Preface

"We don't do it for the glory. We don't do it for the recognition...
We do it because it needs to be done. Because if we don't, no one else will.
And we do it even if no one knows what we've done.
Even if no one knows we exist. Even if no one remembers we ever existed."

Kara, "Christmas with the Superheroes" 2,
"Should Auld Acquaintance Be Forgot", 1989, DC Comics

MOTIVATION

The idea of writing a book on interest-rate modeling crossed our minds a couple of years ago, after a few months of work at Banca IMI, where, as quantitative analysts, we have been supporting traders, dealing with mathematical modeling for financial engineering. We had reached this position after a look at the academic environments and after a subsequent first period as quantitative analysts in a different Italian bank. In Banca IMI we were given the task of studying and developing financial models for the pricing and hedging of a broad range of derivatives, and we have been involved in medium/long-term projects. The farsighted policy of Aleardo Adotti, Head of the Product and Business Development Group, allowed us to work in a serene and stimulating environment, also benefiting from the continuous and valuable feedback of several of our bank's traders.

We began writing technical reports on what we were studying and implementing, and soon we realized that our material was covering some gaps in the existing financial literature. We looked at the material collected and wondered whether it be possible and worthwhile to reorganize it in a coherent structure with the ultimate aim of achieving some circulation. "Why not?", we answered ourselves, and our book-writing project started to be outlined.

We had to tread carefully in this new "book-writing territory", since our writing experience had been limited in the main to academic papers for journals and our Ph.D. theses.

We also asked for the advice of our Head, who immediately gave us his enthusiastic approval and granted us enough time and freedom to devote to our writing.

A year and a half later we had reached this final version and eventually decided to stop thinking of further modifications. Hopefully, we will have the chance of implementing any suggestion for improvements in later editions of the book.

In conclusion of this short initial chronicle, let us say a final word about the language used. We would venture to say it looks like American English, if we could be sure that no one would turn his/her nose up. Indeed, our text editor featured an American spell-checking program.

We are quite aware that many sentences could have been written in a more stylish and classical English, yet we resort to the excuse that our English reflects our Latin background, and at times the temptation to transliterate from Italian has simply been irresistible. So we hope you will find our *Italian English* intelligible if amusing!

Let us now describe our motivations and the content of the book in a more detailed manner, starting by answering the following fundamental, inevitable and obligatory question.

Why a New Book?

In years where every month a new book on financial modeling or on mathematical finance comes out, one of the first questions inevitably is: why one more, and why one on interest-rate modeling in particular?

The answer springs directly from our job experience as quantitative analysts in financial institutions. Indeed, one of the major challenges any financial engineer has to cope with is the practical implementation of mathematical models for pricing derivative securities.

When pricing market financial products, one has to address a number of theoretical and practical issues that are often neglected in the classical, general basic theory: the choice of a satisfactory model, the derivation of specific analytical formulas and approximations, the calibration of the selected model to a set of market data, the implementation of efficient routines for speeding up the whole calibration procedure, and so on. In other words, the general understanding of the theoretical paradigms in which specific models operate does not lead to their complete understanding and immediate implementation and use for concrete pricing. This is an area that is rarely covered by books on mathematical finance.

Undoubtedly, there exist excellent books covering the basic theoretical paradigms, but they do not provide enough instructions and insights for tackling concrete pricing problems. We have therefore thought of writing this book in order to cover this gap between theory and practice.

The Gap between Theory and Practice

A gap, indeed. And a fundamental one. The interplay between theory and practice has proved to be an extremely fecund ingredient in the progress of science and modeling in particular. We believe that practice can help to appreciate theory, thus generating a feedback that is one of the most important and intriguing aspects of modeling and more generally of scientific investigation.

If theory becomes deaf to the feedback of practice or vice versa, great opportunities can be missed. It may be a pity to restrict one's interest only to extremely abstract problems that have little relevance for the scientists or quantitative analysts working in "real life".

Now, it is obvious that everyone working in the field owes a lot to the basic fundamental theory from which such extremely abstract problems stem. It would be foolish to deny the importance of a well developed and consistent theory as a fundamental support for any practice involving mathematical models. Indeed, practice that is deaf to theory or that employs a sloppy mathematical apparatus is quite dangerous.

However, besides the extremely abstract refinement of the basic paradigms, which are certainly worth studying but that interest mostly an academic audience, there are other fundamental and more specific aspects of the theory that are often neglected in books and in the literature, and that interest a larger public.

Is This Book about Theory? What kind of Theory?

In the book, we are not dealing with the fundamental no-arbitrage paradigms with great detail. We resume and adopt the basic well-established theory of Harrison and Pliska, and avoid the debate on the several possible definitions of no-arbitrage and on their mutual relationships. Indeed, we will raise problems that can be faced in the basic framework above. Insisting on the subtle aspects and developments of no-arbitrage theory more than is necessary would take space from the other theory we need to address in the book and that is more important for our purposes.

Besides, there already exist several books dealing with the most abstract theory of no-arbitrage. On the theory that we deal with, on the contrary, there exist only few books. What is this theory? To give a flavor of it, let us select a few questions at random:

- How can the market interest-rate curves be defined in mathematical terms?
- What kind of interest rates does one select when writing the dynamics? Instantaneous spot rates? Forward rates? Forward swap rates? Are there models for each such basic variable?
- What is a sufficiently general framework for expressing no-arbitrage in interest-rate modeling?
- Is there a definition of cap or caplet volatility (and of their term structures) in terms of interest-rate dynamics that is consistent with market practice?
- What kinds of diffusion coefficients in the rate dynamics are compatible with different qualitative evolutions of the term structure of volatilities over time?

- How is “humped volatility shape” translated in mathematical terms and what kind of mathematical models allow for it?
- What is the most convenient probability measure under which one can price a specific product, and how can one derive concretely the related interest-rate dynamics?
- Are different market models of interest-rate dynamics mutually compatible?
- Is it possible to restore some compatibility between incompatible market models via some kind of approximation? How is this done precisely? What are the related formulas?
- What does it mean to calibrate a model to the market in terms of the chosen mathematical model? Is this always possible? Or is there a degree of approximation involved?
- Does terminal correlation among rates depend on instantaneous volatilities or only on instantaneous correlations? Can we analyze this dependence?
- Can terminal correlations among rates be computed with an approximation independent of the particular probability measure chosen? Under which kind of models is this possible? If so, what are the related formulas?
- What is the volatility smile, how can it be expressed in terms of mathematical models and of forward-rate dynamics in particular?
- What is the link between dynamics of rates and their distributions?
- What kind of model is more apt to model correlated interest-rate curves of different currencies, and how does one compute the related dynamics under the relevant probability measures?
- When does a model imply the Markov property for the short rate and why is this important?
- ......

We could go on for a while with questions of this kind. Our point is, however, that the theory dealt with in a book on interest-rate models should consider this kind of question.

We sympathize with anyone who has gone to a bookstore (or perhaps to a library) looking for answers to some of the above questions with little fortune. We have done the same, several times, and we were able to find only limited material and few reference works. We hope this book is a successful step forward towards addressing such questions.

We also sympathize with the reader who has just finished his studies or with the academic who is trying a life-change to work in industry or who is considering some close cooperation with market participants. Being used to precise statements and rigorous theory, this person might find answers to the above questions expressed in contradictory or unclear mathematical language. This is something else we too have been through, and we are trying not to be disappointing on this side either.

Is This Book about Practice? What kind of Practice?

We try and answer some questions on practice that are again overlooked in most of the existing books in mathematical finance, and on interest-rate models in particular. Again, here are some typical questions selected at random:

- What are accrual conventions and how do they impact on the definition of rates?
- Can you give a few examples of how time is measured in connection with some aspects of contracts? What are "day-count conventions"?
- What is the interpretation of most liquid market contracts such as caps and swaptions? What is their main purpose?
- What kind of data structures are observed in the market? Are all data equally significant?
- How is a specific model calibrated to market data in practice? Is a joint calibration to different market structures always possible or even desirable?
- What are the dangers of calibrating a model to data that are not equally important, or reliable, or updated with poor frequency?
- What are the requirements of a trader as far as a calibration results are concerned?
- How can one handle path-dependent or early-exercise products numerically? And products with both features simultaneously?
- What numerical methods can be used for implementing a model that is not analytically tractable? How are trees built for specific models? Can instantaneous correlation be a problem when building a tree in practice?
- What kind of products are suited to evaluation through Monte Carlo simulation? How can Monte Carlo simulation be applied in practice? Under which probability measure is it convenient to simulate?
- Is there a model flexible enough to be calibrated to the market smile for caps?
- What typical qualitative shapes of the volatility term structure are observed in the market?
- What is the impact of the parameters of a chosen model on the market volatility structures that are relevant to the trader?
- What is the accuracy of analytical approximations derived for swaptions volatilities and terminal correlations?
- Does there exist an interest-rate model that can be considered "central" nowadays, in practice? What do traders think about it?
- Can you present some concrete examples of calibration to market data?
- How can we express mathematically the payoffs of some typical market products?
- How do you handle in practice products depending on more than one interest-rate curve at the same time?
- ......

Again, we could go on for a while, and it is hard to find a single book answering these questions with a rigorous theoretical background. Also, answering some of these questions (and others that are similar in spirit) motivates new theoretical developments, maintaining the fundamental feedback between theory and practice we hinted at above.

AIMS, READERSHIP AND BOOK STRUCTURE

Contrary to the equity-derivatives area, interest-rate modeling is a branch of mathematical finance where no general model has been yet accepted as "standard" for the whole sector, although the LIBOR market model is emerging as a possible candidate for this role. Indeed, there exist market standard models for both main interest-rate derivatives "sub-markets", namely the caps and swaptions markets. However, such models are theoretically incompatible and cannot be used jointly to price other interest-rate derivatives.

Because of this lack of a standard, the choice of a model for pricing and hedging interest-rate derivatives has to be carefully dealt with. In this book, therefore, we do not just concentrate on a specific model leaving all implementation issues aside. We instead develop several types of models and show how to use them in practice for pricing a number of specific products.

The main models are illustrated in different aspects ranging from theoretical formulation to a possible implementation on a computer, always keeping in mind the concrete questions one has to cope with. We also stress that different models are suited to different situations and products, pointing out that there does not exist a single model that is uniformly better than all the others.

Thus our aim in writing this book is two-fold. First, we would like to help quantitative analysts and advanced traders handle interest-rate derivatives with a sound theoretical apparatus. We try explicitly to explain which models can be used in practice for some major concrete problems. Secondly, we would also like to help academics develop a feeling for the practical problems in the market that can be solved with the use of relatively advanced tools of mathematics and stochastic calculus in particular. Advanced undergraduate students, graduate students and researchers should benefit as well, from seeing how some sophisticated mathematics can be used in concrete financial problems.

The Prerequisites

The prerequisites are some basic knowledge of stochastic calculus and the theory of stochastic differential equations in particular. The main tools from stochastic calculus are Ito's formula and Girsanov's theorem, which are, however, briefly reviewed in an appendix.

The Book is Structured in Two Parts

The first part is more academic and develops the theoretical basis that is needed for tackling the concrete pricing problems we want to solve.

We start by reviewing some basic concepts and definitions and briefly explain the fundamental theory of no-arbitrage and its implications as far as pricing derivatives is concerned.

We then review some of the basic short-rate models, forward-rate models and market models describing their distributional properties, discussing their analytical tractability and proposing numerical procedures for approximating the interest-rate dynamics. We will make extensive use of the "change-of-numeraire" technique, which is explained in details in a initial section.

The second part in contrast is devoted to concrete applications. We in fact list a series of market financial products that are usually traded over the counter and for which there exists no uniquely consolidated pricing model. We consider some typical interest-rate derivatives dividing them into two classes: i) derivatives depending on a single interest-rate curve; ii) derivatives depending on two interest-rate curves. We also propose a section where we hint at pricing equity derivatives under stochastic interest rates.

Appendices

It is sometimes said that no one ever reads appendices. This book ends with four appendices, and the fourth one is an interview with a quantitative trader, which should be interesting enough to convince the reader to have a look at the appendices, for a change.

The first appendix briefly reviews some basic results from stochastic calculus that are mentioned and applied in the book. The second appendix reports a useful calculation, whereas the third one deals with a general approximation of a diffusion process with a tree.

FINAL WORD AND ACKNOWLEDGMENTS

Whether our treatment of the theory fulfills the targets we have set ourselves, is up to the reader to judge. A disclaimer is necessary though. Assembling a book in the middle of the "battlefield" that is any trading room, while quite stimulating, leaves however little space for planned organization. Indeed, the book is not homogeneous, some topics are more developed than others.

We have tried to follow a logical path in assembling the final manuscript, but we are aware that the book is not optimal in respect of homogeneity and linearity of exposition. Hopefully, the explicit contribution of our work will emerge over these inevitable little misalignments.

Acknowledgments

A book is always the product not only of its authors, but also of their colleagues, of the environment where the authors work, of the encouragements and critics gathered from conferences, referee reports for journal publications, conversations after seminars, and many analogous events. While we cannot do justice to all the above, we thank explicitly our colleagues Gianvittorio "Tree and Optimization Master" Mauri and Francesco "Monte Carlo" Rapisarda, for their help and continuous interaction concerning both modeling and concrete implementations on computers. Francesco also helped by proofreading the manuscript and by suggesting modifications.

We express a thought of gratitude for our Head Aleardo Adotti in Banca IMI, for his farsightedness in allowing us to write this book and more generally to work on the frontiers of mathematical finance inside a bank.

The feedback from the interest-rate-derivatives desk has been fundamental, in the figures of Luca Mengoni and later on also of Antonio Castagna, who have stimulated many developments with their objections, requirements and discussions. Their feeling for market behavior has guided us in cases where mere mathematics and textbook finance could not help us that much.

Two trainees, Giulio Sartorelli and Cristina Capitani, have also contributed to the manuscript: both helped in developing numerical tests and Cristina also proofread part of the book manuscript.

The staff at Springer-Verlag has been active and supportive in the figures of Catriona Byrne and Susanne Denskus, whose help in our first "non-thesis" book has been very valuable.

All mistakes are, needless to say, ours.

Last but not least, we are grateful to our families and friends. Also, Damiano is grateful to the "historical" comics/anime/role-playing/motorcycle friends in Venice, and Fabio to his girlfriend, for their supportive enthusiasm.

Finally, our ultimate gratitude is towards transcendence and is always impossible to express with words. We just say that we are grateful for the Word of the Gospel and the Silence of Zen.

A Special Final Word for Young Readers and Beginners

We close this long preface with a particular thought and encouragement for young readers. Clearly, if you are a professional or academic experienced in interest-rate modeling, we believe you will not be scared by a first quick look at the table of contents and at the chapters.

However, even at a first glance when flipping through the book, some young readers might feel discouraged by the variety of models, by the difference in approaches, and might indeed acquire the impression of a chaotic sequence of models that arose in mathematical finance without a particular order or purpose. Yet, we assure you that this subject is interesting, relevant, and that it can (and should) be fun, however "cliched" this may sound to you. We have tried at times to be colloquial in the book, in an attempt to avoid writing a book on formal mathematical finance from A to Zzzzzzzzzz... (where have you heard this one before?).

We are trying to avoid the two apparent extremes of either scaring or boring our readers. Thus you will find at times opinions from market participants, guided tours, intuition and discussion on things as they are seen in the market. We would like you to give it at least a try. So, if you are one of the above young readers, and be you a student or a practitioner, we suggest you take it easy. This book might be able to help you a little in entering this exciting field of research. This is why we close this preface with the by-now classic recommendations...

"May fear and dread not conquer me"
Majjhima Nikaya VIII.6

"Do not let your hearts be troubled and do not be afraid"
St. John XIV.27

Indeed, *Don't Panic!*

Venice and Milan, April 1, 2001

Damiano Brigo and Fabio Mercurio

DESCRIPTION OF CONTENTS BY CHAPTER

We herewith provide a detailed description of the contents of each chapter.

Part I: MODELS: THEORY AND IMPLEMENTATION

Chapter 1: Definitions and Notation. The chapter is devoted to standard definitions and concepts in the interest-rate world, mainly from a static point of view. We define several interest-rate curves, such as the LIBOR, swap, forward-LIBOR and forward-swap curves, and the zero-coupon curve.

We explain the different possible choices of rates in the market. Some fundamental products, whose evaluation depends only on the initially given curves and not on volatilities, such as bonds and interest-rate swaps, are introduced. A quick and informal account of fundamental derivatives depending on volatility such as caps and swaptions is also presented, mainly for motivating the following developments.

Chapter 2: No-Arbitrage Pricing and Numeraire Change. The chapter introduces the theoretical issues a model should deal with, namely the no-arbitrage condition and the change of numeraire technique. The change of numeraire is reviewed as a general and powerful theoretical tool that can be used in several situations, and, indeed, it will be often used in the book.

We remark how the standard Black models for either the cap or swaption markets, the two main markets of interest-rate derivatives, can be given a rigorous interpretation via suitable numeraires, as we will do later on in Chapter 6.

We finally hint at products involving more than one interest-rate curve at the same time, typically quanto-like products, and illustrate the no-arbitrage condition in this case.

Chapter 3: One-Factor Short-Rate Models. In this chapter, we begin to consider the dynamics of interest rates. The chapter is devoted to the short-rate world. In this context, one models the instantaneous spot interest rate via a possibly multi-dimensional driving diffusion process depending on some parameters. The whole yield-curve evolution is then characterized by the driving diffusion.

If the diffusion is one dimensional, with this approach one is directly modeling the short rate, and the model is said to be "one-factor". In this chapter, we focus on such models, leaving the development of the multi-dimensional (two-dimensional in particular) case to the next chapter.

As far as the dynamics of one-factor models is concerned, we observe the following. Since the short rate represents at each instant the initial point of the yield curve, one-factor short-rate models assume the evolution of the whole yield curve to be completely determined by the evolution of its initial point. This is clearly a dangerous assumption, especially when pricing products depending on the correlation between different rates of the yield curve at a certain time (this limitation is explicitly pointed out in the guided tour of the subsequent chapter).

We then illustrate the no-arbitrage condition for one-factor models and the fundamental notion of market price of risk connecting the objective world, where rates are observed, and the risk-neutral world, where expectations leading to prices occur. We also show how choosing particular forms for the market price of risk can lead to models to which one can apply both econometric techniques (in the objective world) and calibration to market prices (risk-neutral world). We briefly hint at this kind of approach and subsequently leave the econometric part, focusing on the market calibration.

A short-rate model is usually calibrated to some initial structures in the market, typically the initial yield curve, the caps volatility surface, the swaptions volatility surface, and possibly other products, thus determining the model parameters. We introduce the historical one-factor time-homogeneous models of Vasicek, Cox Ingersoll Ross (CIR), Dothan, and the Exponential Vasicek (EV) model. We hint at the fact that such models used to be calibrated only to the initial yield curve, without taking into account market volatility structures, and that the calibration can be very poor in many situations.

We then move to extensions of the above one-factor models to models including "time-varying coefficients", or described by inhomogeneous diffusions. In such a case, calibration to the initial yield curve can be made perfect, and the remaining model parameters can be used to calibrate the volatility structures. We examine classic one-factor extensions of this kind such as Hull and White's extended Vasicek (HW) model, classic extensions of the CIR model, Black and Karasinski's (BK) extended EV model and a few more.

We discuss the volatility structures that are relevant in the market and explain how they are related to short-rate models. We discuss the issue of a humped volatility structure for short-rate models and give the relevant definitions. We also present the Mercurio-Moraleda short-rate model, which allows for a parametric humped-volatility structure while exactly calibrating the initial yield curve, and briefly hint at the Moraleda-Vorst model.

We then present a method of ours for extending pre-existing time-homogeneous models to models that perfectly calibrate the initial yield curve while keeping free parameters for calibrating volatility structures. Our method preserves the possible analytical tractability of the basic model. Our extension is shown to be equivalent to HW for the Vasicek model, whereas it is original in case of the CIR model. We call CIR++ the CIR model being extended through our procedure. We also show how to extend the Dothan and EV models, as possible alternatives to the use of the popular BK model.

We explain how to price coupon-bearing bond options and swaptions with models that satisfy a specific tractability assumption, and give general comments and a few specific instructions on Monte Carlo pricing with short-rate models.

We finally analyze how the market volatility structures implied by some of the presented models change when varying the models parameters. We conclude with an example of calibration of different models to market data.

Chapter 4: Two-Factor Short-Rate Models. If the short rate is obtained as a function of all the driving diffusion components (typically a summation, leading to an additive multi-factor model), the model is said to be "multi-factor".

We start by explaining the importance of the multi-factor setting as far as more realistic correlation and volatility structures in the evolution of the interest-rate curve are concerned.

We then move to analyze two specific two-factor models.

First,we apply our above deterministic-shift method for extending pre-existing time-homogeneous models to the two-factor additive Gaussian case (G2). In doing so, we calibrate perfectly the initial yield curve while keeping five free parameters for calibrating volatility structures. As usual, our method preserves the analytical tractability of the basic model. Our extension G2++ is shown to be equivalent to the classic two-factor Hull and White model. We develop several formulas for the G2++ model and also explain how both a binomial and a trinomial tree for the two-dimensional dynamics can be obtained. We discuss the implications of the chosen dynamics as far as volatility and correlation structures are concerned, and finally present an example of calibration to market data.

The second two-factor model we consider is a deterministic-shift extension of the classic two-factor CIR (CIR2) model, which is essentially the same as extending the Longstaff and Schwartz (LS) models. Indeed, we show that CIR2 and LS are essentially the same model, as is well known. We call CIR2++ the CIR2/LS model being extended through our deterministic-shift procedure, and provide a few analytical formulas. We do not consider this model with the same level of detail devoted to the G2++ model, due to the fact that its volatility structures are less flexible than the G2++'s, at least in case one wishes to preserve analytical tractability.

Chapter 5: The Heath-Jarrow-Morton Framework. In this chapter we consider the Heath-Jarrow-Morton (HJM) framework. We introduce the general framework and point out how it can be considered the right theoretical framework for developing interest-rate theory and especially no-arbitrage. However, we also point out that the most significant models coming out concretely from such a framework are the same models we met in the short-rate approach.

We report conditions on volatilities leading to a Markovian process for the short rate. This is important for implementation of lattices, since one then obtains (linearly-growing) recombining trees, instead of exponentially-growing ones. We show that in the one-factor case, a general condition leading to Markovianity of the short rate yields the Hull-White model with all time-varying coefficients, thus confirming that, in practice, short-rate models already contained some of the most interesting and tractable cases.

We then introduce the Ritchken and Sankarasubramanian framework, which allows for Markovianity of an enlarged process, of which the short rate is a component. The related tree (Li, Ritchken and Sankarasubramanian) is presented. Finally, we present a different version of the Mercurio-Moraleda model obtained through a specification of the HJM volatility structure, pointing out its advantages for realistic volatility behavior and its analytical formula for bond options.

Chapter 6: The LIBOR and Swap Market Models (LFM and LSM). This chapter presents one of the most popular families of interest-rate models: the market models. A paramount fact is that the lognormal forward-LIBOR model (LFM) prices caps with Black's cap formula, which is the standard formula employed in the cap market. Moreover, the lognormal forward-swap model (LSM) prices swaptions with Black's swaption formula, which is the standard formula employed in the swaption market. Now, the cap and swaption markets are the two main markets in the interest-rate-derivatives world, so that compatibility with the related market formulas is a very desirable property. However, even with rigorous separate compatibility with the caps and swaptions classic formulas, the LFM and LSM are not compatible with each other. Still, the separate compatibility above is so important that these models, and especially the LFM, are nowadays seen as the most promising area in interest-rate modeling.

We start the chapter with a guided tour presenting intuitively the main issues concerning the LFM and the LSM, and giving motivation for the developments to come.

We then introduce the LFM, the "natural" model for caps, modeling forward-LIBOR rates. We give several possible instantaneous-volatility structures for this model, and derive its dynamics under different measures. We explain how the model can be calibrated to the cap market, examining the impact of the different structures of instantaneous volatility on the calibration. We introduce rigorously the term structure of volatility, and again check the impact of the different parameterizations of instantaneous volatilities on its evolution in time. We point out the difference between instantaneous and terminal correlation, the latter depending also on instantaneous volatilities.

We then introduce the LSM, the "natural" model for swaptions, modeling forward-swap rates. We show that the LSM is distributionally incompatible with the LFM. We discuss possible parametric forms for instantaneous correlations in the LFM, their impact on swaptions prices, and how, in general, Monte Carlo simulation should be used to price swaptions with the LFM instead of the LSM. We derive several approximated analytical formulas for swaption prices in the LFM (Brace's, Rebonato's and Hull-White's).

We point out that terminal correlation depends on the particular measure chosen for the joint dynamics in the LFM. We derive two analytical formulas based on “freezing the drift” for terminal correlation. These formulas clarify the relationship between instantaneous correlations and volatilities on one side and terminal correlations on the other side. We develop a similar formula for transforming volatility data of semi-annual or quarterly forward rates in volatility data of annual forward rates, and test it against Monte Carlo simulation of the true quantities. This is useful for joint calibration to caps and swaptions, allowing one to consider only annual data.

We present two methods for obtaining forward LIBOR rates in the LFM over non-standard periods, i.e. over expiry/maturity pairs that are not in the family of rates modeled in the chosen LFM.

We conclude the first chapter devoted to the market models with smile modeling. We introduce the smile problem with a guided tour. We provide a little history and a few references on simile modeling, and then present three models for the caplets smile.

The first model is a shifted lognormal dynamics. Having only two parameters for each fixed maturity, the curves along the strike dimension that can be reproduced by such a model are rather poor.

Then we introduce the constant-elasticity-of-variance extension of the LFM as from Andersen and Andreasen. This is an improvement, since this too can be shifted and now there are three parameters for fitting a caplet smile along the strike dimension for a given maturity. However, three parameters are still too few to give a flexible enough model.

We finally introduce the (possibly shifted) lognormal-mixture dynamics of Brigo and Mercurio for the LFM model. This model is always tractable, resulting in prices (and Greeks) that are linear combinations of Black’s prices (Greeks), and can include any number of parameters. Therefore, its fitting capabilities surpass those of the previous models.

Chapter 7: Cases of Calibration of the LIBOR Market Model. In this chapter, we start from a set of market data including zero-coupon curve, caps volatilities and swaptions volatilities, and calibrate the LFM by resorting to several parameterizations of instantaneous volatilities and by several constraints on instantaneous correlations. Swaptions are evaluated through the analytical approximations derived in the previous chapter. We examine the evolution of the term structure of volatilities and the ten-year terminal correlation coming out from each calibration session, in order to assess advantages and drawbacks of every parameterization.

We finally present a particular parameterization establishing a one-to-one correspondence between LFM parameters and swaption volatilities, such that the calibration is immediate by solving a cascade of algebraic second-order equations. No optimization is necessary in general and the calibration is instantaneous. However, if the initial swaptions data are misaligned due to illiquidity or other reasons, the calibration can lead to negative or imaginary volatilities. We show that smoothing the initial data leads again to positive real volatilities.

Chapter 8: Monte Carlo Tests for LFM Analytical Approximations. In this chapter we test Rebonato's and Hull-White's analytical formulas for swaptions prices in the LFM, presented earlier in Chapter 6, by means of a Monte Carlo simulation of the true LFM dynamics. This is done under different parametric assumptions for instantaneous volatilities and under different instantaneous correlations. We conclude that the above formulas are accurate in non-pathological situations.

We also plot the real swap-rate distribution obtained by simulation against the lognormal distribution with variance obtained by the analytical approximation. The two distributions are close in most cases, showing that the previously remarked theoretical incompatibility between LFM and LSM (where swap rates are lognormal) does not transfer to practice in most cases.

We also test our approximated formulas for terminal correlations, and see that these too are accurate in non-pathological situations.

Chapter 9 Other Interest-Rate Models. We present a few interest-rate models that are particular in their assumptions or in the quantities they model, and that have not been treated elsewhere in the book. We do not give a detailed presentation of these models but point out their particular features, compared to the models examined earlier in the book.

Part II: PRICING DERIVATIVES IN PRACTICE

Chapter 10: Pricing Derivatives on a Single Interest-Rate Curve. This chapter deals with pricing specific derivatives on a single interest-rate curve. Most of these are products that are found in the market and for which no standard pricing technique is available. The model choice is made on a case by case basis, since different products motivate different models. The differences are based on realistic behaviour, ease of implementation, analytical tractability and so on. For each product we present at least one model based on a compromise between the above features, and in some cases we present more models and compare their strong and weak points. We try and understand which model parameters affect prices with a large or small influence. The financial products we consider are: in-advance swaps, in-advance caps, autocaps, caps with deferred caplets, ratchets (one-way floaters), constant-maturity swaps (introducing also the convexity-adjustment technique), captions and floortions, Eurodollar futures, accrual swaps, trigger swaps and Bermudan-style swaptions.

Chapter 11: Pricing Derivatives on Two Interest-Rate Curves. The chapter deals with pricing specific derivatives involving two interest-rate curves. Again, most of these are products that are found in the market and for which no standard pricing technique is available. As before, the model choice is made on a case by case basis, since different products motivate different models. The used models reduce to the LFM and the G2++ shifted two-factor Gaussian short-rate model. Under the G2++ model, we are able to model correlation between the interest rate curves of the two currencies. The financial products we consider include differential swaps, quanto caps, quanto swaptions, quanto constant-maturity swaps. A market quanto adjustment and market formulas for basic quanto derivatives are also introduced.

Chapter 12: Pricing Equity Derivatives under Stochastic Interest Rates. The chapter treats equity-derivatives valuation under stochastic interest rates, presenting us with the challenging task of modeling stock prices and interest rates at the same time. Precisely, we consider a continuous-time economy where asset prices evolve according to a geometric Brownian motion and interest rates are either normally or lognormally distributed. Explicit formulas for European options on a given asset are provided when the instantaneous spot rate follows the Hull-White one-factor process. It is also shown how to build approximating trees for the pricing of more complex derivatives, under a more general short-rate process.

Part III: APPENDICES

Appendix A: a Crash Introduction to Stochastic Differential Equations. This appendix is devoted to a quick intuitive introduction on SDE's. We start from deterministic differential equation and gradually introduce randomness. We introduce intuitively Brownian motion and explain how it can be used to model the "random noise" in the differential equation. We observe that Brownian motion is not differentiable, and explain that SDE's must be understood in integral form. We quickly introduce the related Ito and Stratonovich integrals, and introduce the fundamental Ito formula.

We then introduce the Euler and Milstein schemes for the time-discretization of an SDE. These schemes are essential when in need of Monte Carlo simulating the trajectories of an Ito process whose transition density is not explicitly known.

We include two important theorems: the Feynman-Kac theorem and the Girsanov theorem. The former connects PDE's to SDE's, while the latter permits to change the drift coefficient in an SDE by changing the basic probability measure. The Girsanov theorem in particular is used in the book to derive the change of numeraire toolkit.

Appendix B: a Useful Calculation. This appendix reports the calculation of a particular integral against a standard normal density, which is useful when dealing with Gaussian models.

Appendix C: Approximating Diffusions with Trees. This appendix explains a general method to obtain a trinomial tree approximating the dynamics of a general diffusion process. This is then generalized to a two-dimensional diffusion process, which is approximated via a two-dimensional trinomial tree.

Appendix D: Talking to the Traders. This is the ideal conclusion of the book, consisting of an interview with a quantitative trader. Several issues are discussed, also to put the book in a larger perspective.