


OXFORD

An Introduction to QUANTITATIVE FINANCE

Stephen Blyth



An Introduction to Quantitative Finance

STEPHEN BLYTH

Professor of the Practice of Statistics, Harvard University
Managing Director, Harvard Management Company



OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Stephen Blyth 2014

The moral rights of the author have been asserted

First Edition published in 2014

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2013941968

ISBN 978-0-19-966658-4 (hbk.)

ISBN 978-0-19-966659-1 (pbk.)

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

PREFACE

This book is based on *Applied Quantitative Finance on Wall Street*, an upper level undergraduate statistics course that I have taught for several years at Harvard University. The students taking the class are typically undergraduate concentrators in mathematics, applied mathematics, statistics, physics, economics and computer science; master's students in statistics; or PhD students in quantitative disciplines. No prior exposure to finance or financial terminology is assumed.

Many students who are considering a career in finance study this material to gain an insight into the machinery of financial engineering. However, some of my students have no interest in a financial career, but simply enjoy probability and are motivated to explore one of its compelling applications and investigate a new way of thinking about uncertainty. Others may be curious about quantitative finance due to the regulatory and policy-making attention the industry has received since the financial crisis.

The book concerns financial derivatives, a derivative being a contract or trade (or bet, depending on your prejudices) between two entities or counterparties whose value is a function of—derives from—the price of an underlying financial asset. We define various derivative contracts and describe the quantitative and probabilistic tools that were developed to address issues encountered by practitioners as markets developed. The book is steeped in practice, as the methods we explore only developed to such an extent because the markets themselves grew exponentially fast. Whilst we develop theory here, this book is not theoretical in the 'existing only in theory' sense. The products we consider are traded in significant size and have meaningful economic impact.

Probability provides the key tools for analysing and valuing derivatives. The price of a stock or a bond at some future time is a random variable. The payout of a derivative contract is a function of a random variable, and thus itself another random variable. We show that the price one should pay for a derivative is closely linked to the expected value of its payout, and that suitably scaled derivative prices are martingales, which are fundamentally important objects in probability theory. The book focuses largely on interest rate derivatives—where the underlying financial variable is an interest rate—for three reasons. First, they constitute by far the largest and most economically important derivative market in the world. Secondly, they are typically the most challenging mathematically. Thirdly, they have generally been less well addressed by finance textbooks.

Background and motivation

The structure and content of the book are shaped by the experiences of my career in both academia and on Wall Street. I read mathematics as an undergraduate at Cambridge University where I was fortunate to be taught probability by both Frank Kelly and David

Williams. I obtained a PhD in Statistics at Harvard University with Kjell Doksum before teaching for two years at Imperial College London, where I lectured the introductory probability course for mathematics undergraduates.

My move to Wall Street was catalysed indirectly by President Clinton's decision to cancel the superconducting super collider in 1993. Two Harvard housemates, both theoretical physicists, saw the academic job market for physicists collapse overnight and both quickly found employment at Goldman Sachs in New York. Their assertion that derivative markets (whatever in fact they were) were 'pretty interesting' and mathematically challenging convinced me to contemplate a move from my nascent academic career.

I started at HSBC in London, hired directly by a Cambridge PhD mathematician. On my first day at work I was tasked to calculate the expected value of a function of bivariate normal random variables. I moved subsequently to Morgan Stanley in New York where I became managing director and the market maker for US interest rate options. Throughout my stint at Morgan Stanley I worked with another Cambridge mathematician. Together we developed a training program for new analysts, motivated by mathematical and probabilistic challenges and subtleties we encountered in the rapidly developing interest rate options markets of the late 1990s. Several of the problems from that mini-course have found their way into this book. Indeed much of the theory I develop is motivated by a desire to provide a suitably rigorous yet accessible foundation to tackle problems I encountered whilst trading derivatives. This fundamental motivation led me to create the course at Harvard and subsequently develop this book, which I believe combines an unusual blend of derivatives trading experience and rigorous academic background.

After a spell in London running an interest rate proprietary trading group at Deutsche Bank, I rejoined Harvard in 2006 where I currently wear two hats: head of public markets at the Harvard Management Company, the subsidiary of the university responsible for managing the endowment; and Professor of the Practice of statistics at Harvard University, where I teach the course. My role at the endowment includes responsibility for Harvard's investment activities across public bond, equity and commodity markets. This ongoing presence at the heart of financial markets—I continue to trade each day derivative contracts explained in the text—means that the course, and thus I hope the book, retains immediacy to its field. In particular, the financial crisis of 2008–2009 challenged many basic foundational assumptions both academics and practitioners had taken for granted and I detail several of these in the exposition, some surprisingly early.

The Harvard statistician Arthur Dempster made the distinction between 'procedural' statistics, involving the rote application of procedures and methods, and 'logistic' statistics in which reasoning about the problem at hand—using subject-matter knowledge and judgment in addition to quantitative expertise—plays a central role. This book is written in the logistic tradition, interweaving experience and judgment from the real world of finance with the development of appropriate quantitative techniques.

It is often not obvious to an outsider which issues are particularly important to traders and investors in the Wall Street derivative markets. Whilst everything in this book is to some extent motivated by the practical world, we highlight areas where the degree of market relevance is not immediately apparent. We also include market anecdotes and tales from the Street to provide insights and richness into the material.

This book has been motivated by and distilled from my experience on Wall Street, and also by my love for probability. In addition to wanting to understand elements of quantitative finance in practice, students may wish to read the book simply because they have enjoyed their introduction to probability and have a desire to see how it can be applied. Whilst there are many ways to approach financial mathematics, I find the probabilistic route both appealing and fulfilling. There are highly intuitive elements to the theory, for example prices being expected values. Furthermore, we introduce in a straightforward manner powerful concepts such as numeraires and martingales that drive much of modern theory, and discover that there are subtleties in various unexpected places that demand precise and non-trivial understanding.

Prerequisites

The sole prerequisite for mastering the material in this book is a solid introductory undergraduate course in probability (represented by Statistics 110 at Harvard, Part IA Probability at Cambridge University, or Probability & Statistics I at Imperial College). Familiarity is required with: discrete and continuous random variables and distributions; expectation, in particular expectation of a function of a random variable; conditional expectation; the binomial and normal distributions; and an elementary version of the central limit theorem. The single-variable calculus typically associated with such a course—integration by parts, chain rule for differentiation and elementary Taylor series—is used at several points. Exercises at the end of this preface give a sense of the probability prerequisite.

The book is otherwise self-contained and in particular requires no additional preparation or exposure to finance. The necessary financial terminology is introduced and explained as required. I tell my students that the edifice of quantitative finance was largely built by pioneers from academic physics or mathematics who entered Wall Street with no knowledge of finance and having read no finance books, and who still have not.

This is not a book on economics. There are many excellent economics texts on capital markets, corporate finance, portfolio theory and related matters, which complement this material. However, derivative markets are a world largely populated by traders and mathematicians, not economists. Nor is this a book on time series methods or prediction. As we see early on, derivative pricing—even that concerning the price of a forward contract which is an exchange of an asset in the future—is largely unrelated to prediction.

The mathematical prerequisite is modest and no more extensive than that for an introductory undergraduate probability course, although adeptness at logical quantitative reasoning is important. Whilst we build an appropriately rigorous base for our results, this is not a theoretical mathematics book. By developing most of our theory using discrete-time methods, we can address important issues in finance without exposure to stochastic processes, Ito calculus, partial differential equations, Monte Carlo or other numerical methods, all of which play important roles in continuous-time financial mathematics. We raise signposts to areas of further study in the rough guide to continuous time that concludes the book.

I have necessarily had to be ruthless with material I have excluded in order to keep the book appropriate to a one-semester course of approximately 33 lecture hours. The book is selective rather than encyclopaedic and moves at a lively pace. This enables students to be exposed to powerful theory and substantive problems in one semester. Throughout the book I raise signposts to more advanced topics and to different approaches to the material which naturally arise from the exposition. Even an encyclopaedic tome cannot cover all dimensions of derivative markets and associated theory.

Outline

The book is organized in five parts. The first, short part contains some preliminaries regarding interest rates and zero coupon bonds, and a very brief sketch of assets.

The bulk of the material is contained in Parts II–IV. In Part II, we define and explore basic derivative products including forwards, swaps and options. Our first derivative—the forward contract—is introduced and we consider methods for valuation and pricing which we discover do not depend on any probabilistic modelling. Recent violations of traditional assumptions during the financial crisis are soon encountered. We introduce elementary interest rate derivatives, such as forward rate agreements (FRAs) and swaps, and define futures, the first contract whose price depends on distributional assumptions of the underlying asset. We state formally the no-arbitrage principle, a foundation of quantitative finance, which makes precise the arguments to which we have already appealed. Recent empirical challenges to the principle of no-arbitrage manifested during the financial crisis are discussed. We then introduce and define call and put options, key building blocks of modern derivative markets. We investigate option properties and establish model-independent bounds on their price. We also review options on forwards, an area that can still cause puzzlement on Wall Street.

Part III focuses on the key pricing arguments of replication and risk-neutrality. We formalize arguments regarding replication and hedging on a binomial tree, and define risk-neutral valuation in this setting. We show that prices are appropriately scaled expected values under the risk-neutral probability distribution, a result which leads to the fundamental theorem of asset pricing for the binomial tree, the most powerful result of the book. Introducing in a natural manner the concepts of numeraire and martingale, we give a general form of the fundamental theorem, which states that no-arbitrage is equivalent to the ratios of prices to a numeraire being martingales. Bridging in a simple manner to continuous time, we take limits using the central limit theorem and move to a continuous case where expectation under a lognormal density immediately gives the Black–Scholes formula. We review its properties and introduce the concepts of delta and vega. The key duality between option prices and probability distributions is then explored in several ways.

Part IV develops the understanding of interest rate options, the largest, most important and most mathematically challenging of options markets. No book covers the material particularly well and the fabric of these markets is typically only seen clearly by its practitioners. We introduce and value interest rate caps, floors and swaptions, using an elegant choice of numeraire. We then define Bermudan swaptions and derive arbitrage bounds for their value,

and construct trades using cancellable swaps. Two further topics in interest rate options are discussed: *libor-in-arrears* and *general convexity corrections*, in particular the details of one of the great trades in derivatives history; and a brief description of the BGM model and its volatility surfaces, which underlies much of current interest rate modelling.

In the concluding Part V, we provide a brief introductory sketch of the continuous-time analogue of the theory developed in the book, and raise signposts to the partial differential equation approach to mathematical finance.

The book includes substantive homework problems in each chapter (other than Part V) which illustrate and build on the material. The aim is to make the book fulfilling, challenging and entertaining, and to convey the immediacy and practical context of the material we cover.

EXERCISES

The following exercises provide an indication of the level of probability proficiency that is required to master the material in this book.

1. The normal distribution and expectation

Let $X \sim N(\mu, \psi^2)$ and K be a constant. Calculate

- (a) $E(I\{X > K\})$ where I is the indicator function.
- (b) $E(\max\{K - X, 0\})$.

Hint Use $(K - X) = (K - \mu) - (X - \mu)$.

- (c) $E(e^{tX})$.

Leave answers where appropriate in terms of the normal cumulative distribution function $\Phi(\cdot)$.

Note In the language of derivatives (a) is essentially the price of a *digital call option* and (b) is the price of a *put option*. We encounter these in Chapter 7.

2. The lognormal distribution

A random variable Y is said to be *lognormally distributed* with parameters μ and σ if $\log Y \sim N(\mu, \sigma^2)$. We sometimes write $Y \sim \text{lognormal}(\mu, \sigma^2)$.

- (a) Using your answer to Question 1(c), calculate $E(Y)$ and $\text{Var}(Y)$.
- (b) Suppose $\mu = \log f - \frac{1}{2}\sigma^2$ for a constant $f > 0$. Calculate $E(Y)$ and show that $\text{Var}(Y) = f^2\sigma^2(1 + \frac{\sigma^2}{2} + \frac{\sigma^4}{6} + \text{higher order terms})$.

3. Conditional expectation

A sequence of random variables $X_0, X_1, \dots, X_n, \dots$ is defined by $X_0 = 1$ and $X_n = X_{n-1}\xi_n$, where $\xi_i, i = 0, 1, \dots$, are independent and identically distributed with

$$\xi_i = \begin{cases} 1 + u & \text{with probability } p \\ 1 + d & \text{with probability } 1 - p, \end{cases}$$

where $u > d$.

(a) Calculate $E(X_n | X_{n-1})$.

(b) Let $Y_n = \frac{X_n}{(1+r)^n}$ for a constant $r > 0$. Find in terms of p, u and d the value of r such that $E(Y_n | Y_{n-1}) = Y_{n-1}$. Hence find in terms of u and d the range of such possible values for r .

(c) Suppose $E(Y_n | Y_{n-1}) = Y_{n-1}$ for all n . Use any result concerning conditional expectation to prove that $E(Y_n | Y_m) = Y_m$ for all $n > m \geq 0$.

Note The result in (c) is essentially the definition that Y_n is a *martingale*, which we encounter in Chapter 9.

4. The normal cumulative distribution function

Show that the first three non-zero terms of the Taylor series for the normal cumulative distribution function $\Phi(x)$ around zero are

$$\Phi(x) = \frac{1}{2} + \frac{x}{\sqrt{2\pi}} - \frac{x^3}{6\sqrt{2\pi}}.$$

ACKNOWLEDGMENTS

I am grateful to the many colleagues, fellow practitioners and students who read earlier versions of this book, including Raj Bhuptani, Joseph Blitzstein, John Campbell, Howard Corb, Kjell Doksum, Mohamed El-Erian, Jessy Hwang, Amol Jain, Bo Jiang, Frank Kelly, David Land, David Moore and Andy Morton, plus the entire Tufnell Park team. They contributed many helpful and insightful comments.

I was fortunate to have been taught probability by Frank Kelly and David Williams at Cambridge University, both of whom inspired in me a love for the subject that has never faded, and statistical theory and reasoning by Kjell Doksum and Arthur Dempster at Harvard University. These four professors provided me the foundational underpinnings from which I was able to tackle the world of quantitative finance. I learnt finance largely on the job in London and New York, in particular from two fine Cambridge mathematicians, Dan George and David Moore. I sat next to David for six years and gained from him an appreciation of the subtleties of financial markets and the theory required to understand them.

I am particularly grateful to Jane Mendillo who kindly allowed me to spend time in the classroom amid our responsibilities at the endowment, and who has been an unwavering supporter of all my endeavours at the University, to Xiao-Li Meng for welcoming me so warmly to the faculty of the Statistics department, and to Mohamed El-Erian for bringing me back to Harvard in the first place.

Keith Mansfield at OUP has been a strong advocate for this project from the start, and has guided me wisely through its evolution. Willis Ho, Kelly Sampson and Alex Zhu expertly helped produce the figures in the book, and Johnny Roche provided technical support throughout.

Finally, I thank all the students who have taken my class at Harvard, out of which this book grew. They have made teaching an unqualified joy.

CONTENTS

PART I PRELIMINARIES

1 Preliminaries	3
1.1 Interest rates and compounding	3
1.2 Zero coupon bonds and discounting	5
1.3 Annuities	6
1.4 Daycount conventions	7
1.5 An abridged guide to stocks, bonds and FX	8
1.6 Exercises	9

PART II FORWARDS, SWAPS AND OPTIONS

2 Forward contracts and forward prices	13
2.1 Derivative contracts	13
2.2 Forward contracts	14
2.3 Forward on asset paying no income	14
2.4 Forward on asset paying known income	16
2.5 Review of assumptions	17
2.6 Value of forward contract	19
2.7 Forward on stock paying dividends and on currency	20
2.8 Physical versus cash settlement	20
2.9 Summary	21
2.10 Exercises	22
3 Forward rates and libor	25
3.1 Forward zero coupon bond prices	25
3.2 Forward interest rates	26
3.3 Libor	27
3.4 Forward rate agreements and forward libor	27
3.5 Valuing floating and fixed cashflows	29
3.6 Exercises	30
4 Interest rate swaps	33
4.1 Swap definition	33
4.2 Forward swap rate and swap value	34
4.3 Spot-starting swaps	35
4.4 Swaps as difference between bonds	36
4.5 Exercises	37

5 Futures contracts.	41
5.1 Futures definition	41
5.2 Futures versus forward prices	42
5.3 Futures on libor rates	43
5.4 Exercises	45
6 No-arbitrage principle	47
6.1 Assumption of no-arbitrage	47
6.2 Monotonicity theorem	48
6.3 Arbitrage violations	49
6.4 Exercises	50
7 Options.	53
7.1 Option definitions	53
7.2 Put-call parity	56
7.3 Bounds on call prices	57
7.4 Call and put spreads	58
7.5 Butterflies and convexity of option prices	60
7.6 Digital options	61
7.7 Options on forward contracts	62
7.8 Exercises	64
PART III REPLICATION, RISK-NEUTRALITY AND THE FUNDAMENTAL THEOREM	
8 Replication and risk-neutrality on the binomial tree	71
8.1 Hedging and replication in the two-state world	71
8.2 Risk-neutral probabilities	73
8.3 Multiple time steps	74
8.4 General no-arbitrage condition	76
8.5 Exercises	78
9 Martingales, numeraires and the fundamental theorem	81
9.1 Definition of martingales	81
9.2 Numeraires and fundamental theorem	83
9.3 Change of numeraire on binomial tree	85
9.4 Fundamental theorem: a pragmatic example	86
9.5 Fundamental theorem: summary	87
9.6 Exercises	88
10 Continuous-time limit and Black–Scholes formula	93
10.1 Lognormal limit	93
10.2 Risk-neutral limit	95
10.3 Black–Scholes formula	97
10.4 Properties of Black–Scholes formula	98

10.5	Delta and vega	101
10.6	Incorporating random interest rates	105
10.7	Exercises	106
11	Option price and probability duality	111
11.1	Digitals and cumulative distribution function	111
11.2	Butterflies and risk-neutral density	113
11.3	Calls as spanning set	114
11.4	Implied volatility	115
11.5	Exercises	116
 PART IV INTEREST RATE OPTIONS		
12	Caps, floors and swaptions.	119
12.1	Caplets	119
12.2	Caplet valuation and forward numeraire	120
12.3	Swaptions and swap numeraire	122
12.4	Summary	124
12.5	Exercises	124
13	Cancellable swaps and Bermudan swaptions	127
13.1	European cancellable swaps	127
13.2	Callable bonds	128
13.3	Bermudan swaptions	130
13.4	Bermudan swaption exercise criteria	132
13.5	Bermudan cancellable swaps and callable bonds	134
13.6	Exercises	136
14	Libor-in-arrears and constant maturity swap contracts.	141
14.1	Libor-in-arrears	141
14.2	Libor-in-arrears convexity correction	142
14.3	Classic libor-in-arrears trade	144
14.4	Constant maturity swap contracts	145
14.5	Exercises	146
15	The Brace–Gatarek–Musiela framework.	149
15.1	BGM volatility surface	149
15.2	Option price dependence on BGM volatility surface	150
15.3	Exercises	153
 PART V TOWARDS CONTINUOUS TIME		
16	Rough guide to continuous time	157
16.1	Brownian motion as random walk limit	157
16.2	Stochastic differential equations and geometric Brownian motion	158

16.3	Ito's lemma	160
16.4	Black–Scholes equation	161
16.5	Ito and change of numeraire	162
	Glossary of notation	167
	References	171
	Index	173

PART I

Preliminaries

1

Preliminaries

One of the interview questions I was asked when applying for my first job in finance was, ‘Would you rather have one dollar today or one dollar in one year’s time?’ The question was asked abruptly without any polite preamble, as if to throw me off guard. However, the answer is clearly that one would rather have the dollar today since one can, amongst other things, deposit the dollar at a bank and receive interest for the year. This is preferable to receiving a dollar in a year’s time, provided interest rates are not negative. The bulk of this preliminary chapter is concerned with the mechanics of interest rates, compounding and computing the value today of receiving a dollar or other unit of currency at some date in the future.

Note Almost universally throughout quantitative finance, interest rates are assumed to be non-negative. However, finance in practice has the habit of throwing up unexpected complexity even in the simplest of settings, and interest rates are no exception. Figure 1.1 shows the graph of the two-year Swiss interest rate (precisely, the two-year government bond yield, which we encounter in Chapter 4). We observe that this rate has been negative for non-negligible periods. How can this be so? Surely we can simply keep hold of our Swiss francs, and not deposit them at negative rates, thus establishing an effective floor of 0% on interest rates? In practice, it is not easy to keep billions of Swiss francs in safes (or under the mattress). The Swiss national bank—as part of its aim to reduce the strength of its currency during a particularly intense episode of euro concerns in late 2010—penalized all deposits of Swiss francs by way of negative interest rates. However, throughout this book we assume, unless stated otherwise, that all interest rates are non-negative, keeping in mind that markets have a tendency to challenge even fundamentally sound assumptions.

1.1 Interest rates and compounding

Suppose we deposit (or invest) an amount $\$N$ at a rate r per annum, compounded annually. Then we have after one year $\$(1 + r)$, and after T years an amount $\$(1 + r)^T$. Throughout the book we will express interest rates as decimals, so r is a rate of $100r\%$.