ at the left-hand point is equivalent to imposing the *non-anticipativity* of the trading strategy, a property that we always wish to impose.

For a more general process, $X_t(\omega)$, we have

$$\int_0^T X_t(\omega) \; dW_t(\omega) \;:=\; \lim_{n \to \infty} \int_0^T X_t^{(n)}(\omega) \; dW_t(\omega)$$

where $X_t^{(n)}$ is a sequence of elementary processes that converges (in an appropriate manner) to X_t.

---

[4]A function $f(\omega)$ is $\mathcal{F}_t$ measurable if its value is known by time t.

Example 2 We want to compute $\int_0^T W_t \, dW_t$. Towards this end, let $0 = t_0^n < t_1^n < t_2^n < \ldots < t_n^n = T$ be a partition of $[0, T]$ and define

$$X_t^n := \sum_{i=0}^{n-1} W_{t_i^n} I_{[t_i^n, t_{i+1}^n)}(t)$$

where $I_{[t_i^n, t_{i+1}^n)}(t) = 1$ if $t \in [t_i^n, t_{i+1}^n)$ and is 0 otherwise. Then X_t^n is an adapted elementary process and, by continuity of Brownian motion, satisfies $\lim_{n \to \infty} X_t^n = W_t$ almost surely as $\max_i |t_{i+1}^n - t_i^n| \to 0$. The Itô integral of X_t^n is given by

$$
\begin{aligned}
\int_0^T X_t^n \, dW_t &= \sum_{i=0}^{n-1} W_{t_i^n} (W_{t_{i+1}^n} - W_{t_i^n}) \\
&= \frac{1}{2} \sum_{i=0}^{n-1} \left(W_{t_{i+1}^n}^2 - W_{t_i^n}^2 - (W_{t_{i+1}^n} - W_{t_i^n})^2 \right) \\
&= \frac{1}{2} W_T^2 - \frac{1}{2} W_0^2 - \frac{1}{2} \sum_{i=0}^{n-1} (W_{t_{i+1}^n} - W_{t_i^n})^2.
\end{aligned}
\tag{4}
$$

By the definition of quadratic variation the sum on the right-hand-side of (4) converges in probability to T. And since $W_0 = 0$ we obtain

$$\int_0^T W_t \, dW_t = \lim_{n \to \infty} \int_0^T X_t^n \, dW_t = \frac{1}{2} W_T^2 - \frac{1}{2} T.$$

■

Note that we will generally evaluate stochastic integrals using Itô's Lemma (to be discussed later) without having to take limits of elementary processes as we did in Example 2.

Definition 8 *We define the space $L^2[0, T]$ to be the space of processes, $X_t(\omega)$, such that*

$$\mathrm{E} \left[\int_0^T X_t(\omega)^2 \, dt \right] < \infty.$$

Theorem 4 (Itô's Isometry) *For any $X_t(\omega) \in L^2[0, T]$ we have*

$$\mathrm{E} \left[\left(\int_0^T X_t(\omega) \, dW_t(\omega) \right)^2 \right] = \mathrm{E} \left[\int_0^T X_t(\omega)^2 \, dt \right].$$

Proof: (For the case where X_t is an elementary process)
Let $X_t = \sum_i e_i(\omega) I_{[t_i, t_{i+1})}(t)$ be an elementary process where the $e_i(\omega)$'s and t_i's are as defined in Definition 6. We therefore have $\int_0^T X_t(\omega) \, dW_t(\omega) := \sum_{i=0}^{n-1} e_i(\omega) \left(W_{t_{i+1}}(\omega) - W_{t_i}(\omega) \right)$ We then have

$$
\begin{aligned}
\mathrm{E} \left[\left(\int_0^T X_t(\omega) \, dW_t(\omega) \right)^2 \right] &= \mathrm{E} \left[\left(\sum_{i=0}^{n-1} e_i(\omega) \left(W_{t_{i+1}}(\omega) - W_{t_i}(\omega) \right) \right)^2 \right] \\
&= \sum_{i=0}^{n-1} \mathrm{E} \left[e_i^2(\omega) \left(W_{t_{i+1}}(\omega) - W_{t_i}(\omega) \right)^2 \right] \\
&\quad + 2 \sum_{0 \le i < j \le n-1} \mathrm{E} \left[e_i \, e_j(\omega) \left(W_{t_{i+1}}(\omega) - W_{t_i}(\omega) \right) \left(W_{t_{j+1}}(\omega) - W_{t_j}(\omega) \right) \right]
\end{aligned}
$$

$$= \sum_{i=0}^{n-1} \mathrm{E}\left[e_i^2(\omega) \underbrace{\mathrm{E}_{t_i}\left[\left(W_{t_{i+1}}(\omega) - W_{t_i}(\omega)\right)^2\right]}_{= t_{i+1}-t_i} \right]$$

$$+ 2 \sum_{0 \le i < j \le n-1} \mathrm{E}\left[e_i\, e_j(\omega)\left(W_{t_{i+1}}(\omega) - W_{t_i}(\omega)\right) \underbrace{\mathrm{E}_{t_j}\left[\left(W_{t_{j+1}}(\omega) - W_{t_j}(\omega)\right)\right]}_{=0} \right]$$

$$= \mathrm{E}\left[\sum_{i=0}^{n-1} e_i^2(\omega)(t_{i+1} - t_i) \right]$$

$$= \mathrm{E}\left[\int_0^T X_t(\omega)^2\, dt \right]$$

which is what we had to show. ∎

Theorem 5 (Martingale Property of Stochastic Integrals) *The stochastic integral,*
$Y_t := \int_0^t X_s(\omega)\, dW_s(\omega)$, *is a martingale for any* $X_t(\omega) \in L^2[0,T]$.

Exercise 2 *Check that* $\int_0^t X_s(\omega)\, dW_t(\omega)$ *is indeed a martingale when* X_t *is an elementary process. (*Hint: *Follow the steps we took in our proof of Theorem 4.)*

---

4   Stochastic Differential Equations

Definition 9 *An* n*-dimensional Itô process,* X_t*, is a process that can be represented as*

$$X_t = X_0 + \int_0^t a_s\, ds + \int_0^t b_s\, dW_s \tag{5}$$

where W *is an* m*-dimensional standard Brownian motion, and* a *and* b *are* n*-dimensional and* $n \times m$*-dimensional* $\mathcal{F}_t$*-adapted[5] processes, respectively[6].*

We often use the notation

$$dX_t = a_t\, dt + b_t dW_t$$

as shorthand for (5). An n-dimensional *stochastic differential equation* (SDE) has the form

$$dX_t = a(X_t, t)\, dt + b(X_t, t)\, dW_t; \quad X_0 = x \tag{6}$$

where as before, W_t is an m-dimensional standard Brownian motion, and a and b are n-dimensional and $n \times m$-dimensional adapted processes, respectively. Once again, (6) is shorthand for

$$X_t = x + \int_0^t a(X_s, s)\, dt + \int_0^t b(X_s, t)\, dW_s. \tag{7}$$

While we do not discuss the issue here, various conditions exist to guarantee existence and uniqueness of solutions to (7). A useful tool for solving SDE's is Itô's Lemma which we now discuss.

---

[5] a_t and b_t are $\mathcal{F}_t$-'adapted' if a_t and b_t are $\mathcal{F}_t$-measurable for all t. We always assume that our processes are $\mathcal{F}_t$-adapted.
[6] Additional technical conditions on a_t and b_t are also necessary.

5 Itô's Lemma

Itô's Lemma is the most important result in stochastic calculus, the "sine qua non" of the field. We first state and give an outline proof of a basic form of the result.

Theorem 6 (Itô's Lemma for 1-dimensional Brownian Motion)
Let W_t be a Brownian motion on $[0, T]$ and suppose $f(x)$ is a twice continuously differentiable function on $\mathbf{R}$. Then for any $t \leq T$ we have

$$f(W_t) \;=\; f(0) \;+\; \frac{1}{2} \int_0^t f''(W_s) \, ds \;+\; \int_0^t f'(W_s) \, dW_s. \tag{8}$$

Proof: (Sketch) Let $0 = t_0 < t_1 < t_2 < \ldots < t_n = t$ be a partition of $[0, t]$. Clearly

$$f(W_t) \;=\; f(0) \;+\; \sum_{i=0}^{n-1} \left(f(W_{t_{i+1}}) - f(W_{t_i}) \right). \tag{9}$$

Taylor's Theorem implies

$$f(W_{t_{i+1}}) - f(W_{t_i}) \;=\; f'(W_{t_i})(W_{t_{i+1}} - W_{t_i}) \;+\; \frac{1}{2} f''(\theta_i)(W_{t_{i+1}} - W_{t_i})^2 \tag{10}$$

for some $\theta_i \in (W_{t_i}, W_{t_{i+1}})$. Substituting (10) into (9) we obtain

$$f(W_t) \;=\; f(0) \;+\; \sum_{i=0}^{n-1} f'(W_{t_i})(W_{t_{i+1}} - W_{t_i}) \;+\; \frac{1}{2} \sum_{i=0}^{n-1} f''(\theta_i)(W_{t_{i+1}} - W_{t_i})^2. \tag{11}$$

If we let $\delta := \max_i |t_{i+1} - t_i| \to 0$ then it can be shown that the terms on the right-hand-side of (11) converge to the corresponding terms on the right-hand-side of (8) as desired. (This should not be surprising as we know the quadratic variation of Brownian motion on $[0, t]$ is equal to t.) ∎

A more general version of Itô's Lemma can be stated for Itô processes.

Theorem 7 (Itô's Lemma for 1-dimensional Itô process)
Let X_t be a 1-dimensional Itô process satisfying the SDE

$$dX_t \;=\; \mu_t \, dt + \sigma_t \, dW_t.$$

If $f(t, x) : [0, \infty) \times R \to R$ is a $C^{1,2}$ function and $Z_t := f(t, X_t)$ then

$$
\begin{aligned}
dZ_t \;&=\; \frac{\partial f}{\partial t}(t, X_t) \, dt + \frac{\partial f}{\partial x}(t, X_t) \, dX_t + \frac{1}{2} \frac{\partial^2 f}{\partial x^2}(t, X_t) \, (dX_t)^2 \\[2mm]
&=\; \left(\frac{\partial f}{\partial t}(t, X_t) + \frac{\partial f}{\partial x}(t, X_t) \, \mu_t + \frac{1}{2} \frac{\partial^2 f}{\partial x^2}(t, X_t) \, \sigma_t^2 \right) dt \;+\; \frac{\partial f}{\partial x}(t, X_t) \, \sigma_t \, dW_t
\end{aligned}
$$

The "Box" Calculus

In the statement of Itô's Lemma, we implicitly assumed that $(dX_t)^2 = \sigma_t^s \, dt$. The "box calculus" is a series of simple rules for calculating such quantities. In particular, we use the rules

$$
\begin{aligned}
dt \times dt \;=\; dt \times dW_t \;&=\; 0 \quad \text{and} \\
dW_t \times dW_t \;&=\; dt
\end{aligned}
$$

when determining quantities such as $(dX_t)^2$ in the statement of Itô's Lemma above. Note that these rules are consistent with Theorem 1. When we have two correlated Brownian motions, $W_t^{(1)}$ and $W_t^{(2)}$, with correlation coefficient, ρ_t, then we easily obtain that $dW_t^{(1)} \times dW_t^{(2)} = \rho_t \, dt$. We use the box calculus for computing the quadratic variation of Itô processes.

Exercise 3 *Let $W_t^{(1)}$ and $W_t^{(2)}$ be two independent Brownian motions. Use Levy's Theorem to show that*

$$W_t := \rho \, W_t^{(1)} + \sqrt{1 - \rho^2} \, W_t^{(2)}$$

is also a Brownian motion for a given constant, ρ.

Example 3

Suppose a stock price, S_t, satisfies the SDE

$$dS_t = \mu_t S_t \, dt + \sigma_t S_t \, dW_t.$$

Then we can use the substitution, $Y_t = \log(S_t)$ and Itô's Lemma applied to the function[7] $f(x) := \log(x)$ to obtain

$$S_t = S_0 \exp\left(\int_0^t (\mu_s - \sigma_s^2/2) \, ds + \int_0^t \sigma_s \, dW_s \right). \tag{12}$$

Note that S_t does not appear on the right-hand-side of (12) so that we have indeed solved the SDE. When $\mu_s = \mu$ and $\sigma_s \sigma$ are constants we obtain

$$S_t = S_0 \exp\left((\mu - \sigma^2/2) \, t + \sigma \, dW_t \right) \tag{13}$$

so that $\log(S_t) \sim \mathsf{N}\left((\mu - \sigma^2/2)t, \ \sigma^2 t \right)$. ∎

Example 4 (Ornstein-Uhlenbeck Process)

Let S_t be a security price and suppose $X_t = \log(S_t)$ satisfies the SDE

$$dX_t = [-\gamma(X_t - \mu t) + \mu] \, dt + \sigma dW_t.$$

Then we can apply Itô's Lemma to $Y_t := \exp(\gamma t) X_t$ to obtain

$$\begin{aligned}
dY_t &= \exp(\gamma t) \, dX_t + X_t \, d\left(\exp(\gamma t)\right) \\
&= \exp(\gamma t) \left([-\gamma(X_t - \mu t) + \mu] \, dt + \sigma dW_t\right) + X_t \gamma \exp(\gamma t) \, dt \\
&= \exp(\gamma t) \left([\gamma \mu t + \mu] \, dt + \sigma dW_t\right)
\end{aligned}$$

so that

$$Y_t = Y_0 + \mu \int_0^t e^{\gamma s} (\gamma s + 1) \, ds + \sigma \int_0^t e^{\gamma s} \, dW_s \tag{14}$$

or alternatively (after simplifying the Riemann integral in (14))

$$X_t = X_0 e^{-\gamma t} + \mu t + \sigma e^{-\gamma t} \int_0^t e^{\gamma s} \, dW_s. \tag{15}$$

Once again, note that X_t does not appear on the right-hand-side of (15) so that we have indeed solved the SDE. We also obtain $\mathrm{E}[X_t] = X_0 e^{-\gamma t} + \mu t$ and

$$\begin{aligned}
\mathrm{Var}(X_t) = \mathrm{Var}\left(\sigma e^{-\gamma t} \int_0^t e^{\gamma s} \, dW_s \right) &= \sigma^2 e^{-2\gamma t} \, \mathrm{E}\left[\left(\int_0^t e^{\gamma s} \, dW_s \right)^2 \right] \\
&= \sigma^2 e^{-2\gamma t} \int_0^t e^{2\gamma s} \, ds \qquad \text{(by Itô's Isometry)} \tag{16} \\
&= \frac{\sigma^2}{2\gamma} \left(1 - e^{-2\gamma t} \right). \tag{17}
\end{aligned}$$

---

[7] Note that $f(\cdot)$ is not a function of t.

These moments should be compared with the corresponding moments for $\log(S_t)$ in the previous example. ∎

Theorem 8 (Itô's Lemma for n-dimensional Itô process) *Let X_t be an n-dimensional Itô process satisfying the SDE*

$$dX_t = \mu_t \, dt + \sigma_t \, dW_t.$$

where $X_t \in \mathbf{R}^n$, $\mu_t \in \mathbf{R}^n$, $\sigma_t \in \mathbf{R}^{n \times m}$ and W_t is a standard m-dimensional Brownian motion. If $f(t, x) : [0, \infty) \times \mathbf{R}^n \to \mathbf{R}$ is a $C^{1,2}$ function and $Z_t := f(t, X_t)$ then

$$dZ_t = \frac{\partial f}{\partial t}(t, X_t) \, dt + \sum_i \frac{\partial f}{\partial x_i}(t, X_t) \, dX_t^{(i)} + \frac{1}{2} \sum_{i,j} \frac{\partial^2 f}{\partial x_i \partial x_j}(t, X_t) \, dX_t^{(i)} dX_t^{(j)}$$

where $dW_t^{(i)} dW_t^{(j)} = dt \, dW_t^{(i)} = 0$ for $i \neq j$ and $dW_t^{(i)} dW_t^{(i)} = dt$.

Exercise 4 *Let X_t and Y_t satisfy*

$$
\begin{aligned}
dX_t &= \mu_t^{(1)} \, dt + \sigma_t^{(1,1)} \, dW_t^{(1)} \\
dY_t &= \mu_t^{(2)} \, dt + \sigma_t^{(2,1)} \, dW_t^{(1)} + \sigma_t^{(2,2)} \, dW_t^{(2)}
\end{aligned}
$$

and define $Z_t := X_t Y_t$. Apply the multi-dimensional version of Itô's Lemma to find the SDE satisfied by Z_t.

---

6 The Martingale Representation Theorem

The martingale representation theorem is an important result that is particularly useful for constructing replicating portfolios in complete financial models.

Theorem 9 *Suppose M_t is an $\mathcal{F}_t$-martingale where $\{\mathcal{F}_t\}_{t \geq 0}$ is the filtration generated by the n-dimensional standard Brownian motion, $W_t = (W_t^{(1)}, \ldots, W_t^{(n)})$. If $\mathrm{E}[M_t^2] < \infty$ for all t then there exists a unique[8] n-dimensional adapted stochastic process, ϕ_t, such that*

$$M_t = M_0 + \int_0^t \phi_s^T \, dW_t \quad \text{for all } t \geq 0$$

where ϕ_s^T denotes the transpose of the vector, ϕ_s.

Example 5 Let $F = W_T^3$ and define $M_t = \mathrm{E}_t[F]$. We will show that

$$M_t = 3 \int_0^t (T - s + W_s^2) \, dW_s \tag{18}$$

which is consistent with the Martingale Representation theorem. First we must calculate M_t. We do this using the independent increments property of Brownian motion and obtain

$$
\begin{aligned}
M_t = \mathrm{E}_t[W_T^3] &= \mathrm{E}_t[(W_T - W_t + W_t)^3] \\
&= \underbrace{\mathrm{E}_t[(W_T - W_t)^3]}_{=\,0} + \mathrm{E}_t[W_t^3] + 3\mathrm{E}_t[W_t(W_T - W_t)^2] + 3\underbrace{\mathrm{E}_t[W_t^2(W_T - W_t)]}_{=\,0} \\
&= W_t^3 + 3W_t(T - t). \tag{19}
\end{aligned}
$$

---

[8]To be precise, additional integrability conditions are required of ϕ_s in order to claim that it is unique.

We can now apply Itô's Lemma to (19) to obtain

$$
\begin{aligned}
dM_t &= 3W_t^2 \, dW_t \; + \; \frac{1}{2} \, 6W_t \, dt \; + \; 3(T-t) \, dW_t \; - \; 3W_t \, dt \\
&= 3(W_t^2 + T - t) \, dW_t
\end{aligned}
$$

which, noting $M_0 = 0$, is (18). ∎

---

7 Gaussian Processes

Definition 10 *A process X_t, $t \geq 0$, is a Gaussian process if $(X_{t_1}, \ldots, X_{t_n})$ is jointly normally distributed for every n and every set of times $0 \leq t_1 \leq t_2 \leq \ldots \leq t_n$.*

If X_t is a Gaussian process, then it is determined by its *mean function*, $m(t)$, and its covariance function, $\rho(s,t)$, where

$$
\begin{aligned}
m(t) &= \mathrm{E}\left[X_t\right] \\
\rho(s,t) &= \mathrm{E}\left[(X_s - m(s))(X_t - m(t))\right].
\end{aligned}
$$

In particular, the joint moment generating function (MGF) of $(X_{t_1}, \ldots, X_{t_n})$ is given by

$$
M_{t_1,\ldots,t_n}(\theta_1,\ldots,\theta_n) \;=\; \exp\left(\theta^T m(\mathbf{t}) + \frac{1}{2} \, \theta^T \Sigma \theta\right) \tag{20}
$$

where $m(\mathbf{t}) = (m(t_1) \; \ldots \; m(t_n))^T$ and $\Sigma_{i,j} = \rho(t_i, t_j)$.

Example 6 (**Brownian motion**)

Brownian motion is a Gaussian process with $m(t) = 0$ and $\rho(s,t) = \min(s,t)$ for all $s, t \geq 0$. ∎

Theorem 10 (**Integration of a deterministic function w.r.t. a Brownian motion**) *Let W_t be a Brownian motion and suppose*

$$
X_t \;=\; \int_0^t \delta_s \, dW_s
$$

where δ_s is a deterministic function. Then X_t is a Gaussian process with $m(t) = 0$ and $\rho(s,t) = \int_0^{\min(s,t)} \delta_s^2 \, ds$.

Proof: (Sketch)

(i) First use Itô's Lemma to show that

$$
\mathrm{E}\left[e^{uX_t}\right] \;=\; 1 + \frac{1}{2}u^2 \int_0^t \delta_s^2 \, \mathrm{E}\left[e^{uX_s}\right] \, ds. \tag{21}
$$

If we set $y_t := \mathrm{E}\left[e^{uX_t}\right]$ then we can differentiate across (21) to obtain the ODE

$$
\frac{dy}{dt} \;=\; \frac{1}{2}u^2 \delta_t^2 y.
$$

This is easily solved to obtain the MGF for X_t,

$$
\mathrm{E}\left[e^{uX_t}\right] \;=\; \exp\left(\frac{1}{2}u^2 \int_0^t \delta_s^2 \, ds\right) \tag{22}
$$

which, as expected, is the MGF of a normal random variable with mean 0 and variance $\int_0^t \delta_s^2\ ds$.

(ii) We now use (22) and similar computations to show that the joint MGF of $(X_{t_1}, \ldots, X_{t_n})$ has the form given in (20) with $m(\mathbf{t}) = 0$ and $\rho(s,t) = \int_0^{\min(s,t)} \delta_s^2\ ds$. (See Shreve's *Stochastic Calculus for Finance II* for further details.) ∎

The next theorem again concerns Gaussian processes and is often of interest[9] when studying short-rate term structure models.

Theorem 11 *Let W_t be a Brownian motion and suppose δ_t and ϕ_t are deterministic functions. If*

$$X_t := \int_0^t \delta_u\ dW_u \quad and \quad Y_t := \int_0^t \phi_u X_u\ du$$

then Y_t is a Gaussian process with $m(t) = 0$ and

$$\rho(s,t) = \int_0^{\min(s,t)} \delta_v^2 \left(\int_v^s \phi_y\ dy \right) \left(\int_v^t \phi_y\ dy \right)\ dv.$$

Proof: The proof is tedious but straightforward. (Again, see Shreve's *Stochastic Calculus for Finance II* for further details.) ∎

Note for example that Brownian motion with drift and the Ornstein-Uhlenbeck process are both Gaussian processes. Nonetheless, we saw in Example 4 that these two processes behave very differently: the Ornstein-Uhlenbeck process mean-reverts and its variance tends to a finite limit as $T \to \infty$. This is not true of Brownain motion with or without drift.

---

8   The Feynman-Kac Formula

Suppose X_t is a stochastic process satisfying the SDE $dX_t = \mu(t, X_t)\ dt + \sigma(t, X_t)\ dW_t$. Now consider the function, $f(x, t)$, given by

$$f(t, x) = \mathrm{E}_t^x \left[\int_t^T \phi_s^{(t)} h(X_s, s)\ ds + \phi_T^{(t)} g(X_T) \right]$$

where

$$\phi_s^{(t)} = \exp\left(-\int_t^s r(X_u, u)\ du \right)$$

and the notation $\mathrm{E}_t^x[\cdot]$ implies that the expectation should be taken conditional on time t information with $X_t = x$. Note that $f(x, t)$ may be interpreted as the time t price of a security that pays dividends at a continuous rate, $h(X_s, s)$ for $s \geq t$, and with a terminal payoff $g(X_T)$ at time T. Of course $\mathrm{E}[\cdot]$ should then be interpreted as an expectation under an equivalent martingale measure with the cash account as the corresponding numeraire and $r(\cdot, \cdot)$ as the instantaneous risk-free rate.

The Feynman-Kac Theorem[10] states that $f(\cdot, \cdot)$ satisfies the following PDE

$$\frac{\partial f}{\partial t}(t, X_t) + \frac{\partial f}{\partial x}(t, X_t)\ \mu_t(t, X_t) + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}(t, X_t)\ \sigma_t^2(t, X_t) - r(x, t) f(x, t) + h(x, t) = 0, \quad (x, t) \in \mathbf{R} \times [0, T)$$

$$f(x, T) = g(x), \quad x \in \mathbf{R}$$

---

[9]See, for example, Hull and White's one-factor model.

[10]Additional technical conditions on μ, σ, r, h, g and f are required.

Proof: (Sketch proof)
We will use the martingale property of conditional expectations which states that if $M_t := \mathrm{E}_t[F]$ where F is a given fixed random variable then M_t is a martingale. First let

$$f(t, X_t) \ := \ \mathrm{E}_t \left[\int_t^T \phi_s^{(t)} h(X_s, s) \ ds + \phi_T^{(t)} g(X_T) \right] \qquad (23)$$

and note that the random variable inside the expectation on the right-hand-side of (23) is a function of t. This means we cannot yet apply the martingale property of conditional expectations. However by adding $\int_0^t \phi_s^{(t)} h(X_s, s) \ ds$ to both sides of (23) and then multiplying both sides by $\phi_t^{(0)}$ we obtain

$$\phi_t^{(0)} \left(f(t, X_t) + \int_0^t \phi_s^{(t)} h(X_s, s) \ ds \right) \ = \ \mathrm{E}_t \left[\int_0^T \phi_s^{(0)} h(X_s, s) \ ds + \phi_T^{(0)} g(X_T) \right] \qquad (24)$$

$$= \ \mathrm{E}_t \left[Z \right], \quad \text{say.}$$

Note that Z does not depend on t and so we can now apply the martingale property of conditional expectations to conclude that the left-hand-side of (24) is a martingale. We therefore apply Itô's Lemma to the left-hand-side and set the coefficient of dt to zero. This results in the Feynman-Kac PDE given above. ∎

Remark 1 *The Feynman-Kac result generalizes easily to the case where X_t is an n-dimensional Itô process driven by an m-dimensional standard Brownian motion.*

The Feynman-Kac Theorem plays an important role in financial engineering as it enables us to compute the prices of derivative securities (which can be expressed as expectations according to martingale pricing theory) by solving a PDE instead. It is worth mentioning, however, that in the development of derivatives pricing theory, the PDE approach preceded the martingale approach. Indeed by constructing a replicating[11] strategy for a given derivative security, it can be shown *directly* that the price of the derivative security satisfies the Feynman-Kac PDE.

---

9 Change of Probability Measure

The majority of financial engineering models price securities using the EMM, Q, that corresponds to taking the cash account, B_t, as numeraire. Sometimes, however, it is particularly useful to work with another numeraire, N_t, and its corresponding EMM, P_N say. We now describe how to create new probability measures and how to switch back and forth between these measures.

Let Q be a given probability measure and M_t a strictly positive Q-martingale such that $\mathrm{E}_0^Q[M_t] = 1$ for all $t \in [0, T]$. We may then define a new equivalent probability measure, P^M, by defining

$$P_M(A) \ = \ \mathrm{E}^Q \left[M_T 1_A \right].$$

Note that

(i) $P_M(\Omega) = 1$

(ii) $P_M(A) \geq 0$ for every event A and

(iii) It can easily be shown that

$$P_M \left(\bigcup_i A_i \right) \ = \ \sum_i P_M(A_i)$$

whenever the A_i's form a sequence of disjoint sets.

---

[11]We are implicitly assuming that the security can be replicated.

Points (i), (ii) and (iii) above imply that P_M is indeed a probability measure. Because the *null-sets* of Q and P_M coincide we can conclude that P_M is indeed an equivalent probability measure (to Q). Expectations with respect to P_M then satisfy

$$\mathrm{E}_0^{P_M}[X] \ = \ \mathrm{E}_0^Q[M_T X]. \tag{25}$$

Exercise 5 *Verify (25) in the case where $X(\omega) = \sum_{i=1}^n c_i I_{\{\omega \in A_i\}}$, M_T is constant on each A_i, and where $A_1, \ldots, A_n$ form a partition[12] of Ω.*

When we define a measure change this way, we use the notation dP_M/dQ to refer to M_T so that we often write

$$\mathrm{E}_0^{P_M}[X] \ = \ \mathrm{E}_0^Q\left[\frac{dP_M}{dQ} X\right].$$

The following result explains how to switch between Q and P_M when we are taking conditional expectations. In particular, we have

$$\mathrm{E}_t^{P_M}[X] \ = \ \frac{\mathrm{E}_t^Q\left[\frac{dP_M}{dQ} X\right]}{\mathrm{E}_t^Q\left[\frac{dP_M}{dQ}\right]} \tag{26}$$

$$= \ \frac{\mathrm{E}_t^Q\left[\frac{dP_M}{dQ} X\right]}{M_t}. \tag{27}$$

Since M_t is a Q-martingale (27) follows easily from (26). We shall not prove (26), however, as it is more difficult and requires the measure-theoretic definition of a conditional expectation.

Exercise 6 *Show that if X is $\mathcal{F}_t$-measurable, i.e. X is known by time t, then $\mathrm{E}_0^{P_M}[X] = \mathrm{E}_0^Q[M_t X]$.*

Remark 2 *Since M_T is strictly positive we can set $X = I_A/M_T$ in (25) where I_A is the indicator function of the event A. We then obtain $\mathrm{E}_0^{P_M}[I_A/M_T] = \mathrm{E}_0^Q[I_A] = Q(A)$. In particular, we see that dQ/dP_M is given by $1/M_T$.*

Remark 3 *In the context of security pricing, we can take M_t to be the deflated time t price of a security with strictly positive payoff, normalized so that its expectation under Q is equal to 1. For example, let Z_t^T be the time t price of a zero-coupon bond maturing at time T, and let[13] B_t denote the time t value of the cash account. We could then set[14] $M_T := 1/(B_T Z_0^T)$ so that*

$$M_t = \mathrm{E}_t^Q\left[\frac{1}{B_T Z_0^T}\right]$$

is indeed a Q-martingale. The resulting measure, denoted by P_T, is sometimes called the T-forward measure. Note that we have implicitly assumed (why?!) that in this context, Q refers to the EMM when we take the cash account as numeraire. We discuss P_T in further detail in Section 11.

---

[12] The A_i's form a partition of Ω if $A_i \cap A_j = \emptyset$ whenever $i \neq j$ and $\bigcup_i A_i = \Omega$.

[13] We assume the zero-coupon bond has face value \$1 and $B_0 = \$1$.

[14] Note that it is **not** the case that $M_t = 1/(B_t Z_0^t)$ which is not a Q-martingale.

10 Girsanov's Theorem

Girsanov's Theorem is one of the most important results for financial engineering applications. When working with models driven by Brownian motions[15] it enables us to (i) identify the equivalent martingale measure(s) corresponding to a given numeraire and (ii) to move back and forth between different EMM-numeraire pairs. Consider then the process

$$L_t := \exp\left(-\int_0^t \eta_s \, dW_s - \frac{1}{2}\int_0^t \eta_s^2 \, ds\right) \tag{28}$$

where η_s is an adapted process and W_s is a P-Brownian motion. Using Itô's Lemma we can check that $dL_t = -L_t\eta_t \, dW_t$ so L_t is a positive martingale[16] with $E_0^P[L_t] = 1$ for all t.

Theorem 12 (Girsanov's Theorem) *Define an equivalent probability measure, Q^η, by setting*

$$Q^\eta(A) := E_0^P[L_T 1_A]. \tag{29}$$

Then $\widehat{W}_t := W_t + \int_0^t \eta_s ds$ is a standard Q^η-Brownian motion. Moreover, $\widehat{W}_t$ has the martingale representation property under Q^η.

Remark 4 *Suppose $\eta_s = \eta$, a constant. Then since $\widehat{W}_t := W_t + \eta t$ is a standard Q^η-Brownian motion, it implies that $W_t = \widehat{W}_t - \eta t$ is a Q-Brownian motion with drift $-\eta$.*

Example 7 Let $dX_t = \mu_t \, dt + \sigma_t \, dW_t$ and suppose we wish to find a process, η_s, such that X_t is a Q^η-martingale. This is easily achieved as follows:

$$\begin{aligned} dX_t &= \mu_t \, dt + \sigma_t \, dW_t \\ &= \mu_t \, dt + \sigma_t \, (d\widehat{W}_t - \eta_t) \, dt \\ &= \sigma_t \, d\widehat{W}_t \end{aligned}$$

if we set $\eta_t = \mu_t/\sigma_t$ in which case X_t is a Q^η-martingale. We can obtain some intuition for this result: suppose for example that μ_t and σ_t are both positive so that η_t is also positive. Then we can see from the definition of L_t in (28) that L_t places less weight on paths where W_t drifts upwards than it does on paths where it drifts downwards. This relative weighting adjusts X_t for the positive drift induced by μ_t with the result that under Q^η, X_t is a martingale. ∎

Remark 5 *Note that Girsanov's Theorem enables us to compute Q^η-expectations directly without having to switch back to the original measure, P.*

We can get some additional intuition for the Girsanov Theorem by considering a random walk, $\mathbf{X} = \{X_0 = 0, X_1, \ldots, X_n\}$ with the interpretation that X_i is the value of the walk at time iT/n. In particular, X_n corresponds to the value of the random walk at time T. We assume that $\Delta X_i := X_i - X_{i-1} \sim N(0, T/n)$ under P and is independent of $X_0, \ldots, X_{i-1}$ for $i = 1, \ldots n$.

Suppose now that we want to compute $\theta := E_0^Q[h(\mathbf{X})]$ where Q denotes the probability measure under which $\Delta X_i \sim N(\mu, T/n)$, again independently of $X_0, \ldots, X_{i-1}$. In particular, if we set $\mu := -T\eta/n$ then X_t approximates standard Brownian motion under P and X_t approximates Brownian motion with drift $-\eta$ under Q.

---

[15] Versions of Girsanov's Theorem are also available for jump-diffusion and other processes.

[16] In fact we need η_s to have some additional properties before we can claim L_t is a martingale. A sufficient condition is Novikov's Condition which requires $E_0^P\left[\exp\left(\frac{1}{2}\int_0^T \eta_s^2 \, ds\right)\right] < \infty$.

If we let $f(\cdot)$ and $g(\cdot)$ denote the PDF's of $N(\mu, T/n)$ and $N(0, T/n)$ random variables, respectively, then we obtain

$$
\begin{aligned}
\theta \;=\; \mathrm{E}_0^Q[h(\mathbf{X})] \;&=\; \int_{R^n} h(x_1, \ldots, x_n) \left(\prod_{i=1}^n f(\Delta x_i) \right) d\Delta x_1 \ldots d\Delta x_n \\[2mm]
&=\; \int_{R^n} h(x_1, \ldots, x_n) \prod_i \left(\frac{f(\Delta x_i)}{g(\Delta x_i)} g(\Delta x_i) \right) d\Delta x_1 \ldots d\Delta x_n \\[2mm]
&=\; \int_{R^n} h(x_1, \ldots, x_n) \prod_i \left(\frac{f(\Delta x_i)}{g(\Delta x_i)} \right) \left(\prod_i g(\Delta x_i) \right) d\Delta x_1 \ldots d\Delta x_n \\[2mm]
&=\; \mathrm{E}_0^P \left[h(X_1, \ldots, X_n) \prod_i \left(\frac{f(\Delta X_i)}{g(\Delta X_i)} \right) \right] \\[2mm]
&=\; \mathrm{E}_0^P \left[h(X_1, \ldots, X_n) \, \exp\left(-\eta \sum_i \Delta X_i - \frac{\eta^2 T}{2} \right) \right] \\[2mm]
&=\; \mathrm{E}_0^P \left[h(X_1, \ldots, X_n) \, \exp\left(-\eta X_n - \frac{\eta^2 T}{2} \right) \right]
\end{aligned}
$$

which is consistent with our statement of Girsanov's Theorem in (28) and (29) above. (See Remark 4.)

Remark 6 (i) *As in the statement of the Girsanov Theorem itself, we could have chosen μ (and therefore η) to be adapted, i.e. to depend on prior events, in the random walk.*

(ii) *Note that Girsanov's Theorem allows the drift, but not the volatility of the Brownian motion, to change under the new measure, Q^η. It is interesting to see that we are not so constrained in the case of the random walk. Have you any intuition for why this is so?*

Multidimensional Girsanov's Theorem

The multidimensional version of Girsanov's Theorem is a straightforward generalization of the one-dimensional version. In particular let W_t be an n-dimensional standard P-Brownian motion and define

$$
L_t \;:=\; \exp\left(-\int_0^t \eta_s^T \, dW_s - \frac{1}{2} \int_0^t \eta_s^T \, \eta_s \, ds \right)
$$

for $t \in [0, T]$. Then[17] $\widehat{W}_t := W_t + \int_0^t \eta_s ds$ is a standard Q^η-Brownian motion where $dQ^\eta/dP = L_T$.

---

[17] Again it is necessary to make some further assumptions in order to guarantee that L_t is a martingale. Novikov's condition is sufficient.

11 A Useful EMM: The Forward Measure

We now discuss the τ-*forward measure*, P^τ, which is the EMM that corresponds to taking the zero-coupon bond with maturity τ as the numeraire. We therefore let Z_t^τ denote the time t price of a zero-coupon bond maturing at time $\tau \geq t$ and with face value \$1. As usual, we let Q denote the EMM corresponding to taking the cash account, B_t, as numeraire. We assume without loss of generality that $B_0 = \$1$ and now use Z_t^τ to define our new measure. To do this, set

$$\frac{dP^\tau}{dQ} = \frac{1}{B_\tau Z_0^\tau}. \tag{30}$$

Exercise 7 *Check that (30) does indeed define an equivalent probability measure.*

Now let C_t denote the time t price of a contingent claim that expires at time τ. We then have

$$
\begin{aligned}
C_t &= B_t \mathrm{E}_t^Q \left[\frac{C_\tau}{B_\tau} \right] & \text{(by martingale pricing with EMM } Q\text{)} & \tag{31} \\[2mm]
&= \frac{B_t \mathrm{E}_t^{P^\tau} \left[\frac{C_\tau}{B_\tau} B_\tau Z_0^\tau \right]}{\mathrm{E}_t^{P^\tau} [B_\tau Z_0^\tau]} & \text{(going from } Q \text{ to } P^\tau \text{ using (26))} & \tag{32} \\[2mm]
&= \frac{B_t Z_0^\tau \mathrm{E}_t^{P^\tau} [C_\tau]}{\mathrm{E}_t^Q [1] / \mathrm{E}_t^Q [1/(B_\tau Z_0^\tau)]} & \text{(going from } P^\tau \text{ to } Q \text{in the denominator of (32) using (26))} & \\[2mm]
&= Z_t^\tau \mathrm{E}_t^{P^\tau} [C_\tau] & \text{(by martingale pricing with EMM } Q\text{).} & \tag{33}
\end{aligned}
$$

We can now find C_t, either through equation (33) or through equation (31) where we use the cash account as numeraire. Computing C_t through (31) is our "usual method" and is often very convenient. When pricing equity derivatives, for example, we usually take interest rates, and hence the cash account, to be deterministic. This means that the factor $1/B_\tau$ in (31) can be taken outside the expectation so only the Q-distribution of C_τ is needed to compute C_t.

When interest rates are stochastic we cannot take the factor $1/B_\tau$ outside the expectation in (31) and we therefore need to find the *joint* Q-distribution of (B_τ, C_τ) in order to compute C_t. On the other hand, if we use equation (33) to compute C_t, then we only need the P^τ-distribution of C_τ, regardless of whether or not interest rates are stochastic. Working with a univariate-distribution is generally much easier than working with a bivariate-distribution so if we can easily find the P^τ-distribution of C_τ, then it can often be very advantageous to work with this distribution. The forward measure is therefore particularly useful when studying term-structure models.

Switching to a different numeraire can also be advantageous in other circumstances, even when interest rates are deterministic. For example, we will find it convenient to do so when pricing quanto-options and options on multiple underlying securities. Moreover, we will see how we can use Girsanov's Theorem to move back and forth between Q and other EMMs corresponding to other numeraires.

12 Applying Stochastic Calculus: Martingale Pricing Theory

We now briefly outline how the martingale pricing theory that we have seen in a discrete-time, discrete-space framework translates to a continuous-time setting. We use S_t to denote the time t price of a risky asset and B_t to denote the time t value of the cash account. For ease of exposition we assume the risky asset does not[18] pay dividends. Let $\phi_t^{(s)}$ and $\phi_t^{(b)}$ denote the number of units of the security and cash account, respectively, that is held in a portfolio at time t. Then the value of the portfolio at time t is given by $V_t = \phi_t^{(s)} S_t + \phi_t^{(b)} B_t$.

Definition 11 *We say* $\phi_t := (\phi_t^{(s)}, \phi_t^{(b)})$ *is* self-financing *if*

$$dV_t \;=\; \phi_t^{(s)} \; dS_t \;+\; \phi_t^{(b)} \; dB_t.$$

Note that this definition is consistent with our definition for discrete-time models. It is also worth emphasizing the mathematical content of this definition. Note that Itô's Lemma implies

$$dV_t = \phi_t^{(s)} \; dS_t \;+\; \phi_t^{(b)} \; dB_t \;+\; \underbrace{S_t \; d\phi_t^{(s)} \;+\; B_t \; d\phi_t^{(b)} \;+\; dS_t \; d\phi_t^{(s)} \;+\; dB_t \; d\phi_t^{(b)}}_{A}$$

so that the self-financing assumption amounts to assuming that the sum of the terms in A are identically zero.

Our definitions of *arbitrage*, *numeraire securities*, *equivalent martingale measures* and *complete* markets is unchanged from the discrete-time setup. We now state[19] without proof the two fundamental theorems of asset pricing. These results mirror those from the discrete-time theory.

Theorem 13 (The First Fundamental Theorem of Asset Pricing) *There is no arbitrage if and only if there exists an EMM, Q.*

A consequence of Theorem 13 is that in the absence of arbitrage, the *deflated* value process, V_t/N_t, of any self-financing trading strategy is a Q-martingale. This implies that the deflated price of any attainable security can be computed as the Q-expectation of the terminal deflated value of the security.

Theorem 14 (The Second Fundamental Theorem of Asset Pricing) *Assume there exists a security with strictly positive price process and that there are no arbitrage opportunities. Then the market is complete if and only if there exists exactly one risk-neutral martingale measure, Q.*

Beyond their theoretical significance, these theorems are also important in practice. If our model is complete then we can price securities with the unique EMM that accompanies it, assuming of course the absence of arbitrage. This is the case for the Black-Scholes model as well as local volatility models. If our model is incomplete then we generally work directly under an EMM, Q, which is then calibrated to market data. In particular, the true data-generating probability measure, P, is often completely ignored when working with incomplete models. In some sense the issue of completeness then only arises when we discuss replicating or hedging strategies.

Determining Replicating Strategies

The following example is particularly useful in many financial engineering applications. While we focus on the case of a single risky security the result generalizes easily to multiple risky securities. We assume the market is complete of course, since only in that case can we guarantee that a replicating strategy actually exists.

---

[18] We can easily adapt our definition of a self-financing trading strategy to accommodate securities that pay dividends.

[19] Additional technical conditions are generally required to actually prove these results.

Example 8 (Wealth Dynamics and Hedging)

We know the value of the cash-account, B_t, satisfies $dB_t = r_t B_t\ dt$ and suppose in addition that S_t satisfies

$$dS_t\ =\ \mu_t S_t\ dt\ +\ \sigma_t S_t dW_t. \tag{34}$$

Then for a portfolio $(\phi_t^{(s)}, \phi_t^{(b)})$, the portfolio value, V_t, at time t satisfies $V_t := \phi_t^{(s)} S_t + \phi_t^{(b)} B_t$. If $\phi_t := (\phi_t^{(s)}, \phi_t^{(b)})$ is self-financing, then we have

$$
\begin{aligned}
dV_t\ &=\ \phi_t^{(s)}\ dS_t\ +\ \phi_t^{(b)}\ dB_t \\[2mm]
&=\ \phi_t^{(s)} \mu_t S_t\ dt + \phi_t^{(s)} \sigma_t S_t\ dW_t\ +\ \phi_t^{(b)} r_t B_t\ dt \\[2mm]
&=\ V_t \left[\frac{\phi_t^{(s)} S_t}{V_t} \mu_t\ +\ \frac{\phi_t^{(b)} B_t}{V_t} r_t \right]\ dt\ +\ \frac{\phi_t^{(s)} S_t}{V_t} \sigma_t V_t\ dW_t \\[2mm]
&=\ V_t \left[r_t\ +\ \theta_t (\mu_t - r_t) \right]\ dt\ +\ \theta_t \sigma_t V_t\ dW_t
\end{aligned}
\tag{35}
$$

where θ_t and $(1 - \theta_t)$ are the *fractions* of time t wealth, V_t, invested in the risky asset and cash account, respectively, at time t. ∎

Suppose now that V_t is the time t value of some derivative security and that we want to determine the self-financing trading strategy, $(\phi_t^{(s)}, \phi_t^{(b)})$, that replicates it. By writing the dynamics of V_t as in (35), we can immediately determine θ_t, the fraction of time t wealth that is invested in the risky asset. In particular, θ_t is the coefficient of $\sigma_t V_t\ dW_t$ in the dynamics of V_t. Note also that we can work under any probability measure we choose since different probability measures will only change the drift of V_t and not the volatility. Once we have θ_t we immediately have $1 - \theta_t$ and then also $(\phi_t^{(s)}, \phi_t^{(b)})$.

Remark 7 *If a security pays dividends or coupons, then in the above statements we should replace the security price with the total gain process from holding the security. For example, if a stock pays a continuous dividend yield of q, then $Y_t := e^{qt} S_t$ is the total gain process. It is what your portfolio would be worth at time t if you purchased one unit of the security at time $t = 0$ and reinvested all dividends immediately back into the stock. The portfolio $(\phi_t^{(s)}, \phi_t^{(b)})$ is then said to be self-financing if*

$$
\begin{aligned}
dV_t\ &=\ \phi_t^{(s)}\ dY_t\ +\ \phi_t^{(b)}\ dB_t \\
&=\ \phi_t^{(s)}\ d(e^{qt} S_t)\ +\ \phi_t^{(b)}\ dB_t \\
&=\ \phi_t^{(s)}\ \left(e^{qt}\ dS_t\ +\ qY_t\ dt \right)\ +\ \phi_t^{(b)}\ dB_t.
\end{aligned}
\tag{36}
$$

We shall make use of (36) later in the course when we price options on dividend-paying securities.

The Role of Girsanov's Theorem

Suppose we again adopt the setting of Example 8 and that the risky security does not pay dividends. If we deflate as usual by the cash account then we can let $Z_t := S_t / B_t$ denote the deflated time t value of the risky asset. Assume also that W_t is a P-Brownian motion where P is the true data-generating probability measure. Then it is easy to check using Itô's Lemma that

$$dZ_t\ =\ Z_t (\mu_t - r)\ dt\ +\ Z_t \sigma_t\ dW_t. \tag{37}$$

But we know from the first fundamental theorem that Z_t must be a Q-martingale. Girsanov's Theorem, however, implies that Z_t is a Q-martingale only if $\widehat{W}_t := W_t + \int_0^t \eta_s\ ds$ is a Q-Brownian motion with $\eta_s := (\mu_s - r)/\sigma_s$. Indeed, because there is only one such process, η_s, this Q is unique[20] and the market is therefore complete by the second fundamental theorem. We then obtain that the Q-dynamics for S_t satisfy

$$dS_t\ =\ rS_t\ dt\ +\ \sigma_t S_t\ d\widehat{W}_t.$$

---

[20] If, for example, there were two Brownian motions and only one risky security then there would be infinitely many processes, η_s, that we could use to make Z_t a martingale. In this case we could therefore conclude that the market was incomplete.

Exercise 8 *Suppose now that the risky asset pays a continuous dividend yield, q. Show that the Q-dynamics of S_t now satisfy*

$$dS_t \;=\; (r-q)S_t \; dt \; + \sigma_t S_t \; d\widehat{W}_t.$$

(Hint: Remember that it is now the deflated total gains process that is now a Q-martingale.)

We can use these observations now to derive the Black-Scholes formula. Suppose the stock-price, S_t, has P-dynamics

$$dS_t \;=\; \mu_t S_t \; dt \; + \sigma S_t \; dW_t$$

where W_t is a P-Brownian motion. Note that we have now assumed a constant volatility, σ, as Black and Scholes originally assumed. If the stock pays a continuous dividend yield of q then martingale pricing implies that the price of a call option on the stock with maturity T and strike K is given by

$$C_0 \;=\; \mathsf{E}_0^Q \left[e^{-rT}(S_T - K)^+ \right] \tag{38}$$

where $\log(S_T) \sim \mathsf{N}\left((r - q - \sigma^2/2)T, \; \sigma T \right)$. We can therefore compute the right-hand-side of (38) analytically to obtain the Black-Scholes formula.

The following exercise is perhaps too difficult at this stage of the course. Later on, however, when we have had plenty of practice with Girsanov's Theorem and changing probability measures, we will see how straightforward such questions are.

Exercise 9 *Consider an equity model with two securities, A and B, whose price processes, $S_t^{(a)}$ and $S_t^{(b)}$ respectively, satisfy the following SDE's*

$$\begin{aligned}
dS_t^{(a)} &= r S_t^{(a)} \; dt + \sigma_1 S_t^{(a)} \; dW_t^{(1)} \\
dS_t^{(b)} &= r S_t^{(b)} \; dt + \sigma_2 S_t^{(b)} \left(\rho \; dW_t^{(1)} + \sqrt{1 - \rho^2} \; dW_t^{(2)} \right)
\end{aligned}$$

where $(W_t^{(1)}, \, W_t^{(2)})$ is a 2-dimensional Q-standard Brownian motion. We assume the cash account is the numeraire security corresponding to Q (which is consistent with the Q-dynamics of $S_t^{(a)}$ and $S_t^{(b)}$) and that the continuously compounded interest rate, r, is constant. Use the change of numeraire technique to compute the time 0 price, C_0 of a European option that expires at time T with payoff $\max\left(0, \; S_T^{(a)} - S_T^{(b)}\right)$.