Preface

Why We Wrote the Book

The subject of financial markets is fascinating to many people: to those who care about money and investments, to those who care about the well-being of modern society, to those who like gambling, to those who like applications of mathematics, and so on. We, the authors of this book, care about many of these things (no, not the gambling), but what we care about most is teaching. The main reason for writing this book has been our belief that we can successfully teach the fundamentals of the economic and mathematical aspects of financial markets to almost everyone (again, we are not sure about gamblers). Why are we in this teaching business instead of following the path of many of our former students, the path of making money by pursuing a career in the financial industry? Well, they don't have the pleasure of writing a book for the enthusiastic reader like yourself!

Prerequisites

This text is written in such a way that it can be used at different levels and for different groups of undergraduate and graduate students. After the first, introductory chapter, each chapter starts with sections on the single-period model, goes to multiperiod models, and finishes with continuous-time models. The single-period and multiperiod models require only basic calculus and an elementary introductory probability/statistics course. Those sections can be taught to third- and fourth-year undergraduate students in economics, business, and similar fields. They could be taught to mathematics and engineering students at an even earlier stage. In order to be able to read continuous-time sections, it is helpful to have been exposed to an advanced undergraduate course in probability. Some material needed from such a probability course is briefly reviewed in chapter 16.

Who Is It For?

The book can also serve as an introductory text for graduate students in finance, financial economics, financial engineering, and mathematical finance. Some material from continuous-time sections is, indeed, usually considered to be graduate material. We try to explain much of that material in an intuitive way, while providing some of the proofs in appendixes to the chapters. The book is not meant to compete with numerous excellent graduate-level books in financial mathematics and financial economics, which are typically written in a mathematically more formal way, using a theorem-proof type of structure. Some of those more advanced books are mentioned in the references, and they present a natural next step in getting to know the subject on a more theoretical and advanced level.

Structure of the Book

We have divided the book into three parts. Part I goes over the basic securities, organization of financial markets, the concept of interest rates, the main mathematical models, and ways to measure in a quantitative way the risk and the reward of trading in the market. Part II deals with option pricing and hedging, and similar material is present in virtually every recent book on financial markets. We choose to emphasize the so-called martingale, probabilistic approach consistently throughout the book, as opposed to the differential-equations approach or other existing approaches. For example, the one proof of the Black-Scholes formula that we provide is done calculating the corresponding expected value. Part III is devoted to one of the favorite subjects of financial economics, the equilibrium approach to asset pricing. This part is often omitted from books in the field of financial mathematics, having fewer direct applications to option pricing and hedging. However, it is this theory that gives a qualitative insight into the behavior of market participants and how the prices are formed in the market.

What Can a Course Cover?

We have used parts of the material from the book for teaching various courses at the University of Southern California: undergraduate courses in economics and business, a masters-level course in mathematical finance, and option and investment courses for MBA students. For example, an undergraduate course for economics/business students that emphasizes option pricing could cover the following (in this order):

• The first three chapters without continuous-time sections; chapter 10 on bond hedging could also be done immediately after chapter 2 on interest rates

• The first two chapters of part II on no-arbitrage pricing and option pricing, without most of the continuous-time sections, but including basic Black-Scholes theory

• Chapters on hedging in part II, with or without continuous-time sections

• The mean-variance section in chapter 5 on risk; chapter 13 on CAPM could also be done immediately after that section

If time remains, or if this is an undergraduate economics course that emphasizes equilibrium/asset pricing as opposed to option pricing, or if this is a two-semester course, one could also cover

• discrete-time sections in chapter 4 on utility.

• discrete-time sections in part III on equilibrium models.

Courses aimed at more mathematically oriented students could go very quickly through the discrete-time sections, and instead spend more time on continuous-time sections. A one-semester course would likely have to make a choice: to focus on no-arbitrage option pricing methods in part II or to focus on equilibrium models in part III.

Web Page for This Book, Excel Files

The web page **http://math.usc.edu/∼cvitanic/book.html** will be regularly updated with material related to the book, such as corrections of typos. It also contains Microsoft Excel files, with names like ch1.xls. That particular file has all the figures from chapter 1, along with all the computations needed to produce them. We use Excel because we want the reader to be able to reproduce and modify all the figures in the book. A slight disadvantage of this choice is that our figures sometimes look less professional than if they had been done by a specialized drawing software. We use only basic features of Excel, except for Monte Carlo simulation for which we use the Visual Basic programming language, incorporated in Excel. The readers are expected to learn the basic features of Excel on their own, if they are not already familiar with them. At a few places in the book we give "Excel Tips" that point out the trickier commands that have been used for creating a figure. Other, more mathematically oriented software may be more efficient for longer computations such as Monte Carlo, and we leave the choice of the software to be used with some of the homework problems to the instructor or the reader. In particular, we do not use any optimization software or differential equations software, even though the instructor could think of projects using those.

Notation

Asterisk Sections and problems designated by an asterisk are more sophisticated in mathematical terms, require extensive use of computer software, or are otherwise somewhat unusual and outside of the main thread of the book. These sections and problems could be skipped, although we suggest that students do most of the problems that require use of computers.

Dagger End-of-chapter problems that are solved in the student's manual are preceded by a dagger.

Greek Letters We use many letters from the Greek alphabet, sometimes both lowercase and uppercase, and we list them here with appropriate pronunciation: α (alpha), β (beta), γ, Γ (gamma), δ, Δ (delta), ε (epsilon), ζ (zeta), η (eta), θ (theta), λ (lambda), μ (mu), ξ (xi), π, Π (pi), ω, Ω (omega), ρ (rho), σ, Σ (sigma), τ (tau), φ, Φ (phi).

Acknowledgments

First and foremost, we are immensely grateful to our families for the support they provided us while working on the book. We have received great help and support from the staff of our publisher, MIT Press, and, in particular, we have enjoyed working with Elizabeth Murry, who helped us go through the writing and production process in a smooth and efficient manner. J. C.'s research and the writing of this book have been partially supported by National Science Foundation grant DMS-00-99549. Some of the continuous-time sections in parts I and II originated from the lecture notes prepared in summer 2000 while J. C. was visiting the University of the Witwatersrand in Johannesburg, and he is very thankful to his host, David Rod Taylor, the director of the Mathematical Finance Programme at Wits. Numerous colleagues have made useful comments and suggestions, including Krzysztof Burdzy, Paul Dufresne, Neil Gretzky, Assad Jalali, Dmitry Kramkov, Ali Lazrak, Lionel Martellini, Adam Ostaszewski, Kaushik Ronnie Sircar, Costis Skiadas, Halil Mete Soner, Adam Speight, David Rod Taylor, and Mihail Zervos. In particular, D. Kramkov provided us with proofs in the appendix of chapter 6. Some material on continuous-time utility maximization with incomplete information is taken from a joint work with A. Lazrak and L. Martellini, and on continuous-time mean-variance optimization from a joint work with A. Lazrak. Moreover, the following students provided their comments and pointed out errors in the working manuscript: Paula Guedes, Frank Denis Hiebsch, and Chulhee Lee. Of course, we are solely responsible for any remaining errors.

A Prevailing Theme: Pricing by Expected Values

Before we start with the book's material, we would like to give a quick illustration here in the preface of a connection between a price of a security and the optimal trading strategy of an investor investing in that security. We present it in a simple model, but this connection is present in most market models, and, in fact, the resulting pricing formula is of the form that will follow us through all three parts of this book. We will repeat this type of argument later in more detail, and we present it early here only to give the reader a general taste of what the book is about. The reader may want to skip the following derivation, and go directly to equation (0.3).

Consider a security S with today's price $S(0)$, and at a future time 1 its price $S(1)$ either has value s^u with probability p, or value s^d with probability $1 - p$. There is also a risk-free security that returns $1 + r$ dollars at time 1 for every dollar invested today. We assume that $s^d < (1 + r)S(0) < s^u$. Suppose an investor has initial capital x, and has to decide how many shares δ of security S to hold, while depositing the rest of his wealth in the bank

account with interest rate r. In other words, his wealth $X(1)$ at time 1 is

$$X(1) = \delta S(1) + [x - \delta S(0)](1 + r)$$

The investor wants to maximize his expected utility

$$E[U(X(1))] = pU(X^u) + (1 - p)U(X^d)$$

where U is a so-called utility function, while X^u, X^d is his final wealth in the case $S(1) = s^u$, $S(1) = s^d$, respectively. Substituting for these values, taking the derivative with respect to δ and setting it equal to zero, we get

$$pU'(X^u)[s^u - S(0)(1 + r)] + (1 - p)U'(X^d)[s^d - S(0)(1 + r)] = 0$$

The left-hand side can be written as $E[U'(X(1))\{S(1) - S(0)(1 + r)\}]$, which, when made equal to zero, implies, with arbitrary wealth X replaced by optimal wealth $\hat{X}$,

$$S(0) = E\left[\frac{U'(\hat{X}(1))}{E(U'[\hat{X}(1)])} \frac{S(1)}{1 + r}\right] \tag{0.1}$$

If we denote

$$Z(1) := \frac{U'(\hat{X}(1))}{E\{U'(\hat{X}(1))\}} \tag{0.2}$$

we see that the today's price of our security S is given by

$$S(0) = E\left[Z(1)\frac{S(1)}{1 + r}\right] \tag{0.3}$$

We will see that prices of most securities (with some exceptions, like American options) in the models of this book are of this form: the today's price $S(0)$ is an expected value of the future price $S(1)$, multiplied ("discounted") by a certain random factor. Effectively, we get the today's price as a weighted average of the discounted future price, but with weights that depend on the outcomes of the random variable $Z(1)$. Moreover, in standard option-pricing models (having a so-called completeness property) we will not need to use utility functions, since $Z(1)$ will be independent of the investor's utility. The random variable $Z(1)$ is sometimes called **change of measure,** while the ratio $Z(1)/(1 + r)$ is called **state-price density, stochastic discount factor, pricing kernel,** or **marginal rate of substitution,** depending on the context and interpretation. There is another interpretation of this formula, using a new probability; hence the name "change of (probability) measure." For example, if, as in our preceding example, $Z(1)$ takes two possible values $Z^u(1)$ and $Z^d(1)$ with

probabilities p, $1 - p$, respectively, we can define

$$p^* := p Z^u(1), \qquad 1 - p^* = (1 - p) Z^d(1)$$

The values of $Z(1)$ are such that p^* is a probability, and we interpret p^* and $1 - p^*$ as modified probabilities of the movements of asset S. Then, we can write equation (0.3) as

$$S(0) = E^* \left[\frac{S(1)}{1 + r} \right] \tag{0.4}$$

where E^* denotes the expectation under the new probabilities, p^*, $1 - p^*$. Thus the price today is the expected value of the discounted future value, where the expected value is computed under a special, so-called **risk-neutral** probability, usually different from the **real-world probability.**

Final Word

We hope that we have aroused your interest about the subject of this book. If you turn out to be a very careful reader, we would be thankful if you could inform us of any remaining typos and errors that you find by sending an e-mail to our current e-mail addresses. Enjoy the book!

Jakša Cvitanić and Fernando Zapatero

E-mail addresses: cvitanic@math.usc.edu, zapatero@usc.edu