

COPYRIGHT NOTICE:

Riccardo Rebonato: Modern Pricing of Interest-Rate Derivatives

is published by Princeton University Press and copyrighted, © 2002, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

For COURSE PACK and other PERMISSIONS, refer to entry on previous page. For more information, send e-mail to permissions@pupress.princeton.edu

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.

poor theoretical standing (see, e.g., Schaefer 1977), but it was close to impossible to produce a satisfactory financial justification in terms of risk-neutral valuation for its use in a Black-like formula. In all likelihood, however, the trading community abandoned the log-normal yield model for reasons other than its lack of sound theoretical standing. As discussed below, in fact, the cap and swaption pricing formulas, which used log-normal forward and swap rates, were at the very same time gaining general market acceptance. Prima facie, these formulas appear deceptively similar to the log-normal yield approach. It so happens that, unlike the log-normal yield approach, these latter formulas can be justified theoretically. In the early days, however, this was rarely recognized, and papers showing the ‘correctness’ of the Black swaption formula were still appearing as late as the mid-1990s (see Neuberger 1990, Gustavsson 1997). Traders, although with a somewhat guilty conscience, were nonetheless still using the Black formula for caplets and European swaptions well before its theoretical justification had become common knowledge.

The reasons for the abandonment of the log-normal yield model must therefore have been rooted in causes other than its lack of theoretical justifiability. This constitutes more than a piece of historical curiosity: the fact that traders carried on using the Black formula for caplets and swaptions despite its then-perceived lack of theoretical standing turned the approach into a market standard. The subsequent realization that this market standard could be put on solid theoretical ground then provided the impetus behind the development of more general approaches (such as the LIBOR market model) capable of pricing complex interest-rate derivatives consistently with the (Black) caplet market. Therefore, without the somewhat fortuitous choices made during the establishment of a market standard, the current ‘modern’ pricing approach would probably not exist in its present form.

Going back to the pull-to-par problem, its ‘correct’ solution was, of course, to use the Black – rather than the Black and Scholes – formula, with the forward (as opposed to spot) price and the volatility of the forward as inputs. This route was often overlooked, and for different reasons, both by naive and by sophisticated market players. The naive traders simply did not appreciate the subtle, but fundamental, difference between the Black and the Black and Scholes formulas and the volatilities used as input for both, and believed the pull-to-par phenomenon to be relevant to the Black formula as well. The sophisticated traders understood that the Black formula, with the appropriate inputs, is perfectly correct and justifiable, and that no pull-to-par effect applies to a forward bond price (which remains of the same residual maturity throughout its life). They also realized, however, that something important was still missing. Much as the Black formula can give a perfectly correct answer for a series of options considered independently of each other, there is no way of telling whether these options inhabit a plausible, or even logically consistent, universe. In reality, despite the fact that different forward bond prices correspond to different assets, we also know that these prices are significantly correlated, but there is no mechanism within the Black formula to incorpo-

rate views about this joint dynamics. To make this important point clearer, let us assume that a trader were to give two different forward bond price volatilities as inputs for the Black formula for two options with different expiries on forward bonds with the same residual maturities. By so doing the trader is implicitly saying ‘something’ about the volatilities of the forward bonds spanning the period between the first and the second expiry, and the first and second maturity. Notice, however, the vagueness of the term ‘something’. The relationship between the various volatilities is not strong enough to be of deterministic nature: given the volatilities for two same-residual-maturity, different-expiry forward bonds, the volatility of the forward bond spanning, say, the two expiries is by no means uniquely determined. At the same time, the relationship between the various volatilities is not weak enough for the various underlying ‘assets’ to be considered completely independent.

Sure enough, if an exogenous model existed, capable of providing self-consistent inputs (i.e., volatilities) to the Black formula, then all problems of consistency would be solved. Indeed, I have shown elsewhere (Rebonato 1998) that the Black formula can be rigorously extended to a surprising variety of path-dependent and compound-option cases (at least, as long as one is prepared to work with higher-and-higher-dimension cumulative normal distributions). However, paraphrasing what I wrote in a slightly different context (Rebonato 1998):

... what is needed is some strong structure to be imposed on the co-movements of the financial quantities of interest; ... this structure can be provided by specifying the dynamics of a small number of variables. ... Once the process for all these driving factors has been chosen, the variances of and correlations among all the financial observables can be obtained, analytically or numerically as the case might be, as a by-product of the model itself. The implied co-dynamics of these quantities might turn out to be simplified to the point of becoming simplistic, but, at least, the pricing of different options can be undertaken on a consistent basis. ...

These reasons for the early dissatisfaction with the Black approach are of more than historical interest because, paradoxically, the modern approach to pricing derivatives can suffer from the same over-abundance of degrees of freedom. I will argue that the need to impose the strong structure on the co-movements of the yield curve mentioned above is absolutely central to the modern implementation of the pricing approaches derived from the HJM family. But more about this later.

In parallel with the Treasury/repo market, the LIBOR-related market was also providing interesting pricing challenges. The demand for LIBOR-based derivative products was coming on the one hand from liability managers (borrowers) seeking interest-rate insurance via the purchase of caps, and on the other from issuers and investors seeking better returns or funding rates, respectively, via callable or puttable bonds. The coupons from the latter would

then be swapped with an investment firm, who would also receive the optionality (typically a swaption) embedded in the structure. The mechanics of, and financial rationale for, these trades is an interesting topic in itself: the story of the search for funding ‘advantage’ that drove the evolution of financial instruments from the early once-callable coupon bonds to the 30-year multi-callable zero-coupon swaptions of late 1997/early 1998 deserves to be told, but elsewhere. For the present discussion it is sufficient to point out two salient features.

The first is that the consensus of the investment houses that were most active in the field was crystallizing around the use of the Black formula with log-normal forward and swap rates. This need not have been the case, since a caplet, for instance, can be equivalently regarded as a call on a (log-normal) rate or a put on a (log-normal) bond. Once again, despite the (rather unconvincing) rationalizations for the choice given at the time, the choice of the Black formula with the respective forward rates for caplets and swaptions turned out to be based on little more than historical accident. Without this choice, the development of the HJM approach would have taken a significantly different turn, and would have actually led to the market acceptance and popularity of formulations much easier to implement (such as the ones described in Section 17.3 of Rebonato (1998); see also Carverhill (1993) or Hull (1993)). Perhaps more importantly, had the choice of the put on a (log-normal) forward bond been made, today’s term structure of volatilities for caplets would display a radically different smile structure (or, perhaps, no smile at all). And as for Chapter 11 of this book, which shows how to extend the modern pricing approach in the presence of monotonically decaying smiles, it might have been absent altogether, or relegated to a footnote.

The second important consequence of the market supply and demand for optionality was that investment houses found themselves as the natural buyers of swaptions and sellers of caps. Of course, it was readily recognized that these instruments were intimately linked: a swap rate, after all, is simply a linear combination (albeit with stochastic weights) of the underlying forward rates. However, the Black formulas by themselves once again failed to give any indication as to what this link should be: each individual caplet would be perfectly priced by the Black formula with the ‘correct’ input volatility, but it would inhabit its own universe (in later-to-be-introduced terminology, it would be priced in its own forward measure). Even more so, there would be no systematic way, within a Black framework, to make a judgement about the mutual consistency between the volatility of a swaption and the volatilities of the underlying caplets.

Much as in the Treasury/repo market, the need was felt for a comprehensive approach that would bring unity to the pricing of these different and loosely connected instruments. The search for a model with the ability to price convincingly and simultaneously the caplet and the swaption markets was actually to be frustrated until the late 1990s, and the progress made in this direction actually constitutes one of topics of Part III of this book. Once again,

the story of the search for this particular ‘Holy Grail’ is in itself very interesting, and even has an almost dramatic twist, linked as it is to the ill-fated trades of a well-known investment house of the late summer of 1998. For the purpose of the present discussion, however, the salient feature of this early phase was the perceived need for a more comprehensive modelling approach that would bring unity and simplicity to what appeared to be hopelessly disjointed markets.

In moving from this earliest stage in interest-rate derivatives pricing to its adolescence, one last important factor should be mentioned: the caplet and swaption markets soon acquired such importance and liquidity that they became the new ‘underlyings’. In other terms, the gamma and vega of the more complex products that were beginning to appear would be routinely hedged using caplets and European swaptions.³ This placed another important constraint on the ‘comprehensive’ model that was being sought: since the price of an option is, after all, intimately linked to the cost of its hedges, the market prices of caplets or European swaptions (seen now not as options in their own right, but as the new hedging instruments) would have to be correctly recovered. From the point of view of the trader who had to make a price in an exotic product, it would simply not be good enough to have a model that implied that caplets prices ‘should’ be different from what the (Black-driven) market implied. This feature, in the early days, was to be found rather down the wish list of the desiderata for the model the trader would dream of receiving for Christmas. It was nonetheless to become one of the most compelling reasons for the development and the ready market acceptance of the modern pricing approach (namely, in the incarnation usually referred to as the ‘market model’).

1.1.3 The First Yield-Curve Models

The modelling lead was still taken at this stage by academics (namely, Vasicek (1977) and Cox, Ingersoll and Ross (CIR) (Cox et al. 1985)). Faced with the multitude of instruments (bond options, caplets, swaptions, not to mention the bonds themselves) in search for a coherent and self-consistent description, Vasicek and Cox et al.⁴ made the sweeping assumption that the dynamics of the whole yield curve would be driven by the instantaneous short rate. The evolution of the latter was then assumed to be described by a stochastic differential equation made up of a deterministic mean-reverting compo-

³For a distinction between and discussion of ‘in-model’ (delta and gamma) hedging and ‘outside-the-model’ (vega) hedging, see Section 11.1.

⁴The brief account presented in this chapter barely does justice to the CIR approach, which is actually considerably farther-reaching, ensuring as it does that the resulting dynamics for the short rate is consistent with an equilibrium economy. Given the rather draconian nature of the assumptions the model has to enforce, this feature was perceived to be more of an intellectual bonus than of practical relevance. As a consequence, the trading community never placed much importance on this theoretically important feature, and regarded the CIR approach as a purely no-arbitrage model. The same applies to the Longstaff and Schwartz (1992) (LS) model.

ment and a stochastic part, with a diffusion coefficient either constant or proportional to the square root of the short rate itself. Given the existence of a single source of uncertainty (as described by the evolution of the short rate) the stochastic movements of any bond could be perfectly hedged by taking a suitable position in a bond of different maturity. The appropriate hedge ratio was obtainable, via Ito's lemma, on the basis of the assumed evolution of the single driver of the yield-curve dynamics (i.e., the short rate). Given the existence of a single stochastic driver, the prescription was therefore given to build a purely deterministic portfolio made up of just two bonds of different maturities (in a suitable ratio). Once this step was accomplished, the solution of the problem could therefore be conceptually reduced to the classic Black-and-Scholes framework, where a stock and an option are combined in such an amount as to give rise to a risk-less portfolio.

Despite the fact that the practical success of these first models was rather limited, their influence was enormous, and, with exceedingly few exceptions (such as Duffie and Kan's (1996) approach), all the models that were to be developed up until the HJM approach was introduced were part of the same, short-rate-based, research program. With hindsight, it might seem bizarre that so much modelling effort should have been concentrated in the same direction. It might appear even more surprising, given the seemingly arbitrary and unpromising choice for the driving factor, that the models spawned by this approach should turn out to be, after all, as useful as they have been. As usual, a rather complex combination of factors conjured to produce this state of affairs – some of computational, some of conceptual and some of purely accidental nature. Cutting a long story to a reasonable length, a few salient features of these early models are worth pointing out in order to appreciate how and why the modern approach came to be. First of all, the Vasicek/CIR models had both a prescriptive and a descriptive dimension. More precisely, if one took as given a certain set of parameters, they showed what the shape of the yield curve should be. If, on the other hand, one left these parameters as free-fitting degrees of freedom, the model would then show what shapes the yield curve could assume. In other words, if applied in a 'fundamental' way (e.g., by estimating econometrically the values of parameters – including the market price of risk!), the model would indicate what the yield curve should look like today; if used in an 'implied' manner, a cross-sectional analysis of prices would give the best combination of parameters capable of producing the closest fit to the observed yield curve. In either case, for a believer in the model, any discrepancy between a market and model bond price signals a cheap or expensive bond, and, therefore, a trading opportunity. For the bond- or LIBOR-option trader, on the other hand, the failure to reproduce the market price would simply mean that the underlying was being mis-priced. In a Black-and-Scholes framework, this was tantamount to being forced to price an option using the wrong input, of all things, for the spot price of the underlying. The assessment of the usefulness of the model was therefore sharply different for the relative-value cash trader and for the option trader: for the former the model,

however crude, had some explanatory power, and could, at least in theory, identify portions of the yield curve as cheap or dear; for the latter, the failure to recover the underlying correctly was too high a price to pay for the much sought-after self-consistency among prices of different options.

Relative-value traders, incidentally, tended to be somewhat intrigued by the fact that models with such a simple (or simplistic?) structure and sweeping assumptions did, after all, a very acceptable job at describing the yield curve. This partial, but encouraging, success was often taken as an indication that ‘at the bottom there must be something right in choosing the short rate as the driving factor for the yield curve’. With the usual benefit of hindsight, one can venture a more prosaic explanation for this ‘intriguing’ degree of success: copious econometric research has shown that the stochastic evolution of the yield curve is explained to a very large extent by its first principal component. The loadings of the different forward rates onto this first principal component are also well known to be approximately constant. Therefore, virtually any rate, and, therefore, in particular, the short rate, could have been taken as a reasonable proxy for the first component, and hence for the yield curve as a whole. This way of explaining the partial success of the early short-rate-based models is not just another curious item in the history of derivatives pricing, but has a direct influence on many of the implementations of the modern approach.

1.1.4 The Second-Generation Yield-Curve Models

The third phase in this brief account of the path that led to the modern methods that are the topic of this book was ushered in by Black, Derman and Toy (BDT) (Black et al. 1990) and by Hull and White (1990) (HW) with their extended Vasicek and extended CIR models. The most salient feature of this class of models was the addition of a purely deterministic (time-dependent) term to the mean-reverting component in the drift of the short rate.⁵ Minor as this feature might appear, it allowed the introduction of a simple but powerful *deus ex machina* capable of disposing of whatever discrepancy the stochastic and mean-reverting components of the short-rate dynamics would leave between market and model bond prices. Therefore, given an arbitrary market yield curve, however twisted and bizarre, the second-generation yield-curve models could always augment the mean-reverting drift with a deterministic ‘correction’ term capable of reproducing the market prices.

Needless to say, the relative-value bond traders and the option traders had at this point to part company, and all the model developments that were to

⁵The BDT model was originally introduced in algorithmic form, and its continuous-time equivalent presented only later (Jamshidian 1991, Hull and White 1990). Given the conceptual equivalence between the algorithmic, discrete-time version and its continuous-time counterpart, and for the sake of brevity, in the discussion of this section I have decided to treat all the yield-curve fitting approaches together, irrespective of whether the fitting itself must be obtained numerically or can be arrived at analytically. The Black and Karasinski (1991) (BK) model therefore also falls in this category.

follow have, if anything, widened their paths. Obviously, it could not be otherwise: a model that could account for any conceivable input shape of the yield curve automatically loses all explanatory power. This selfsame feature, however, makes it attractive for the option trader, who will be able to carry out her model-suggested delta hedging (by buying or selling bonds) at the prices actually encountered in the market.

At the same time, a new important type of relative-value trader became keenly interested in the models of the BDT/HW/BK family. These second-generation models might well have been able to reproduce any yield curve, but could not automatically account for the prices of all plain-vanilla options (caplets and European swaptions) – actually, if implemented as their inventors (correctly) recommended, that is, with constant volatility parameters, they failed to account exactly even for the prices of all the caplets. The explanatory mandate of these models was therefore transferred from accounting for the shape of the yield curve to assessing the reasonableness of the market term structure of volatilities. As a consequence, plain-vanilla LIBOR traders were put for the first time in a position to speculate that, *if* the yield-curve evolution was truly driven by the short rate, *if* its (risk-adjusted) drift was indeed of the prescribed mean-reverting form⁶ and *if* its volatility had been chosen appropriately, then the model could give an indication that the relative prices of, say, two caplets could not be what was observable in the market. In other words, the second-generation models brought about for the first time the possibility of model-driven option-based ‘arbitrage’ trades. The story of these trades, and of the application of models to discover ‘arbitrage’ between plain-vanilla options, is still unfolding, and has been profoundly influenced on the one hand by the development of the modern pricing approach and on the other by the market events that followed the Russia crisis of 1998.

The class of market professionals who were still unhappy with second-generation models were, of course, the exotic traders. By the early 1990s – when the models of the BDT/HW family were mainstream in the financial community, and those houses boasting an ‘HJM model’ often actually had little more than a glorified version of the HW model in place – a variety of more and more complex trades were regularly appearing in the market. Beside the ubiquitous Bermudan swaptions, indexed-principal swaps, ratchet caps, callable inverse floaters, knock-out caps, index accruals, digitals and many other products were continuously being introduced. From personal experience, I feel confident to say that the years between 1990 and 1994 probably saw the highest pace in the introduction of new types of product. For those exotic traders who, like myself, had to carry out their hedges using caplets and European swaptions dealt with their colleagues from the plain-vanilla desks, the greatest desideratum that could at the time be requested of a model was

⁶In the case of the BDT model the mean reversion is rather *sui generis*, and is an artifact of its algorithmic description. See Rebonato (1998, 1999c) for a discussion. Since this distinction is not very relevant to the thrust of the discussion, and for the sake of simplicity, in this exposition I simply refer to ‘mean reversion’ *tout court*, without any further qualifications.

its ability to price at least the required option hedges for each individual trade in line with the plain-vanilla market. Much as plain-vanilla option traders wanted their hedges (bonds and swaps) correctly priced by the model, so exotic traders would have liked the prices of their hedges (caplets and/or swaptions) reproduced by the model in line with the market. And as for the latter, the caplet and European swaption markets back in the early to mid-1990s were still solidly anchored to the Black framework, with no smiles to speak of in all currencies apart from the Japanese Yen. The ‘target to hit’ for exotic traders was therefore both reasonably well defined and tantalizingly close to achieve.

1.1.5 The Modern Pricing Approach

It was approximately at this time that the first non-trivial applications of the HJM model began to be developed by front-office quants and to appear on their trading desks. This is the fourth and latest (but certainly not last) phase in the evolution of interest-rate derivatives pricing that I am describing. A puzzling fact is that the original HJM working paper began to be circulated as early as 1987, yet exceedingly few bona fide implementations were to appear before, approximately, 1993–94. There were several reasons for this delay. First of all the paper was cast in a relatively new language (or, more precisely, in a language that was new for the for the option community): set theory, measure theory and relatively advanced stochastic calculus. Techniques to solve analytically or numerically parabolic linear partial differential equations, which had been the staple diet of the pre-HJM quants, had little application for non-trivial implementations of the new approach. Similarly, recombining-tree techniques, so thoroughly and profitably explored in the wake of the Cox et al. (1979) paper, became of little use to cope with the non-Markovian nature of the log-normal forward-rate processes.

More generally, a whole new vocabulary was introduced to the trading community:⁷ filtrations, martingales, sigma fields, etc., were terms more likely to be familiar to pure and applied statisticians than to the young physicists who had been recruited as rocket scientists at the main investment houses. So, not-so-old dogs suddenly had to learn not-so-new tricks.

Even as the intellectual barrier to entry was being surmounted, a technological impasse was becoming apparent: for meaningful and interesting implementations of the HJM model, closed-form solutions and recombining-tree-based techniques were of very little use. The obvious tool was Monte Carlo simulation. Despite the fact that the application to finance of this technique had been known since the late 1970s (Boyle 1977), the method enjoyed very little popularity, and tended to be regarded as a ‘tool of last resort’, to be used when everything else failed. The fact that the ‘advanced’ variance

⁷The introduction to the financial world of the concepts of martingales and stochastic integrals to characterize conditions of no-arbitrage pre-dates significantly the HJM paper (see Harrison and Kreps 1979, Harrison and Pliska 1981). Before HJM, however, this important work had remained mainly read and appreciated by the academic community.

reduction techniques of the day boiled down to little more than the drawing of antithetic variates and the use of a contravariate variable gives an idea of the rather primitive state of the financial applications of Monte Carlo techniques in the early 1990s. Coincidentally, and luckily for the acceptance of the HJM approach, a significant breakthrough for financial applications occurred in those very same years in the form of high-dimensional low-discrepancy sequences of quasi-random numbers. By their ability to reduce, sometimes by orders of magnitude, the computational time required to perform a simulation, high-dimensional quasi-random numbers contributed to making the HJM approach a practical proposition.

The next big problem to be tackled was the so-called ‘calibration issue’, that is, how to make the HJM model reproduce the prices of the desired plain-vanilla options (i.e., caplets). In those early days the HJM model was perceived to be difficult to calibrate, and a small cottage industry quickly emerged, which began to spin rather cumbersome, and by and large ineffectual, numerical procedures to ensure that the market caplet prices could be recovered by the model.

This might appear extraordinary today, since, after all, one of the greatest advantages of the LIBOR-market-model approach is that it can be made to reproduce the market prices of plain-vanilla options virtually by inspection. If anything, it is the excessive ease with which this calibration can be accomplished that raises problems today (see Parts II and III of this book). The HJM model, however, was originally cast in terms of instantaneous forward rates, which had no obvious equivalent in traded market instruments. Furthermore, the HJM paper mentioned the fact that, in the continuous-time limit and for truly instantaneous and log-normal forward rates, their process explodes with positive probability. Whilst perfectly true and correct, this statement, often repeated and rarely understood, acted as a powerful deterrent against the development of an HJM-based log-normal market model. The fact that log-normal forward rates were considered to be a ‘no-fly zone’ obviously made recovery of the (Black) log-normal plain-vanilla option prices difficult, to say the least.

These fears were actually misplaced: as soon as the process is discretized and the forward rates become of finite tenor, the log-normal explosion disappears. Actually, however one numerically implemented a log-normal forward-rate HJM model (by Monte Carlo simulation, using a bushy tree, or in any other way), one could not have observed the dreaded explosion, hard as one might have tried. The discomfort in moving down the log-normal route was nonetheless both palpable and widespread, and encouraged a line of research devoted to the study of log-normal bond-price-based HJM approaches (which imply approximately normal forward rates). This was to prove a dead alley, but significantly slowed down the development of the standard version of the modern LIBOR-market-based approach.

As for the latter, it is fair to say that, before any of the now-canonical papers appeared, it was simultaneously and independently ‘discovered’ by analysts and practitioners who, undaunted by the expected occurrence of the

log-normal explosion, went ahead and discretized a log-normal forward-rate-based HJM implementation. These practitioners were more slaves of necessity than endowed with visionary foresight. The path they were forced to follow if they wanted to price a LIBOR-based derivative security using the HJM approach, can be approximately reconstructed as follows:

- Given a LIBOR-based path-dependent exotic product, they first of all had to specify the discrete points on the yield curve where the price-sensitive events occurred.
- This partition of the yield curve automatically defined a set of discrete-tenor forward rates. The next logical step was then to assume, plausibly if not rigorously, that these discrete forward rates could be regarded as the finite-time equivalent of the HJM instantaneous forward rates.
- At this point the no-arbitrage drifts for the latter (expressed in the original HJM paper in the form of integrals of volatility functions) had to be translated into their discrete-time equivalent. Not surprisingly, by so doing, integrals turned into summations, and continuous-time stochastic integrals had to be simulated using discrete-time-step stochastic quadrature.
- A numeraire also had to be chosen, and, by a natural extension of HJM's approach, the discretely compounded money-market account was 'invented'.
- Finally, a distributional assumption had to be made for the discrete forward rates. Given that so much in the approach lacked, in those first implementations, rigor and mathematical justifiability, why not try the forbidden explosive log-normal assumption? *If* it worked, after all, log-normality was still the standard for the plain-vanilla market, and the goal of pricing exotic and plain-vanilla options consistently would have been fulfilled.

This 'heuristic' approach did work. No explosion wrecked the computers of those who went ahead and tried the discrete-tenor log-normal implementation, and for the first time one could price *simultaneously* a series of caplets exactly in line with the market *using the same numeraire*.

Sure enough, pricing a series of caplets correctly had been a feat well within the reach of any trader capable of programming the Black formula on her pocket calculator for almost 15 years. But, as discussed above, these Black caplets inhabited separate universes (were priced in different measures), and therefore no other LIBOR security could be priced at the same time in an internally consistent manner. The exotic option trader could now for the first time use her Frankenstein model to price an exotic product (trigger swaps were the flavor of the time) and rest assured that, at the same time, the implied prices of all the 'underlying' caplets would be correct; and that they

would be so, not approximately and by virtue of a possibly dubious choice of model parameters, but virtually by construction.

This account is somewhat stylized and simplified, but captures in a fundamentally correct manner both the genesis and the impetus behind the ‘discovery’ of the modern pricing approach. Needless to say, a lot of work still remained to be done to justify in a rigorous manner the procedure outlined above, and this is the area where the papers by Brace et al. (1995), Jamshidian (1997), Musiela and Rutkowski (1997a) and many others made a very important contribution. The importance of this body of work was much more fundamental than dotting the mathematical ‘i’s and crossing the financial ‘t’s. These papers showed with clarity that any discrete-time implementation of the HJM model, and, therefore, in particular, also the log-normal one, was fully and uniquely specified by the instantaneous volatilities of, and the instantaneous correlations among, the discrete forward rates. More precisely, any discrete-tenor implementation was shown to be fully and uniquely specified by a series of integrals of time-dependent covariance terms. For each time step, that is, for each price-sensitive event, there would correspond one matrix with elements given by integrals of covariance terms; each matrix, in turn, would contain a number of entries proportional to the square of the forward rates still ‘alive’. In its maximum generality, an implementation of the modern approach would therefore require the specification of a number of ‘parameters’ (i.e., the covariance integrals) proportional to the cube of the forward rates in the problem. It was no wonder that the market prices of benchmark plain-vanilla options could, if one wanted, be exactly reproduced!

Actually, this recovery of market prices, so cumbersome with the first- and second-generation models, could now be achieved with disturbing ease, and in an infinity of ways. Unfortunately, each of the possible choices for the instantaneous volatility functions (or for the above-mentioned integrals) would, in general, give rise to different prices for exotic products. Furthermore, if this fitting was injudiciously carried out, it could produce implausible, or even positively pathological, evolutions for such quantities as the term structure of volatilities or the swaption matrix. Much as discussed in the first part of this chapter, what was needed was once again some strong structure that would reduce the degrees of freedom in a systematic and financially transparent way, and yet still preserved the ability to achieve a market fit in a quick and efficient manner. Imposing what I have called above ‘strong structural constraints’ and enforcing internal consistency between different sets of state variables might not bear a great resemblance to choosing a traditional model, but in reality comes extremely close to fulfilling exactly the same function. This task is still an exciting ongoing research program, and constitutes one of the main topics of this book.

A second useful by-product of the formalizations provided in the papers referred to above was the realization that the modern approach, and, a fortiori, the HJM ‘model’, are not truly ‘models’ in the same sense as the HW, the CIR, the Vasicek or the BDT are. Rather, the approaches derived from the

HJM root simply provide conditions on the drifts of the forward rates if arbitrage is to be prevented, and, given a set of chosen instantaneous volatility and correlation functions, express these drifts purely in terms of these functions. It is actually the choice of the parametric form and functional dependence on the state variables of these volatility and correlation inputs that more closely resembles the choice of a ‘model’ in the traditional sense. After all, if any of the ‘traditional’ models were indeed arbitrage-free, they had to be a subset of the admissible HJM models. Therefore, there had to be a suitable specification of the volatility of the forward rates⁸ that would reproduce exactly the traditional model.

This equivalence between the ‘old’ and the ‘new’ approach is actually extremely important. First of all, by embracing the new approach, one can rest assured that, at least in principle, no financially significant feature of a traditional model (such as, for instance, its mean reversion) will have to be abandoned. In other words, all the previously known models, as long as arbitrage-free, had to be just a subset of the HJM family. The modern approach simply provides a new, and often more flexible, vocabulary to describe the characteristics of a given ‘model’.

The second, and arguably far more important, consequence of this equivalence is that the new volatility-based vocabulary forces the trader to express her trading views directly in terms of tradable, or at least market-related, quantities. With the modern approach, the option trader no longer has to ‘guess’ values for, say, the volatility of the consol yield, the mean-reversion level of the instantaneous short rate or the variance of the short-rate volatility, and hope that the opaque mechanism by means of which these inputs are turned into market observables will produce something reasonable. Admittedly, it is true that, with the modern methodology, the user still has to specify quantities not directly observable from the market, such as the instantaneous volatilities of forward rates. It is however possible to translate directly and transparently these choices into trading views about the future evolution of market observables, such as the term structure of volatilities or the swaption matrix. If sufficient liquid instruments were traded in the market⁹ (i.e., if the market in forward-rate volatilities were complete), traders would not be forced to express such views, and could make use of any implementation of the LIBOR market model capable of correctly reproducing all the market prices¹⁰. The market prices of the plain-vanilla instruments (together with the associated serial options) would therefore give rise either to an effectively unique self-consistent parametrization of the LIBOR market model, or to no solution at

⁸Note carefully, however, that this volatility would not in general be deterministic. See the discussion in Section 2.3.

⁹Serial options complete the instantaneous-volatility market, but, depending on the currency, they are either illiquid or confined to the very short end of the yield curve. Serial options, and how they can be used to complete the market in instantaneous forward-rate volatilities, are discussed in Section 2.2 and in Rebonato (1999c).

¹⁰See the discussion in Section 11.1.

all. If any one market price could not be fitted (and assuming, of course, the distributional assumptions of the model to be correct), in this universe endowed with complete-instantaneous-volatility markets the trader could then put in place a replicating strategy capable of arbitraging away the offending plain-vanilla prices.¹¹

Unfortunately, the instruments (serial options) that would be required to complete the market imperfectly spanned by caplets and European swaptions are far too illiquid and sparsely quoted to constitute a reliable market benchmark. Given this predicament, expressing a view about the plausibility of the model-implied behavior for such market observables as the term structure of volatilities becomes the main tool in the hands of the trader to establish whether the model parametrization being used is reasonable. This congruence between the trader's views and the model-implied evolution of the market is far more fundamental than a simple 'sanity check': any arbitrage-free model implicitly presupposes a self-financing dynamic trading strategy capable of reproducing the terminal payoff of the derivative product. This strategy, in turn, implies future transactions in the 'underlying instruments' (which, in the case of an exotic derivative product, are caplets and swaptions) with no net injection or withdrawal of cash. If a poorly calibrated model assumes an unrealistic future term structure of volatility or swaption matrix, it will produce the wrong set-up cost for the initial replicating portfolio. This, after all, is just a fancy way of saying that, by assuming the wrong future re-hedging costs, it will predict the 'wrong' price today for the exotic product.

Another important bonus of the modern pricing approach is that it can be equivalently cast in terms of (discrete) forward rates or of swap rates. This has two important positive consequences: first of all, given the ease with which the market-implied volatilities of the state variables can be recovered, the user can rest assured that at least one set of hedging instruments (caplets or swaptions) will be exactly priced by the model. It will then be up to user to decide whether recovery of the volatility of forward or swap rates is more relevant for the pricing of a given particular exotic product (see Jamshidian (1997) for a nice discussion of this point in the context of trigger swaps). At the same time, most complex exotic products require hedging positions both in caplets and in swaptions. It is therefore extremely useful to ascertain what swaption matrix is implied by a given forward-rate-based application, or what term structure of caplet volatilities is produced by the chosen swap-rate-based implementation.

More fundamentally, the modern approach provides the trader with the perfect tools to analyze the congruence between the two benchmark plain-vanilla option markets (caplets and swaptions). Needless to say, for a religious

¹¹This statement requires careful qualification, since a finite number of instruments cannot complete the market in instantaneous volatilities, but can at most complete the market in instruments whose value depends on integrals over finite time intervals of the instantaneous volatility. In other words, if serial options, caplets and swaptions were simultaneously traded, one could conceptually create instruments that paid a unit certain payoff if one particular realization of each finite-time-interval integral of the instantaneous volatility occurred. This aspect is covered in Chapter 4.

believer in informationally perfectly efficient prices, markets are always congruent, and the inability to price simultaneously caplets and swaptions simply points to an inadequacy of the model used for the task. Alternatively, if the model is thought to be correct and it does manage to reproduce all the prices, the resulting volatility and correlations must reflect the market consensus about these quantities, no matter how implausible these might appear if compared, for instance, with historical and statistical evidence.

In reality, the picture is considerably more complex: on the one hand, there are the natural market flows and the actions of agents with preferred habitats who create an 'excess' demand or supply for certain products. The most obvious examples of this imbalance are, to name just a few: the 'natural' intrinsic mismatch between demand for caplet optionality (from liability managers) and the supply of swaption optionality (from issuers and investors in search of better-than-plain-vanilla funding rates or investment yields);¹² the recent and exceptional demand from British pension funds for long-dated GBP swaptions due to their need to provide fixed-rate annuities; the behavior of US investors who want to hedge prepayment risk inherent in their mortgage-backed securities; etc.

On the other side of the equation, such imbalances of supply and demand should in theory be ironed away by arbitrageurs, proprietary and relative-value traders who do not have a preferred habitat and can indifferently take either side of the market. However, I will argue in the next section that any inefficiency and market constraint reduces the ability of the arbitrageurs to exploit price discrepancies caused by an unbalanced supply or demand and to bring with their position-taking related markets in line with each other. In the caplet and swaption markets these imperfections and constraints are still abundant (and, if anything, they are likely to have increased in the past few years). For instance, a generalized reduction in the scale of the activities of proprietary relative-value arbitrage desks after the Russia events has had a detrimental effect on market liquidity. With poorer liquidity, it has become relatively easier to 'move the market' with a few well-placed large trades. This, in turn, has discouraged the activity of traders who, on the basis of a model, perceive a certain portion of, say, the swaption matrix cheap or dear relative to the underlying caplets. In this landscape of relatively poor liquidity, the rigorous discipline of marking positions to market has made the weathering of P&L storms in relative-value trades sometimes extremely painfully. The magnitude of 'temporary'¹³ losses has often forced the closing out of positions, even when models, common-sense and trading experience would indicate that, given enough time, the mark-to-market losses would 'eventually' be reversed.

The bottom line of this digression is that a blind and unquestioning belief

¹²See, in this respect, the discussion in Section 9.1.

¹³If relative-value trades were put on in, say, long-dated swaptions, the 'temporary' period by the end of which the losses will, in all likelihood, be reversed (if the positions were held and model-hedged) could be as long as a decade or more. This, on the time-scale of a typical investment bank, can easily correspond to several generations of traders.

in the congruence of the caplet and swaption market requires a significant act of faith. This state of affairs creates for the trader a difficult situation, since the information from the two sister markets can neither be ignored nor fully accepted at face value. The issue of the joint analysis of these two markets, and of its impact on the calibration of the market model, is one of the main topics treated in this book, and is revisited in the final section of this chapter.

This somewhat stylized, but fundamentally accurate, account of the developments in interest-rate derivative pricing has brought us virtually to the current state of affairs. As in every good plot, there is a twist at the end: just when the market model has achieved the ability to value exotic products while, at the same time, correctly pricing the Black-model-driven plain-vanilla options, the market standard has begun to move resolutely away from the log-normal paradigm. This departure has been signalled by the appearance of smirks and smiles in the implied volatility curves. These plots, which, for a given maturity, should be exact straight lines as a function of strike if the log-normal assumption held true, first began to assume a monotonically decreasing shape as early as 1995–96 (see, e.g., Rebonato 1999c); after the market events that followed the summer/autumn of 1998, a hockey-stick smile shape then began to appear. This distinction is important, because I will make the point, in Part IV of this book, that different financial mechanisms are at play in producing these two features. I shall therefore argue that, if one wants to account for these distinct mechanisms in a financially convincing manner, different and simultaneous modelling routes are necessary. The first requires only relatively minor tinkering at the edges of the market model; the second calls for far more radical surgery.

In closing this introductory section, I would like to add two more remarks. First, it is essential to point out that the LIBOR market model as it is known today is much more than a set of equations for the no-arbitrage evolution of forward or swap rates: it includes a very rich body of calibration procedures and of approximate but very accurate numerical techniques for the evolution of the forward rates that have turned the approach into today's most popular pricing tool for complex interest-rate derivatives. Second, these calibration and approximation techniques have turned out to be, by and large, extremely simple.

Subtle as the reasons for its success might be, the fact therefore remains that, once the modern approach is properly implemented and calibrated, very complex computational tasks can be carried out with ease and in real trading time. (These are indeed the topics that I cover in Parts I–III of this book.) This state of affairs, however, makes extending the LIBOR market model in such a way that the observed smiles can be accounted for in a financially convincing way a very tall order. Not only must the resulting equations make 'financial sense', but the established calibration and evolution results must also be recovered, if not totally, at least to a significant extent. The treatment to be found in Part IV of the book has therefore been informed by these joint requirements of financial plausibility and ease of practical implementation for

the calibration and evolution techniques that will be presented in the chapters to follow.

These topics arguably constitute the most exciting areas of current development in derivatives pricing. Rather than labelling this activity as the ‘fourth phase’ of model evolution, I prefer to post a brightly painted sign with the words ‘Work in progress’, and to wait and see which route the financial community will take in the years to come.

1.2 Some Important Remarks

The LIBOR market model allows the user to obtain *simultaneously* the correct prices of exogenous sets of (i) discount bonds, (ii) caplets and (iii) European swaptions. The recovery of the discount curve is virtually built into the construction of the model, and therefore comes, as it were, ‘for free’. The ability to fit almost exactly to the two other asset classes, however, would appear to be a very important and desirable feature. I shall argue in the remainder of this book it will be argued that forcing the chosen model implementation to yield *simultaneously* the market prices of caplets and European swaptions is in most cases not desirable. I shall also repeatedly invite the reader to check the financial ‘reasonableness’ of the chosen parametrization.

Since caplet and swaptions are made up of the same building blocks (forward rates), both recommendations smack of heresy: Doesn’t failure to recover market prices expose the trader to the possibility of being at the receiving end of arbitrage trades? Isn’t the ‘reasonableness’ of a financial quantity irrelevant in derivatives pricing, given that the trader can ‘lock in’ via dynamic trading the implied values, however ‘unreasonable’?

The answer to both questions does not so much lie in the ‘in theory/in practice’ dichotomy, as in modelling a given financial phenomenon within theoretical frameworks of different scope and generality. The bedrock of the whole analysis is to be found in the concept of no-arbitrage and in the efficient market hypothesis and its corollaries. Arbitrage, which can be defined in this context as the ‘simultaneous purchase and sale of the same, *or essentially similar*, security ... for advantageously different prices’ (Sharpe and Alexander 1990; my emphasis), ensures that prices remain anchored to fundamentals, and, ultimately, provides the strongest argument for market efficiency. Recall, in fact, that the validity of the efficient market hypothesis does not require investors’ rationality (see, e.g., Shleifer 2000). It does not even require that the actions of the ‘noise’ traders should be uncorrelated. It does require, however, the existence of arbitrageurs capable of exploiting the moves away from fundamental value brought about by the (possibly correlated) noise traders. By virtue of the actions of the arbitrageurs, the values of financial instruments can never stray too far from fundamentals, and redundant instruments (as either caplets or swaptions appear to be in the LIBOR market model) must display congruous prices. When it comes to these LIBOR derivatives products, how-

ever, I shall argue that these quasi-arbitrage trades are in reality difficult and risky, *even if the trader could assume the underlying deterministic volatilities and correlations to be known with certainty*. In other words, even neglecting the market incompleteness that arises when these functions are imperfectly known, locking in the prices implied by the complementary sets of plain-vanilla options is far from easy. The ‘accidental’ factors that make this quasi-arbitrage difficult are several. Let me recall a few:

- *Agency problems* – The standard analyses in asset pricing assume that investors allocate their own wealth among the various available securities on the basis of their risk aversion. However, traders in general, and arbitrage traders in particular, typically do not act as principals, but as agents of investment houses (banks, hedge funds, etc.). As such, they are evaluated and compensated by managers or investors who do not share their market knowledge, and must base the appraisal of the traders’ performance on their results over a relatively short period of time (see, e.g., Shleifer and Vishny 1997). A fundamental feature of this activity is therefore that ‘brains and resources are separated by an agency relationship’ (Shleifer 2000). If the strategy an arbitrageur has put in place is likely to show the market wrong only over a time-scale longer than the next bonus date, the trader might have to unwind the ‘correct’ positions before time.
- *Limit structures* – Arbitrage trades of this type typically reside on the trading book of a financial institution. Before the market can be shown the errors of its ways, large mark-to-market swings, possibly in the ‘wrong’ direction, can occur. Internal or regulatory limit losses may be hit, Value-at-Risk (VaR) limits breached and the positions might have to be prematurely unwound.
- *Market liquidity* – All of the above is more likely to happen in not very liquid and/or deep markets. The temptations for competing market participants to ‘corner’ the arbitrageur, once her positions are known or suspected, is increased by market illiquidity. This feature is compounded by the fact that these pseudo-arbitrage trades, to be ‘worth the trouble’, typically require very large notionals.

Indeed, the events alluded to in the previous section, connected with the caplet/swaption arbitrage trades of 1998, can be explained by a combination of all these three sources of market imperfections (or, one might say, of market reality). In my opinion, and to the extent that I have been able to ascertain the details of those trades, I can say that they were ‘right’. In this context the word ‘right’ means that they would have probably brought the trader substantial gains if ‘only’ they could have been kept on for the best part of 10 years; if the market had not so violently swung against the holders of the positions causing very large mark-to-market losses; if the market had not been at the

time so illiquid as to lend itself to manipulation; if the arbitrageurs had not been under the constraints of VaR and stop-loss limits; if the traders thought that their compensation, and the continuation of their gainful employment, would be decided on the basis of their (very-)long-term performance. As it happened, none of the above applied, the trades had to be unwound (at a substantial loss), and the caplet and swaption markets remained out of line.

So, is the efficient market hypothesis wrong? If the question is meant in the sense ‘Are some of its assumptions incorrect or unrealistic?’, the answer is obviously ‘yes’. But, in this sense, *any* model is wrong – if it were not, it would not be a model in the first place. The more pertinent interpretation of the question above is instead: ‘Are some of the assumptions so crude as to invalidate the results in some interesting cases?’ If the question is understood in this sense, the answer is probably: ‘It depends.’ When speaking of a model I find it more profitable to think in terms of its appropriateness to a problem than of its correctness. The efficient-market framework has been proven to be extremely powerful and useful, and it provides a very robust tool of analysis in a wide range of applications. This does not imply, however, that the conceptualizations it affords, and the conclusions that can be derived from it, should be applied without question to every asset class and in every market condition. This is exactly what is done, however, when the model parametrization of the LIBOR market model is so chosen as to enforce perfect simultaneous pricing of caplets and swaptions. Given the difficulty of carrying out the ‘arbitrage’ trade that would bring their values in line with each other, I believe that assuming that the efficient market hypothesis should hold in the case of these instruments is unwarranted.

Does it matter? I believe it does, because the ‘contortions’ imposed onto the model by this over-perfect parametrization can produce very undesirable pricing and hedging effects. So, if a financially convincing parametrization of the LIBOR market model can be found that ‘naturally’ prices swaptions and caplets well, so much the better. If this is not possible, the trader will have to ask herself what future hedging instruments she is most likely to use during the course of the trade, and to ensure that their future prices (as predicted by the model) are not too much at variance with the real ones.

Isn’t the whole approach, then, based as it is on dynamic replication of arbitrary payoffs, self-contradictory? It is only so for the believer in ‘true’ and ‘false’ models. Perfect payoff replication might be impossible, but a considerable degree of hedging can certainly be achieved. The trader who uses models as tools to aid the analysis will not be upset by this, and will know that using a model can be a useful crutch up to a point, but can become a hindrance if pushed too far.

In this book I shall analyze in detail additional and more technical causes of the near-impossibility to enforce via quasi-arbitrage trades the exact congruence between the two markets. These are linked to the fact that, even if the volatility and correlation functions that describe the model were deterministic, they are not perfectly known by the trader. Because of this lack of

perfect information, the market is not complete. The difficulties are then compounded by the fact that, in reality, volatilities are unlikely to be deterministic, and, if stochastic, do not even appear to be describable in terms of simple diffusions (see Chapter 13). If this is true, supply and demand can, and in my opinion do, drive a wedge between the theoretical prices of caplets and swaptions. The main message that I tried to convey in this section will be expanded upon and will reappear as a recurrent theme throughout the book.