

Preface to the Second Edition

The second edition of *Mathematical Techniques in Finance* appears at a very turbulent time in the global financial markets. The collapse of the U.S. subprime mortgage market and the ensuing payouts on insurance contracts known as credit default swaps have caused a massive tightening of credit supply around the world and left many apparently healthy financial institutions reeling, some taken over by their rivals and still others nationalized by their respective governments.

Against this background the subtitle of this textbook, *Tools for Incomplete Markets*, seems ever more timely. It reminds us that no amount of financial engineering can protect investors from all financial risk. It urges us to acknowledge this risk and to model it realistically, rather than assuming it away as a mathematical inconvenience.

There are three substantial changes in the second edition. First, the computing environment supported by the textbook has changed from GAUSS to MATLAB. Second, I have incorporated a new chapter on finite-difference methods, developed for students at Cass Business School, City University London. This becomes Chapter 12 and the material on mean-variance hedging for incomplete markets finds a new home in Chapter 13. Finally, I have made available a set of classroom-tested slides, divided into two-hour blocks. These are aimed primarily at instructors, but can also be used as an aid to self-study. All supporting materials for the second edition can be found at

<http://press.princeton.edu/titles/9079.html>

Bibliographic references have been updated throughout, with particular emphasis on Chapters 3 and 13, which map recent research in the area of incomplete markets. I am grateful to my coauthors Chris Brooks, Jan Kallsen, Fabio Maccheroni, Massimo Marinacci, Joelle Miffre and Aldo Rustichini for their substantial contributions to my understanding of the finance of incomplete markets. I would like to thank Sam Clark of T&T Productions Ltd for his editorial guidance and careful implementation of the many changes in this edition. My biggest thanks go to Richard Baggeley for a decade of continuing support and encouragement.

Study Guide

Before you start reading the book take a look at the book's website (URL above) to find out what resources are available.

Not all the material in this book is suitable for all students. There are essentially two long coherent themes appropriate for Master's 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 Master's 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 MATLAB 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 MATLAB 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 minutes.

The discrete-time numerical explorations are a natural springboard to more theoretical calculations on the continuous-time complete market trail (Chapters 6, 10 and 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 Section 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. Chapter 12 is devoted to numerical solutions of PDEs via finite-difference methods. The above trails on discrete and continuous-time complete markets are suitable for a core Master's 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 13).

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 advice 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 13 transports the reader into a multi-period model where option hedging is risky. In Section 13.1 we describe the optimal hedging strategy and the minimum hedging error, and compute these quantities in a spreadsheet. Section 13.2 discusses the option pricing business in incomplete markets. Section 13.3 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 13 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 Master’s 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 13, and particularly Section 13.4, to helicopter students into the issues of dynamic programming, its advantages, challenges, principles, and the mathematical language. Chapter 13 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 13.3.2.

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 the 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 (2005) 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), Joshi (2003), Neftci (1996), Pliska (1997), Shreve (2004a) and Luenberger (1998); *Mathematical Techniques in Finance* belongs to this category.

Wilmott (1998) gives a practitioner's perspective on financial engineering mathematics, biased towards partial differential equations, but with plenty of numerical examples and many important topics. Björk (1998), Duffie (1996), Hunt and Kennedy (2000) and Shreve (2004b) 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.