Introduction

This book is an introduction to probabilistic methods in Finance. It is intended for graduate students in mathematics, and it may also be useful for mathematicians in academia and in the financial industry. Our focus is on stochastic models in discrete time. This limitation has two immediate benefits. First, the probabilistic machinery is simpler, and we can discuss right away some of the key problems in the theory of pricing and hedging of financial derivatives. Second, the paradigm of a complete financial market, where all derivatives admit a perfect hedge, becomes the exception rather than the rule. Thus, the discrete-time setting provides a shortcut to some of the more recent literature on incomplete financial market models.

As a textbook for mathematicians, it is an introduction at an intermediate level, with special emphasis on martingale methods. Since it does not use the continuous-time methods of Itô calculus, it needs less preparation than more advanced texts such as [63], [64], [72], [109], [161]. On the other hand, it is technically more demanding than textbooks such as [138]: We work on general probability spaces, and so the text captures the interplay between probability theory and functional analysis which has been crucial for some of the recent advances in mathematical finance.

The book is based on our notes for first courses in Mathematical Finance which both of us are teaching in Berlin at Humboldt University and at Technical University. These courses are designed for students in mathematics with some background in probability. Sometimes, they are given in parallel to a systematic course on stochastic processes. At other times, martingale methods in discrete time are developed in the course, as they are in this book. Usually the course is followed by a second course on Mathematical Finance in continuous time. There it turns out to be useful that students are already familiar with some of the key ideas of Mathematical Finance.

The core of this book is the dynamic arbitrage theory in the first chapters of Part II. When teaching a course, we found it useful to explain some of the main arguments in the more transparent one-period model before using them in the dynamical setting. So one approach would be to start immediately in the multi-period framework of Chapter 5, and to go back to selected sections of Part I as the need arises. As an alternative, one could first focus on the one-period model, and then move on to Part II.

We include in Chapter 2 a brief introduction to the mathematical theory of expected utility, even though this is a classical topic, and there is no shortage of excellent expositions; see, for instance, [117] which happens to be our favorite. We have three reasons for including this chapter. Our focus in this book is on incompleteness, and incompleteness involves, in one form or another, preferences in the face of risk and uncertainty. We feel that mathematicians working in this area should be aware, at least to some extent, of the long line of thought which leads from Daniel Bernoulli via von Neumann–Morgenstern and Savage to some more recent developments which are motivated by shortcomings of the classical paradigm. This is our first reason. Second, the analysis of risk measures has emerged as a major topic in mathematical finance, and this is closely related to a robust version of the Savage theory. Third, but not least, our experience is that this part of the course was found particularly enjoyable, both by the students and by ourselves.

We acknowledge our debt and express our thanks to all colleagues who have contributed, directly or indirectly, through their publications and through informal discussions, to our understanding of the topics discussed in this book. Ideas and methods developed by Freddy Delbaen, Darrell Duffie, Nicole El Karoui, David Heath, Yuri Kabanov, Ioannis Karatzas, Dimitri Kramkov, David Kreps, Stanley Pliska, Chris Rogers, Steve Ross, Walter Schachermayer, Martin Schweizer, Dieter Sondermann and Christophe Stricker play a key role in our exposition. We are obliged to many others; for instance the textbooks [47], [63], [64], [97], and [121] were a great help when we started to teach courses on the subject.

We are grateful to all those who read parts of the manuscript and made useful suggestions, in particular to Dirk Becherer, Ulrich Horst, Steffen Krüger, Irina Penner, and to Alexander Giese who designed some of the figures. Special thanks are due to Peter Bank for a large number of constructive comments. We also express our thanks to Erhan Çinlar, Adam Monahan, and Philip Protter for improving some of the language, and to the Department of Operations Research and Financial Engineering at Princeton University for its hospitality during the weeks when we finished the manuscript.

Berlin, June 2002

Hans Föllmer
Alexander Schied