

Financial calculus
An introduction to derivative pricing

Martin Baxter
University of Cambridge

Andrew Rennie
Union Bank of Switzerland

 **CAMBRIDGE**
UNIVERSITY PRESS

Contents

Published by the Press Syndicate of the University of Cambridge
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© M.W. Baxter and A.J.O. Rennie 1996

First published 1996

Printed in Great Britain at the University Press, Cambridge

Typeset in Monotype Bembo by the authors using \TeX

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

Library of Congress cataloguing in publication data available

ISBN 0 521 55289 3 hardback

	Preface	vii
	The parable of the bookmaker	1
Chapter 1	Introduction	3
	1.1 Expectation pricing	4
	1.2 Arbitrage pricing	7
	1.3 Expectation vs arbitrage	9
Chapter 2	Discrete processes	10
	2.1 The binomial branch model	10
	2.2 The binomial tree model	17
	2.3 Binomial representation theorem	28
	2.4 Overture to continuous models	41
Chapter 3	Continuous processes	44
	3.1 Continuous processes	45
	3.2 Stochastic calculus	51
	3.3 Itô calculus	57
	3.4 Change of measure – the C-M-G theorem	63
	3.5 Martingale representation theorem	76
	3.6 Construction strategies	80
	3.7 Black–Scholes model	83
	3.8 Black–Scholes in action	92

Contents

Chapter 4	Pricing market securities	99
	4.1 Foreign exchange	99
	4.2 Equities and dividends	106
	4.3 Bonds	112
	4.4 Market price of risk	116
	4.5 Quantos	122
Chapter 5	Interest rates	128
	5.1 The interest rate market	129
	5.2 A simple model	135
	5.3 Single-factor HJM	142
	5.4 Short-rate models	149
	5.5 Multi-factor HJM	158
	5.6 Interest rate products	163
	5.7 Multi-factor models	172
Chapter 6	Bigger models	178
	6.1 General stock model	178
	6.2 Log-normal models	181
	6.3 Multiple stock models	183
	6.4 Numeraires	189
	6.5 Foreign currency interest-rate models	193
	6.6 Arbitrage-free complete models	196
Appendices	A1 Further reading	201
	A2 Notation	205
	A3 Answers to exercises	209
	A4 Glossary of technical terms	216
Index		228

Preface

Notoriously, works of mathematical finance can be precise, and they can be comprehensible. Sadly, as Dr Johnson might have put it, the ones which are precise are not necessarily comprehensible, and those comprehensible are not necessarily precise.

But both are needed. The mathematics of finance is not easy, and much market practice is based on a soft understanding of what is actually going on. This is usually enough for experienced practitioners to price existing contracts, but often insufficient for innovative new products. Novices, managers and regulators can be left to stumble around in literature which is ill suited to their need for a clear explanation of the basic principles. Such 'seat of the pants' practices are more suited to the pioneering days of an industry, rather than the mature \$15 trillion market which the derivatives business has become.

On the academic side, effort is too often expended on finding precise answers to the wrong questions. When working in isolation from the market, the temptation is to find analytic answers for their own sake with no reference to the concerns of practitioners. In particular, the importance of hedging both as a justification for the price and as an important end in itself is often underplayed. Scholars need to be aware of such financial issues, if only because some of the very best work has arisen in answering the questions of industry rather than academe.

Guide to the chapters

Chapter one is a brief warning, especially to beginners, that the expected

worth of something is not a good guide to its price. That idea has to be shaken off and arbitrage pricing take its place.

Chapter two develops the idea of hedging and pricing by arbitrage in the discrete-time setting of binary trees. The key probabilistic concepts of conditional expectation, martingales, change of measure, and representation are all introduced in this simple framework, accompanied by illustrative examples.

Chapter three repeats all the work of its predecessor in the continuous-time setting. Brownian motion is brought out, as well as the Itô calculus needed to manipulate it, culminating in a derivation of the Black–Scholes formula.

Chapter four runs through a variety of actual financial instruments, such as dividend paying equities, currencies and coupon paying bonds, and adapts the Black–Scholes approach to each in turn. A general pattern of the distinction between tradable and non-tradable quantities leads to the definition the market price of risk, as well as a warning not to take that name too seriously. A section on quanto products provides a showcase of examples.

Chapter five is about the interest rate market. In spirit, a market of bonds is much like a market of stocks, but the richness of this market makes it more than just a special case of Black–Scholes. Market models are discussed with a joint short-rate/HJM approach, which lies within the general continuous framework set up in chapter three. One section details a few of the many possible interest rate contracts, including swaps, caps/floors and swaptions. This is a substantial chapter reflecting the depth of financial and technical knowledge that has to be introduced in an understandable way. The aim is to tell one basic story of the market, which all approaches can slot into.

Chapter six concludes with some technical results about larger and more general models, including multiple stock n -factor models, stochastic numeraires, and foreign exchange interest-rate models. The running link between the existence of equivalent martingale measures and the ability to price and hedge is finally formalised.

A short bibliography, complete answers to the (small) number of exercises, a full glossary of technical terms and an index are in the appendices.

How to read this book

The book can be read either sequentially as an unfolding story, or by random access to the self-contained sections. The occasional questions are to allow

practice of the requisite skills, and are never essential to the development of the material.

A reader is not expected to have any particular prior body of knowledge, except for some (classical) differential calculus and experience with symbolic notation. Some basic probability definitions are contained in the glossary, whereas more advanced readers will find technical asides in the text from time to time.

Acknowledgements

We would like to thank David Tranah at CUP for politely never mentioning the number of deadlines we missed, as well as his much more invaluable positive assistance; the many readers in London, New York and various universities who have been subjected to writing far worse than anything remaining in the finished edition. Special thanks to Lorne Whiteway for his help and encouragement.

Martin Baxter
Andrew Rennie

June 1996

The parable of the bookmaker

A bookmaker is taking bets on a two-horse race. Choosing to be scientific, he studies the form of both horses over various distances and goings as well as considering such factors as training, diet and choice of jockey. Eventually he correctly calculates that one horse has a 25% chance of winning, and the other a 75% chance. Accordingly the odds are set at 3-1 against and 3-1 on respectively.

But there is a degree of popular sentiment reflected in the bets made, adding up to \$5 000 for the first and \$10 000 for the second. Were the second horse to win, the bookmaker would make a net profit of \$1667, but if the first wins he suffers a loss of \$5000. The expected value of his profit is $25\% \times (-\$5000) + 75\% \times (\$1667) = \$0$, or exactly even. In the long term, over a number of similar but independent races, the law of averages would allow the bookmaker to break even. Until the long term comes, there is a chance of making a large loss.

Suppose however that he had set odds according to the money wagered – that is, not 3-1 but 2-1 against and 2-1 on respectively. Whichever horse wins, the bookmaker exactly breaks even. The outcome is irrelevant.

In practice the bookmaker sells more than 100% of the race and the odds are shortened to allow for profit (see table). However, the same pattern emerges. Using the actual probabilities can lead to long-term gain but there is always the chance of a substantial short-term loss. For the bookmaker to earn a steady riskless income, he is best advised to assume the horses' probabilities are something different. That done, he is in the surprising

position of being disinterested in the outcome of the race, his income being assured.

A note on odds

When a price is quoted in the form $n-m$ against, such as 3-1 against, it means that a successful bet of $\$m$ will be rewarded with $\$n$ plus stake returned. The implied probability of victory (were the price fair) is $m/(m+n)$. Usually the probability is less than half a chance so the first number is larger than the second. Otherwise, what one might write as 1-3 is often called odds of 3-1 on.

Actual probability	25%	75%	
Bets placed	\$5000	\$10 000	
1. Quoted odds	13-5 against	15-4 on	
Implied probability	28%	79%	Total = 107%
Profit if horse wins	-\$3000	\$2333	Expected profit = \$1000
2. Quoted odds	9-5 against	5-2 on	
Implied probability	36%	71%	Total = 107%
Profit if horse wins	\$1000	\$1000	Expected profit = \$1000

Allowing the bookmaker to make a profit, the odds change slightly. In the first case, the odds relate to the actual probabilities of a horse winning the race. In the second, the odds are derived from the amounts of money wagered.

Chapter 1
Introduction

Financial market instruments can be divided into two distinct species. 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, for example – 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 stock without needing the capital to buy it outright.

In form, one species depends on the other – without the underlying (stock) there could be no future claims – but the connection between the two is sufficiently complex and uncertain for both to trade fiercely in the same market. The apparently random nature of stocks filters through to the claims – they appear random too.

Yet mathematicians have known for a while that to be random is not necessarily to be without some internal structure – put crudely, things are often random in non-random ways. The study of probability and expectation shows one way of coping with randomness and this book will build on probabilistic foundations to find the strongest possible links between claims and their random underlying stocks. The current state of truth is, however, unfortunately complex and there are many false trails through this zoo of the new. Of these, one is particularly tempting.

1.1 Expectation pricing

Consider playing the following game – someone tosses a coin and pays you one dollar for heads and nothing for tails. What price should you pay for this prize? If the coin is fair, then heads and tails are equally likely – about half the time you should win the dollar and the rest of the time you should receive nothing. Over enough plays, then, you expect to make about fifty cents a go. So paying more than fifty cents seems extravagant and less than fifty cents looks extravagant for the person offering the game. Fifty cents, then, seems about right.

Fifty cents is also the expected profit from the game under a more formal, mathematical definition of expectation. A probabilistic analysis of the game would observe that although the outcome of each coin toss is essentially random, this is not inconsistent with a deeper non-random structure to the game. We could posit that there was a fixed measure of likelihood attached to the coin tossing, a *probability* of the coin landing heads or tails of $\frac{1}{2}$. And along with a probability ascription comes the idea of expectation, in this discrete case, the total of each outcome's value weighted by its attached probability. The expected payoff in the game is $\frac{1}{2} \times \$1 + \frac{1}{2} \times \$0 = \$0.50$.

This formal expectation can then be linked to a 'price' for the game via something like the following:

Kolmogorov's strong law of large numbers

Suppose we have a sequence of independent random numbers X_1, X_2, X_3 , and so on, all sampled from the same distribution, which has mean (expectation) μ , and we let S_n be the arithmetical average of the sequence up to the n th term, that is $S_n = (X_1 + X_2 + \dots + X_n)/n$. Then, with probability one, as n gets larger the value of S_n tends towards the mean μ of the distribution.

If the arithmetical average of outcomes tends towards the mathematical expectation with certainty, then the average profit/loss per game tends towards the mathematical expectation less the price paid to play the game. If this difference is positive, then in the long run it is certain that you will end up in profit. And if it is negative, then you will approach an overall loss with certainty. In the *short term* of course, nothing can be guaranteed, but over time, expectation will out. Fifty cents is a fair price in this sense.

But is it an enforceable price? Suppose someone offered you a play of the game for 40 cents in the dollar, but instead of allowing you a number of plays, gave you just one for an arbitrarily large payoff. The strong law lets you take advantage of them over repeated plays: 40 cents a dollar would then be financial suicide, but it does nothing if you are allowed just one play. Mortgaging your house, selling off all your belongings and taking out loans to the limit of your credit rating would not be a rational way to take advantage of this source of free money.

So the 'market' in this game could trade away from an expectation justified price. Any price might actually be charged for the game in the short term, and the number of 'buyers' or 'sellers' happy with that price might have nothing to do with the mathematical expectation of the game's outcome. But as a guide to a starting price for the game, a ball-park amount to charge, the strong law coupled with expectation seems to have something going for it.

Time value of money

We have ignored one important detail – the time value of money. Our analysis of the coin game was simplified by the payment for and the payoff from the game occurring at the same time. Suppose instead that the coin game took place at the end of a year, but payment to play had to be made at the beginning – in effect we had to find the value of the coin game's contingent payoff not as of the future date of play, but as of now.

If we are in January, then one dollar in December is not worth one dollar now, but something less. Interest rates are the formal acknowledgement of this, and bonds are the market derived from this. We could assume the existence of a market for these future promises, the prices quoted for these bonds being structured, derivable from some interest rate. Specifically:

Time value of money

We assume that for any time T less than some time horizon τ , the value now of a dollar promised at time T is given by $\exp(-rT)$ for some constant $r > 0$. The rate r is then the *continuously compounded* interest rate for this period.

The interest rate market doesn't have to be this simple; r doesn't have to be constant. And indeed in real markets it isn't. But here we assume it is. We can derive a strong-law price for the game played at time T . Paying 50 cents at time T is the same as paying $50 \exp(-rT)$ cents now. Why? Because the payment of 50 cents at time T can be guaranteed by buying half a unit of the appropriate bond (that is, promise) now, for cost $50 \exp(-rT)$ cents. Thus the strong-law price must be not 50 cents but $50 \exp(-rT)$ cents.

Stocks, not coins

What about real stock prices in a real financial market? One widely accepted model holds that stock prices are *log-normally* distributed. As with the time value of money above, we should formalise this belief.

Stock model

We assume the existence of a random variable X , which is normally distributed with mean μ and standard deviation σ , such that the change in the logarithm of the stock price over some time period T is given by X . That is

$$\log S_T = \log S_0 + X, \quad \text{or} \quad S_T = S_0 \exp(X).$$

Suppose, now, that we have some claim on this stock, some contract that agrees to pay certain amounts of money in certain situations – just as the coin game did. The oldest and possibly most natural claim on a stock is the *forward*: two parties enter into a contract whereby one agrees to give the other the stock at some agreed point in the future in exchange for an amount agreed now. The stock is being *sold forward*. The 'pricing question' for the forward stock 'game' is: what amount should be written into the contract now to pay for the stock one year in the future?

We can dress this up in formal notation – the stock price at time T is given by S_T , and the forward payment written into the contract is K , thus the value of the contract at its expiry, that is when the stock transfer actually takes place, is $S_T - K$. The time value of money tells us that the value of this claim as of now is $\exp(-rT)(S_T - K)$. The strong law suggests that the expected value of this random amount, $\mathbb{E}(\exp(-rT)(S_T - K))$, should

be zero. If it is positive or negative, then long-term use of that pricing should lead to one side's profit. Thus one apparently reasonable answer to the pricing question says K should be set so that $\mathbb{E}(\exp(-rT)(S_T - K)) = 0$, which happens when $K = \mathbb{E}(S_T)$.

What is $\mathbb{E}(S_T)$? We have assumed that $\log(S_T) - \log(S_0)$ is normally distributed with mean μ and variance σ^2 – thus we want to find $\mathbb{E}(S_0 \exp(X))$, where X is normally distributed with mean μ and standard deviation σ . For that, we can use a result such as:

The law of the unconscious statistician

Given a real-valued random variable X with probability density function $f(x)$ then for any integrable real function h , the expectation of $h(X)$ is

$$\mathbb{E}(h(X)) = \int_{-\infty}^{\infty} h(x)f(x) dx.$$

Since X is normally distributed, the probability density function for X is

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-(x - \mu)^2}{2\sigma^2}\right).$$

Integration and the law of the unconscious statistician then tells us that the expected stock price at time T is $S_0 \exp(\mu + \frac{1}{2}\sigma^2)$. This is the strong-law justified price for the forward contract; just as with the coin game, it can only be a suggestion as to the market's trading level. But the technique will clearly work for more than just forwards. Many claims are capable of translation into functional form, $h(X)$, and the law of the unconscious statistician should be able to deliver an expected value for them. Discounting this expectation then gives a theoretical value which the strong law tempts us into linking with economic reality.

1.2 Arbitrage pricing

So far, so plausible – but seductive though the strong law is, it is also completely useless. The price we have just determined for the forward could only be the market price by an unfortunate coincidence. With markets where

the stock can be bought and sold freely and arbitrary positive and negative amounts of stock can be maintained without cost, trying to trade forward using the strong law would lead to disaster – in most cases there would be *unlimited* interest in selling forward to you at that price.

Why does the strong law fail so badly with forwards? As mentioned above in the context of the coin game, the strong law cannot enforce a price, it only suggests. And in this case, *another completely different mechanism does enforce a price*. The fair price of the contract is $S_0 \exp(rT)$. It doesn't depend on the expected value of the stock, it doesn't even depend on the stock price having some particular distribution. Either counterparty to the contract can in fact *construct* the claim at the *start* of the contract period and then just wait patiently for expiry to exchange as appropriate.

Construction strategy

Consider the seller of the contract, obliged to deliver the stock at time T in exchange for some agreed amount. They could borrow S_0 now, buy the stock with it, put the stock in a drawer and just wait. When the contract expires, they have to pay back the loan – which if the continuously compounded rate is r means paying back $S_0 \exp(rT)$, but they have the stock ready to deliver. If they wrote less than $S_0 \exp(rT)$ into the contract as the amount for forward payment, then they would lose money *with certainty*.

So the forward price is bounded below by $S_0 \exp(rT)$. But of course, the buyer of the contract can run the scheme in reverse, thus writing *more* than $S_0 \exp(rT)$ into the contract would guarantee them a loss. The forward price is bounded above by $S_0 \exp(rT)$ as well.

Thus there is an *enforced* price, not of $S_0 \exp(\mu + \frac{1}{2}\sigma^2)$ but $S_0 \exp(rT)$. Any attempt to strike a different price and offer it into a market would inevitably lead to someone taking advantage of the free money available via the construction procedure. And unlike the coin game, mortgaging the house *would* now be a rational action. This type of market opportunism is old enough to be ennobled with a name – *arbitrage*. The price of $S_0 \exp(rT)$ is an *arbitrage price* – it is justified because any other price could lead to unlimited riskless profits for one party. The strong law wasn't wrong – if $S_0 \exp(\mu + \frac{1}{2}\sigma^2)$ is greater than $S_0 \exp(rT)$, then a buyer of a forward contract *expects* to make money. (But then of course, if the stock is expected to grow faster than the riskless interest rate r , so would buyers of the stock itself.) But the existence of an *arbitrage price*, however surprising, overrides the strong

law. To put it simply, if there is an *arbitrage price*, any other price is too dangerous to quote.

1.3 Expectation vs arbitrage

The strong law and expectation give the wrong price for forwards. But in a certain sense, the forward is a special case. The construction strategy – buying the stock and holding it – certainly wouldn't work for more complex claims. The standard call option which offers the buyer the right but not the obligation to receive the stock for some strike price agreed in advance certainly couldn't be constructed this way. If the stock price ends up above the strike, then the buyer would exercise the option and ask to receive the stock – having it salted away in a drawer would then be useful to the seller. But if the stock price ends up below the strike, the buyer will abandon the option and any stock owned by the seller would have incurred a pointless loss.

Thus maybe a strong-law price would be appropriate for a call option, and until 1973, many people would have agreed. Almost everything appeared safe to price via expectation and the strong law, and only forwards and close relations seemed to have an *arbitrage price*. Since 1973, however, and the infamous Black–Scholes paper, just how wrong this is has slowly come out. Nowhere in this book will we use the strong law again. Just to muddy the waters, though, *expectation* will be used repeatedly, but it will be as a tool for risk-free *construction*. All derivatives can be built from the underlying – *arbitrage* lurks everywhere.