Preface

For quantitative researchers working in an investment bank, the process of writing a fixed income model usually has two stages. First, a theoretical framework for yield curve dynamics is specified, using the language of mathematics (especially stochastic calculus) to ensure that the underlying model is well-specified and internally consistent. Second, in order to use the model in practice, the equations arising from the first step need to be turned into a working implementation on a computer. While specification of the theoretical model may be seen as the difficult part, in quantitative finance applications the second step is technically and intellectually often more challenging than the first. In the implementation phase, not only does one need to translate abstract ideas into computer code, one also needs to ensure that the resulting numbers being produced are meaningful to a trading desk, are stable and robust, are in line with market observations, and are produced in a timely manner. Many of these requirements are, as it turns out, extremely challenging, and not only demand a strong knowledge of actual market practices (which tend to deviate in significant ways from "textbook" theory), but also require application of a large arsenal of techniques from applied mathematics, chiefly approximation methods and numerical techniques.

While there are many good introductory books on fixed income derivatives on the market, when we hire people who have read them we find that they still require significant training before they become productive members of our quantitative research teams. For one, while existing literature covers some aspects of the first step above, advanced approaches to specifying yield curve dynamics are typically not covered in sufficient detail. More importantly, there is simply too little said in the literature about the process of getting the theory to work in the real world of trading and risk management. An important goal of our book series is to close these gaps in the literature.

As we write this in early 2010, financial markets are still reeling from a severe crisis that has, at least in part, been blamed on over-the-counter (OTC) options markets, the venue where complex derivative securities are transacted. Stricter regulation of some types of OTC derivatives currently seems all but inevitable, and many common OTC securities may in the future either be outlawed or traded only on public exchanges. In the wake of the crisis, opinion of financial engineers and bankers has hit an all-time low, with many in the public convinced that they are peddlers of toxic waste or "weapons of financial destruction". All things considered, the present may therefore seem like an inauspicious moment to launch a series of monographs on the pricing and risk management of interest rate derivatives. We disagree, for several reasons. First, in defense of OTC derivatives we note that although they certainly can be used inappropriately to create excessive leverage and risk, many complex (or "exotic") derivatives serve as innovative and cost-effective vehicles for bank clients to reduce their financial risk. Second, irrespective of what will ultimately transpire on the regulatory front, it has become obvious that going forward both regulators and market participants need a better grasp of the management and characterization of complex financial risk. This is perhaps particularly true for the quantitative research professionals (the "quants", in common parlance) who recently have been taken to task by the press for the failure of their models and their inability to predict the credit crisis. While this simplistic characterization is actually quite unfair, there is no doubt that many derivatives models that worked well enough before the credit crisis are no longer adequate. Indeed, even the simple task of pricing a basic interest rate swap — possibly the simplest of all interest rate derivatives — has recently required major methodology revisions[1]. If nothing else, a severe crisis serves to expose weaknesses in the foundation on which models are built, allowing one to reinforce it for future storms. In this light, we feel that the time is just about right for a comprehensive, practical, and up-to-date exposition of interest rate modeling and risk management[2].

The three volumes of *Interest Rate Modeling* are aimed primarily at practitioners working in the area of interest rate derivatives, but much of the material is quite general and, we believe, will also hold significant appeal to researchers working in other asset classes. Students and academics interested in financial engineering and applied work will find the material particularly useful for its description of real-life model usage and for its expansive discussion of model calibration, approximation theory, and numerical methods. In preparing the books we have drawn on nearly 30 years of combined industry experience, and much of the material has never been exposed in book form before.

---

[1] We cover this in Chapter 6.

[2] We ought to note that interest rate derivatives (unlike *credit* derivatives) so far have not been directly implicated in the financial crisis.

Quantitative finance attracts students and practitioners from many different academic fields, and with varying levels of preparation in mathematics and computation. (Case in point: L.B.G.A was originally a robotics engineer and V.V.P a probabilist.) To cater to a broad audience, we have kept the exposition fairly informal; graduate students in applied fields such as engineering and physics should feel at home with the level (or lack) of rigor used in the book. We have relied on a proposition-proof format throughout, largely because this facilitates easier cross-referencing in a long text, but acknowledge that the format is occasionally more formal than the results themselves. For instance, we tend to skip over technical regularity conditions in our proofs and also frequently list approximate results in propositions without explicitly specifying the sense in which they approximate true values. Although the exposition is largely self-contained, some previous knowledge of basic option pricing principles (e.g., at the level of Hull [2006]) may be useful.

Interest Rate Modeling divides into three separate volumes. *Volume I* provides the theoretical and computational foundations for the series, emphasizing the construction of efficient grid- and simulation-based methods for contingent claims pricing. Numerical methods serve an extremely important role in the text, so we develop this topic to an advanced level suitable for professional-quality model implementations. Placing this material early in the text allows us to incorporate it into our discussion of individual models in subsequent chapters. The second part of Volume I is dedicated to local-stochastic volatility modeling and to the construction of vanilla models for individual swap and Libor rates. Although the focus is eventually turned toward fixed income securities, much of the material in this volume applies to a broad capital market setting and will be of interest to anybody working in the general area of asset pricing.

Volume II is dedicated to in-depth study of term structure models of interest rates. While providing a thorough analysis of classical short rate models, the primary focus of the volume is on multi-factor stochastic volatility dynamics, in the setups of both the separable HJM and Libor market models. Implementation techniques are covered in detail, as are strategies for model parameterization and calibration to market data.

The first half of *Volume III* contains a detailed study of several classes of fixed income securities, ranging from simple vanilla options to highly exotic cancelable and path-dependent trades. The analysis is done in product-specific fashion, covering, among other subjects, risk characterization, calibration strategies, and valuation methods. In its second half, Volume III studies the general topic of derivative portfolio risk management, with a particular emphasis on the challenging problem of computing smooth price sensitivities to market input perturbations.

Although much of the material in *Interest Rate Modeling* is focused on the technical and theoretical issues surrounding model implementation on a computer, it is impractical for us to delve into the exercise of writing actual computer routines. Fortunately, there are several specialized books on how to write good quant code, see, e.g., Hyer [2010] and Joshi [2004]. Both of these books work with C++ which is still the most common computer language used in professional quant libraries. For those that choose to work with C++, we wholeheartedly endorse books by Scott Meyers (see, e.g., Meyers [2005]) and Andrei Alexandrescu (see, e.g., Sutter and Alexandrescu [2004]) as guides to sound and maintainable code.

During the six year process of writing this book series, we have received encouragement and constructive criticism from many people. We particularly wish to thank Peter Carr, Peter Forsyth, Alexandre Antonov, Peter Jäckel, Dominique Bang, Martin Dahlgren, Neil Oliver, Patrick Roome, Regis van Steenkiste, Natasha Bushueva and many members of the research teams at Barclays Capital and Bank of America Merrill Lynch. Natalia Kryzhanovskaya meticulously proofread our first draft, and contributed greatly to the harmonization of notation across what turned out to be a very long manuscript. All remaining errors are, of course, entirely our own. Speaking of errors: with nearly 20,000 equations, it is probable that a few typos remain, despite our best efforts to weed them out. A list of errata will be maintained on `www.andersen-piterbarg-book.com` where supplemental material and news will also be posted on a running basis. We greatly appreciate reporting of typos or factual errors to our web address, and will list the names of all those who contribute to error spotting in future editions of *Interest Rate Modeling*.

Lastly, we owe a great debt of gratitude to our families for their support and patience, even when our initial plans for a brief book on tips and tricks for working quants ballooned into something more ambitious that consumed many evenings and weekends over the last six years.

London, New York,
June 2004 — August 2010

Leif B.G Andersen
Vladimir V. Piterbarg