


MODERN PRICING OF INTEREST-RATE DERIVATIVES

The LIBOR Market Model and Beyond RICCARDO REBONATO



Modern Pricing of Interest-Rate Derivatives

THE LIBOR MARKET MODEL AND BEYOND

Riccardo Rebonato

Princeton University Press

Princeton and Oxford

Copyright © 2002 by Princeton University Press

Published by Princeton University Press,
41 William Street, Princeton, New Jersey 08540

In the United Kingdom: Princeton University Press,
3 Market Place, Woodstock, Oxfordshire OX20 1SY

All Rights Reserved

Library of Congress Cataloging-in-Publication Data applied for.

Rebonato, Riccardo

Modern Pricing of Interest-Rate Derivatives: The LIBOR Market and Beyond /

Riccardo Rebonato

p.cm.

Includes bibliographical references and index.

ISBN 0-691-08973-6

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

This book has been composed in New Baskerville

Printed on acid-free paper ∞

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Introduction

Rationale for the Book

The aim of this book is to present my views as to the most satisfactory approach to pricing a wide class of interest-rate derivatives. This approach falls squarely within the framework of the LIBOR market model. However, many competing versions, and even more modes of implementation, exist. I have not attempted to present a comprehensive, unbiassed review of all these possible approaches. Rather, I have chosen the particular version and the overall calibration philosophy that I have found most conceptually satisfying and practically useful. I have not been shy to express my opinions, but having strong views on a subject is, per se, no great virtue: I have therefore provided the reader with detailed reasons and explanations for my choices and preferences. Especially when it comes to its more recent treatments, I have consciously tried to cover an important version of the subject in depth, rather than its entirety in breadth.

Since 1973, the year of the Black and Scholes paper, both practice and theory have moved very rapidly in the derivatives area in general, and in interest-rate modelling in particular. Despite, or perhaps because of, this rapid development, the academic and practitioners' communities have not always communicated as productively as might have been desirable. Since my professional and academic experience straddles both fields, I have tried to bring theory and the practical trading experience together with as much 'constructive interference' as possible.

I imagine a rather diverse readership. All my readers should have in common a basic but solid knowledge of derivatives pricing, of the Black and Scholes model, of interest-rate products, of differential calculus and of elementary stochastic calculus. Some knowledge of stochastic calculus beyond Ito's lemma and some familiarity with modern probability theory would be helpful, but neither is a prerequisite. As long as they share this common ground, I hope that different types of reader will find the book useful. I have in mind, for instance: quantitative analysts, stronger in mathematical training than in trading experience, who want to deepen their understanding of how the modern interest-rate models work; practitioners interested in understanding the conceptual implications of the models they are using, up to what point

these models can be pushed and trusted, and what aspects of financial reality they leave out; students who want to understand how asset pricing is currently applied to interest-rate derivatives; and researchers with an interest in improving upon and expanding the current approaches. The pace of the book might be a bit slow for some, but should not be too fast for any of these readers.

Plan and Structure of the Book

The book is organized in four parts, as shown in Figure 1. Part I opens with a historical introduction that traces the development of interest-rate derivatives modelling over the last 30 years or so. This introductory chapter aims to provide an indication as to why certain modelling choices have been adopted by the industry, to explain how the perspective of what constitutes an ‘underlying’ and a ‘hedging’ instrument has changed over time, and to illustrate how the required explanatory power of a model has evolved.

I then move on to laying out the mathematical and modelling framework of the standard LIBOR market model. The latter is shown to be characterized by imposing that the volatilities of forward rates should be deterministic functions of time. From this the observation follows that, given a pricing measure, forward rates cannot all be simultaneously log-normally distributed. This has deep consequences, because the lack of simultaneous log-normality for the forward rates would seem to indicate that their future conditional distributions are not known a priori, and therefore must be sampled via computationally intensive numerical methods (e.g., short-stepped Monte Carlo simulations). I show, however, that it is possible to introduce some very simple techniques that retain in an approximate but very accurate way the joint log-normality. Once these approximations have been enforced, I show (Chapter 5) that the drifts can be expressed purely as a function of the covariances. Therefore, within the validity of these approximations, the covariance elements are truly all that matters for pricing in a LIBOR market model context. This is the first important ‘message’ of Part I.

The attention is then shifted to the types of product that can be handled by the modern pricing approach. I show that their payoff functions must satisfy some, practically rather mild, conditions (related to the measurability and homogeneity of the payoffs). More importantly, I show that for all these products the Monte Carlo simulations that are typically used for valuation purposes can be carried out by evolving the yield curve either from one price-sensitive event to the next, or directly to the very final maturity. I have called these two modes of evolution of the yield curve the long- and the very-long-jump technique, to emphasize that, once the drift approximations referred to above are enforced, the computationally expensive short-stepped Monte Carlo evolution is never required (in the deterministic-volatility case). This is the second important result of Part I.

Part I is concluded by a chapter that shows how these numeraire- and measure-dependent drift terms can be expressed in terms of market-related

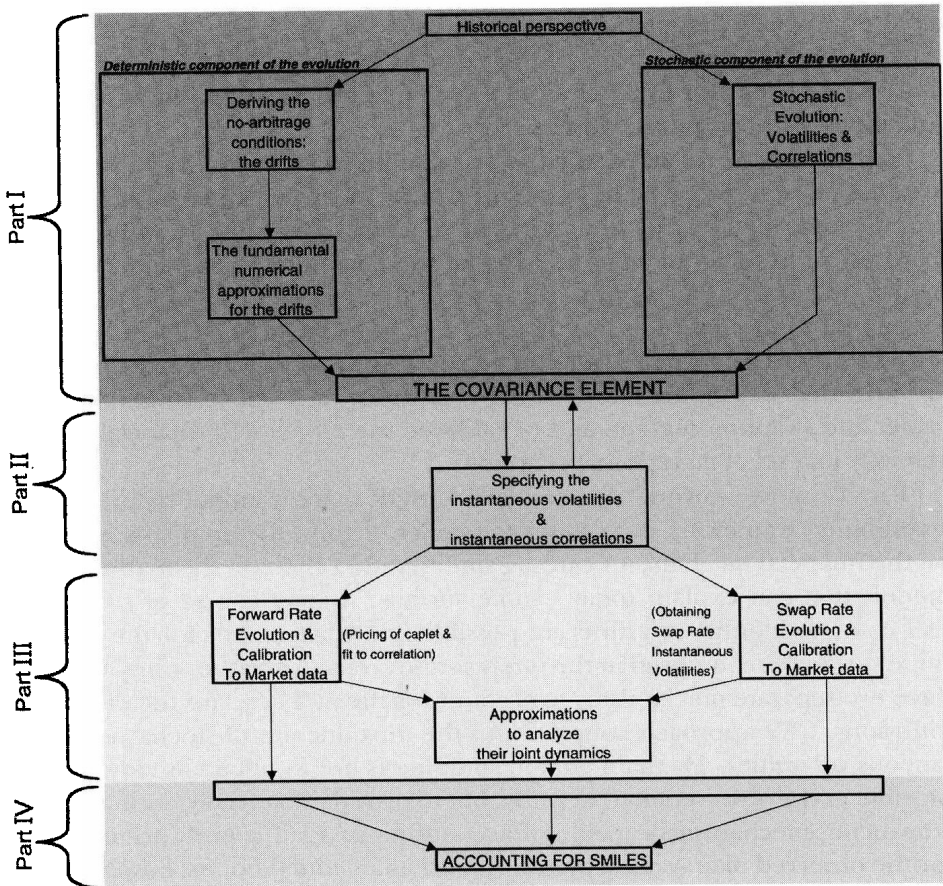


Figure 1 Plan of the book.

forward-rate volatilities and correlations. The focus is therefore naturally shifted to the analysis of these latter quantities, which constitutes the topic of Part II.

The logical development up to this point has highlighted the centrality of the marginal or of the total forward-rate covariance matrix to the whole modern pricing approach. The elements of this covariance matrix are, in turn, made up of instantaneous volatilities and correlations, and the LIBOR market model, which ultimately simply expresses no-arbitrage conditions for discrete-tenor forward rates, only and truly becomes a 'model' when these quantities are chosen. The quality of the whole pricing approach therefore hinges on the quality of these inputs, and Part II is devoted to exploring in considerable detail what the desiderata should be of plausible volatility and correlation functions. The task is made challenging (and interesting) because the market data available from the plain-vanilla instruments do not uniquely determine

either the instantaneous correlation or volatility functions, and a strong element of financial judgement therefore necessarily enters the analysis. In this respect I explain why I place a great deal of emphasis on the condition of time homogeneity for the evolution of such quantities as the term structure of volatilities or the swaption matrix.

Part III assumes that, on the basis of the criteria presented in Part II, the user has chosen to her heart's content a satisfactory set of instantaneous correlation and volatility functions. In general, these chosen inputs will not automatically reproduce the market prices of the benchmark plain-vanilla instruments. The task undertaken in this part of the book is therefore to calibrate the model to the desired set of market data not only in an efficient way, but also, and more importantly, in a financially desirable way. In particular, tools are provided to explore the internal coherence and self-consistency of the caplet and swaption market, and to calibrate the model in a financially desirable way to a set of co-terminal swaptions.

Part IV moves beyond the standard LIBOR market model by presenting a combined framework to account for smiles. Again, the emphasis is placed on the financial mechanisms causing smiles, rather than on the mathematical models that can explain today's smile surface. In the course of this analysis I come to identify two different possible financial reasons for the types of smile surface encountered in the interest-rate arena. Correspondingly, I introduce two separate mechanisms to account for them. These are the displaced-diffusions/CEV approach coupled with the introduction of stochastic instantaneous volatilities. My main goal in so doing is not to obtain as close a fit as possible to the observed market smile, but to provide financially justifiable and convincing mechanisms capable of accounting in a sufficiently accurate way for the observed plain-vanilla prices. Almost as an added bonus, however, it so turns out that the quality of the obtained fit is very good indeed. Finally, I argue in the last three chapters that the standard LIBOR market model as it has evolved in the last few years is far more than a set of no-arbitrage conditions for forward rates. In its recent developments, it has become intimately enmeshed with the set of approximation and calibration techniques that have turned it into a practical and powerful pricing tool, and that I present and discuss in Parts I to III. I show in Part IV that the surprisingly simple, but very powerful, techniques that allow the efficient calibration of the model and evolution of the yield curve can be translated to the displaced-diffusion/stochastic-volatility setting in a natural and straightforward way. The last chapter presents and discusses empirical evidence concerning the swaption implied-volatility matrix, and its relevance for the stochastic volatility modelling. A general systematic procedure is introduced to assess the quality of a model when the term structure of volatilities it produces is stochastic.

As for the presentation style of the book, it has been strongly influenced by the use to which I have assumed it will be put by its presumably diverse readership. My ideal reader would begin on page 1, and would read the book sequentially (although, possibly, with a few breaks!) all the way to the end. To

such a reader I must offer my thanks and some apologies. I have had in the back of my mind, in fact, the existence of a 'sub-optimal' reader who might prefer to read some sections out of their natural order, skip some chapters altogether, re-read others over and over, etc. A reader, in other words, rather like myself. For the benefit of this erratic reader, but possibly to the annoyance of the ideal one, I have therefore provided frequent indicators of the progress of the logical flow and of the direction ahead, and liberally posted reminders of the main results obtained on the way. I am aware that this does not make for elegant writing. I would at least like the ideal reader to know that I have not inflicted unnecessary repetitions out of sheer stylistic sloppiness or mental laziness, and that, in the end, it is all the fault of readers like me.

Acknowledgements

It is a pleasure to acknowledge the help, encouragement and support that I have received from colleagues and friends. In particular, I owe a debt of gratitude to Dr Mark Joshi and Dr Chris Hunter for stimulating discussions, useful suggestions and inspiring bantering. I would also like to thank Mark and Chris again, and Dr Peter Jaeckel, for invaluable computational help, and Dr Jochen Theis for careful reading of the manuscript and useful comments. Needless to say, I am solely responsible for all the errors in the text.

As I have acknowledged in my previous books, the seeds for many of the ideas that appear in my work were sown many years ago by Mike Sherring: my heartfelt thanks to him.

I am grateful to the staff at Princeton University Press, and to Richard Baggaley in particular, for the enthusiasm they have shown for the project and the help they have provided.

I have written this book in a private capacity, and the opinions expressed in it should not be attributed to my employer, the Royal Bank of Scotland, which, however, I wish to thank for providing a stimulating environment in which the ideas here presented have found intellectual challenge and practical application.

Finally, I hope I will be as supportive to my wife in her current and forthcoming book projects as she has been to me during the course of this work.

Contents

Introduction	xi
Acknowledgements	xvii
I The Structure of the LIBOR Market Model	1
1 Putting the Modern Pricing Approach in Perspective	3
1.1 Historical Developments	3
1.2 Some Important Remarks	21
2 The Mathematical and Financial Set-up	25
2.1 The Modelling Framework	25
2.2 Definition and Valuation of the Underlying Plain-Vanilla Instruments	28
2.3 The Mathematical and Financial Description of the Securities Market	40
3 Describing the Dynamics of Forward Rates	57
3.1 A Working Framework for the Modern Pricing Approach	57
3.2 Equivalent Descriptions of the Dynamics of Forward Rates	65
3.3 Generalization of the Approach	79
3.4 The Swap-Rate-Based LIBOR Market Model	83
4 Characterizing and Valuing Complex LIBOR Products	85
4.1 The Types of Product That Can be Handled Using the LIBOR Market Model	85
4.2 Case Study: Pricing in a Three-Forward-Rate, Two-Factor World	96
4.3 Overview of the Results So Far	107
5 Determining the No-Arbitrage Drifts of Forward Rates	111
5.1 General Derivation of the Drift Terms	112
5.2 Expressing the No-Arbitrage Conditions in Terms of Market-Related Quantities	118
5.3 Approximations of the Drift Terms	123

5.4	Conclusions	131
II	The Inputs to the General Framework	133
6	Instantaneous Volatilities	135
6.1	Introduction and Motivation	135
6.2	Instantaneous Volatility Functions: General Results	141
6.3	Functional Forms for the Instantaneous Volatility Function – Financial Implications	153
6.4	Analysis of Specific Functional Forms for the Instantaneous Volatility Functions	167
6.5	Appendix I – Why Specification (6.11c) Fails to Satisfy Joint Conditions	171
6.6	Appendix II – Indefinite Integral of the Instantaneous Covariance	171
7	Specifying the Instantaneous Correlation Function	173
7.1	General Considerations	173
7.2	Empirical Data and Financial Plausibility	180
7.3	Intrinsic Limitations of Low-Dimensionality Approaches	185
7.4	Proposed Functional Forms for the Instantaneous Correlation Function	189
7.5	Conditions for the Occurrence of Exponential Correlation Surfaces	196
7.6	A Semi-Parametric Specification of the Correlation Surface	204
III	Calibration of the LIBOR Market Model	209
8	Fitting the Instantaneous Volatility Functions	211
8.1	General Calibration Philosophy and Plan of Part III	211
8.2	A First Approach to Fitting the Caplet Market: Imposing Time-Homogeneity	214
8.3	A Second Approach to Fitting the Caplet Market: Using Information from the Swaption Matrix	218
8.4	A Third Approach to Fitting the Caplet Market: Assigning a Future Term Structure of Volatilities	226
8.5	Results	231
8.6	Conclusions	248
9	Simultaneous Calibration to Market Caplet Prices and to an Exogenous Correlation Matrix	249
9.1	Introduction and Motivation	249
9.2	An Optimal Procedure to Recover an Exogenous Target Correlation Matrix	254
9.3	Results and Discussion	260
9.4	Conclusions	274

10 Calibrating a Forward-Rate-Based LIBOR Market Model to Swaption Prices	276
10.1 The General Context	276
10.2 The Need for a Joint Description of the Forward- and Swap-Rate Dynamics	280
10.3 Approximating the Swap-Rate Instantaneous Volatility	294
10.4 Computational Results on European Swaptions	306
10.5 Calibration to Co-Terminal European Swaption Prices	312
10.6 An Application: Using an FRA-Based LIBOR Market Model for Bermudan Swaptions	318
10.7 Quality of the Numerical Approximation in Realistic Market Cases	326
IV Beyond the Standard Approach: Accounting for Smiles	331
11 Extending the Standard Approach – I: CEV and Displaced Diffusion	333
11.1 Practical and Conceptual Implications of Non-Flat Volatility Smiles	333
11.2 Calculating Deltas and Other Risk Derivatives in the Presence of Smiles	342
11.3 Accounting for Monotonically Decreasing Smiles	349
11.4 Time-Homogeneity in the Context of Displaced Diffusions	363
12 Extending the Standard Approach – II: Stochastic Instantaneous Volatilities	367
12.1 Introduction and Motivation	367
12.2 The Modelling Framework	372
12.3 Numerical Techniques	382
12.4 Numerical Results	397
12.5 Conclusions and Suggestions for Future Work	413
13 A Joint Empirical and Theoretical Analysis of the Stochastic-Volatility LIBOR Market Model	415
13.1 Motivation and Plan of the Chapter	415
13.2 The Empirical Analysis	420
13.3 The Computer Experiments	437
13.4 Conclusions and Suggestions for Future Work	442
Bibliography	445
Index	453

Part I

The Structure of the LIBOR Market Model

1

Putting the Modern Pricing Approach in Perspective

1.1 Historical Developments

1.1.1 Introduction

The set of techniques to price interest-rate derivatives that stemmed from the original work of Heath, Jarrow and Morton (HJM) in the late 1980s (HJM 1989) are referred to in this book as the ‘modern’ or the ‘LIBOR-market-model’ approach. At a superficial glance, the differences between the various ‘incarnations’ of the approach might appear greater than what they have in common. The state variables could be instantaneous or discretely compounded rates; they could be swap rates or forward rates; they might be normally or log-normally (or otherwise) distributed; the associated numeraire could be a zero-coupon bond, a swap annuity or the money-market account; and so on. Despite these non-trivial differences, these approaches share one essential common feature: the recognition that the no-arbitrage evolution of the state variables (however chosen) can be expressed purely as a function of the volatilities of, and of the correlations among, the state variables themselves. Different choices of numeraires will give rise to different combinations for these covariance elements, but this fundamental result, which goes back to the original insight of HJM, is shared by all the approaches that will be dealt with in this book. This result and its implications are sufficiently fundamental and far-reaching to justify a self-contained and unified treatment.

Given the various ‘versions’, ‘implementations’ and choices of numeraires, no general agreement exists in the financial community on how to call this set of approaches: the terms ‘BGM (Brace, Gatarek and Musiela) model’ and ‘Jamshidian approach’ are often used, but ‘pricing in the forward measure’, the ‘LIBOR market model’ and other terms are also frequently encountered. Some purists insist on calling the approach simply the ‘HJM model’. The difficulty in establishing intellectual priority is compounded by the fact that

many of the key results were first obtained (and used) by practitioners but, for obvious reasons, not published in the academic press. I have therefore avoided any identification of the approach with names of academics or market professionals, and used the more neutral terms ‘LIBOR market model’ or ‘modern pricing approach’ – much as I am sure the latter may read rather quaint in a few years’ time.

This introductory chapter is meant to provide a brief map of the development of interest-rate derivative pricing from its earliest (modern) days to the present. I have chosen to present such an introduction not only for its intrinsic historical interest, but also because it illustrates rather clearly that an uneven combination of market events, ‘right choices made for the wrong reasons’, computational expediency and sound judgement have conspired to produce the market standard that the later, more sophisticated, models have striven to recover. In other words, the modern approach is, justifiably, so loved by practitioners because of its ability to price exotic products while at the same time recovering exactly the prices of the relevant plain-vanilla options (caplets or European swaptions). I shall explain below how the market consensus has crystallized around the Black framework, partly for sound financial reasons, but partly also by historical accident. If this analysis is correct, there is nothing ‘inevitable’ about the current market standard, and it is quite possible that the target the modern approach has been trying to hit might in the near future turn out to be a rather rapidly moving one.

Indeed, this phenomenon is already becoming apparent: as discussed in Part IV of this book, in the last few years the prices of plain-vanilla options have been able to be strait-jacketed into their log-normal-rate Black framework only by increasingly complex ad hoc adjustments.¹ As a consequence, just when the pricing of exotic products had finally been successfully tuned onto the log-normal-rate wavelength, the prices of the underlying vanilla instruments have ceased to inhabit the same (Black) world. The brief account of the developments that brought about this state of affairs is presented below, and should give a clear indication of the fact that the ‘modern’ approach is virtually certain to be anything but the last step in interest-rate derivative pricing. The reader keen to delve into the quantitative aspects of the pricing can safely skip these pages. She would miss, however, not only a good story, but also some perspective useful in appreciating what aspects of today’s market consensus are more likely to be challenged and superseded tomorrow.

¹The malaise first became apparent with the appearance of implied volatilities monotonically decreasing as a function of strike. This phenomenon has acquired increasing importance since the mid-to-late 1990s. In addition, after the 1998 market turmoil, a distinct hockey-stick or U-shape for the implied volatility curve has appeared both for caplets and for swaptions. See the discussion in Section 11.1.

1.1.2 The Early Days

The relatively brief history of the evolution of the pricing of interest-rate derivatives can be subdivided into four distinct periods. The very first one corresponds to the use of the Black and Scholes (1973) (BS), Black (1976) and Merton (1973) approaches. In all these cases, the same distributional assumption (namely, the log-normal) was made for the underlying variable, and the resulting expiry-time distribution was integrated over the terminal payoff of a European option. For all of the three above-mentioned models, the solutions have a very (and deceptively) similar appearance, with the integration over the appropriate log-normal probability densities giving rise to the familiar cumulative-normal-distribution terms. The quantities that were typically assumed to be log-normal were bond prices (spot or forward), forward rates, forward swap rates, or bond yields. As for the products priced using these modelling approaches, they belonged to two rather distinct markets and yield curves: the Treasury/repo world, on the one hand, was the relevant environment for the pricing of plain-vanilla bond options, either directly, or embedded in once-callable or extendable-maturity structures; and the LIBOR environment, on the other, which provided caps and European swaptions.

The most straightforward approach (i.e., the use of the Black and Scholes formula with the spot bond price as the underlying) was popular, but it also came under early criticism because of the so-called pull-to-par phenomenon: in its original form the Black and Scholes formula requires a constant percentage volatility of the underlying. For a coupon or a discount bond, however, the volatility is certainly not constant (since the price has to converge to par at maturity). This fact was considered to be cause for little worry if the expiry of the bond option was much shorter than the maturity of the underlying bond (e.g., for a few weeks' or months' option on a, say, 10-year bond); but it would create discomfort when the times to expiry and maturity were comparable.

The 'easy-fix' solution was to consider a non-traded quantity (the bond yield) as the underlying log-normal variable. The main advantage of this approach was that a yield does not exhibit a deterministic pull to par. The most obvious drawback, on the other hand, was that the yield is not a traded asset, and, therefore, the Black and Scholes reasoning behind the self-financing dynamic trading strategy that reproduces the final payoff of the option could not be easily adapted. Despite being theoretically not justifiable, the approach was for a period widely used because it allowed the trader to think in terms of a volatility (the volatility of the yield) that was more independent of the maturity of the underlying instrument than the volatility of the spot price. With hindsight, and much to the relief of the academic community, this route was to prove a blind alley, and very little more will be said in this book about it.² Academics, of course and correctly, 'hated' the log-normal yield approach: not only did it make use as its driving variable of a quantity (the yield itself) of

²Yield-based approaches are still used to some extent in the bond option world. This asset class is not dealt with in this book, which mainly addresses the pricing of LIBOR derivatives.