

A Course in Financial Calculus

Alison Etheridge

University of Oxford



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK

40 West 20th Street, New York, NY 10011-4211, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

Ruiz de Alarcón 13, 28014 Madrid, Spain

Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Cambridge University Press 2002

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2002

Printed in the United Kingdom at the University Press, Cambridge

Typeface Times 10/13pt. *System* L^AT_EX 2_ε [DBD]

A catalogue record of this book is available from the British Library

Library of Congress Cataloguing in Publication data

ISBN 0 521 81385 9 hardback

ISBN 0 521 89077 2 paperback

Contents

<i>Preface</i>	<i>page vii</i>
1 Single period models	1
<i>Summary</i>	1
1.1 Some definitions from finance	1
1.2 Pricing a forward	4
1.3 The one-step binary model	6
1.4 A ternary model	8
1.5 A characterisation of no arbitrage	9
1.6 The risk-neutral probability measure	13
<i>Exercises</i>	18
2 Binomial trees and discrete parameter martingales	21
<i>Summary</i>	21
2.1 The multiperiod binary model	21
2.2 American options	26
2.3 Discrete parameter martingales and Markov processes	28
2.4 Some important martingale theorems	38
2.5 The Binomial Representation Theorem	43
2.6 Overture to continuous models	45
<i>Exercises</i>	47
3 Brownian motion	51
<i>Summary</i>	51
3.1 Definition of the process	51
3.2 Lévy's construction of Brownian motion	56
3.3 The reflection principle and scaling	59
3.4 Martingales in continuous time	63
<i>Exercises</i>	67
4 Stochastic calculus	71
<i>Summary</i>	71

4.1	Stock prices are not differentiable	72
4.2	Stochastic integration	74
4.3	Itô's formula	85
4.4	Integration by parts and a stochastic Fubini Theorem	93
4.5	The Girsanov Theorem	96
4.6	The Brownian Martingale Representation Theorem	100
4.7	Why geometric Brownian motion?	102
4.8	The Feynman–Kac representation	102
	<i>Exercises</i>	107
5	The Black–Scholes model	112
	<i>Summary</i>	112
5.1	The basic Black–Scholes model	112
5.2	Black–Scholes price and hedge for European options	118
5.3	Foreign exchange	122
5.4	Dividends	126
5.5	Bonds	131
5.6	Market price of risk	132
	<i>Exercises</i>	134
6	Different payoffs	139
	<i>Summary</i>	139
6.1	European options with discontinuous payoffs	139
6.2	Multistage options	141
6.3	Lookbacks and barriers	144
6.4	Asian options	149
6.5	American options	150
	<i>Exercises</i>	154
7	Bigger models	159
	<i>Summary</i>	159
7.1	General stock model	160
7.2	Multiple stock models	163
7.3	Asset prices with jumps	175
7.4	Model error	181
	<i>Exercises</i>	185
	<i>Bibliography</i>	189
	<i>Notation</i>	191
	<i>Index</i>	193

1 Single period models

Summary

In this chapter we introduce some basic definitions from finance and investigate the problem of pricing financial instruments in the context of a very crude model. We suppose the market to be observed at just two times: zero, when we enter into a financial contract; and T , the time at which the contract expires. We further suppose that the market can only be in one of a finite number of states at time T . Although simplistic, this model reveals the importance of the central paradigm of modern finance: the idea of a perfect hedge. It is also adequate for a preliminary discussion of the notion of ‘complete market’ and its importance if we are to find a ‘fair’ price for our financial contract.

The proofs in §1.5 can safely be omitted, although we shall from time to time refer back to the statements of the results.

1.1 Some definitions from finance

Financial market instruments can be divided into two types. There are the *underlying* stocks – shares, bonds, commodities, foreign currencies; and their *derivatives*, claims that promise some payment or delivery in the future contingent on an underlying stock’s behaviour. Derivatives can reduce risk – by enabling a player to fix a price for a future transaction now – or they can magnify it. A costless contract agreeing to pay off the difference between a stock and some agreed future price lets both sides ride the risk inherent in owning a stock, without needing the capital to buy it outright.

The connection between the two types of instrument is sufficiently complex and uncertain that both trade fiercely in the same market. The apparently random nature of the underlying stocks filters through to the derivatives – they appear random too.

Derivatives Our central purpose is to determine how much one should be willing to pay for a derivative security. But first we need to learn a little more of the language of finance.

Definition 1.1.1 A forward contract is an agreement to buy (or sell) an asset on a specified future date, T , for a specified price, K . The buyer is said to hold the long position, the seller the short position.

Forwards are not generally traded on exchanges. It costs nothing to enter into a forward contract. The ‘pricing problem’ for a forward is to determine what value of K should be written into the contract. A *futures contract* is the same as a forward except that futures *are* normally traded on exchanges and the exchange specifies certain standard features of the contract and a particular form of settlement.

Forwards provide the simplest examples of derivative securities and the mathematics of the corresponding pricing problem will also be simple. A much richer theory surrounds the pricing of *options*. An option gives the holder the *right*, but not the *obligation*, to do something. Options come in many different guises. Black and Scholes gained fame for pricing a European call option.

Definition 1.1.2 A European call option gives the holder the right, but not the obligation, to buy an asset at a specified time, T , for a specified price, K .

A European put option gives the holder the right to sell an asset for a specified price, K , at time T .

In general *call* refers to buying and *put* to selling. The term *European* is reserved for options whose value to the holder at the time, T , when the contract expires depends on the state of the market only at time T . There are other options, for example American options or Asian options, whose payoff is contingent on the behaviour of the underlying over the whole time interval $[0, T]$, but the technology of this chapter will only allow meaningful discussion of European options.

Definition 1.1.3 The time, T , at which the derivative contract expires is called the exercise date or the maturity. The price K is called the strike price.

The pricing problem

So what is the pricing problem for a European call option? Suppose that a company has to deal habitually in an intrinsically risky asset such as oil. They may for example know that in three months time they will need a thousand barrels of crude oil. Oil prices can fluctuate wildly, but by purchasing European call options, with strike K say, the company knows the *maximum* amount of money that it will need (in three months time) in order to buy a thousand barrels. One can think of the option as insurance against increasing oil prices. The pricing problem is now to determine, for given T and K , how much the company should be willing to pay for such insurance.

For this example there is an extra complication: it costs money to store oil. To simplify our task we are first going to price derivatives based on assets that can be held without additional cost, typically company shares. Equally we suppose that there is no additional benefit to holding the shares, that is no dividends are paid.




Figure 1.1 Payoff at maturity of (a) a forward purchase, (b) a European call and (c) a European put with strike K as a function of S_T .

Assumption Unless otherwise stated, the underlying asset can be held without additional cost or benefit.

This assumption will be relaxed in Chapter 5.

Suppose then that our company enters into a contract that gives them the right, but not the obligation, to buy one unit of stock for price K in three months time. How much should they pay for this contract?

Payoffs

As a first step, we need to know what the contract will be worth at the expiry date. If at the time when the option expires (three months hence) the actual price of the underlying stock is S_T and $S_T > K$ then the option will be exercised. The option is then said to be *in the money*: an asset worth S_T can be purchased for just K . The value to the company of the option is then $(S_T - K)$. If, on the other hand, $S_T < K$, then it will be cheaper to buy the underlying stock on the open market and so the option will not be exercised. (It is this freedom *not* to exercise that distinguishes options from futures.) The option is then worthless and is said to be *out of the money*. (If $S_T = K$ the option is said to be *at the money*.) The *payoff* of the option at time T is thus

$$(S_T - K)_+ \triangleq \max \{(S_T - K), 0\}.$$

Figure 1.1 shows the payoff at maturity of three derivative securities: a forward purchase, a European call and a European put, each as a function of stock price at maturity. Before embarking on the valuation *at time zero* of derivative contracts, we allow ourselves a short aside.

Packages

We have presented the European call option as a means of reducing risk. Of course it can also be used by a speculator as a bet on an increase in the stock price. In fact by holding *packages*, that is combinations of the ‘vanilla’ options that we have described so far, we can take rather complicated bets. We present just one example; more can be found in Exercise 1.

Example 1.1.4 (A straddle) Suppose that a speculator is expecting a large move in a stock price, but does not know in which direction that move will be. Then a possible combination is a straddle. This involves holding a European call and a European put with the same strike price and maturity.

Explanation: The payoff of this straddle is $(S_T - K)_+$ (from the call) plus $(K - S_T)_+$ (from the put), that is, $|S_T - K|$. Although the payoff of this combination is always positive, if, at the expiry time, the stock price is too close to the strike price then the payoff will not be sufficient to offset the cost of purchasing the options and the investor makes a loss. On the other hand, large movements in price can lead to substantial profits. $\square$

1.2 Pricing a forward

In order to solve our pricing problems, we are going to have to make some assumptions about the way in which markets operate. To formulate these we begin by discussing forward contracts in more detail.

Recall that a forward contract is an agreement to buy (or sell) an asset on a specified future date for a specified price. Suppose then that I agree to buy an asset for price K at time T . The payoff at time T is just $S_T - K$, where S_T is the actual asset price at time T . The payoff could be positive or it could be negative and, since the cost of entering into a forward contract is zero, this is also my total gain (or loss) from the contract. Our problem is to determine the fair value of K .

Expectation
pricing

At the time when the contract is written, we don't know S_T , we can only guess at it, or, more formally, assign a probability distribution to it. A widely used model (which underlies the Black–Scholes analysis of Chapter 5) is that stock prices are *lognormally distributed*. That is, there are constants ν and σ such that the *logarithm* of S_T/S_0 (the stock price at time T divided by that at time zero, usually called the *return*) is normally distributed with mean ν and variance σ^2 . In symbols:

$$\begin{aligned} \mathbb{P}\left[\frac{S_T}{S_0} \in [a, b]\right] &= \mathbb{P}\left[\log\left(\frac{S_T}{S_0}\right) \in [\log a, \log b]\right] \\ &= \int_{\log a}^{\log b} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x - \nu)^2}{2\sigma^2}\right) dx. \end{aligned}$$

Notice that stock prices, and therefore a and b , should be positive, so that the integral on the right hand side is well defined.

Our first guess might be that $\mathbb{E}[S_T]$ should represent a fair price to write into our contract. However, it would be a rare coincidence for this to be the market price. In fact we'll show that the cost of borrowing is the key to our pricing problem.

The risk-free
rate

We need a model for the *time value of money*: a dollar now is worth more than a dollar promised at some later time. We assume a market for these future promises (the *bond* market) in which prices are derivable from some interest rate. Specifically:

Time value of money We assume that for any time T less than some horizon τ the value now of a dollar promised at T is e^{-rT} for some constant $r > 0$. The rate r is then the *continuously compounded* interest rate for this period.

Such a market, derived from say US Government bonds, carries no risk of default – the promise of a future dollar will always be honoured. To emphasise this we will often refer to r as the *risk-free interest rate*. In this model, by buying or selling cash bonds, investors can borrow money for the same risk-free rate of interest as they can lend money.

Interest rate markets are not this simple in practice, but that is an issue that we shall defer.

Arbitrage
pricing

We now show that it is the *risk-free interest rate*, or equivalently the price of a cash bond, and not our lognormal model that forces the choice of the strike price, K , upon us in our forward contract.

Interest rates will be different for different currencies and so, for definiteness, suppose that we are operating in the dollar market, where the (risk-free) interest rate is r .

- Suppose first that $K > S_0 e^{rT}$. The seller, obliged to deliver a unit of stock for $\$K$ at time T , adopts the following strategy: she borrows $\$S_0$ at time zero (i.e. sells bonds to the value $\$S_0$) and buys one unit of stock. At time T , she must repay $\$S_0 e^{rT}$, but she has the stock to sell for $\$K$, leaving her a *certain* profit of $\$(K - S_0 e^{rT})$.
- If $K < S_0 e^{rT}$, then the buyer reverses the strategy. She sells a unit of stock at time zero for $\$S_0$ and buys cash bonds. At time T , the bonds deliver $\$S_0 e^{rT}$ of which she uses $\$K$ to buy back a unit of stock leaving her with a *certain* profit of $\$(S_0 e^{rT} - K)$.

Unless $K = S_0 e^{rT}$, one party is guaranteed to make a profit.

Definition 1.2.1 *An opportunity to lock into a risk-free profit is called an arbitrage opportunity.*

The starting point in establishing a model in modern finance theory is to specify that there is no arbitrage. (In fact there are people who make their living entirely from exploiting arbitrage opportunities, but such opportunities do not exist for a significant length of time before market prices move to eliminate them.) We have proved the following lemma.

Lemma 1.2.2 *In the absence of arbitrage, the strike price in a forward contract with expiry date T on a stock whose value at time zero is S_0 is $K = S_0 e^{rT}$, where r is the risk-free rate of interest.*

The price $S_0 e^{rT}$ is sometimes called the *arbitrage price*. It is also known as the *forward price* of the stock.

Remark: In our proof of Lemma 1.2.2, the buyer sold stock that she may not own. This is known as *short selling*. This can, and does, happen: investors can ‘borrow’ stock as well as money. $\square$

Of course forwards are a very special sort of derivative. The argument above won’t tell us how to value an option, but the strategy of seeking a price that does not provide either party with a risk-free profit will be fundamental in what follows.

Let us recap what we have done. In order to price the forward, we constructed a portfolio, comprising one unit of underlying stock and $-S_0$ cash bonds, whose value at the maturity time T is *exactly* that of the forward contract itself. Such a portfolio is said to be a *perfect hedge* or *replicating portfolio*. This idea is the central paradigm of modern mathematical finance and will recur again and again in what follows. Ironically we shall use expectation repeatedly, but as a tool in the construction of a perfect hedge.

1.3 The one-step binary model

We are now going to turn to establishing the fair price for European call options, but in order to do so we first move to a *simpler* model for the movement of market prices. Once again we suppose that the market is observed at just two times, that at which the contract is struck and the expiry date of the contract. Now, however, we shall suppose that there are just two possible values for the stock price at time T . We begin with a simple example.

Pricing a
European
call

Example 1.3.1 Suppose that the current price in Japanese Yen of a certain stock is ¥2500. A European call option, maturing in six months time, has strike price ¥3000. An investor believes that with probability one half the stock price in six months time will be ¥4000 and with probability one half it will be ¥2000. He therefore calculates the expected value of the option (when it expires) to be ¥500. The riskless borrowing rate in Japan is currently zero and so he agrees to pay ¥500 for the option. Is this a fair price?

Solution: In the light of the previous section, the reader will probably have guessed that the answer to this question is no. Once again, we show that one party to this contract can make a risk-free profit. In this case it is the seller of the contract. Here is just one of the many possible strategies that she could adopt.

Strategy: At time zero, sell the option, borrow ¥2000 and buy a unit of stock.

- Suppose first that at expiry the price of the stock is ¥4000; then the contract will be exercised and so she must sell her stock for ¥3000. She then holds $¥(-2000 + 3000)$. That is ¥1000.
- If, on the other hand, at expiry the price of the stock is ¥2000, then the option will not be exercised and so she sells her stock on the open market for just ¥2000. Her




Figure 1.2 The seller of the contract in Example 1.3.1 is guaranteed a risk-free profit if she can buy any portfolio in the shaded region.

net cash holding is then $\yen(-2000 + 2000)$. That is, she exactly breaks even.

Either way, our seller has a positive chance of making a profit with *no risk* of making a loss. The price of the option is too high.

So what is the right price for the option?

Let's think of things from the point of view of the seller. Writing S_T for the price of the stock when the contract expires, she knows that at time T she needs $\yen(S_T - 3000)_+$ in order to meet the claim against her. The idea is to calculate how much money she needs at time zero, to be held in a combination of stocks and cash, to guarantee this.

Suppose then that she uses the money that she receives for the option to buy a portfolio comprising x_1 Yen and x_2 stocks. If the price of the stock is $\yen4000$ at expiry, then the time T value of the portfolio is $x_1 e^{rT} + 4000x_2$. The seller of the option requires this to be at least $\yen1000$. That is, since interest rates are zero,

$$x_1 + 4000x_2 \geq 1000.$$

If the price is $\yen2000$ she just requires the value of the portfolio to be non-negative,

$$x_1 + 2000x_2 \geq 0.$$

A profit is guaranteed (without risk) for the seller if (x_1, x_2) lies in the interior of the shaded region in Figure 1.2. On the boundary of the region, there is a positive probability of profit and no probability of loss at all points other than the intersection of the two lines. The portfolio represented by the point $(\bar{x}_1, \bar{x}_2)$ will provide *exactly* the wealth required to meet the claim against her at time T .

Solving the simultaneous equations gives that the seller can exactly meet the claim if $\bar{x}_1 = -1000$ and $\bar{x}_2 = 1/2$. The cost of building this portfolio at time zero is $\yen(-1000 + 2500/2)$, that is $\yen250$. For any price higher than $\yen250$, the seller can make a risk-free profit.

If the option price is *less* than ¥250, then the *buyer* can make a risk-free profit by ‘borrowing’ the portfolio $(\bar{x}_1, \bar{x}_2)$ and buying the option. In the absence of arbitrage then, the fair price for the option is ¥250. $\square$

Notice that just as for our forward contract, we did not use the probabilities that we assigned to the possible market movements to arrive at the fair price. We just needed the fact that we could *replicate* the claim by this simple portfolio. The seller can *hedge* the *contingent claim* $\mathbb{Y}(S_T - 3000)_+$ using the *portfolio* consisting of $\mathbb{Y}x_1$ and x_2 units of stock.

Pricing
formula for
European
options

One can use exactly the same argument to prove the following result.

Lemma 1.3.2 Suppose that the risk-free dollar interest rate (to a time horizon $\tau > T$) is r . Denote the time zero (dollar) value of a certain asset by S_0 . Suppose that the motion of stock prices is such that the value of the asset at time T will be either S_0u or S_0d . Assume further that

$$d < e^{rT} < u.$$

At time zero, the market price of a European option with payoff $C(S_T)$ at the maturity T is

$$\left(\frac{1 - de^{-rT}}{u - d} \right) C(S_0u) + \left(\frac{ue^{-rT} - 1}{u - d} \right) C(S_0d).$$

Moreover, the seller of the option can construct a portfolio whose value at time T is exactly $(S_T - K)_+$ by using the money received for the option to buy

$$\phi \triangleq \frac{C(S_0u) - C(S_0d)}{S_0u - S_0d} \quad (1.1)$$

units of stock at time zero and holding the remainder in bonds.

The proof is Exercise 4(a).

1.4 A ternary model

There were several things about the binary model that were very special. In particular we assumed that we knew that the asset price would be one of just two specified values at time T . What if we allow *three* values?

We can try to repeat the analysis of §1.3. Again the seller would like to replicate the claim at time T by a portfolio consisting of $\mathbb{Y}x_1$ and x_2 stocks. This time there will be three scenarios to consider, corresponding to the three possible values of S_T . If interest rates are zero, this gives rise to the three inequalities

$$x_1 + S_T^i x_2 \geq (S_T^i - 3000)_+, \quad i = 1, 2, 3,$$

where S_T^i are the possible values of S_T . The picture is now something like that in Figure 1.3.




Figure 1.3 If the stock price takes three possible values at time T , then at any point where the seller of the option has no risk of making a loss, she has a *strictly positive* chance of making a profit.

In order to be *guaranteed* to meet the claim at time T , the seller requires (x_1, x_2) to lie in the shaded region, but at any point in that region, she has a strictly positive probability of making a profit and zero probability of making a loss. Any portfolio from outside the shaded region carries a risk of a loss. There is no portfolio that *exactly* replicates the claim and there is no unique ‘fair’ price for the option.

Our market is not *complete*. That is, there are contingent claims that cannot be perfectly hedged.

Bigger
models

Of course we are tying our hands in our efforts to hedge a claim. First, we are only allowing ourselves portfolios consisting of the underlying stock and cash bonds. Real markets are bigger than this. If we allow ourselves to trade in a third ‘independent’ asset, then our analysis leads to three non-parallel planes in $\mathbb{R}^3$. These *will* intersect in a single point representing a portfolio that exactly replicates the claim. This then raises a question: when is there arbitrage in larger market models? We shall answer this question for a single period model in the next section. The second constraint that we have placed upon ourselves is that we are not allowed to adjust our portfolio between the time of the selling of the contract and its maturity. In fact, as we see in Chapter 2, if we consider the market to be observable at intermediate times between zero and T , and allow our seller to rebalance her portfolio at such times (without changing its value), then we *can* allow any number of possible values for the stock price at time T and yet still replicate each claim at time T by a portfolio consisting of just the underlying and cash bonds.

1.5 A characterisation of no arbitrage

In our binary setting it was easy to find the right price for an option simply by solving a pair of simultaneous equations. However, the binary model is very special and, after our experience with the ternary model, alarm bells may be ringing. The binary model describes the evolution of just one stock (and one bond). One solution to our

difficulties with the ternary model was to allow trade in another ‘independent’ asset. In this section we extend this idea to larger market models and characterise those models for which there are a sufficient number of independent assets that any option has a fair price. Other than Definition 1.5.1 and the statement of Theorem 1.5.2, this section can safely be omitted.

A market
with N
assets

Our market will now consist of a finite (but possibly large) number of tradable assets. Again we restrict ourselves to single period models, in which the market is observable only at time zero and a fixed future time T . However, the extension to multiple time periods exactly mirrors that for binary models that we describe in §2.1.

Suppose then that there are N tradable assets in the market. Their prices at time zero are given by the column vector

$$S_0 = (S_0^1, S_0^2, \dots, S_0^N)^t \triangleq \begin{pmatrix} S_0^1 \\ S_0^2 \\ \vdots \\ S_0^N \end{pmatrix}.$$

Notation For vectors and matrices we shall use the superscript ‘ t ’ to denote transpose.

Uncertainty about the market is represented by a finite number of possible states in which the market might be at time T that we label $1, 2, \dots, n$. The security values at time T are given by an $N \times n$ matrix $D = (D_{ij})$, where the coefficient D_{ij} is the value of the i th security at time T if the market is in state j . Our binary model corresponds to $N = 2$ (the stock and a riskless cash bond) and $n = 2$ (the two states being determined by the two possible values of S_T).

In this notation, a portfolio can be thought of as a vector $\theta = (\theta_1, \theta_2, \dots, \theta_n)^t \in \mathbb{R}^n$, whose market value at time zero is the scalar product $S_0 \cdot \theta = S_0^1\theta_1 + S_0^2\theta_2 + \dots + S_0^N\theta_N$. The value of the portfolio at time T is a vector in $\mathbb{R}^n$ whose i th entry is the value of the portfolio if the market is in state i . We can write the value at time T as

$$\begin{pmatrix} D_{11}\theta_1 + D_{21}\theta_2 + \dots + D_{N1}\theta_N \\ D_{12}\theta_1 + D_{22}\theta_2 + \dots + D_{N2}\theta_N \\ \vdots \\ D_{1n}\theta_1 + D_{2n}\theta_2 + \dots + D_{Nn}\theta_N \end{pmatrix} = D^t \theta.$$

Notation For a vector $x \in \mathbb{R}^n$ we write $x \geq 0$, or $x \in \mathbb{R}_+^n$, if $x = (x_1, \dots, x_n)$ and $x_i \geq 0$ for all $i = 1, \dots, n$. We write $x > 0$ to mean $x \geq 0, x \neq 0$. Notice that $x > 0$ does not require x to be strictly positive in *all* its coordinates. We write $x \gg 0$, or $x \in \mathbb{R}_{++}^n$, for vectors in $\mathbb{R}^n$ that are *strictly* positive in *all* coordinates.

In this notation, an *arbitrage* is a portfolio $\theta \in \mathbb{R}^N$ with either

$$S_0 \cdot \theta \leq 0, \quad D^t \theta > 0 \quad \text{or} \quad S_0 \cdot \theta < 0, \quad D^t \theta \geq 0.$$

Arbitrage
pricing

The key to arbitrage pricing in this model is the notion of a state price vector.

Definition 1.5.1 A state price vector is a vector $\psi \in \mathbb{R}_{++}^n$ such that $S_0 = D\psi$.

To see why this terminology is natural, we first expand this to obtain

$$\begin{pmatrix} S_0^1 \\ S_0^2 \\ \vdots \\ S_0^N \end{pmatrix} = \psi_1 \begin{pmatrix} D_{11} \\ D_{21} \\ \vdots \\ D_{N1} \end{pmatrix} + \psi_2 \begin{pmatrix} D_{12} \\ D_{22} \\ \vdots \\ D_{N2} \end{pmatrix} + \cdots + \psi_n \begin{pmatrix} D_{1n} \\ D_{2n} \\ \vdots \\ D_{Nn} \end{pmatrix}. \quad (1.2)$$

The vector, $D^{(i)}$, multiplying ψ_i is the security price vector if the market is in state i . We can think of ψ_i as the marginal cost at time zero of obtaining an additional unit of wealth at the end of the time period if the system is in state i . In other words, if at the end of the time period, the market is in state i , then the value of our portfolio increases by one for each additional ψ_i of investment at time zero. To see this, suppose that we can find vectors $\{\theta^{(i)} \in \mathbb{R}^N\}_{1 \leq i \leq n}$ such that

$$\theta^{(i)} \cdot D^{(j)} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

That is, the value of the portfolio $\theta^{(i)}$ at time T is the indicator function that the market is in state i . Then, using equation (1.2), the cost of purchasing $\theta^{(i)}$ at time zero is precisely $S_0 \cdot \theta^{(i)} = (\sum_{j=1}^n \psi_j D^{(j)}) \cdot \theta^{(i)} = \psi_i$. Such portfolios $\{\theta^{(i)}\}_{1 \leq i \leq n}$ are called *Arrow–Debreu securities*.

We shall find a convenient way to think about the state price vector in §1.6, but first, here is the key result.

Theorem 1.5.2 For the market model described above there is no arbitrage if and only if there is a state price vector.




Figure 1.4 There is no arbitrage if and only if the regions K and M of Theorem 1.5.2 intersect only at the origin.

This result, due to Harrison & Kreps (1979), is the simplest form of what is often known as the Fundamental Theorem of Asset Pricing. The proof is an application of a Hahn–Banach Separation Theorem, sometimes called the *Separating Hyperplane Theorem*. We shall also need the *Riesz Representation Theorem*. Recall that $M \subseteq \mathbb{R}^d$ is a *cone* if $x \in M$ implies $\lambda x \in M$ for all strictly positive scalars λ and that a *linear functional* on $\mathbb{R}^d$ is a linear mapping $F: \mathbb{R}^d \rightarrow \mathbb{R}$.

Theorem 1.5.3 (Separating Hyperplane Theorem) Suppose M and K are closed convex cones in $\mathbb{R}^d$ that intersect precisely at the origin. If K is not a linear subspace, then there is a non-zero linear functional F such that $F(x) < F(y)$ for each $x \in M$ and each non-zero $y \in K$.

This version of the Separating Hyperplane Theorem can be found in Duffie (1992).

Theorem 1.5.4 (Riesz Representation Theorem) Any bounded linear functional on $\mathbb{R}^d$ can be written as $F(x) = v_0 \cdot x$. That is $F(x)$ is the scalar product of some fixed vector $v_0 \in \mathbb{R}^d$ with x .

Proof of Theorem 1.5.2: We take $d = 1 + n$ in Theorem 1.5.3 and set

$$M = \left\{ (-S_0 \cdot \theta, D^t \theta) : \theta \in \mathbb{R}^N \right\} \subseteq \mathbb{R} \times \mathbb{R}^n = \mathbb{R}^{1+n},$$

$$K = \mathbb{R}_+ \times \mathbb{R}_+^n.$$

Note that K is a cone and not a linear space, M is a linear space. Evidently, there is no arbitrage if and only if K and M intersect precisely at the origin as shown in

Figure 1.4. We must prove that $K \cap M = \{0\}$ if and only if there is a state price vector.

(i) *Suppose first that $K \cap M = \{0\}$.* From Theorem 1.5.3, there is a linear functional $F: \mathbb{R}^d \rightarrow \mathbb{R}$ such that $F(z) < F(x)$ for all $z \in M$ and non-zero $x \in K$.

The first step is to show that F must vanish on M . We exploit the fact that M is a linear space. First observe that $F(0) = 0$ (by linearity of F) and $0 \in M$, so $F(x) \geq 0$ for $x \in K$ and $F(x) > 0$ for $x \in K \setminus \{0\}$. Fix $x_0 \in K$ with $x_0 \neq 0$. Now take an arbitrary $z \in M$. Then $F(z) < F(x_0)$, but also, since M is a linear space, $\lambda F(z) = F(\lambda z) < F(x_0)$ for all $\lambda \in \mathbb{R}$. This can only hold if $F(z) = 0$. $z \in M$ was arbitrary and so F vanishes on M as required.

We now use this actually to construct explicitly the state price vector from F . First we use the Riesz Representation Theorem to write F as $F(x) = v_0 \cdot x$ for some $v_0 \in \mathbb{R}^d$. It is convenient to write $v_0 = (\alpha, \phi)$ where $\alpha \in \mathbb{R}$ and $\phi \in \mathbb{R}^n$. Then

$$F(v, c) = \alpha v + \phi \cdot c \quad \text{for any } (v, c) \in \mathbb{R} \times \mathbb{R}^n = \mathbb{R}^d.$$

Since $F(x) > 0$ for all non-zero $x \in K$, we must have $\alpha > 0$ and $\phi \gg 0$ (consider a vector along each of the coordinate axes). Finally, since F vanishes on M ,

$$-\alpha S_0 \cdot \theta + \phi \cdot D^t \theta = 0 \quad \text{for all } \theta \in \mathbb{R}^N.$$

Observing that $\phi \cdot D^t \theta = (D\phi) \cdot \theta$, this becomes

$$-\alpha S_0 \cdot \theta + (D\phi) \cdot \theta = 0 \quad \text{for all } \theta \in \mathbb{R}^N,$$

which implies that $-\alpha S_0 + D\phi = 0$. In other words, $S_0 = D(\phi/\alpha)$. The vector $\psi = \phi/\alpha$ is a state price vector.

(ii) *Suppose now that there is a state price vector, ψ .* We must prove that $K \cap M = \{0\}$. By definition, $S_0 = D\psi$ and so for any portfolio θ ,

$$S_0 \cdot \theta = (D\psi) \cdot \theta = \psi \cdot (D^t \theta). \quad (1.3)$$

Suppose that for some portfolio θ , $(-S_0 \cdot \theta, D^t \theta) \in K$. Then $D^t \theta \in \mathbb{R}_+^n$ and $-S_0 \cdot \theta \geq 0$. But since $\psi \gg 0$, if $D^t \theta \in \mathbb{R}_+^n$, then $\psi \cdot (D^t \theta) \geq 0$ which, by equation (1.3), tells us that $S_0 \cdot \theta \geq 0$. Thus it must be that $S_0 \cdot \theta = 0$ and $D^t \theta = 0$. That is, $K \cap M = \{0\}$, as required. $\square$

1.6 The risk-neutral probability measure

The state price vector then is the key to arbitrage pricing for our multiasset market models. Although we have an economic interpretation for it, in order to pave the way for the full machinery of probability and martingales we must think about it in a different way.

Recall that all the entries of ψ are strictly positive.

State prices
and
probability

Writing $\psi_0 = \sum_{i=1}^n \psi_i$, we can think of

$$\underline{\psi} \triangleq \left(\frac{\psi_1}{\psi_0}, \frac{\psi_2}{\psi_0}, \dots, \frac{\psi_n}{\psi_0} \right)^t \quad (1.4)$$

as a vector of *probabilities* for being in different states. It is important to emphasise that they may have nothing to do with our view of how the markets will move. First of all,

What is ψ_0 ?

Suppose that as in our binary model (where we had a risk-free cash bond) the market allows *positive riskless borrowing*. In this general setting we just suppose that we can replicate such a bond by a portfolio $\bar{\theta}$ for which

$$D^t \bar{\theta} = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix},$$

i.e. the value of the portfolio at time T is one, no matter what state the market is in. Using the fact that ψ is a state price vector, we calculate that the cost of such a portfolio at time zero is

$$S_0 \cdot \bar{\theta} = (D\psi) \cdot \bar{\theta} = \psi \cdot (D^t \bar{\theta}) = \sum_{i=1}^n \psi_i = \psi_0.$$

That is ψ_0 represents the discount on riskless borrowing. In our notation of §1.2, $\psi_0 = e^{-rT}$.

Expectation
recovered

Now under the probability distribution given by the vector (1.4), the expected value of the i th security at time T is

$$\mathbb{E}[S_T^i] = \sum_{j=1}^n D_{ij} \frac{\psi_j}{\psi_0} = \frac{1}{\psi_0} \sum_{j=1}^n D_{ij} \psi_j = \frac{1}{\psi_0} S_0^i,$$

where in the last equality we have used $S_0 = D\psi$. That is

$$S_0^i = \psi_0 \mathbb{E}[S_T^i], \quad i = 1, \dots, n. \quad (1.5)$$

Any security's price is its discounted expected payoff under the probability distribution (1.4). The same must be true of any portfolio. This observation gives us a new way to think about the pricing of contingent claims.

Definition 1.6.1 We shall say that a claim, C , at time T is attainable if it can be hedged. That is, if there is a portfolio whose value at time T is exactly C .

Notation When we wish to emphasise the underlying probability measure, $\mathbb{Q}$, we write $\mathbb{E}^{\mathbb{Q}}$ for the expectation operator.

Theorem 1.6.2 *If there is no arbitrage, the unique time zero price of an attainable claim C at time T is $\psi_0 \mathbb{E}^{\mathbb{Q}}[C]$ where the expectation is with respect to any probability measure $\mathbb{Q}$ for which $S_0^i = \psi_0 \mathbb{E}^{\mathbb{Q}}[S_T^i]$ for all i and ψ_0 is the discount on riskless borrowing.*

Remark: Notice that it is crucial that the claim is attainable (see Exercise 11). $\square$

Proof of Theorem 1.6.2: By Theorem 1.5.2 there is a state price vector and this leads to the probability measure (1.4) satisfying $S_0^i = \psi_0 \mathbb{E}[S_T^i]$ for all i . Since the claim can be hedged, there is a portfolio θ such that $\theta \cdot S_T = C$. In the absence of arbitrage, the time zero price of the claim is the cost of this portfolio at time zero,

$$\theta \cdot S_0 = \theta \cdot (\psi_0 \mathbb{E}[S_T]) = \psi_0 \sum_{i=1}^N \theta_i \mathbb{E}[S_T^i] = \psi_0 \mathbb{E}[\theta \cdot S_T].$$

The same value is obtained if the expectation is calculated for any vector of probabilities, $\mathbb{Q}$, such that $S_0^i = \psi_0 \mathbb{E}^{\mathbb{Q}}[S_T^i]$ since, in the absence of arbitrage, there is only one riskless borrowing rate and this completes the proof. $\square$

Risk-neutral
pricing

In this language, our arbitrage pricing result says that if we can find a probability vector for which the time zero value of each underlying security is its discounted expected value at time T then we can find the time zero value of any *attainable* contingent claim by calculating its discounted expectation. Notice that we use the *same* probability vector, whatever the claim.

Definition 1.6.3 *If our market can be in one of n possible states at time T , then any vector, $\underline{p} = (p_1, p_2, \dots, p_n) \gg 0$, of probabilities for which each security's price is its discounted expected payoff is called a risk-neutral probability measure or equivalent martingale measure.*

The term *equivalent* reflects the condition that $p \gg 0$; cf. Definition 2.3.12. Our simple form of the Fundamental Theorem of Asset Pricing (Theorem 1.5.2) says that in a market with positive riskless borrowing there is no arbitrage if and only if there is an equivalent martingale measure. We shall refer to the process of pricing by taking expectations with respect to a risk-neutral probability measure as *risk-neutral pricing*.

Example 1.3.1 revisited Let us return to our very first example of pricing a European call option and confirm that the above formula really does give us the arbitrage price.

Here we have just two securities, a cash bond and the underlying stock. The discount on borrowing is $\psi_0 = e^{-rT}$, but we are assuming that the Yen interest rate is zero, so $\psi_0 = 1$. The matrix of security values at time T is given by

$$D = \begin{pmatrix} 1 & 1 \\ 4000 & 2000 \end{pmatrix}.$$

Writing p for the risk-neutral probability that the security price vector is $(1, 4000)^t$, if the stock price is to be equal to its discounted expected payoff, p must solve

$$4000p + 2000(1 - p) = 2500,$$

which gives $p = 0.25$. The contingent claim is ¥1000 if the stock price at expiry is ¥4000 and zero otherwise. The expected value of the claim under the risk-neutral probability, and therefore (since interest rates are zero) the price of the option, is then $¥0.25 \times 1000 = ¥250$, as before.

An advantage of this approach is that, armed with the probability p , it is now a trivial matter to price all European options on this stock with the same expiry date (six months time) by taking expectations with respect to the *same* probability measure. For example, for a European put option with strike price ¥3500, the price is

$$¥\mathbb{E}[(K - S_T)_+] = ¥0.75 \times 1500 = ¥1125.$$

Our original argument would lead to a new set of simultaneous equations for each new claim. □

Complete
markets

We now have a prescription for the arbitrage price of a claim if one exists, that is if the claim is attainable. But we must be a little cautious. Arbitrage prices only exist for attainable claims – even though the prescription may continue to make sense.

Definition 1.6.4 *A market is said to be complete if every contingent claim is attainable, i.e. if every possible derivative claim can be hedged.*

Proposition 1.6.5 *A market consisting of N tradable assets, evolving according to a single period model in which at the end of the time period the market is one of n possible states, is complete if and only if $N \geq n$ and the rank of the matrix, D , of security prices is n .*

Proof: Any claim in our market can be expressed as a vector $v \in \mathbb{R}^n$. A hedge for that claim will be a portfolio $\theta = \theta(v) \in \mathbb{R}^N$ for which $D^t \theta = v$. Finding such a θ amounts to solving n equations in N unknowns. Thus a hedging portfolio exists for every choice of v if and only if $N \geq n$ and the rank of D is n , as required. □

Notice in particular that our single period binary model is complete.

Suppose that our market is complete and arbitrage-free and let $\mathbb{Q}$ and $\mathbb{Q}'$ be any two equivalent martingale measures. By completeness every claim is attainable, so for every random variable X , using that there is only one risk-free rate,

$$\mathbb{E}^{\mathbb{Q}}[X] = \mathbb{E}^{\mathbb{Q}'}[X].$$

In other words $\mathbb{Q} = \mathbb{Q}'$. So in a complete arbitrage-free market the equivalent martingale measure is *unique*.

The main results so far Let us summarise the results for our single period markets. They will be reflected again and again in what follows.

Results for single period models

- The market is arbitrage-free if and only if there exists a martingale measure, $\mathbb{Q}$.
- The market is complete if and only if $\mathbb{Q}$ is unique.
- The arbitrage price of an attainable claim C is $e^{-rT} \mathbb{E}^{\mathbb{Q}}[C]$.

Martingale measures are a powerful tool. However, in an incomplete market, if a claim C is not attainable different martingale measures can give different prices. The arbitrage-free notion of *fair price* only makes sense if we can *hedge*.

Trading in two different markets We must sound just one more note of caution. It is important in calculating the risk-neutral probabilities that all the assets being modelled are tradable in the same market. We illustrate with an example.

Example 1.6.6 Suppose that in the US dollar markets the current Sterling exchange rate is 1.5 (so that £100 costs \$150). Consider a European call option that offers the holder the right to buy £100 for \$150 at time T . The riskless borrowing rate in the UK is u and that in the US is r . Assuming a single period binary model in which the exchange rate at the expiry time is either 1.65 or 1.45, find the fair price of this option.

Solution: Now we have a problem. The exchange rate is *not* tradable. Nor, in *dollar* markets, is a Sterling cash bond – it is a tradable instrument, but in *Sterling* markets. However, the product of the two is a dollar tradable and we shall denote the value of this product by S_t at time t .

Now, since the riskless interest rate in the UK is u , the time zero price of a Sterling cash bond, promising to pay £1 at time T , is e^{-uT} and, of course, at time T the bond price is one. Thus we have $S_0 = e^{-uT} 150$ and $S_T = 165$ or $S_T = 145$.

Let p be the risk-neutral probability that $S_T = 165$. Then, since the discounted price (in the *dollar* market) of our ‘asset’ at time T must have expectation S_0 , we obtain

$$150e^{-uT} = e^{-rT} (165p + 145(1 - p)),$$

which yields

$$p = \frac{150e^{(r-u)T} - 145}{20}.$$

The price of the option is the discounted expected payoff with respect to this


probability which gives

$$V_0 = e^{-rT} 15p = \frac{3}{4} (150e^{-uT} - 145e^{-rT}).$$

□

Exercises

- 1 What view about the market is reflected in each of the following strategies?
 - (a) *Bullish vertical spread*: Buy one European call and sell a second one with the same expiry date, but a larger strike price.
 - (b) *Bearish vertical spread*: Buy one European call and sell a second one with the same expiry date but a smaller strike price.
 - (c) *Strip*: Buy one European call and two European puts with the same exercise date and strike price.
 - (d) *Strap*: Buy two European calls and one European put with the same exercise date and strike price.
 - (e) *Strangle*: Buy a European call and a European put with the same expiry date but different strike prices (consider all possible cases).
- 2 A *butterfly spread* represents the complementary bet to the straddle. It has the following payoff at expiry:



Find a portfolio consisting of European calls and puts, all with the same expiry date, that has this payoff.

- 3 Suppose that the price of a certain asset has the lognormal distribution. That is $\log(S_T/S_0)$ is normally distributed with mean ν and variance σ^2 . Calculate $\mathbb{E}[S_T]$.
- 4 (a) Prove Lemma 1.3.2.
(b) What happens if we drop the assumption that $d < e^{rT} < u$?
- 5 Suppose that at current exchange rates, £100 is worth €160. A speculator believes that by the end of the year there is a probability of 1/2 that the pound will have fallen to €1.40, and a 1/2 chance that it will have gained to be worth €2.00. He therefore buys a European put option that will give him the right (but not the obligation) to

sell £100 for €1.80 at the end of the year. He pays €20 for this option. Assume that the risk-free interest rate is zero across the Euro-zone. Using a single period binary model, either construct a strategy whereby one party is certain to make a profit or prove that this is the fair price.

- 6 How should we modify the analysis of Example 1.3.1 if we are pricing an option based on a commodity such as oil?
- 7 Show that if there is no arbitrage in the market, then any portfolio constructed at time zero that exactly replicates a claim C at time T has the same value at time zero.
- 8 *Put-call parity*: Denote by C_t and P_t respectively the prices at time t of a European call and a European put option, each with maturity T and strike K . Assume that the risk-free rate of interest is constant, r , and that there is no arbitrage in the market. Show that for each $t \leq T$,

$$C_t - P_t = S_t - Ke^{-r(T-t)}.$$

- 9 Use risk-neutral pricing to value the option in Exercise 5. Check your answer by constructing a portfolio that exactly replicates the claim at the expiry of the contract.
- 10 What is the payoff of a forward at expiry? Use risk-neutral pricing to solve the pricing problem for a forward contract.
- 11 Consider the *ternary* model for the underlying of §1.4. How many equivalent martingale measures are there? If there are two different martingale measures, do they give the same price for a claim? Are there arbitrage opportunities?
- 12 Suppose that the value of a certain stock at time T is a random variable with distribution $\mathbb{P}$. Note we are *not* assuming a binary model. An option written on this stock has payoff C at time T . Consider a portfolio consisting of ϕ units of the underlying and ψ units of bond, held until time T , and write V_0 for its value at time zero. Assuming that interest rates are zero, show that the extra cash required by the holder of this portfolio to meet the claim C at time T is

$$\Psi \triangleq C - V_0 - \phi(S_T - S_0).$$

Find expressions for the values of V_0 and ϕ (in terms of $\mathbb{E}[S_T]$, $\mathbb{E}[C]$, $\text{var}[S_T]$ and $\text{cov}(S_T, C)$) that minimise

$$\mathbb{E}[\Psi^2],$$

and check that for these values $\mathbb{E}[\Psi] = 0$.

Prove that for a binary model, any claim C depends *linearly* on $S_T - S_0$. Deduce that in this case we can find V_0 and ϕ such that $\Psi = 0$.

When the model is *not* complete, the parameters that minimise $\mathbb{E}[\Psi^2]$ correspond to finding the best linear approximation to C (based on $S_T - S_0$). The corresponding value of the expectation is a measure of the *intrinsic risk* in the option.

- 13 *Exchange rate forward:* Suppose that the riskless borrowing rate in the UK is u and that in the USA is r . A dollar investor wishes to set the exchange rate, C_T , in a forward contract in which the two parties agree to exchange C_T dollars for one pound at time T . If a pound is currently C_0 dollars, what is the fair value of C_T ?
- 14 The option writer in Example 1.6.6 sells a *digital* option to a speculator. This amounts to a bet that the asset price will go up. The payoff is a fixed amount of cash if the exchange rate goes to \$165 per £100, and nothing if it goes down. If the speculator pays \$10 for this bet, what cash payout should the option writer be willing to write into the option? You may assume that interest rates are zero.
- 15 Suppose now that the seller of the option in Example 1.6.6 operates in the Sterling markets. Reexpress the market in terms of Sterling tradables and find the corresponding risk-neutral probabilities. Are they the same as the risk-neutral probabilities calculated by the dollar trader? What is the dollar cost at time zero of the option as valued by the Sterling trader?
- This is an example of *change of numeraire*. The dollar trader uses the dollar bond as the reference risk-free asset whereas the Sterling trader uses a Sterling bond.