Preface

Modern finance overlaps with many fields of mathematics, in particular, probability theory, linear algebra, calculus, partial differential equations, stochastic calculus, numerical mathematics, and not least programming. The diversity of mathematical skills makes finance a very challenging subject, putting a lot of strain on its prospective students. *Mathematical Techniques in Finance* offers an introduction to the mathematical tools which are needed to price uncertain income streams such as derivative securities. It is primarily intended as a textbook for Masters in Finance courses with a significant quantitative element, although it has also been popular with Finance PhD students here at Imperial College London, and it has found its way on to the desks of financial analysts.

This book is about the *active and practical use* of mathematics with the main focus on three interrelated financial topics: asset pricing, portfolio allocation, and risk measurement. The book contains a mix of applications and theory working together in a happy union; theory underpins the applications and the applications illustrate the theory. Working out the exercises is more important than trying to memorize the financial and mathematical theory contained in the text.

Study Guide

Before you start reading the book take a look at the book's website

http://pup.princeton.edu/titles/7606.html

to find out what resources are available. Secondly, set up your computer. The computer programs in this book are written in GAUSS, which is a simple language based on vectors and matrices. GAUSS LIGHT is an inexpensive student version that has all the functionality of GAUSS, except it will only handle matrices with no more than 10 000 elements, which is appropriate for most applications in the book. Many economics and finance departments already have a GAUSS license; on the other hand engineering departments seem to prefer MATLAB. This book may provide support to MATLAB and Visual Basic users in the future—watch the website. There is no need to worry about using C++ at this stage; it runs fast but the set-up costs are high. To use an analogy, GAUSS is a car with an automatic gearbox, whereas C++ has a manual one. At this stage we are quite happy just turning the steering wheel and admiring the view without having to worry about the gear shifts. There is always time to learn C++ later if you decide that you really like programming.

Not all the material in this book is suitable for all students. There are essentially two long coherent themes appropriate for Masters programmes, and several digressions intended for short courses aimed at doctoral students. The difficulty is largely conceptual, not mathematical. The book uses linear algebra at the level of Anton (2000), calculus at the level of Binmore and Davies (2001), and probability at the level of Mood et al. (1974); all three are standard undergraduate textbooks. The background to Itô calculus is self-contained and the applications of Itô calculus require little more than partial differentiation and ordinary integration.

Longer Masters Courses

The *discrete-time complete market trail* (Chapters 1, 2, 5 and 6) has a number of exciting computer simulations looking into dynamic asset pricing. Here one can get away with very little mathematics, especially if one is willing to take a few crucial results on trust.

Chapter 1 establishes the basics of the one-period model, shows how securities can be represented by vectors and matrices, and introduces the concept of hedging. It also provides a simple context in which to explore the GAUSS commands. Chapter 2 introduces important financial notions such as returns, arbitrage and state prices, and gives examples of asset pricing both in complete and incomplete markets. Sections 2.1–2.4 are not essential for the complete market modelling and can be skipped.

Chapter 5 introduces the multi-period binomial model for stock prices and computes a dynamic hedging strategy that replicates a given option. We observe how the risk-neutral probabilities arise within the multi-period framework and how the option price can be expressed as a risk-neutral expectation. The calculations are implemented in a spreadsheet.

Chapter 6 takes the binomial modelling one step further by introducing more/shorter time periods. To achieve consistency across models one must make sure that the mean and variance of annual returns match the empirical data, which brings up the basic properties of mean and variance. At this stage it may be desirable to revise the elementary concepts in Appendix B on probability. Once the model is calibrated we realize that with many periods it is extremely time-consuming to implement it in a spreadsheet. This difficulty is overcome by a simple GAUSS program where we can use some of the matrix algebra of Chapters 1 and 2. Once the model is up and running it is natural to explore the continuous-time limit; on a computer one can consider hedging as frequently as every 10 min.

The discrete-time numerical explorations are a natural springboard to more theoretical calculations on the *continuous-time complete market trail* (Chapters 6, 10, 11). The numerical simulations show that the option price settles down as the rehedging intervals shorten; the real challenge is to work out the limit with pen and paper. This brings up the notions of the central limit theorem and continuous random variables, in particular the normal distribution. The optional (hard) calculations needed to work out the risk-neutral mean and variance of log returns are in Section 6.2.4; it is a good exercise in Taylor expansions and limits. The Black–Scholes integral (Section 6.2.5) is easier and likely to be compulsory in most finance courses. Chapter 6 demonstrates an important point: there are computations one can do with pen and paper that even the fastest computers cannot perform. Here, our productivity tool is standard calculus.

The second half of Chapter 6 deals with the Poisson jump limit of the binomial model. Some courses may wish to discuss the jumps there and then to show that Brownian motion is *not the only* continuous-time limit logically possible. An alternative is to leave jumps as an optional reading and stay on the Brownian motion path moving straight to Chapter 10, where we introduce continuous-time Brownian motion, Itô processes, and most importantly Itô calculus. Itô calculus is another great productivity tool, and it receives plenty of attention in Chapters 10 and 11. In my experience it is hard to *understand* the Itô calculus, but it is possible to *get used to* it and to apply it quickly and consistently; the main focus is therefore on practice. There is a large number of worked examples in Chapter 10, and the end-of-chapter exercises offer yet more opportunities to practise.

With Itô calculus under the belt, Section 11.2 explains the martingale approach to pricing; it represents the condensed wisdom of continuous-time asset pricing. Section 11.2 draws heavily on the martingale properties discussed in Chapter 9; these can be taken for granted if time is at a premium. For a good understanding one will also need the notion of state variable, Markov process and information filtration, which can be found in Chapter 8.

Section 11.3 discusses the Girsanov Theorem (required in Section 11.2) and its use in investment evaluation. Section 11.4 extends 11.2 to several risky assets. Sections 11.3 and 11.4 are more advanced and can be skipped on a first read. Section 11.5 talks about the relationship between martingales and partial differential equations, which is central to most finance applications. Section 11.6 surveys numerical methods used in continuous-time pricing.

The above trails on discrete and continuous-time complete markets are suitable for a core Masters course and can be covered in approximately 40 hours of lectures and 20 hours of tutorials.

Complete market pricing is remarkable by the conspicuous absence of risk, which is mathematically convenient but clearly at odds with reality. Risk is omnipresent in financial markets, as documented by the fate of Long Term Capital Management. Where there is risk one must, first of all, be able to measure it and only then one can come up with a price. Hence the other major theme in this book is *risk measurement and asset pricing in incomplete markets* (Chapters 3 and 4, and the first half of Chapter 12).

Chapter 3 starts by explaining how risky investment opportunities are ranked by the expected utility paradigm. Expected utility is often criticized for being ad hoc, for using meaningless units, for its results being dependent on initial wealth, etc., in short, for being worlds apart from mean–variance analysis. Chapter 3 dispels this *dangerous myth*. When correct measurement units are used all utility functions look exactly the same for small risks, and their investment advice is consistent with mean–variance analysis. When the risks are large and/or asymmetric the mean–variance analysis may lead to investment decisions that are logically inconsistent, whereas increasing utility functions will give consistent advice, albeit one that depends on the investor's attitude to large risks. Formally, this is shown by examining the scaling properties of the HARA class of utility functions. We will see that the risk–return trade-off of utility functions can be measured in terms of generalized Sharpe ratios similar to the standard Sharpe ratio of mean-to-standard deviation.

Naturally, one wishes to achieve the best risk–return trade-off, which leads to the maximization of expected utility. Chapter 4 discusses the numerical techniques that are needed for this task because, sadly, closed-form formulae are not available in incomplete markets. On the other hand the algorithms are quite simple and intuitive.

The use of numerical techniques in Chapter 4 is not an attempt to be innovative at all costs, rather, this chapter follows a trend that is increasingly apparent in financial economics as it relies more and more on numerical analysis to provide answers to pressing practical problems that are beyond the reach of closed-form solutions. As these developments take root financial economics will soon need a large number of professionals who are confident and competent users of numerical techniques. Chapter 4 is an accessible introduction to the economic and mathematical issues of numerical optimization that will prepare the reader for the road ahead.

Chapters 3 and 4 are set in a one-period environment. Chapter 12 transports the reader into a multi-period model where option hedging is risky. In Section 12.1 we describe the optimal hedging strategy and the minimum hedging error, and compute these quantities in a spreadsheet. Section 12.2 discusses the option pricing business in incomplete markets. Section 12.3 then talks about the continuous-time limit, where we will see that continuous hedging is *not* riskless, after all.

Chapter 3, with small digressions to Chapter 4, and the first two sections of Chapter 12 will need at least 15 hours of lectures and tutorials. Ideally, students should be given plenty of space to experiment with the programs and to feed the programs with real market data. This material is suitable for an elective Masters course.

Shorter PhD Courses

The book offers opportunities for short courses targeted at doctoral students.

In the absence of introductory textbooks on dynamic programming one can use Chapter 12, and particularly Section 12.4 to helicopter students into the issues of dynamic programming, its advantages, challenges, principles, and the mathematical language. Chapter 12 is the simplest multi-period optimization problem one will ever encounter (quadratic target function, linear controls) and therefore it is an ideal pedagogical tool. It is the only set-up that does not require iterative numerical optimization. Dynamic programming highlights the importance of the information set, Markov property and state variables covered in Chapter 8.

To complement the dynamic programming one may wish to introduce the martingale duality approach that appears in Section 9.4. This naturally leads to the connection between pricing kernels and the best investment opportunities (Hansen–Jagannathan duality) in Section 9.4.6 and via the extension theorem leads to the equilibrium price kernel restrictions used in the diagnostics of asset-pricing models (Cochrane 2001).

Chapter 7 gives an introduction to the fast Fourier transform (FFT) in finance, and it will appeal mainly to students specializing in derivative pricing. Chapter 7 offers the best of discrete and continuous-time worlds, fast pricing in combination with rich structure (affine models). Motivation for the FFT can be given quickly by referring to the numerical examples in Chapter 6. Complex numbers are introduced with minimum fuss by appealing to their geometric properties. The FFT naturally leads to the continuous-time limit, continuous Fourier transforms and characteristic functions, and it opens a new world of opportunities for numerical and theoretical explorations. The practical usefulness of the FFT can be seen, for example, in Section 12.3.2.

Exercises

The book is about empowering students and helping them to become confident users of the techniques they have seen in the lectures. For this purpose each chapter is accompanied by a tutorial that gives students an opportunity to practise the material just covered. Exercises are an integral part of the book, and solutions are freely available on this book's website (see p. xiii). If the reader can solve the exercises, then he or she can be pretty sure to have understood the theoretical concepts, and vice versa.

Related Reading

Hull (1997) is a classical all-round finance text with accessible mathematics, plenty of institutional details and many different types of financial instruments. There are several *intermediate* texts that concentrate more on the valuation methodology and less on the market practicalities, namely Baxter and Rennie (1996), Neftci (1996), Pliska (1997) and Luenberger (1998); *Mathematical Techniques* belongs to this category. Wilmott (1998) gives a practitioner's prospective of financial engineering mathematics, biased towards partial differential equations, but with plenty of numerical examples and many important topics. Duffie (1996) and Hunt and Kennedy (2000) represent *advanced* textbooks that start almost directly with continuous-time stochastic processes and martingale pricing. Further to this general list of textbooks each chapter provides references to sources and suggested reading.

Acknowledgments

My fascination with applied mathematics goes back a long way and I would like to take this opportunity to thank those who made my journey through the mathematical world both rewarding and enjoyable. First and foremost are my loving parents, who were my first maths and physics teachers. Next comes Bohuš Sivák and the staff at the Pedagogical University in Banská Bystrica who ran a phenomenally entertaining series of summer camps for young mathematicians. Maja Sadloňová and the staff at GVOZA (Maths High School in Žilina) have provided refuge for many generations of budding mathematicians. The introduction to complex numbers in Chapter 7 goes back to Žilina days and Ivo Čáp's afternoon lectures in physics. I would like to thank the staff at the Mathematics Department of the Faculty of Nuclear Sciences and Physical Engineering for the support, enthusiasm and inspiration they provided me with during my undergraduate studies.

My interest in financial economics took longer to develop. I am grateful to the staff at CERGE in Prague for providing an opportunity to study solid theoretical

economics. Aydin Hayri is in many ways responsible for the rest of the story; he helped me to enroll as a PhD student at Warwick and, some time later, sent me to a talk by Darrell Duffie. This was a talk that I did not understand, but it nevertheless changed the course of my professional life. It contained heaps of fascinating mathematics, mentioning *martingales* in every other sentence and, best of all, seemed to say something about the real world around us. I have been exploring the fascinating world of continuous-time finance ever since and this book is an opportunity to share parts of that journey with the reader.

In economics and finance no amount of mathematical sophistication can compensate for the ability to identify important problems and then, in spite of complicated analysis needed to generate the solution, communicate the results in simple terms. It has been my privilege to learn the craft from experienced financial economists, Stewart Hodges and David Miles. Stewart is one of the fathers of the incomplete market pricing paradigm and the exposition in Chapter 3 owes much to his pioneering work on the generalized Sharpe ratio. David, my colleague and co-author at Imperial, has taught me the value of numerical analysis in solving real world economic questions. He is the inspiration behind much of the numerical exploration in the book.

This book reflects numerous research discussions I have had with professional colleagues and co-authors. My thanks go to Karim Abadir, David Barr, Simon Benninga, Antonio Bernardo, Tomas Björk, Francis Breedon, Ian Buckley, Lara Cathcart, Martin Cripps, John Cochrane, Keith Cuthbertson, Bernard Dumas, Lina El-Jahel, Bob Flood, Stefano Galluccio, Hélyette Geman, Aydin Hayri, Berthold Herrendorf, Stewart Hodges, Lane Hughston, Tomáš Jandík, Stefan Jaschke, Jan Kallsen, Turalay Kenc, Jo Kennedy, Dima Kramkov, Harold Kushner, David Miles, Marcus Miller, Dirk Nitzsche, Huyên Pham, Nick Schmidt, James Sefton, Euan Sinclair, Jeremy Smith, Andrea Roncoroni, Walter Schachermayer, Martin Schweizer and Martin Weale.

This textbook has taken more than three years to write, and it would have been impossible to complete without the support of people around me. I would like to thank my wife, Barbora, for her care, support and understanding during this campaign. My thanks go to the Business School at Imperial College London, and particularly to Keith Cuthbertson for organizing my teaching and administrative workload efficiently, which freed up the time for the book. I am grateful to Ian Buckley, who provided me with his own lecture notes when I started at Imperial five years ago and who has since contributed many excellent ideas, in particular his trademark multiple choice questions. David Miles has read large parts of the typescript and suggested important clarifications. I have received valuable feedback from four anonymous referees. The book is a testament to the professionalism and perseverance of the PUP team. I am indebted to Jonathan Wainwright for an extremely meticulous editing and typesetting of the manuscript and for the many improvements he has suggested. A big thanks goes to Richard Baggaley, who initiated this project in November 1999 and guided it through the perilous waters of peer reviews, editorial boards and missed deadlines right to the end.

Last, but not least, I have received valuable feedback from five generations of students at Imperial, and I hope that we will join forces in the years to come.