

Term-Structure Models: a Review

Riccardo Rebonato
QUARC (QUAntitative Research Centre) - Royal Bank of Scotland
Oxford University - OCIAM

February 25, 2003

1 Introduction

1.1 Justification for Another Review of Interest-Rate Models

The topic of term-structure modelling for derivatives pricing has been covered in recent years, at least in book form (see, eg, [Rebonato (1998)], [James and Webber (2000)], [Hughston (2000)], [Brigo and Mercurio (2001)]). One obvious justification for updating these works is that the modelling of interest rates is still rapidly evolving. There are, however, deeper reasons why a fresh examination of the current state of research in this area deserves attention.

The first is that recently a qualitatively new dimension has been added to the modelling complexity, because of the appearance of pronounced and complex smiles in the implied volatility surfaces of caplets and swaptions. Sure enough, smiles of sorts were produced by some of the early modelling approaches (eg, by the [Hull and White (1993)] model), but these were at the time obtained as an 'afterthought', and by and large regarded as an unpleasant feature to be ignored or excused away. Currently, the modelling of smiles is possibly the most active area of research in interest-rate derivatives pricing, and a sufficiently substantial body of work has accumulated to warrant a review of its achievements and of the problems still to be tackled.

The second reason why a review of term structure modelling is timely is connected with two related market developments, namely the compression in profit margins for the well-established complex products and the simultaneous increase in complexity of many of the new-generation interest-rate derivatives. These joint market developments are putting under severe strain some of the basic underpinning theoretical concepts of the classic pricing approach, which ultimately derives from the [Black and Scholes (1973)] (BS in the following) paradigm (although, often, via the [Harrison and Pliska (1981)] route). In particular, given the robustness of the BS model, market completeness and exact payoff replicability (both reasonable but inaccurate descriptions of financial reality) can constitute a valid framework when one is pricing, say, a simple cap

(and the profit margin is generous). However, the wisdom of adopting the same pricing approach (which implies irrelevance of attitudes towards risk) should at least be questioned if the trader is pricing a product with payoff depending on, say, the joint realizations of two currencies and of a non-particularly liquid currency pair over a 30-year period¹.

More generally, for most problems in asset pricing a mixture of the relative and absolute approaches (see, eg, [Cochrane (2000)]) tends to be more profitable. Traditionally, derivatives pricing has been the area of asset pricing where the relative approach has been most successful, and its conceptual corollaries (such as irrelevance of attitudes towards risk) have therefore been extended to more and more complex products. This trend, however, should not be regarded as inevitable or invariably desirable, and I intend to provide in this review an opportunity for a debate.

There is, I believe, a third reason why a review of interest-rate models is timely. There is an implicit but heavy reliance of current pricing practice on market efficiency. This reliance becomes apparent in the estimation of the quantities to which a model has to be calibrated (eg, instantaneous volatilities or correlations). The prevailing current philosophy appears to be that it is better to 'imply' the inputs to a model from the prices of plain-vanilla traded instruments (eg, the correlation among forward rates from the prices of European swaptions), than using statistical estimation. In this approach, the simultaneous recovery of as many market prices (of bonds, caplets, European swaptions, etc) as possible is therefore intrinsically desirable. The validity of relying purely on market-implied quantities, however, hinges either on exact payoff replicability, and, therefore, on market completeness, or on the informational efficiency of derivatives markets. The validity of this approach, however, should not be accepted as self-evident. I therefore provide in the last section a re-assessment along these lines of the current prevailing market practice.

To avoid repetition, I shall treat relatively quickly material already well covered elsewhere, and place more emphasis on the new angles mentioned above. This, however, should not provide a skewed view of interest-rate derivatives pricing: the efficient market hypothesis, market completeness, payoff replication and deterministic volatilities are still an excellent and reliable starting point for derivatives pricing. However, they are not its beginning and end.

1.2 Plan of the Survey

Review papers typically either take a historical approach, or prefer to analyze the status of the research in a given topic in a more abstract and 'de-contextualized' manner. I have chosen to present a strongly historically-based

¹The product referred to is a power reverse dual swap, which can be seen as the compound option to exchange a series of FX calls for a stream of floating payments (see, eg, [Del Bano (2002)] for a description). It is interesting to note that the market in the underlying (very-long-dated FX options, in US\$/Yen for the first trades, and then AUD/Yen) was virtually created by power reverse duals. The liquidity implication of this should not be underestimated (see, again, [Del Bano (2002)]), yet the standard pricing approach assumes availability and perfect liquidity for the FX options.

account of term structure modelling because I believe that the current state of research has been largely motivated by its history, and by almost 'accidental' industry and technological developments.

Furthermore, some of the practices of the users of models (the traders) might well be theoretically questionable or unsound (eg, the day-by-day recalibration of a model before re-hedging). Nonetheless, prices, that models strive not only to produce but also to recover, are created by traders who (correctly or incorrectly) use models. As a consequence, understanding this complex dynamic process is essential in order to appreciate the evolution of modelling, and to understand why certain models have been more successful than other.

Given the chosen approach, the second section of this survey puts in perspective those aspects of trading practice and of the industry developments that are of relevance to current interest-rate derivatives modelling.

Despite this historically-based approach, I have not presented the derivations following their original formulations (much as today one would probably not present classical mechanics using Newton's *Principia*). In order to deal with the various models within a conceptually coherent analytical framework, I present instead in Section 3 a unified treatment of the 'classic' pricing approach by payoff replication in complete markets. In so doing I have drawn extensively from [Hughston (2000)] and [Hughston and Brody (2000)]. Many of the ideas can also be found in the earlier work by [Gustafsson (1992)]. As the various models then appear on the scene, I shall then refer to Section 3 for their theoretical underpinning, and comment on their specific features (reasons for introduction, calibration issues, etc) as I introduce them.

This historical account is subdivided into seven phases (Sections 4 to 10), the latest of which deals with the models specifically designed to deal with volatility smiles.

Section 11 of my survey deals with the issue of how, and to what extent, models should be calibrated to market prices of plain-vanilla derivatives. This debate is still very open, and, to argue my point, which is somewhat at odds with the prevalent market and academic practice, I discuss market completeness and the difficulties that pseudo-arbitrageurs might encounter in bringing different sets of prices in line with fundamentals, and coherent with each other.

Finally, the last section provides a summary, perspectives on the likely evolution of modelling and suggests a comparison with areas of asset pricing where different approaches are currently being used.

As for the mathematical treatment, I have kept the mention of the technical and regularity conditions to a minimum, because excellent references already exist that fulfill this task: see, eg, [Jamshidian (1997)] and [Musiela and Rutkowski (1997)] for the LIBOR market model, [Hunt and Kennedy (2000)] for Markov-functional interest-rate models, [Rogers (1997)] for the potential approach and [Kennedy (1997)] for Gaussian models.

2 Market Practice

2.1 Different Users of Term-Structure Models

To understand pricing models it is essential to grasp how they are put to use. In general, term-structure models can be of interest to relative-value bond traders, to plain-vanilla-option traders and to complex-derivatives traders. Relative-value (eg, coupon-to-coupon) bond traders are interested in models because, if a given description of the yield curve dynamics were correct, failure of the model to fit the market bond prices exactly could indicate trading opportunities. Similarly, plain-vanilla option traders are interested in understanding how a given model 'sees' the implied volatility of different swaptions and/or caplets, and compare these predictions with the market data. For these users, the inability to recover the observed market price of a caplet or swaption is not necessarily a failure of the model, but could indicate a trading opportunity. For both these classes of user, models should therefore have not just a *descriptive*, but also a *prescriptive* dimension.

The situation is different for complex-derivative traders, who do not have access to readily visible market prices for the structured products they trade in, and therefore require the models to 'create' a price for the exotic product, given the observable market inputs for the bond prices and the plain-vanilla implied volatilities. Since complex traders will, in general, vega hedge their positions (see Section 2.4), exact recovery of the plain-vanilla hedging instruments (caplets and European swaptions²) - the descriptive aspect of a model - becomes paramount.

Term-structure models have therefore been used in different ways by different classes of market participant, and complex traders have become the largest group. As a consequence, these days few traders rely on equilibrium or arbitrage-free models to assess the relative value of traded bonds, and the emphasis has strongly shifted towards the pricing of complex products, which requires an accurate description of (rather than prescription for) the underlying inputs (bond prices and yield volatilities).

This state of affairs is closely linked to the type of econometric information that is deemed to be of relevance to construct and validate the models. This is touched upon in the next section.

2.2 Statistical Information, Hedging Practice and Model Evolution

Recent econometric research in the derivatives area has been more focussed on the statistical behaviour of volatilities (see, eg, [Derman and Kamal (1997)], [Derman (1999)], [Alexander (2000)], [Alexander (2001)], [Rebonato and Joshi (2002)], [Rebonato (2003)]) than on the arguably more fundamental topic of the dynamics of the underlying rates. This can be understood as follows.

²For brevity, unless otherwise stated, the term 'swaptions' will always refer to European swaptions.

In the derivatives area at least two possible strands of empirical analysis, which I will call 'fundamental' and 'derived' or 'relative', coexist. In the fundamental approach, it is recognized that the value of derivative contracts is ultimately derived from the dynamics of the underlying assets, and therefore the focus is on the statistical properties of the latter. The approach should be of great interest to the plain-vanilla trader, for whom the underlying asset is the main hedging instrument. See, eg, [Bouchaud et al. (1998)] or [Cont (2001)] in the equity area, or the references in Section 5.1 for interest rates. Models such as the Variance Gamma, which boasts a better description of the dynamics of the underlying than traditional diffusive approaches (see, eg, [Madan, Carr and Chang (1998)]), or the Levy extension of the LIBOR market model ([Eberlein and Raible (1999)]) fall in this category.

The derived approach is of greater interest to the complex-derivative trader, for whom not only the underlying asset, but also other plain-vanilla options constitute the set of hedging instruments. The dynamics of option prices (or, more commonly, of the implied volatilities) therefore become as, if not more, important, and mis-specification of the dynamics of the underlying asset is often considered to be a 'second-order effect'. This is because, typically, the complex trader will engage in substantial vega and/or gamma hedging of the complex product using plain-vanilla options. Once a suitable hedge has been put in place, it is often found that the delta exposure of the complex instrument and of the plain-vanilla hedging options to a large extent cancel out. When this is the case, the *net* dependence of the whole portfolio on the underlying (the delta) can be relatively modest, and consistent errors in the estimation of the delta exposure of the complex product and of the plain-vanilla hedging options can compensate each other. To some extent, the complex traders who has vega-hedged her positions shifts (part of) the risk of a poorly specified dynamics for the underlying to the book of the plain-vanilla trader³.

The evolution in the use of models from relative-value/plain-vanilla traders to complex-derivatives traders alluded to in the previous section has therefore had several effect of relevance for the topic of this survey: first, it has influenced the type of statistical analysis deemed to be of relevance by most industry practitioners. This is great relevance for the calibration practice, discussed in Section 11. Second, it has shifted the emphasis from models that can better account for the statistical properties of the underlyings to models that better capture important features (such as time homogeneity, smile dynamics, etc) of the smile surface. Finally, it has created a situation whereby it is difficult for a model that 'simply' better accounts for the statistical properties of the underlyings to supplant a model that better recovers (or can be more easily fitted to) the prices of plain-vanilla instruments.

The fundamental reason for this is that, unlike natural phenomena, the plain-

³This statement is clearly very imprecise and difficult to quantify, since a mis-specification of the dynamics for the underlying will also give rise to an incorrect estimate of the vega. These inconsistencies are however integral part of the use of models: the trader who uses a deterministic-volatility model, for instance, would have no theoretical justification for vega-hedging her positions.

vanilla prices that constitute the 'observable inputs' for the complex models are in turn obtained using models that the more complex approaches are made to recover as limiting cases. It is therefore not surprising that a certain degree of 'self-fulfilling prophesising' should be observed in the pricing of derivatives products. This is not to say that, once a modelling practice becomes established, a vicious circle of self-corroboration establishes itself, and that it becomes impossible to dislodge the perhaps erroneous model choice. However a certain degree of positive feedback and circularity certainly exists among plain-vanilla prices, models, calibration procedures and complex prices, which makes the embracing of radically new, and possibly better, modelling approaches rather difficult.

While today's models are indubitably more effective (at least in certain respects) than the early ones, and, therefore, one can certainly speak of an 'evolution' of interest-rate models, this notion of progress does not necessary imply that certain choices, abandoned in the past, might not have ultimately led a more fruitful description of interest-rate derivatives products if they had been pursued more actively. In other words, I do not believe that the 'linear evolution' paradigm, possibly applicable to some areas of physics⁴, and according to which later models are 'better' in the sense of being closer to the 'phenomenon', is necessarily suited to describing the evolution of yield-curve models.

2.3 Reasons for Inertia to Model Changes

A stylized account of the interaction between plain-vanilla prices, traders of structured products and pricing models might help to make the statements above more concrete. Typically, a complex trader will enter a transaction with a non-professional counterparty, and will attempt to hedge her exposure by transacting hedging trades with the plain-vanilla derivatives desk. These hedging transactions will attempt to neutralize both the exposure to interest rates, and to their volatility (vega exposure). In order to arrive at the price of the complex product the exotic trader will typically make use of a suitably calibrated model. 'Suitably calibrated' in this context means that the prices of plain-vanilla options (caplets and swaptions) used for hedging should be correctly recovered. Caplets and swaptions therefore enjoy a special (benchmark) status among the interest-rate products.

What gives rise to the 'model inertia' referred to above? What discourages an exotic trader who believed her model to describe financial reality better than the current market practice from choosing to 'swim against the tide', ie, from accepting an initially unfavourable marking of her positions in the plain-vanilla hedging instruments, in the hope that the theoretical profits will ultimately (by option expiry) accrue to her?

There are several reasons for this. To begin with, if, according to the 'better'

⁴I realize that, with this statement, I am walking into a philosophical thicket, and that some philosophers of science would deny even models in fundamental physics any claim of being 'true' in an absolute sense. I stay clear of this controversy, and the more mundane points I make about the evolution of yield-curve models remain valid irrespective if whether an absolute or a 'social' view of scientific progress is more valid.

model, the trader is delta and vega neutral, she will, in general, not be so according to the market standard practice (and, therefore, to the risk management assessment of her trading house). Therefore she will utilize a higher proportion of the internal and regulatory (VaR) limits. [Mode risk refs here] If the trader works for a regulated institution, the regulatory capital charge will be higher. So will the economic capital utilization, and, if the traders' compensation is based on return on regulatory or economic capital, her performance will be assessed in a less favourable light. Furthermore, a dynamic model reserve might be applied to the marking of her positions, which, for a book of approximately constant size, effectively translates into a 'tax' on her trading activity.

This state of affairs generates a disincentive for the exotic trader to question the prices of plain-vanilla instruments and contributes to the 'self-fulfilling prophecies' I referred to above. I shall make the point in my review that, as a consequence of this, the natural desire for the exotic trader to reproduce almost 'at all costs' the prices of caplets and swaptions has guided a large part of the theoretical and practical development of interest-rate models.

2.4 In-Model and Out-of-Model Hedging

No aspect of derivatives trading has stronger implications on what is required of a desirable model than the joint practices of out-of-model hedging and model recalibration. In-model hedging refers to the practice to hedge a complex option by taking positions in 'delta' amounts of traded instruments to neutralize the uncertainty from the stochastic drivers. Out-of-model hedging is the taking of positions to neutralize the sensitivity of a complex product to variations in input quantities that the model assumes deterministic (eg, volatility). Vega hedging is an example of out-of-model hedging.

Despite the fact that out-of-model hedging is on conceptually rather shaky grounds (if the volatility is deterministic and known, as most models assume, there would be no need to undertake vega hedging), its adoption is universal in the trading community. Similarly universal and difficult to justify theoretically is the practice of re-calibrating a model to the current market plain-vanilla prices throughout the life of the complex trade.

These two practices are closely linked. In a friction-less market, if a model did not have to be re-calibrated during its life, future vega transactions would have no economic impact: contingent on a particular realization of the forward rates, these trades would be transacted at the future conditional prices for the plain-vanilla hedging instruments implicit in the day-0 calibration. This is no longer true, however, if, in order to recover the future spot plain-vanilla prices, the model has to be re-calibrated day after day. In particular, if the future conditional prices predicted at time 0 for the plain-vanilla options used for vega-hedging turn out to be different from the future prices actually encountered, an extra additional cost (not accounted for by the model) may be incurred.

Choosing good inputs to a model therefore means recovering today's prices in such a way that tomorrow's volatilities and correlations, as predicted by the model, will produce future plain-vanilla option prices as similar as possible to

what will be encountered in the market. Given the joint practices of vega-hedging and re-calibration, the 'best' calibration methodology is therefore the one that will require as little future re-estimation of the model parameters as possible.

The question however remains of how the model inputs that will give rise to the more stable calibration and to the smallest re-hedging 'surprises' should be estimated. Answering this fundamental question requires choosing the source of information (statistical analysis or market 'implication') that can best serve the trading practice. This topic is dealt with in Section 11.

3 A Unified Modelling Approach

3.1 The Mathematical Setting

One of the most obvious requirements for the modelling of bond prices or forward rates is that they should be strictly positive in all possible states of the world. Strictly positive (semi-)martingales therefore provide a natural mathematical description. There are deeper reasons, however, for wanting to model bond prices and rates as semi-martingales, namely the fact that (strictly positive) semi-martingales remain invariant under a change between equivalent measures, and that they can be uniquely decomposed (Doob-Meyer decomposition) into the sum of a pure martingale component (the 'innovation') and a previsible process (the 'predictable' part): consider on a standard filtered probability space⁵ $(\Omega, \mathcal{F}_t, P)$ (see, eg, [Dothan (1990)]) the set \mathcal{SM} of strictly positive semi-martingales, sm_t , such that

$$sm_t = sm_0 + m + a \tag{1}$$

with $m \in \mathcal{M}$ (the set of real-valued martingales, and $a \in \mathcal{A}$ (the set of previsible, real-valued, finite-variation processes. ($\mathcal{SM}^+ \subset \mathcal{SM}$ will then denote the set of *strictly positive* semi-martingales). One can then decompose a price or rate change in a previsible component, and an innovation ('surprise') component linked to the arrival of unexpected market information.

As for the innovation part, different choices could be made: the process could be chosen to have continuous or discontinuous paths; it could be given finite first variation and infinity activity (such as, for instance, the Variance Gama process, [Madan, Carr and Chang (1998)]). Given the legacy of the early models, and the desire for analytic tractability, the prevalent practice has been to choose for the innovation continuous martingales, m_X , with finite quadratic variation. This can be satisfied if the innovations are chosen to be Wiener processes. Furthermore, since (Levy theorem) any Wiener process relative to a filtration \mathcal{F}_t is a Brownian motion, one can simply write:

$$dm_{X_i}(t) = \sigma_i dw_i(t) \tag{2}$$

⁵Among the many hypotheses required for the set-up to produce the results derived below one should mention that the filtration \mathcal{F} is assumed to be known by all agents and that markets are perfectly liquid, frictionless, and with free and public information.

where $dm_{X_i}(t)$ is the innovation associated with the i -th forward rate, $dw_i(t)$ the increment of a standard Brownian motion, and the responsiveness of the i -th innovation to the Brownian shock (the volatility) is denoted by σ_i .

3.2 Different Descriptions

Given the above, a natural way to model (in the real world) the price process, P_t^T , for a bond of maturity T is to require that it should be an element of \mathcal{SM}^+ (a *strictly positive* semi-martingale). With the assumptions above regarding the nature of the innovation, one is led to write for its real-world dynamics⁶

$$\frac{dP_t^T}{P_t^T} = \mu_t^T dt + \sum_{k=1,n} v_k(t, T) dw_k^T = \mu_t^T dt + v(t, T) dw^T \quad (3)$$

$$v(t, T)^2 = \sum_{k=1,n} |v_k(t, T)|^2 \quad (4)$$

From a standard replication argument, similar in spirit to BS's, and first presented in [Vasicek (1977)], it is well known that, in the one-factor case, the extra return over the riskless rate per unit volatility must be independent of the security to avoid arbitrage: for any two securities, a and b , either 'primitive' or 'derivatives', of respective volatilities v_a and v_b , it must hold that

$$\frac{\mu_a - r}{v_a} = \frac{\mu_b - r}{v_b} = \lambda \quad \forall a, b \quad (5)$$

Therefore Equation 5 unambiguously *defines* the market price of risk in the one-factor case:

$$\mu_t^T = r + \lambda^t v(t, T) \quad (6)$$

and the decomposition 6 is unique. In order to extend the definition to many factors, let us assume the number of bond maturities to be finite. The market will be said to be complete if the rank of the matrix $v_k(t, T_i)$ is equal to n , ie, if any Brownian shock is reflected in the change in price of at least one bond. Therefore, if we allow no degeneracies in the real symmetric matrix ρ_{ij}

$$\rho_{ij} = \sum_{k=1,n} v_k(t, T_i) v_k(t, T_j) \quad (7)$$

a purely deterministic portfolio can always be created, by taking suitable positions in the underlying bonds, such that the coefficients of *each* Brownian motion should be exactly zero. Since the return from such a deterministic portfolio must, by no-arbitrage, be equal to the riskless rate, r , a vector $\lambda_k^t, k = 1, 2, \dots, n$,

⁶I will always assume that the required integrability and technical conditions for the volatility and drift functions are met. See, for instance, [Bjork (1998)] or [Mikosh (1998)]. For simplicity, I will also assume that a finite number, n , of Brownian processes shock the yield curve, but this need not be the case

must exist (see, eg, [James and Webber (2000)]) such that the drift of each asset must be of the form

$$\mu_t^{T_i} = r + \sum_{k=1,n} \lambda_k^t v_k(t, T_i) \quad (8)$$

If the market is complete this vector is unique. Therefore, following [Hughston (2000)], one can interpret the market price of risk, λ_k , associated with the k -th Brownian process as the extra return over the risk-less rate per unit of the volatility loading, $v_k(t, T)$, onto the k -th factor:

$$\lambda_k^t = \frac{\mu_t^{T_i} - r}{v_k(t, T_i)} \quad \forall T_i \quad (9)$$

(Again, no index T_i has been associated with λ_k^t). Another interpretation of the quantity λ^t is as a Sharpe ratio, ie the ratio of the excess bond return above the risk-less rate to the standard deviation,

$$\frac{E_t[\frac{dP_t^T}{P_t^T} - r dt]}{v(t, T)} = \frac{(\mu_t^T - r) dt}{v(t, T)} \quad (10)$$

(See, eg, [Cochrane (2000)]). For the multi-factor case, λ_k^t can be seen as the Sharpe ratio (excess return per unit 'risk') earned by the bond because of its loading over the k -th factor. Note that in arriving at Equation 9 purely no-arbitrage arguments were invoked. Much stronger general-equilibrium assumptions (eg, the validity of the CAPM model) are needed to prescribe specific forms for the vector λ_k^t , eg its relationship to the market portfolio. This observation will be re-visited in the context of the CIR/Vasicek models.

Going back to Equation 9, using this no-arbitrage condition for the bond price process, one can therefore write ([Hughston and Brody (2000)]) (again, in the real-world)⁷

$$\frac{dP_t^T}{P_t^T} = r dt + \sum_{k=1,n} v_k(t, T) (\lambda_k^t dt + dw_k) \quad (11)$$

Equation 11 can be integrated to give [**First Formulation**]

$$P_t^T = P_0^T \exp\left(\int_0^t r_s ds\right) \exp\left(\int_0^t \sum_{k=1,n} v_k(s, T) [\lambda_k^s ds + dw_k(s)] - \frac{1}{2} \int_0^t v(s, T)^2 ds\right) \quad (12)$$

Recalling that $P(T, T) = 1$ and the definition of the continuously-compounded money-market account, B_t^T :

$$B_t^T = \exp\left(-\int_t^T r_s ds\right) \quad (13)$$

⁷a somewhat similar approach can also be found in [Gustafsson (1992)]

one can 'invert' 12 for the money-market account

$$B_0^T = \frac{1}{P_0^T} \exp \left(\int_0^T \sum_{k=1,n} -v_k(s, T) (\lambda_k^s ds + dw_k^s) + \frac{1}{2} \int_0^T v(s, T)^2 ds \right) \quad (14)$$

and, making use of Equation 13, one can obtain the no-arbitrage process for the short rate implied by the bond price dynamics 11 [**Second Formulation**]:

$$r_T = -\frac{\partial \ln B_0^T}{\partial T} = -\frac{\partial \ln P_0^T}{\partial T} + \int_0^T v(s, T) \frac{\partial v(s, T)}{\partial T} ds - \int_0^T \sum_{k=1,n} \frac{\partial v_k(s, T)}{\partial T} (\lambda_k^s ds + dw_k^s) \quad (15)$$

In order to obtain a third equivalent formulation, one can define the forward bond price process $P_t^{T_1, T_2}$ (with $T_2 > T_1$) as

$$P_t^{T_1, T_2} = \frac{P_t^{T_2}}{P_t^{T_1}} \quad (16)$$

and, from Equation 12 one can write [**Third Formulation**]

$$P_t^{T_1} = P_0^{T_1, T_2} \frac{\exp \left(\int_0^t \sum_{k=1,n} v_k(s, T_2) (\lambda_k^s ds + dw_k^s) - \frac{1}{2} \int_0^t v(s, T_2)^2 ds \right)}{\exp \left(\int_0^t \sum_{k=1,n} v_k(s, T_1) (\lambda_k^s ds + dw_k^s) - \frac{1}{2} \int_0^t v(s, T_1)^2 ds \right)} \quad (17)$$

Equation 17 expresses the bond price process in terms of two exogenous vectors, the volatility vector and the risk premium vector, and the short rate no longer appears explicitly in the formulation (although it is obviously linked to these vectors via Equation 15).

Finally, if we go back to the first formulation, and define the instantaneous forward rate, $f(t, T)$, as

$$f(t, T) = -\frac{\partial \ln P_t^T}{\partial T} \quad (18)$$

by differentiation of Equation 17 it follows that

$$f(t, T) = f(0, T) + \int_0^t v(s, T) \frac{\partial v(s, T)}{\partial T} ds - \int_0^t \sum_{k=1,n} \frac{\partial v_k(s, T)}{\partial T} (\lambda_k^s ds + dw_k^s) \quad (19)$$

This equation provides yet another equivalent 'set of coordinates' for the yield curve evolution. Equation 19 gives the evolution of a forward rate in terms of the volatilities of the *bond price* processes. However, given the definition of

the forward rate, Ito's lemma shows that the volatility, σ_t^T , of the forward rate $f(t, T)$ is given by⁸

$$\sigma_t^T = -\frac{\partial v(t, T)}{\partial T} \quad (20)$$

Using Equation 20 one can therefore write the dynamics for the instantaneous forward rates in terms of forward rate volatilities as [**Fourth Formulation**]

$$df(t, T) = -\sigma_t^T \left(\int_t^T \sigma_t^u du \right) dt + \sum_{k=1, n} (\sigma_t^T)_k (\lambda_k^s ds + dw_k^s) \quad (21)$$

Also the fourth formulation (Equation 21), often known as the HJM condition, was arrived at by pure no-arbitrage arguments, and shows that the real-world no-arbitrage drift of the forward rates is just a function of the chosen forward rate volatilities (and of the risk vector λ_k^s).

All these formulations are very 'rich' in that they allow not only the pricing of derivatives, but also say (prescribe) something very specific about the real-world dynamics of the term structure. This is the conceptual route taken by the absolute pricing approach, (see, eg, [Vasicek (1977)], [Cox, Ingersoll and Ross (1985)], [Longstaff and Schwartz (1992)], but see also the discussion in Section 5.1 regarding their implementation in practice). If we 'only' want to price derivatives, which, in the complete-market setting we used so far, are 'redundant' securities, such a fundamental description is not required. See Section 3.4.

3.3 Equivalence of the Different Approaches

The derivation above highlights the conceptual 'symmetry' among the various different possible formulations of the no-arbitrage yield curve dynamics: one could have taken, for instance, expression 15 as a starting point, used Ito's lemma to obtain the short-rate volatility, σ_r , in terms of the bond price volatilities $v(s, T)$, and obtained a totally equivalent description of the short rate evolution completely in terms of the short-rate-related quantities⁹. What would have changed in moving from one choice to the other would have been the 'natural' input volatilities: σ_t^T in one case, σ_r in another and $v(t, T)$ in the third. If, in particular, one chose one set of input volatilities to be a deterministic function of time, neither of the other two sets would turn out to be deterministic.

All the modelling approaches reviewed in this survey fall in one of these equivalent formulations. Despite this conceptual equivalence, however, any

⁸The alternative and equivalent formulation

$$v(t, T) = \int_t^T \sigma_t^u du$$

makes clear that

$$v(T, T) = 0$$

ie, that the volatility of a bond price must vanish as its maturity approaches.

⁹Needless to say, if one had started from 'simple' bond price volatilities, the dynamics for the short rate would have been rather complex, and, in general, non-Markovian

choice for the preferred set of co-ordinates has implications for the ease of calibration of the model to market-observable prices. The various modelling approaches can be looked at as the requirement that certain set of volatilities should have a 'simple' or 'nice' form (for instance, that they should be deterministic functions of time), or lend themselves to simple calibration.

Given this way of looking at different models, one would have thought that a lot of statistical research must have been devoted to the study of which sets of volatilities have the 'simplest' (eg, more homoskedastic) behaviour. This research has to some extent indeed been carried out (for a review of some of the results see, eg, [Cambell, Lo and MacKinley (1996)], or the references quoted in Section 4), but it has had relatively little influence on derivatives modelling. I discuss below that market practice (and, in particular, the market's sometimes arbitrary choice of which set of volatilities should be regarded as deterministic functions of time) has had a much stronger influence on this process.

3.4 The Relative Risk Density Process and Derivatives Pricing

Following again [Hughston and Brody (2000)], let us go back to Equation 12 and to the definition 13 of the money-market account. Putting them together one can write

$$\frac{P_t^T}{B_0^t} = P_0^T \exp \left[\int_0^t \sum_{k=1,n} v_k(s, T) (\lambda_k^s ds + dw_k^s) - \frac{1}{2} \int_0^t v(s, T)^2 ds \right] \quad (22)$$

After defining

$$\lambda_t^2 = \sum_{k=1,n} |\lambda_k^t|^2 \quad (23)$$

let us create the following strictly positive martingale process, called the relative risk density process:

$$\frac{d\rho_t}{\rho_t} = - \sum_{k=1,n} \lambda_k^t dw_k^t \quad (24)$$

$$\rho_t = \exp \left(- \int_0^t \sum_{k=1,n} \lambda_k^s dw_k^s - \frac{1}{2} \int_0^t \lambda_s^2 ds \right) \quad (25)$$

After multiplying 25 by 22 one obtains (by completing the squares) [**Fifth Formulation**]

$$\frac{P_t^T \rho_t}{B_0^t} = P_0^T \exp \left[\int_0^t \left(\sum_{k=1,n} v_k(s, T) - \lambda_s^s \right) dw_k^s - \frac{1}{2} \int_0^t [v(s, T) - \lambda_s]^2 ds \right] \quad (26)$$

showing that, unlike the quantity $\frac{P_t^T}{B_0^t}$, the ratio $\frac{P_t^T \rho_t}{B_0^t}$ is a (strictly positive) martingale. Furthermore, Equation 26 also shows that multiplication of P_t^T by the ratio $\frac{\rho_t}{B_0^t}$ provides an unbiased forecast for the future (time- t) value of P_t^T ¹⁰. Indeed, if, to lighten notation, we denote by Z_t^T the bond price discounted by the money market account, $Z_t^T \equiv \frac{P_t^T}{B_0^t}$, one can readily obtain

$$E [Z_t^T \rho_t | \mathcal{F}_0] = Z_0^T \rho_0 = Z_0^T \quad (27)$$

showing that the martingale process ρ_t is directly linked (to within the riskless discount factor B_0^t) to the stochastic discount factor used in asset pricing.

It is worthwhile pausing to comment about the relative risk density process (also known as the stochastic discount factor, the pricing kernel or the change or measure'). In the general asset pricing context, the stochastic discount factor (see, eg, [Cochrane (2000)]) appears in the consumption-based pricing equation as the marginal rate of substitution between time periods, and therefore assigning it exogenously requires specifying a particular form for the investor's utility function. Historically, the most 'natural' way to specify investors' aversion to risk has been via the market-price-of-risk route. However, this need not be the case: taking the stochastic discount factor as the primitive building block might appear to provide little intuitive appeal, but it is a theoretically perfectly feasible (and, in some ways, a more fundamental) way to specify a term structure model. This is the route taken, for instance, by [Constantinides (1992)].¹¹ The equivalence of the different formulations can be seen by noticing that all models implicitly or explicitly assign the relative risk density process. What changes is whether this specification is effected via an estimation of the investors' risk aversion (eg, [Cox, Ingersoll and Ross (1985)], [Longstaff and Schwartz (1992)]), or by implicitly enforcing the fulfillment of the martingale condition 26 (eg, [Black, Derman and Toy (1990)], [Black and Karasinski (1991)]), or by the direct modelling of the pricing kernel ([Constantinides (1992)], [Rogers (1997)]).

3.5 From Absolute to Relative Pricing

The approach just sketched suggests how to price bond options *and* bonds consistently in an arbitrage-free way, given the real-world dynamics of the driving factors, and the investors' attitude towards risk. In the case of stock options, the most powerful and useful result of the BS treatment, however, is that the investors' utility function should not affect the price of a replicable option. The logical path that leads to the equations just obtained seem to run counter to this philosophy. Isn't there a way to price 'just' bond options, *if we are given* the prices of bonds? In other terms, one can one move from absolute to relative pricing in the interest-rate context?

¹⁰The quantity $\frac{\rho_t}{B_0^t}$ goes under the name of the 'Long portfolio'. Its properties are discussed in detail in [Long (1990)].

¹¹[Harrison and Kreps (1979)] derive the condition for its uniqueness (essentially, market completeness, but see [Nielsen (1999)] for a more thorough discussion)

This shift in perspective can be presented most simply if one lightens the notation by dealing with the case of a single factor. Following [Hughston (2000)], one can start from the real-world measure Q and define, for any random variable X_t measurable with respect to a filtration \mathcal{F}_t , a new measure Q^ρ under which expectations are linked to the real-world expectations by

$$E_s^{Q^\rho}[X_t] = \frac{E_s[\rho_t X_t]}{\rho_s} \quad (28)$$

Notice that the measure change depends on λ via Equation 25. Recalling that $\frac{P_t^T \rho_t}{B_t^T}$ is a martingale (see Equation 26) and making use of Equation 28 with $X_t = \frac{P_t^T}{B_t^T}$, we obtain the result

$$\frac{E_s[\rho_t \frac{P_t^T}{B_t^T}]}{\rho_s} = E_s^{Q^\rho}[\frac{P_t^T}{B_t^T}] = \frac{P_s^T}{B_s^T} \quad (29)$$

which shows that, under the measure Q^ρ , the bond price P discounted by the money market account is a martingale. When the change of measure ρ is constructed from risk premia λ_k as in Equation 11, this measure is called the *risk-neutral measure*.

Consider now a derivative, C^τ , whose price is adapted to \mathcal{F}_t like the bond price P_t^T , and which is fully characterized by its terminal (time- τ) payoff ($\tau < T$). Since the no-arbitrage relationship 5 holds for any security, and therefore also for C^τ , the underlying bond and the derivative must share the same risk premium λ , and therefore it must also be true that

$$\frac{E_s[\rho_t \frac{C_t^\tau}{B_t^T}]}{\rho_s} = E_s^{Q^\lambda}[\frac{C_t^\tau}{B_t^T}] = \frac{C_s^\tau}{B_s^T} \quad (30)$$

and

$$C_0^\tau = E_0[\rho_\tau \frac{C_\tau^\tau}{B_0^\tau}] = E_0^{Q^\lambda}[\frac{C_\tau^\tau}{B_0^\tau}] \quad (31)$$

ie, *the price of the derivative today is obtained by taking the expectation of its discounted terminal payoff under the risk neutral measure.*

Because again of the commonality of the risk premia across all securities, and reverting to the multi-factor case, the dynamics of the bond price and of the derivative can be written in the form

$$\frac{dP_t^T}{P_t^T} = rdt + \sum_{k=1,n} v_k(t, T) (\lambda_k^t dt + dw_k^t) \quad (32)$$

$$\frac{dC_t^\tau}{C_t^\tau} = rdt + \sum_{k=1,n} \sigma_k^C(t, \tau) (\lambda_k^t dt + dw_k^t) \quad (33)$$

If, starting from each standard Brownian process W_k^t , we define a new vector process

$$W_{\lambda_k}^t = W_k^t + \int_0^t \lambda_k^s ds \quad (34)$$

it can be shown (Girsanov's theorem) that $W_{\lambda_k}^t$ is a standard Brownian process with respect to the risk-neutral measure Q^λ , and therefore

$$\frac{dP_t^T}{P_t^T} = rdt + \sum_{k=1,n} v_k(t, T) dw_{\lambda_k}^t \quad (35)$$

$$\frac{dC_t^\tau}{C_t^\tau} = rdt + \sum_{k=1,n} \sigma_k^C(t, \tau) dw_{\lambda_k}^t \quad (36)$$

Equations 35 and 36 show that, if the value of the derivative is calculated by taking the expectation under the risk neutral measure, the risk premia do not affect the dynamics of the bond or of the derivative, and therefore the price of the latter.

This still leaves open the question of how to carry out an expectation in the a priori unknown risk measure. In theory, one could explicitly obtain the new measure via the Radon-Nikodym derivative, $\frac{dQ}{dQ^p}$. In practice, the route followed is typically to *construct* the risk-neutral measure in such a way to ensure that the exogenously-given bond price processes are correctly recovered. This is done by using the property of the quantity $\frac{\rho_t}{B_0}$ discussed above (the 'Long portfolio'[Long (1990)]), which is security-independent, of producing an unbiased forecast of an asset. A measure is first constructed such that all the marginal and conditional expectations recover all the spot and forward *bond* prices. Given the invariance of the risk premia across securities, once such a measure (in computational terms, perhaps a lattice) has been built, the same measure (lattice) is used to value the derivative. This is the route implicitly taken by those modelling approaches such as, for instance, [Black, Derman and Toy (1990)] or [Black and Karasinski (1991)] which work directly in the risk-neutral measure. Arriving at the value of a derivative following this route is therefore a case of pure relative pricing, and nothing can be said, in this approach, about the real-world behaviour of bond prices, of forward rates or of the short rate.

It is important to point out that, in the derivation of the results above, market completeness was invoked in order to deduce the uniqueness of the market price of risk vector. This uniqueness is lost if the traded securities do not span the space of all the Brownian processes (ie, if there exists a dw_k that does not affect the price of at least one of the hedging bonds). When this is the case, a perfectly replicating portfolio cannot be created with certainty, and one has to resort, at least for the unhedgeable part, to absolute rather than relative pricing. This route is embarked upon in derivatives pricing with great reluctance, and many difficult-to-justify *ad hoc* assumptions are often invoked in order to recover a framework compatible with the complete-market treatment. I shall argue in the last part of this survey that the judgement as to whether invoking market completeness is a profitable working assumption should be made on the basis of the actual ability to replicate a terminal payoff, not of the simplicity of the resulting approach, and that it can be dangerous to ignore in the pricing the practical difficulties in enforcing replication. The choice need not be between full

relative or absolute pricing, and 'compromise solutions' (eg, 'no-too-good-deal' pricing [Cochrane and Saa-Requejo (2000)]) can be invoked.

Returning to the complete-market case, if the short rate and the volatility were constant (see however the discussion below), the relative pricing results 35 and 36 would clearly apply in the one-factor case. The expectation of the payoff of a τ -expiry call on a T -maturity bond struck at K would in the risk-neutral measure would then be given by the familiar BS formula:

$$C_0^\tau = P_0^T \mathcal{N}(h_1) - K \mathcal{N}(h_2) \frac{1}{B_0^\tau} \quad (37)$$

with

$$h_1 = \frac{\ln \frac{P_0^T B_0^\tau}{K} + \frac{1}{2} v^2 \tau}{v \sqrt{\tau}} \quad (38)$$

$$h_2 = \frac{\ln \frac{P_0^T B_0^\tau}{K} - \frac{1}{2} v^2 \tau}{v \sqrt{\tau}} \quad (39)$$

and $\mathcal{N}(\cdot)$ the cumulative normal distribution¹². Notice however, how, in the formulae above the volatility of the bond price has simply been expressed as v , rather than as $v(t, T)$. Looking at expressions ?? and ?? it is not clear which constant volatility could possibly be assigned to a bond price, without implying something very strange about the volatility of the forward rates. Furthermore, we have seen that the processes for bond prices and for the short rate are inextricably linked by the condition of no arbitrage. Therefore one cannot simply assume a process for one set of variables (as one implicitly does by assuming the trivial constant process for the short rate) and exogenously assign the volatilities of the equivalent 'co-ordinates'. This leads us directly to some of the difficulties encountered in the very first applications of the BS reasoning to interest-derivatives pricing, discussed in the next section.

4 Phase 1: The Black-and-Scholes/Merton Approach

The impact of the BS work was strongly felt also in areas, such as interest-rate derivatives pricing, that were considerably removed from the original scope of the BS paper. However the adaptation of their results to the new domain was not always simple, and the (real and fictitious) problems encountered at the time are reviewed in this section.

¹²In the classic Black-and-Scholes derivation the pricing formula is obtained as a solution of a PDE, which is in turn derived by a hedging argument. The hedging argument is still implicit in the approach sketched above, via the relationship 5.

4.1 Adapting the BS/M Framework to Interest-Rate Derivatives

The BS model, in its original formulation, assumes constant interest rates. If interest rates are deterministic (let alone constant) the problem of pricing interest-rate derivatives becomes trivial. Traders who began to use the BS formula to price bond options¹³ were obviously aware of this inconsistency, but, implicitly or explicitly, argued that assuming deterministic interest rates in the discounting and the forwarding would be a second-order effect, compared to the first-order change in the bond price.

As for the log-normal assumption for the bond price, this was partially satisfactory, because it did not constrain the price to be smaller than face value, thereby assigning a non-zero probability to negative interest rates. Since typical market volatilities tended to assign relatively small (risk-adjusted) probabilities to these events, traders were happy to gloss over this blemish as well.

What gave greater discomfort, however, was the so-called pull-to-par problem.

4.2 The Pull-To-Par Problem

The BS paper assumes constant volatility. Since the volatility of a bond must go to zero with its residual maturity, this creates a problem. The crudest answer was to argue that, if one were pricing a 'short-dated' option on a long-maturity bond, this effect could be neglected (see, eg, Hull (1990)). This solution is, however, clearly unsatisfactory, both because it allows no assessment of what 'long' and 'short' should mean, and, more seriously, because for many LIBOR options (caplets) the expiry of the option and the maturity of the associated notional bond are typically as little as six or three months apart. The expiry of the option and the maturity of the underlying (notional) bond are therefore very similar.

However, as long as the instantaneous percentage volatility, σ_t , is a deterministic function of time, it is well known that by inputting in the BS formula the root-mean-squared volatility, $\hat{\sigma}$, defined by

$$\hat{\sigma}^2 = \frac{1}{T} \int_0^T \sigma(u)^2 du \quad (40)$$

all the replication results still apply. In particular, a perfectly replicating portfolio can still be established, based on a delta-holding of stock given, at all points in time, by the BS formula with the residual root-mean-squared volatility. Therefore, if one had to price, say, a 9.75-year option on a 10-year bond, the correct volatility, one might argue, would be given by the root-mean squared volatility of the bond over the life of the option, approximately equal to volatility

¹³The term 'bond options' should be understood in a generalized sense, since a caplet and a floorlet can be seen as a put and a call on a bond. These LIBOR-based applications were in fact far more common than 'straight' bond options.

of a 5-year bond. However plausible, this reasoning gave rise to a paradox, illustrated in the section below.

4.3 From Black-and-Scholes to Black

Three years after BS's work, Black published a paper [Black (1976)] with the deceptively narrow title '*The Pricing of Commodity Contracts*'. Its emphasis was on forward (rather than spot) prices, and obtained a formally very similar pricing formula for a call option. Despite the superficial similarity, however, there was one fundamental difference, in that the volatility of relevance was now the volatility of a *forward* (as opposed to spot) price.

If rates were deterministic there would be no difference between the two volatilities, but if rates are stochastic, and correlated with the process for the underlying the difference can be large, both conceptually and in practice. In the case of a 9.75-year option on a 10-year bond, the argument presented above would indicate that the appropriate volatility would (approximately) be that of a 5-year bond. If one focusses on the forward price, on the other hand, the relevant volatility is that of a (forward) bond with the *unchanging* residual maturity of 0.25 years. Which volatility is the 'correct' one? The Black, approach, by focussing on the forward bond price, ie on the ratio of the spot asset (itself a bond) to the discounting bond, makes clear that it is the volatility of this ratio that matters. It is only when the volatility of the discounting bond is negligible with respect to the volatility of the asset that the volatility of the spot and of the forward process can be approximately considered to be the same. However, for the case of a long-dated caplet, the denominator contributes a volatility approximately as large as the volatility of the spot bond, but the closer the expiry of the caplet to the maturity of the bond, the greater the correlation between the two bonds, reducing the volatility of the ratio. Therefore the assumption of constant rates in the BS formula is poor because, by neglecting the covariance between the asset and the discounting bond, it systematically mis-specifies the relevant volatility. As a corollary, when the trader imputes from a traded price the 'implied volatility' this is always the volatility of the *forward* quantity (and only approximately the *spot* volatility).

So, despite the fact that the use of the Black formula with the forward price would have completely eliminated the pull-to-par problem, the legacy of the BS approach was such that at the time¹⁴ these subtle points sometimes failed to be fully appreciated, the distinction between the volatility of the spot price (considered the 'fundamental' process) and the forward price often glossed over, and the discomfort with the Black(-and-Scholes) approach applied to prices (spot or forward) remained.

¹⁴A mathematical paper showing that the appropriate input to the Black formula is the volatility of the *forward* rate appeared as late as 1995.

4.4 From Price Models to Rate and Yield Models

An alternative solution to the pull-to-par problem was to use as state variable a (log-normal) yield or rate. As an added bonus, this approach would also get rid of the problem of (spot or forward) bond prices greater than 1. So, either the yield to redemption of a bond or the equilibrium rate of a swap or forward rate began to be regularly input in the Black formula to obtain the prices of swaptions or caplets (Hull's book as reference here). Since neither yields nor rates are directly traded assets, this practice was entered upon with some discomfort, and it was generally regarded as intuitively appealing but not theoretically justifiable. In particular, it was not widely appreciated in the trading community that the superficially similar practices of using the Black formula with yields or forward rates are fundamentally different, in that the latter can be theoretically justified, for instance via the introduction of a 'natural payoff' [Doust (1995)], but the former, based on a flawed concept [Schaefer (1977)], remains theoretically unsound. So, despite the fact that the theoretical justification for the use of the Black formula for swaptions had been given as early as 1990 [Neuberger (1990)], papers proving the correctness of the approach were still appearing as late as 1997 (see, eg, [Gustafsson (1997)]). In the early 1990s the rate- and yield-based-models were therefore generally regarded, at least by the less sophisticated traders, as roughly equivalent, and as useful but similarly theoretically unsound.

By luck or by inspired foresight the trading practice crystallized around the use of the Black formula for swap and forward rates. This constitutes an important step in the history of term-structure modelling: the joint facts that traders carried on using the Black formula for caplets and swaptions despite its then-perceived lack of sound theoretical standing, and that this approach would be later theoretically justified contributed to turning the approach into a market standard, and directly led to the modern LIBOR market models. Therefore, without the somewhat fortuitous choices made during the establishment of a market standard, the 'modern' pricing approach might currently have a considerably different structure.

4.5 The Need for a Coherent Model

Going back to the pull-to-par problem, use of forward rather than spot prices could have solved the problem, but this route was often overlooked, and for different reasons, both by naïve and by sophisticated market players. The naïve traders simply did not appreciate the subtle, but fundamental, differences between the Black and the BS formulae 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 the appropriateness of the Black approach. However, the Black formula can give a perfectly correct answer for a series of options considered independently of each other, but there is no way of telling whether these options inhabit a plausible, or even logically consis-

tent, universe (ie, every option is priced in its own measure)¹⁵. Despite the fact that different forward bond prices correspond to different assets, their prices are strongly (albeit imperfectly) correlated. There is however no mechanism within the Black formula to incorporate views about this joint dynamics. The need was therefore increasingly felt for a comprehensive and self-consistent model, capable of providing coherent inputs for the pricing of complex products dependent on the joint realizations of several forward rates.

Indeed, I have shown elsewhere [Rebonato (1998)] extending work by [Doust (1995)] that, by working with suitably-high-dimension cumulative normal distributions, the Black formula *can* be rigorously extended to a large number of path-dependent and compound-option cases. In particular, when this route is taken the very same quadratic covariation terms between forwards

$$\left\langle \frac{df_i}{f_i}, \frac{df_j}{f_j} \right\rangle = \sigma_i \sigma_j \rho_{ij} \quad (41)$$

appear as in the modern LIBOR market model. There is however a fundamental difference: in the Black-based approach all these covariance elements have to be assigned exogenously and independently. In a model-based approach they might, individually, be less accurate or plausible, but they are at least guaranteed to be all internally consistent. Quoting [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 [...] 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 important because they directly led to the first-generation yield-curve models.

¹⁵As it is well known there is also an inconsistency of distributional nature, since forward rates of different tenor, or forward rates and swap rates cannot all be simultaneously log-normal. More fundamentally, if the volatility of one set of variables (say, forward rates) is assumed to be deterministic, the volatilities of the other sets of variables is stochastic [Jaeckel and Rebonato (2000)]. However, the pricing impact of these inconsistencies is in general very small, (see, eg, [Rebonato (1999)]). The lack of internal coherence across measures referred to here is of more fundamental nature.

5 Phase 2: First-Generation Yield-Curve Models

5.1 Vasicek and CIR

Historically, the first internally consistent term structure models were [Vasicek (1977)] and [Cox, Ingersoll and Ross (1985)]. Using as a blueprint the treatment presented in Section 2, their logical structure can be described as follows. The starting points are the joint assumptions that the short rate, r_t , is a diffusion process, with *real-world* drift μ_r , and volatility σ_r .

$$dr_t = \mu_r(r_t, t)dt + \sigma_r(r_t, t)dw_t \quad (42)$$

and that the prices of all bonds purely only depend on the short rate itself:

$$P_t^T = P_t^T(r_t) \quad (43)$$

Applying Ito's lemma,

$$dP_t^T = \left(\frac{\partial P_t^T}{\partial t} + \mu_r \frac{\partial P_t^T}{\partial r_t} + \frac{1}{2} \sigma_r^2 \frac{\partial^2 P_t^T}{\partial r_t^2} \right) dt + \left(\sigma_r \frac{\partial P_t^T}{\partial r_t} \right) dw_t \quad (44)$$

and using the definition of the market price of risk (Equation ??), one obtains

$$\frac{dP_t^T}{P_t^T} = (r_t + \lambda_t v(t, T)) dt + v(t, T) dw_t \quad (45)$$

with

$$v(t, T) = \frac{\sigma_r(r_t, t)}{P_t^T(r_t)} \frac{\partial P_t^T(r_t)}{\partial r_t} \quad (46)$$

Equating the drift in Equation 45 with the drift in Equation 44 one obtains the PDE:

$$\frac{\partial P_t^T}{\partial t} + (\mu_r - \lambda_t \sigma_r) \frac{\partial P_t^T}{\partial r_t} + \frac{1}{2} \sigma_r^2 \frac{\partial^2 P_t^T}{\partial r_t^2} = r_t P_t^T \quad (47)$$

to be solved with the initial condition $P_T^T = 1$. Therefore, in theory, in order to specify the *real-world* dynamics of bond prices in this approach, one would undertake an econometric analysis of the statistical properties of the short rate (to determine μ_r and σ_r), and separately estimate a utility function for the bond investors in order to derive the market price of risk (λ_t). This would tell the trader everything that there is to know about the real-world dynamics of the bond processes, and, in particular, all their prices today, P_0^T .

In practice, however, estimating the real world drift of the short rate and the market price of risk are notoriously difficult tasks¹⁶. And as for the estimation

¹⁶To give an idea of the variety of conclusions reached in recent statistical studies of the yield curve, Ait-Sahalia [Ait-Sahalia (1996)] finds that the short rate is close to a random walk in the middle of its historical range (approximately between 4% and 17%), but mean

from econometric data of the market price of risk (let alone its derivation from first principles) the difficulties are understandably even greater:

In the light of these difficulties, a formally similar, but conceptually very different, approach to using the Vasicek and CIR models was followed in practice by making the following joint assumptions:

1. that the real-world drift of the short rate should be of the form

$$\mu_r = k(\theta - r) \quad [\text{Vasicek and CIR}] \quad (49)$$

2. that the volatility of the short rate should be of the form

$$\sigma_r = \sigma \quad [\text{Vasicek}] \quad (50)$$

$$\sigma_r = \sigma\sqrt{r} \quad [\text{CIR}] \quad (51)$$

3. that the market price of risk should be of the form

$$\lambda_t = \lambda \quad [\text{Vasicek}] \quad (52)$$

$$\lambda_t = \lambda\sqrt{r} \quad [\text{CIR}] \quad (53)$$

(In these equations, the quantities λ, σ, k and θ are all assumed to be constant). These assumptions partly reflected common views about the dynamics of rates (eg, the assumption of mean reversion), but were largely made in order to allow analytical solution of the PDE 47.

If the model had been correctly specified, instead of estimating the real-world quantities directly from the dynamics of the short rate and the investor's attitude towards risk, one could turn the approach upside down, and impute these quantities from the observed prices of bonds¹⁷. This was the approach universally adopted by practitioners.

Was there evidence that the model was correctly specified? The empirical record was, not surprisingly, rather poor. If either model had been correctly specified, cross-sectional estimates of the (combinations of) parameters characterizing the models should be constant over time. Unfortunately, [Brown and Dybvig (1986)]

reverts strongly when outside this range. [Hamilton (1989)], [Garcia and Perron (1996)], [Gray (1996)], [And and Bekaert (1998)], [Naik and Lee (1997)], [Bansal and Zhou (2002)] among others instead argue that a switch behaviour, each with its reversion level and reversion speed, is more appropriate. Then again, a number of authors (see, eg, [Ait-Sahalia (1996)], [Chan et al. (1992)]) find evidence that the exponent β in the diffusion equation

$$dr = \mu_r dt + r^\beta \sigma dw \quad (48)$$

should be greater than 1. As for the estimation of the market price of risk, [Dai and Singleton (2000)], [Jagannathan et al (2000)], [Ahn et al. (2002)], [Bansal and Zhou (2002)] and [Duffee (2002)] suggest that a complex specification may be required to account for the observed yield curve dynamics.

¹⁷In reality, since the term in $\frac{\partial P_t^T}{\partial r_t}$ in Equation 47 contains the expression $[(k\theta - \lambda\sigma) - kr]$ (Vasicek) or $[k\theta - (\lambda\sigma + k)r]$ (CIR), the market price of risk and the reversion parameters k and θ cannot be independently estimated using the implied approach. This has no impact, however, on derivatives pricing, but switches the approach from *absolute* to *relative* pricing.

and [Brown and Schaefer (1994)] found the estimates to be not only non-stationary over long periods of time (which could be compatible with a regime-switch view), but also wildly fluctuating on a day-by-day basis¹⁸. Extensions of the CIR approach to several factors did not achieve better statistical corroboration: see, eg, [Chen and Scott (1993)].

Apart from these difficulties (and from the well-known fact that the Vasicek model allowed for negative interest rates), a more fundamental problem affected these models. Because of their stationary nature (λ, σ, k and θ are all constants), neither the Vasicek nor the CIR approach could recover for an arbitrary observed yield curve, P_0^T . This was both the strength and the limitation of the approach: it provided the *prescriptive* element alluded to in Section 1 ('if this is indeed the dynamics of the short rate, this is what the prices of the bonds and their volatilities *should* be'), thereby making relative-value bond trading applications possible. When used to price derivatives, however, the inability to recover the prices of the underlying bonds was clearly unsatisfactory. The shift in interest in the trading community towards the pricing of interest-rate derivatives was rapidly making the shortcomings to be more acutely felt than the advantages.

5.2 Reasons for the Good Performance of Short-Rate Models

The models presented so far belong, at least in spirit, to the 'absolute pricing school'. Therefore they should have genuine explanatory power for the dynamics of bond prices. Yet, the assumption that all that there is to know about a future yield curve could be 'summarized' in the short rate should have given rise to some puzzlement. With HJM-aided insight, it is easy to see that the 'intriguing' ability of the short rate to account in a satisfactory fashion for the yield curve dynamics can be explained in rather prosaic econometric terms: numerous studies (see, eg, [Martellini and Priaulet (2001)] for a recent review) have shown that the first principal component obtained by orthogonalizing the correlation matrix formed from changes in rates explains a high proportion (around 90%) of the observed yield curve variability; and that the first eigenvector assigns approximately constant loadings to the various rates (ie, the first mode of variation is the change in average level of the yield curve). This being the case, virtually *any* rate could have been chosen as a reasonable proxy for the curve level, and therefore the after-all-satisfactory performance of short-rate-based models had little to do with any specific property of the short rate itself. The main 'virtue' of short-rate-based approaches was simply that, given the implicit choice of the money market account as numeraire enforced by computational trees, the short rate provided the most expedient quantity to effect the discounting between time horizons.

¹⁸It should be noted in passing that more stable results were obtained [Brown and Schaefer (1991)] for the dynamics of the term structure of real rates, as can be estimated from UK index-linked government bonds, suggesting that a one-factor description of the CIR type could be adequate for real rates.

The legacy of this class of models was therefore very strong, and it is therefore not surprising that the HJM framework, in which the short rate plays no privileged role, was at the beginning nonetheless looked at from the rather inappropriate vantage point of 'what does all this entail for the dynamics of the short rate?'. More about this later.

6 Phase 3: Second-Generation Yield Curve Models

6.1 Fitting the Yield Curve

A solution to the problem of how to recover an arbitrary exogenous set of bond prices was offered by the Hull-and-White (HW in the following) approach [Hull and White (1990)], [Hull and White (1993)], [Hull and White (1994a)], [Hull and White (1994a)]. The modification of the Vasicek equation is apparently minor (a similar treatment can be applied to the CIR approach), and consists in making the reversion level in the real-world dynamics, θ , of the short rate become time-dependent:

$$\mu_r = k[\theta_t - r] \tag{54}$$

It is easy to show [Hull and White (1993)] that the introduction of these extra degrees of freedom is exactly what is required in order to fit an arbitrary exogenous yield curve. Interest-rate options could now be priced with the confidence that the underlying bonds would certainly be priced correctly. Furthermore, the still-relatively-simple form for the drift ensured that, conditional on a future realization of the short rate, r_τ , the corresponding future bond prices, P_τ^T , could be obtained analytically. This made the computational burden commensurate with the computing power available on the trading desks in the mid 1990s. For this reason (ie, the need to compute a price in a trading time ranging from a few seconds to a couple of minutes) the extended Vasicek approach was generally preferred to be theoretically more appealing extended CIR model (which ensures non-negative rates). Only a few years later this advantage would no longer have been significant, but, by the time the CPU evolution caught up with the pricing and hedging requirements of the extended CIR model, the trading practice had moved in different directions.

6.2 Fitting the Caplet Volatilities

If an arbitrary yield curve could now be recovered by construction, the same was not true for the implied caplet volatilities: with a constant short-rate volatility σ_r the model would *prescribe* what the term structure of volatilities should look like (and it would therefore constitute a useful, if risky, tool for the plain-vanilla option trader), but could not account for (*describe*) an exogenous set of market volatilities. The main users of term-structure models, however, were becoming traders in complex products, who would transact their vega hedges buying and selling caplets and swaptions from the plain-vanilla desk.

See the discussion in Sections 2.2 and 2.4. Therefore the inability to price correctly the second-order hedging instruments was becoming the new source of modelling dissatisfaction.

Sure enough, the same 'trick' used to recover the market bond prices could be played again, and Hull and White [Ref here] showed that, by also making the volatility time-dependent, and exogenous term structure of volatilities could now be recovered. The price to be paid, however, was to make the model even more strongly non-stationary, thereby implying that future term structure of volatilities would look very different from today's. Despite the fact that Hull and White themselves cautioned against taking this route [Ref here], their advice often went unheeded by practitioners and the all-fitting approach was used in market practice.

Did this lack of time-homogeneity matter, since after all the prices of the hedging instruments today were correctly recovered? As discussed in Section 2 (see also Section EMH) the answer depends on one's views regarding market completeness, replicability, and the ability to 'lock-in' parameters. In making these choices the intellectual influence of the BS framework, and of the no-arbitrage approach to pricing in general, where the realization of (some!) real-world quantities becomes irrelevant, played an important part in making traders focus on the recovery of today's prices, and believe (or hope) that the 'model would take care of the rest'. In other words, traders had made an impressive intellectual leap, and considerable effort, in grasping the counterintuitive concept that one does not have to believe the market value of, say, a forward rate to be econometrically plausible for 'locking it in' by trading in discount bonds (the only instruments needed to complete the market in FRAs). Such had been the power and legacy of this result that the same conclusions were extended to many much more complex situations, where the market completeness should have been questioned much more carefully. Indeed, the over-strong application of this 'irrelevance theorem' still characterizes vast areas of derivatives pricing¹⁹. Its critical reappraisal has been mainly undertaken in the wake of the introduction of the LIBOR market models, that showed that a perfect fit to bond *and* caplet prices can be obtained in an infinity of ways, and thereby forced traders to choose among these alternative on the basis of criteria other than recovery of market prices. This aspect is discussed further in Section 7.2.

6.3 Tree-Based Models

The extended models by Hull and White were chronologically not the first of the yield-curve-fitting type to be used. Around the same time other models, (Ho and Lee, [Black, Derman and Toy (1990)], [Black and Karasinski (1991)]) were introduced, which were also capable of reproducing by construction an arbitrary exogenous set of current bond prices. These models were presented

¹⁹The difficulty of fitting equity smiles with time-stationary parameters, for instance, has not stopped traders from using models (such as Dupire's [Ref here] or Derman and Kani's [Ref here]) that imply very 'implausible' future smile surfaces, but exactly recover (by construction) any exogenous set of plain-vanilla option prices

in computational form, ie as an algorithmic construction to discount cashflows in such a way that the discounted expectations of the terminal payoffs of pure discount bonds would recover correctly their market prices. A common feature of these computational approaches was that the description of the yield-curve dynamics was now directly carried out in the risk-neutral measure, and no attempt was made to address the issue of the real-world evolution of bond prices.

Despite the fact that the work by [Harrison and Kreps (1979)] pre-dated these developments, the martingale approach had not yet become common parlance in the derivatives-pricing area. Therefore the theoretical justification for the algorithmic prescription was adequate (fundamentally, it hinged on the recovery of the observed market prices, assumed to be always arbitrage free), but not particularly profound. Also the link with continuous-time models, for which most of the theoretical results apply, was only made once the similarities between the computational trees and the finite-differences approaches were appreciated: the tri-nomial trees immediately suggested an explicit-finite-differences interpretation of the construction (see, eg, [Hull and White (1990)], building on the work in [Brennan and Schawrtz (1978)] and [Schaefer (1977)]) and it took little time to associate a limiting PDE, and to this an SDE. The fact that this translation to the continuous-time vocabulary was very much an afterthought explains why some important features of these models was initially overlooked. The continuous-time equivalent of the BDT model, for instance, was shown by Hull and White to be

$$d \ln r_t = [\theta_t - f'_t(\psi_t - \ln r_t)] dt + \sigma_r(t) dz_t \quad (55)$$

with

$$f'_t = \frac{\partial \ln \sigma_r(t)}{\partial t} \quad (56)$$

This expression clearly shows that mean-reversion in the BDT model is rather *sui generis*, since it can only be obtained if the volatility of the short rate decays with time. More importantly, [Rebonato (1999b)] shows that this feature is a direct consequence of the requirement that the computational tree should be recombining. This 'technical' constraint, indispensable in order to avoid computationally intractable 'bushy' trees, has no theoretical justification, and actually brings about one of the most unpleasant features of the BDT model: once the market prices of caplets are correctly recovered, the resulting short rate volatility is in general indeed found to be time-decaying. This, however, implies that future yield curves will be less and less volatile, with clear and unpleasant consequences for the implied future hedge ratios.

Given these shortcomings, why was the BDT model so popular, and why did it remain so for so long? Quite simply because it allows speedy and simple calibration to caplet prices (or swaption prices). In a way, the very same reasons that make the calibration so simple (see the discussion in [Rebonato and Kazziha (1997)] or [Rebonato (1999b)]) should have made traders reluctant to use the BDT model for pricing (and even more so for hedging). Nonetheless, the excessive

reliance on the recovery of *today's* prices discussed above often made the BDT the model of choice.

The competitors of the day, mainly the HW and the BK models, implied a much more palatable dynamics for the short rate: the BK model, for instance, incorporated, in common with the BDT approach, a deterministic function of time in the drift for the short rate, thereby allowing the correct pricing of discount bonds; however it displayed a 'genuine' mean-reverting behaviour, ie, a mean reversion that did not depend on the time behaviour of the short-rate volatility, and managed to do so on a recombining (although *tri*-nomial) tree. As a consequence, an exact fit to caplet prices could be obtained while implying a much more time-stationary term structure of volatilities. Despite these important positive features, neither the HW nor the BK models ever became as widely used as the BDT. As for the extended Vasicek model, the ability to obtain analytically the future yield curve emanating for any given node in the tree was a clear computational advantage. However, it was more difficult to calibrate the model, either exactly (using a time dependent σ_r) or approximately (using a constant σ_r), to caplet prices. Also, the possibility of negative interest rates was seen as an undesirable feature of the model²⁰. The BK model did not suffer from negative rates, yet it was probably even less widely used than the extended-Vasicek approach. The reasons for this were the greater calibration difficulties, coupled with the over-strong belief that exact recovery of the current observed market prices was all that mattered for a model (at least in practice, if perhaps not in theory): it would be the LMM to show traders that substantially different prices could be obtained for complex products even when the market prices of bonds and caplets had been correctly recovered ([Sidenius (2000)]).

6.3.1 Forward Induction

Despite the fact that this review is only marginally focussed on computational techniques, it is important to mention briefly the introduction of the forward-induction technique ([Jamshidian (1991)]). The idea allowed to calibrate trees of the BDT, BK family to bond prices without retraversing portions of the tree that had already been computed for bonds of shorter maturity. This apparently minor 'trick' turned the task of fitting a tree with n steps from a $O(n^3)$ to a $O(n^2)$ problem. Related in these terms in a review of the conceptual strides that were being accomplished in modelling the dynamics of interest rates, this might seem not to amount to much. However, given the computing power of the day, it is safe to say that, without forward induction, models such as the BDT (that was to set a trading standard for a long time, at least for the pricing of

²⁰In the mid 1990s, when these models were used, the levels of most yield curves ranged between 5% and 10%. The risk-adjusted probability of negative rates, especially in the presence of mean-reversion, was therefore extremely small, and the models were probably too harshly penalized for what was at the time a rather minor blemish. At the time of writing, the short end of the US\$ yield curve is well below 2%, and these concerns would be better founded today.

Bermudan swaptions) could not have been used for pricing and hedging, trading practice would have been different and models would have developed differently.

6.4 The Dimensionality of the Underlying Drivers

6.4.1 The Industry Needs

Had an interest-rate-option trader been asked in the mid 1990s what the greatest shortcoming of the then-mainstream models was, she would have probably referred to their being almost invariably one-factor models. The reason for this concern was the realization that the value of some of the most popular complex derivatives products of the day was strongly dependent on the ability of different forward rates to move out of lock-step²¹. At the time, little distinction was made between *instantaneous* and *terminal* decorrelation (see. eg. [Rebonato (2002)] for a detailed discussion), and little attention was paid to the ability of a time-dependent forward rate volatility to produce decorrelation among rates, even in the presence of perfect instantaneous correlation. This was probably due to the fact that, for the pre-LIBOR-market-model pricing approaches, the volatility was uniquely determined once a best- or exact fit to the caplet prices had been carried out. As a consequence, after the fitting no 'volatility degrees of freedom' were left to ensure, for instance, a realistic terminal de-correlation between rates.

Despite the fact that a (time-stationary) forward-rate volatility with a dependence on the residual forward-rate maturity can be very effective in producing decorrelation among rates (see, eg, [De Jong, Driessen and Pelsser (1999)]), the natural reaction at the time was therefore to look for models with more than one Brownian driver.

6.4.2 The Modelling Response

[Brennan and Schwartz (1982)], [Brennan and Schwartz (1983)] had presented a two-factor model of the yield curve as early as the mid 1980s. However, the industry needs and the increased computer power did not make their approach a strongly desirable or viable solution until the early-to-mid 1990s. Their approach, conceptually similar to the Vasicek/CIR, was to posit that the term structure is driven by two state variables, the short rate and the consol yield²², L , of respective stochastic real-world dynamics given by

$$dr_t = \mu_r(r, L)dt + \sigma_r(r, L, t)dw_r \quad (57)$$

$$dL_t = \mu_L(r, L)dt + \sigma_L(r, L, t)dw_L \quad (58)$$

²¹An example of these products were indexed-principal swaps, which were aggressively and successfully marketed in the mid 1990s to proxy hedge the exposure arising from the pre-payment risk of US\$ mortgage-backed securities. The notional of this long-dated products was prescribed to be a function of the path of a very short rate (eg, 6-month LIBOR), thereby introducing a dependence of the value of the swap on the correlation between short and long rates.

²²the consol yield, L , is the yield paid by a consol (irredimable) bond. If the coupon paid by the bond is c it is easy to show that the consol yield is given by c/L .

The idea was particularly appealing, because several principal-component-analysis studies had by then shown that the slope of the yield curve was the most important mode of deformation after the level, and that together these two eigenvectors, ξ_1 and ξ_2 , that could be proxied by

$$\xi_1 \simeq \alpha r + L \quad \alpha > 0 \quad (59)$$

$$\xi_2 \simeq \beta L - r \quad \beta > 0 \quad (60)$$

were able to account for more than 90% of the observed yield curve variability.

From Equations 57 and 58 the reasoning then followed the usual route: market completeness and absence of arbitrage were invoked, leading to the introduction of two market prices of risk ($\lambda_r(r, L, t)$ and $\lambda_L(r, L, t)$): see Equations 9 and 8 in Section 3.2. One of these two quantities, $\lambda_r(r, L, t)$, was then eliminated by invoking the fact that one of the securities in the economy (the consol bond) has a price that does not depend on the short rate (it is simply given by $1/L$), and therefore its real-world return cannot depend on r : $\mu_L(r, L) = \mu_L(L)$. Some algebraic manipulations then yielded for the consol market price of risk (that, to avoid arbitrage, must apply to *any* security)

$$\lambda_L(r, L, t) = \frac{rL - L^2 + \mu_L}{\sigma_L} - \frac{\sigma_L}{L} \quad (61)$$

and the resulting associated PDE only contained the market price of risk for the short rate.

Once again at this point a 'fundamental' approach could in theory have been followed, by estimating the real-world dynamics of the driving factors and by deriving from a utility function the required market prices of short-rate risk. In practice, the 'implied' (cross-sectional) approach was employed, by best-fitting the model 'parameters' to the prices of the observed bonds. Since no closed-form solutions were available, the match to bond prices had to be carried out by solving a relatively time-consuming two-dimensional finite-differences method. This is the reason why the approach was not seriously considered for almost a decade after its appearance in the literature. Unfortunately, apart from the numerical burden, other, there were more fundamental problems.

6.4.3 Stability Problems and Solutions

Following the favourite numerical route of the time, the Brennan and Schwartz approach had been computationally presented as an application of the tree/explicit-finite-differences-grid technology, and very few traders or researchers tried to apply a Monte Carlo (forward induction) methodology to it. Had they done so, they would have immediately realized that the coupled dynamic system of Equations 57 and 58 was dynamically unstable. This realization was first made precise by [Hogan (1993)], who showed that, if the dynamics of the yield curve is described by Equations of the type 57 and 58, and the short rate drift has

the Brennan-and-Schwartz functional form, either r or L will reach infinity with probability one in finite time.

This problem was solved by [Rebonato (1997)], who, by working directly in the risk-neutral measure, invoked the fact that the reciprocal of the consol yield, being a traded asset, had to grow at the risk-less rate to eliminate one market price of risk, and then ensured a perfect pricing of the market bonds by introducing, in the spirit of Hull and White, a time-dependent function θ in the drift of the short rate²³.

Other two-factor models were appearing which did not suffer from these drawbacks. In particular, the two-factor model by [?] implied a stable dynamics with desirable features such as stationarity, while allowing for a good and efficient calibration to many observed term structures of volatilities. More pressing practical problems were however appearing, that brought about the demise of this line of research.

6.4.4 Poor Instantaneous Decorrelation

Recall that the motivation behind the growing interest in two-factor models was the perceived need to account for rate decorrelation by increasing the dimensionality of the underlying term structure model. [Rebonato and Cooper (1995)], however, showed that low-factor models display intrinsic limitations in producing rapid de-correlation between contiguous forward rates, and that these limitations are largely independent of the details of the model. In particular, it was shown that, under very general conditions, for a two-factor model the correlation between two instantaneous forward rates would go to zero with the difference in maturities as a cosine law, thereby effectively negating the very effect (*rapid* decorrelation between contiguous forward rates) that the approach was seeking to address.

Sure enough, by adding a sufficiently large number of factors, a more plausible behaviour for the decorrelation could be obtained, but the marginal contribution of what are effectively higher and higher Fourier frequencies was shown to be rather slow, and the numerical approaches of the day were still wedded to lattice-based implementations (at least for compound options). The curse of dimensionality was therefore rendering this route *de facto* impracticable.

²³Despite the fact that the resulting dynamical equations were now stable, objections of a different nature were raised to this type of approach by [Duffie, Ma and Yong (1995)], who questioned the self-consistency of all approaches based on the joint specification of the dynamics for the short rate and the consol yield. For the approach to be self-consistent, in fact, the price of a consol bond, C , paying $\$dt$ every infinitesimal time interval, must be given by

$$C_0 = E_0 \left[\int_0^\infty \exp \left(- \int_0^t r_s ds \right) dt \right]$$

showing that there must be a fundamental link between the price and the dynamics of the short rate and of the consol yield. It is therefore not *a priori* clear to what extent the dynamics of r and L can be specified without taking this link into account.

6.4.5 Longstaff and Schwartz

Despite the fact that the model by [Longstaff and Schwartz (1992)] conceptually fits in the set of the much-earlier equilibrium approaches that start from joint prescription of the market-price-of-risk vector and of the real-world dynamics of the state variables, it chronologically belongs to this period. The Authors, possibly taking up a suggestion by [Dybvig (1998)], chose the short rate and its instantaneous volatility as the driving state variables to describe the dynamics of the yield curve, and were able to obtain closed-form solutions for bond prices (in terms of the non-central chi-squared distribution). This made the calibration and implementation of the model easier to achieve. When used for trading purposes, the 'implied' procedure described for the CIR/Vasicek models was invariably employed.

Very good fits could be obtained to a variety of market yield curves, some of very complex shape, and the fitted short rate and volatility were observed to have the correct order of magnitude and to change rather smoothly with the fitting date (see the discussion in [Rebonato (1998)], suggesting that the model might be better-specified than the CIR/Vasicek models of the same family). There were, however, some problems: to begin with, with the parameters obtained by the best fit to market data the correlation between the short rate and its volatility almost invariably turned out to be very high (eg 90% to 99%), implying that the model was effectively very close to being one-factor. See [Clewlow and Strickland (1994)] and the reply by [Longstaff and Schwarz (1994)]. Furthermore, [Rebonato (1998)] studied the predicted changes in yields of all maturities once the model was calibrated and the observed market changes in (proxies of) the short rate and its volatility were fed as input. (Clearly, good predictive power is necessary for an effective 'in-model' hedging). The results were mixed at best. Finally, and almost paradoxically, the fact that the LS model would produce smiles was at the time regarded as a drawback. After a period of great interest, the approach was therefore almost completely abandoned by traders.

6.4.6 Unresolved Modelling Issues

At the close of the second phase in interest-rate derivatives pricing, the awareness was growing that

- terminal de-correlation is important,
- low-dimensionality models are not effective in producing appreciable instantaneous decorrelation,
- once the market caplet prices have been fitted, short-rate-based approaches leave no degrees of freedom to specify a time dependence for the volatility that could produce a realistic terminal decorrelation.

This feeling of impasse was to be broken by the growing acceptance among the trading community of the HJM approach, reviewed below.

7 Phase 4: Third-Generation Yield Curve Models

With few exceptions (see, eg, [Duffie and Kan (1996)]), most of the term-structure models introduced until the work by [Heath et al (1989)], [Heath et al (1992)] were (Markov) models driven by the short rate (and, possibly, by another state variable). The Markov property had, up to this point, almost been 'taken for granted', the more so because all the recombining-tree-based computational techniques of the day relied on the full yield curve for maturities $T \geq t$ being known (either analytically or numerically) once the short rate process and its realization up to time t was assigned. Actually, such was the reliance on backward induction carried out on a tree (ideally suited to compound-option problems, but ill-designed to deal with path-dependent products), that papers [Hull and White (1994a)] were written showing how to deal with path-dependent payoffs while still using recombining trees²⁴.

Models, until HJM, were almost invariably associated with a recombining computational lattice²⁵, with a Markovian process for the short rate, and with a finite-differences-implied PDE/SDE. Indeed, such had been the reluctance to use forward induction and Monte Carlo techniques that it had hidden for so long the glaringly obvious instabilities of the [Brennan and Schawrtz (1978)] approach. However, the third-generation yield-curve models, of which the HJM was the prototype, were in general intrinsically non-Markovian, and therefore did not lend themselves to mapping onto a low-dimensional recombining tree. This would have made forward-induction the natural computational choice, but the evaluation of the prices of derivatives using Monte-Carlo techniques, which can readily handle high-dimensionality Markov processes, despite having been introduced as early as 1977 by [Boyle (1977)], was still regarded as a 'tool of last resort', to be tried when everything else failed. As for the variance-reduction techniques of the day, there was on offer little more than the use of antithetical or control variates [Hull and White (1988)], and some moment matching.

Almost coincidentally, a number of papers began to appear in this period, which introduced high-dimensional low-discrepancy sequences into the derivatives pricing arena [Ref here]. Despite the fact that these paper were originally introduced with equity/FX applications in mind, the timing could not have been better, because it was soon realized that they provided the tool needed to give a new impetus to interest-rate modelling.

²⁴The need to deal with path-dependent payoffs was strongly felt because of the popularity at the time of indexed-principal swaps mentioned above. In order to mimic burn-out effects of mortgage-backed-securities their notional was made to depend on the path of a reference (index) short rate.

²⁵[Nelson and Ramaswamy (1990)] had in the meantime addressed from a solid mathematical point of view the relevant convergence questions raised by the use of these intuitively appealing tools.

7.1 The HJM Results

The HJM no-arbitrage conditions have been presented in Equation 21. The main results of the approach can be summarized as follows:

1. the instantaneous forward rates can be used as the fundamental building blocks of the yield curve dynamics. Their exogenous specification ensures that a market yield curve can always be recovered;
2. all short-rate-based or bond-price-based models, as long as free of arbitrage, can be regarded as special cases of the HJM approach (see the discussion Section 3.2).
3. the no-arbitrage drifts of the forward rates are uniquely specified once their volatilities and correlations are assigned. Therefore model specification is exactly equivalent to these quantities. The apparent ability to specify independently drifts and volatilities enjoyed by, say, short rate models, is therefore purely a consequence of using a 'set of co-ordinates' that obscures the essential simplicity of the no-arbitrage yield curve dynamics
4. even if a single factor shocks the forward rates the resulting dynamics is no longer Markovian (an up-and-down shock does not give rise to the same yield curve as a down-and-up shock). As a consequence, recombining trees no longer offer a suitable computational tool, and, in particular, the process for the short rate is in general path-dependent²⁶.

At this stage, apart from the conceptual elegance and generality, the advantage of using the HJM was from obvious. In particular, since instantaneous forward rates are not directly observable, nor linked to the price of any traded instrument, the calibration of the HJM model, at least in this first incarnation, was no simpler than the calibration of a bond-price model, and certainly more complex than some of the then-popular short-rate models.

Traders were faced with another difficulty: since the market standard for the pricing of caplets was still linked to the (log-normal-forward-rates) Black model, it would have been natural to impose a log-normal behaviour to the instantaneous forward rates of the HJM approach. Unfortunately, log-normal forward rates were guaranteed to reach infinity with probability one in a finite time (the 'log-normal explosion') and one of the most natural modelling choices therefore appeared to be precluded. Remarkably inelegant and artificial ways to circumvent this problem were offered (see, eg,), but this situation of impasse was to be satisfactorily broken only by the introduction of the LIBOR market model.

²⁶[Carverhill A (1992)] showed which specifications of the forward rate volatilities are compatible with a Markovian short rate.

7.2 The LIBOR Market Models

One of the main goals of the LIBOR market model is to recover exactly the prices of caplets as produced by the Black model. As discussed above, the application of the Black model to the pricing of caplets (and swaptions) was originally done on a 'heuristic' basis, and only subsequently rationalized and justified in financially and mathematically acceptable terms. Doing so was important in order to ensure the logical consistence of the approach, but was not a very demanding task, since each individual caplet (or swaption) was considered in isolation, and priced in its own (terminal) pricing measure. All caplets and swaptions were therefore independently and simultaneously assumed to be log-normally distributed, even when this constituted a logical impossibility²⁷.

Such an approach is adequate to price caplets and swaptions in isolation. However, a complex product is, by definition, an instrument whose payoff depends on the *joint* probability of realization of several state variables. The LIBOR market model was therefore designed so as to recover exactly the market-standard (Black) prices of all the compatibly log-normal caplets while specifying a coherent and desirable dynamics for all the underlying forward rates.

Looked at in this manner, any model that reproduces exactly the prices of all the relevant plain-vanilla instruments can be regarded as a very high dimensional copula function that conjoins the various exogenous (ie market given) marginal distributions associated with the individual forward rates. So, if one works in terms of logarithm of forward rates, the LIBOR market model is approximately equivalent to a Gaussian copula.

The reason why it is not *exactly* equivalent to a Gaussian copula is important. In order to price a product whose payoff depends on the joint realizations of several forward rates, one has to work under a common pricing measure. The price to be paid in order to recover exactly *under the same measure* all the Black market prices of caplets is that, if one wants to prevent the possibility of arbitrage, the joint distribution of the logarithms of the forward rates can no longer be retained to be (exactly) a multi-variate normal²⁸. As a consequence, closed-form solutions for the evolution of the forward rate can no longer be obtained, and the resulting yield-curve dynamics is no longer Markovian.

It must be emphasized that a different strategy could have been employed. One could have arbitrarily chosen a numeraire and an associated common pricing measure. One could have then imposed that, *under this common measure*, all the marginal distributions should be exactly normal and created a model by imposing absence of arbitrage and requiring that the conjoining copula should be exactly Gaussian. This could have been a viable modelling route, but the consequence would have been that the market prices of caplets would no longer have been exactly recovered. Because of the discussion in section 2, however,

²⁷Nothing prevents a set of same-tenor and non-overlapping forward-rates to be simultaneously log-normally distributed. However, if a set of caplets is assumed to be log-normally distributed, the swap rate associated with them cannot be also log-normal, nor can, say, 6-month forward rates be log-normally distributed if 3-month forward rates are.

²⁸The reason for this is that the no-arbitrage drifts of the forward rates contain, for a general choice of numeraire, the state variables themselves. See Equations 64 and 65.

there is an intrinsic advantage for complex traders in their being able to recover the plain-vanilla prices, to some extent almost irrespective of the quality of the market-agreed model, and this logically possible route was not followed.

The strategy employed to develop the LMM approach can be reconstructed as follows. The first step is to recognize that the payoffs of most interest-rate derivatives products depend on the joint realization of a finite number of rates at pre-specified times. Therefore the pricing of LIBOR contingent claims depends on the continuous-time evolution of discrete-tenor forward rates, $F(t, T, T + \tau)$,

$$F(t, T, T + \tau) = F_t^T = \left[\frac{P(t, T)}{P(t, T + \tau)} - 1 \right] \frac{1}{\tau} \quad (62)$$

We saw in Section 3 that the quantities that it is natural to model as strictly positive semi-martingales spot or forward bond prices or spot or forward rates. We also saw, however, that these quantities are inter-related, and, because of Ito's lemma, the volatilities of three sets of variables are fully specified once the volatilities of any other set of variables are given. Furthermore, (again, because of Ito's lemma) if the volatilities of the elements of any one set are purely deterministic functions of time, the volatilities of all the elements of the other three sets will, in general, contain the state variables themselves, and therefore be stochastic. So, this modelling framework prepares the setting for a whole class of no-arbitrage specifications ('models') of the yield curve dynamics, each one associated with a choice for the covariance structure relating the rates or prices in any one set of variables.

This way of looking at the LMM is particularly useful because its different 'incarnations' can be easily and naturally understood in this light. The FRA-based LIBOR market model (see, eg, [Brace, Gatarek and Musiela (1995)], [Brace, Gatarek and Musiela (1996)], [Musiela and Rutkowski (1997)]) for instance, is obtained by imposing that the volatilities of *forward rates* should be deterministic functions of time; the swap-rate version (see [Jamshidian (1997)]) is obtained by identifying the strictly positive semi-martingales with the forward swap rates and imposing that *their* volatility vector should be deterministic; another version of the model (which now however loses the label 'market') would be obtained by identifying the semi-martingales $m_X(t)$ with the forward bond prices (see [Musiela and Rutkowski (1997b)]), and imposing that these should have a deterministic volatility. With the first choice, caplet (but not swaption) prices as produced by the model will be equal to their Black values; with the second choice swaption (but not caplet) prices will be in accordance with the Black formula; with the third choice neither will (but bond option prices will be Black-consistent).

As mentioned above, forward and swap rates cannot be simultaneously log-normal. Once, a forward- or swap-rate-based version of the LMM is chosen, the Black prices of the complementary associated plain-vanilla options (swaptions or caplets, respectively) cannot be simultaneously recovered. The 'luck of the model' stemmed from the fact that the pricing discrepancies arising from lack of simultaneous log-normality are in practice very small (see, eg, Rebonato (1998)), and certainly far too subtle to be arbitrated away. Therefore the logically

inconsistent Black market practice can be retained, and at the same time it can be not too severely at odds with the outputs of the logically coherent LIBOR market model.

7.2.1 Calibration of the LMM

The no-arbitrage dynamics of discrete-tenor forward rates, F , in a multi-factor LIBOR market model in terms of n orthogonal Brownian motions can be described as follows:

$$\frac{dF_t^{T_i}}{F_t^{T_i}} = \mu_i dt + \sigma_t^{T_i} \sum_{k=1,n} b_{ik} dz_k^t \quad (63)$$

with the drift of the i -th forward rate depending on the discount bond, $P_t^{T_j}$, used as common numeraire, and given by

$$\mu_i = \sigma_t^{T_i} \sum_{k=j+1}^i \frac{\sigma_t^{T_k} \rho_{ik}^t F_t^{T_k} \tau_k}{1 + F_t^{T_k} \tau_k} \quad \text{if } i > j \quad (64)$$

$$\mu_i = -\sigma_t^{T_i} \sum_{k=i+1}^j \frac{\sigma_t^{T_k} \rho_{ik}^t F_t^{T_k} \tau_k}{1 + F_t^{T_k} \tau_k} \quad \text{if } i < j \quad (65)$$

$$\mu_i = 0 \quad \text{if } i = j$$

It is easy to see that, if the market caplets are priced on the basis of the Black formula with implied volatility $\sigma_{Black}^{T_i}$, their recovery will always be guaranteed as long as

$$\sigma_{Black}^{T_i} = \sqrt{\frac{1}{T_i} \int_0^{T_i} (\sigma_u^{T_i})^2 du} \quad (66)$$

and

$$\sum_{k=1,n} (b_{ik})^2 = 1 \quad (67)$$

Equations 66 and 67 clearly show that recovery of any exogenous set of caplet prices is not only always achievable, but that one can do so in an infinity of ways. These equations also provide a link between the volatility of and the correlation among the forward rates: since one can easily derive that

$$\rho_{ij} = \sum_{k=1,n} b_{ik} b_{kj} \quad (68)$$

to each caplet-recovering specification of the instantaneous volatilities (ie, to each given matrix $\{b_{ik}\}$ such that Equation 67 is satisfied) there corresponds one and only one correlation matrix. This fact was exploited by [Rebonato (1999a)] in order to obtain the 'best' correlation matrix compatible with a chosen number of factors and with the exact recovery of the caplet prices. More generally, since Equation 66 underdetermines the instantaneous volatility function, the trader

now has additional degrees of freedom to recover other desirable features of the yield curve evolution. For instance, the trader might find desirable that the evolution of the term structure of volatilities should be as time homogeneous as possible. This can be achieved by starting from a function $\sigma_t^{T_k}$ of the form:

$$\sigma_t^{T_k} = g(T_k)h(T_k - t) \quad (69)$$

If the forward-rate-specific term $g(T_k)$ were constant, the evolution of the term structure of volatilities would be exactly time homogeneous. In general, exact pricing of market caplets cannot be achieved by imposing $\sigma_t^{T_k} = h(T_k - t)$, but, given a parametric form $h(T_k - t; \{\alpha\})$, the parameters $\{\alpha\}$ can be chosen in such a way that the function $g(T_k)$ is as close to unity as possible, thereby ensuring that the evolution of the term structure of volatilities should be as time homogeneous as possible given that all the market caplet prices have been recovered. See [Rebonato (2002)] for a detailed discussion²⁹. Useful parametric forms for the function $h(T_k - t)$ are discussed in [Rebonato (1999)], [Rebonato (2002)] and [Brigo and Mercurio (2001)].

Notice the difference with short-rate based approaches, where the fit to the market prices of caplets, if at all possible, tends to exhaust all of the available degrees of freedom, thereby leaving the option trader with no way to express a trading view on the very quantity (the evolution of the term structure of volatilities) in which she is expected to make a market. This feature arguably constitutes one of the strongest points of the LMM. Equally important were the following:

- despite the fact that the drifts of the forward rates contain the forward rates themselves (see Equations 64 and 65), and the latter are therefore *not* log-normally distributed, simple approximations were found (see [Hunter, Jaeckel and Joshi (2001)], based on work by [Kloeden and Platen (1992)], or [Pietersz et al. (2002)], [Kurbanmuradov et al. (2002)]) for their evolution over time periods as long as ten-to-thirty years;
- a simple but accurate approximation was provided to link the model-implied prices of swaptions for a set of forward rate volatilities (??, [Hull and White (2000b)], [Jaeckel and Rebonato (2000)], Latest Geneva SchCoff)
- a systematic methodology was presented to ensure that the best possible fit (given the chosen dimensionality of the LMM) could be achieved to an exogenous correlation matrix, while exactly recovering at the same the prices of caplets [Rebonato (1999a)];
- it was shown ([Rebonato (2000)], [Rebonato (2002)]) how to recover (almost) exactly an exogenous set of prices of co-terminal swaptions, while obtaining at the same time an almost time-homogeneous evolution of the

²⁹Whenever the quantity $(\sigma_{Black}^{T_i})^2 T_i$ is not a strictly increasing function of T_i it is easy to show ([Rebonato (2002)]) that in a Black world the evolution of the term structure of volatilities *cannot* be time homogeneous.

swaption matrix and the best possible fit (again, given the chosen dimensionality of the model);

- [Brace and Womersley (2000)] showed how to recover simultaneously caplet and swaption prices.

7.3 The Early-Exercise Problem

The greatest stumbling block in the acceptance of the LMM was at this stage the evaluation of compound options. This was an industry problem that could not easily be ignored, since Bermudan swaptions are possibly the most common exotic products. More generally, the search for greater and greater 'yield advantage' to issuers and investors naturally leads to the granting to the investment bank who receive the optionality in exchange more and more powerful options. These could be, for instance, the right to call a complex structure not on a single pre-specified date, but on any of a set of possible dates, thereby introducing compound optionality. This is indeed the route that has led from multi-callable zero-coupon swaptions to power reverse dual swaps.

It was the luck of the LMM that, in the very same years when it was introduced, techniques were being developed from independent research outside the interest-rate arena to deal with the problem of American options with many underlyings. These results, often, but not always, based on the estimation of the early-exercise boundary, were quickly adapted to the LIBOR setting. Work in this area include [Broadie and Glasserman (1997)], (which develops lower and upper bounds via simulation), [Broadie and Glasserman (1997b)], [Broadie, Glasserman and Jain (1997)], (a simulated-tree approach which works very well for Bermudan swaptions with a small number of exercise opportunities) [Jaekel (2000)], [Andersen and Broadie (2001)], [Broadie and Cao (2003)] (these papers deal with computing upper bounds for option prices given a sub-optimal exercise policy), and [Joshi and Theis (2002)].

The importance of these computational techniques cannot be over-emphasized: as in the earlier case of forward-induction with trees, they made all the difference between an theoretically-appealing model and a viable market standard. This state of affairs explains why extensions of the existing models to account for smiles has taken as a starting point an approach (the LMM) - which, after all, produces in its original formulation, no smiles at all - rather than any of the many models that naturally give rise to a smile, sometimes of very desirable nature³⁰.

8 Variations on the LMM Theme

Following, or contemporary to, the introduction of the LMM, other modelling approaches were developed. The main 'drawbacks' of the LMM that they at-

³⁰I argue in section 10 that part of the smile can be accounted for by a market-perceived deviation from log-normality of the forward rates. Square-root or mean-reverting Gaussian short rate processes could have provided a good starting point in this direction.

tempted to improve upon were:

- the imperfect simultaneous recovery of the Black cap and swaption prices
- the lack of recombination onto a low-dimensionality lattice of the high-dimensional Markov process required to describe the discrete yield curve.

Of these two perceived shortcomings, the former was much the milder, since it became apparent quite soon (see, eg, [Rebonato (1999)]) that the incompatibility of the simultaneous log-normal assumption for forward and swap rates had a very mild pricing impact. The first computational drawback, however, was more acutely felt, since Bermudan swaptions (which require the evaluation of compound optionality) are one of the most widely traded exotic products. The solutions to these problems are briefly discussed below. It should be said, however, that the increase calibration burden required was in general perceived to be too high a price to pay, especially as the Monte Carlo techniques to estimate the early-exercise boundary (applicable with the easy-to-calibrate LMM) became more and more common and efficient.

8.1 Low-Dimensionality Versions of the HJM Model

The simplest approach to tackling the high dimensionality of the Markov processes that describe the HJM evolution of the yield curve (and the consequent explosion of the associated bushy trees) is perhaps the one proposed by [Li et al (1995)] and [Ritchen et al. (1995)]. They prove that necessary and sufficient condition for the price of any interest-rate derivative to be function of a two-dimensional Markov process, $\chi = \chi(r, \Psi)$, is that the volatility, σ_t^T , of the instantaneous forward rates should have the separable of the form:

$$\sigma_t^T = g_t \exp - \int_t^T h_u du \quad (70)$$

with $g()$ an adapted process and h_t a deterministic function of time. If this is the case, the second factor that, together with the short rate fully describes the dynamics can then be shown to be given by the total variance up to time t of the forward rate of expiry t :

$$\Psi_t = \int_0^t (\sigma_u^t)^2 du \quad (71)$$

Closed-form expressions for the bond prices can then be obtained as a function of r and Ψ . Differentiating Equation 19 with $r_t = f(t, t)$ then shows that the function g that appears in the definition (Equation 70) of the forward-rate volatility σ_t^T is just the volatility of the short rate:

$$dr_t = \mu_r(t)dt + g_t dz_t \quad (72)$$

The approach is interesting, especially if $g()$ is taken to be a function of the short rate (for some choices of the functional dependence on r , for instance, an important component of the interest-rate smile can be modelled).

[Rebonato (2002)], however, discusses why the plausible choice 70 may not always be desirable.

The approach just described ensures that the dynamics of the instantaneous forward rate should be driven by a low-dimensionality Markov process, thereby ensuring recombination of the computational tree, and a reasonably easy evaluation of compound (eg Bermudan) option. Based as it was on instantaneous forward rates, however, this approach still left substantial calibration problems. The challenge was taken up in the discrete-forward rate setting by the more general models described below.

8.2 Low-Dimensionality Markov Market Models

There is no better example of how model development is driven by the joint requirements of calibration (recovery of the prices of plain-vanilla options) and ease of computation than the Markov-Functional (MF) models. The idea [Pelsser (2000)] is to take a low-dimensionality Markov process, x_t , and to define its relationships to the market prices in such a way that they assume the market-implied distribution. More precisely, let Q^{N+1} be the measure under which the last of a set of spanning forward rates is a martingale, (the 'terminal' measure) and let $dw^{Q^{N+1}}$ be the increment of a standard Brownian motion in this measure. A simple choice for x_t could be

$$dx_t = s(t)dw_t^{Q^{N+1}} \quad (73)$$

(with $s(t)$ a deterministic function of time). Therefore, with this choice x_t is both conditionally and unconditionally Gaussian (see [Nielsen (1999)]), with probability density $\phi(\cdot)$ ³¹. A Markov model is then built by requiring that the bond prices, P_t^T should be a *monotonic* function of the Markov process x_t :

$$P_t^T = P_t^T(x_t) \quad (74)$$

Let us denote by $P_t^{T_{N+1}}$ the price of the bond maturing at the payoff time of the last forward rate. Assuming this process to be known, the functional form of all the other bonds, $P_t^{T_i}, i \leq N+1$ is given by

$$\begin{aligned} \frac{P_t^{T_i}}{P_t^{T_{N+1}}} &= E^{Q^{N+1}} \left[\frac{P_{T_i}^{T_i}}{P_{T_i}^{T_{N+1}}} | \mathcal{F}_t \right] = E^{Q^{N+1}} \left[\frac{1}{P_{T_i}^{T_{N+1}}} | \mathcal{F}_t \right] = \\ &= \int_{-\infty}^{\infty} \frac{1}{P_{T_i}^{T_{N+1}}(u)} \phi(u|x_t) du \end{aligned} \quad (75)$$

Given a process for x_t , the model becomes fully specified by choosing a functional form for the discount bond. This choice can be made so as to ensure that the prices of caplets (LIBOR MF model) or swaptions (swap MF

³¹The 'old' short rate models were clearly a special case of a MF model, where the one-dimensional Markov driver is the short rate. In particular, if the process for the short rate was assumed to be of the Ornstein-Uhlenbeck type, as in the Vasicek model, the density $\phi(\cdot)$ would be known to be conditionally Gaussian.

model) are correctly recovered [Pelsser (2000)]. These prices might, but need not, be compatible with the Black formula³². The recovery of their marginal distributions could be achieved for a variety of functional choices for x_t . The quality of a model, however, is characterized not just by its ability to recover prices today (recovery of the marginal distributions), but also by its ability to specify in a desirable way the future prices (specification of conditional distributions). The functional form of the driving Markov process should therefore be chosen in such a way as to exploit these degrees of freedom in the most desirable way (eg, to achieve mean reversion, time homogeneity, etc). A variety of procedures (either parametric, semi- or non-parametric) are possible. See [Hunt and Kennedy (2000)] or [Pelsser (2000)] for details.

The main advantage of the model is the fact that it can be mapped onto a recombining lattice, thereby facilitating the evaluation of compound options (Bermudan swaptions *in primis*, especially in the swap-rate implementation). The price to be paid for this is that calibration to the market prices of caplets requires the solution of a non-linear integral equation, which must be solved numerically. Also, the curse of dimensionality constrains the Markov process to be one- or at most two-dimensional. Whether this limitation is important for the pricing of Bermudan swaptions is a hotly debated topic, briefly touched upon in the following sections.

9 Other Approaches

9.1 Positive-Rate Models

In this approach ([Flesaker and Hughston (1996)], [Rutkowski (1997b)]) the pricing kernel (see Section 2.4) becomes the fundamental quantity that drives the dynamics of bonds and rates. The starting point is the real-world (objective-measure) Equation 26. After defining $V_k(t, T) \equiv v_k(t, T) - \lambda_k^t$ and $V(t, T) \equiv v(t, T) - \lambda^t$, this can be re-written as:

$$\frac{P_t^T \rho_t}{B_0^t} = P_0^T \exp \left[\int_0^t \left(\sum_{k=1, n} V_k(s, T) \right) dw_k^s - \frac{1}{2} \int_0^t [V(s, T)]^2 ds \right] \quad (76)$$

³²In recovering exactly *today's* prices of caplets and swaptions, the MF models achieve something that the LMM cannot. The extra 'degrees of freedom' stem from the fact that the LMM requires not only that today's distribution of forward or swap rates should be log-normal (as the MF model does), but also that, in the associated terminal or swaption measure, forward or swap rates should be log-normal martingales. This is not required by the MF approach, which only recovers a log-normal *marginal* distribution for the forward (swap) rates, and the martingale property. As a consequence, however, future conditional distributions will not, in general, be log-normal, and the future conditional prices of caplets and European swaptions will not be correctly recovered. See Section for a discussion of the importance of this point. Luckily, the deviation from log-normality of swap rates given log-normal forward rates (or vice versa) is sufficiently small that the 'deformation' of the log-normal dynamics implied by the MF model is often small.

Setting $T = t$ allows to solve for $\frac{P_t^T \rho_t}{B_0^t}$ (because $P_t^t = 1$):

$$\frac{\rho_t}{B_0^t} = P_0^t \exp \left[\int_0^t \left(\sum_{k=1, n} V_k(s, t) \right) dw_k^s - \frac{1}{2} \int_0^t [V(s, t)]^2 ds \right] \quad (77)$$

Therefore Equation 76 becomes

$$P_t^T = \frac{P_0^T \exp \left[\int_0^t \left(\sum_{k=1, n} V_k(s, T) \right) dw_k^s - \frac{1}{2} \int_0^t [V(s, T)]^2 ds \right]}{P_0^t \exp \left[\int_0^t \left(\sum_{k=1, n} V_k(s, t) \right) dw_k^s - \frac{1}{2} \int_0^t [V(s, t)]^2 ds \right]} \quad (78)$$

The numerator depends on t and T , and can therefore be designated as Δ_t^T and, for a given T , is an exponential martingale, initialized at $\Delta_0^T = P_0^T$. The denominator only depends on t , and can be written as Δ_t^t . Therefore

$$P_t^T = \frac{\Delta_t^T}{\Delta_t^t} \quad (79)$$

If one makes the (reasonable) assumption that

$$\lim_{T \rightarrow \infty} P_t^T = 0 \quad (80)$$

it follows that

$$\lim_{T \rightarrow \infty} \Delta_t^T = 0 \quad (81)$$

and therefore

$$\Delta_t^T = - \int_T^\infty \frac{\partial \Delta_t^s}{\partial s} ds \quad (82)$$

Given the initialization condition $\Delta_0^t = P_0^t$ and the fact that $\frac{\partial \Delta_t^s}{\partial s}$ is a martingale, one can write

$$M_t^T = \frac{\frac{\partial \Delta_t^s}{\partial s}}{\frac{\partial \Delta_0^s}{\partial s}} = \frac{\frac{\partial \Delta_t^s}{\partial s}}{\frac{\partial P_0^s}{\partial s}} \quad (83)$$

with $M_0^T = 1$. Therefore Δ_t^T , the quantity that characterizes the dynamics of the bond prices (the 'model') can be written as

$$\Delta_t^T = - \int_T^\infty \frac{\partial P_0^s}{\partial s} M_t^s ds \quad (84)$$

Positivity of future rates is ensured if the family of martingales M_t^T are positive (and so are the initial rates). This constitutes a 'positive interest' representation of a class of general interest rate models. Each choice for the family of martingales M_t^T specifies one particular model.

9.1.1 Rational Models

Rational models [Flesaker and Hughston (1996)] are obtained by prescribing that M_t^T should be of the form:

$$M_t^T = \alpha_T + \beta_T M_t \quad (85)$$

with α_t and β_t positive, deterministic functions satisfying $\alpha_t + \beta_t = 1$, and M_t any positive martingale such that $M_0 = 1$. It is then easy to show that

$$P_t^T = \frac{\Delta_t^T}{\Delta_t^t} = \frac{F_T + G_T M_t}{F_t + G_t M_t} \quad (86)$$

for F_T and G_T positive decreasing functions satisfying $F_T + G_T = P_0^T$. For the current term structure to be recovered it is enough to require that the functions F and G satisfy

$$P_0^T = \frac{F_T + G_T}{F_0 + G_0} \quad (87)$$

The rational-model route is appealing for its relative simplicity, but suffers from the drawback that it can be shown ([Babbs (1997)]) that the bond prices and the short rate are constrained to lie between

$$\frac{F_T}{F_t} \geq P_t^T \geq \frac{G_T}{G_t} \quad (88)$$

and

$$-\frac{G'_t}{G_t} \geq r_t \geq -\frac{F'_t}{F_t} \quad (89)$$

This can easily make calibration to market data impossible. This shortcoming is not shared, however, by the potential approach.

9.1.2 Potential Models

The potential model adapts the approach by [Constantinides (1992)], which takes the pricing kernel (relative risk price density) positive super-martingale ρ_t (Equation 25)

$$\rho_t = \exp \left(- \int_0^t \sum_{k=1,n} \lambda_k^s dw_k^s - \frac{1}{2} \int_0^t \lambda_s^2 ds \right)$$

as the fundamental state variable for the bond price dynamics) to the positive-interest framework. The conceptual link with the positive-interest model is achieved by noticing [Hughston and Brody (2000)] that $\Delta_t^t = \rho_t$. If the extra condition

$$\lim_{t \rightarrow \infty} E[\rho_t] = 0 \quad (90)$$

is imposed (and technical conditions are satisfied), the quantity ρ_t is a potential (see [Rogers (1997)]), whence the name or the approach.

9.1.3 Link Between MF and Positive-Interest Models

The MF approach requires bond prices to be a functional of a low-dimensionality Markov process. Recalling ([Hunt, Kennedy and Pellser (2000)]) that the state-density ρ_t is associated with a numeraire $(1/\rho_t)$, and that the positive-interest models make the requirement that the numeraire should depend on a low-dimensionality Markov process, one can see the conceptual link between the two. The main difference is the flexibility to choose functional forms for the Markov functionals to match market option prices. These are (too?) easily recovered with the MF model, but with much greater difficulty with the positive-interest approach. Whatever the relative fundamental explanatory appeal, the former have been used by traders much more than the latter.

10 Phase 5: Accounting for Smiles

In a way, the timing of the appearance and acceptance of the LMM turned out to be almost ironic. No sooner had the approach described in Section 7 earned the 'market' label and won near-universal acceptance than progressively marked smiles began to appear³³. Clearly, these smiles were indicating that the log-normal-diffusion paradigm that had won widespread acceptance in the plain-vanilla option market was no longer adequate. At the beginning the deviations were seen as a relatively small 'perturbation' of a basically correct approach, but following the market turmoil of 1998, it became increasingly difficult to ignore smiles, or to treat them in a perfunctory manner. As a result, a series of approaches were (and are being) introduced in order to account for this phenomenon.

The challenge was not so much to provide modelling approaches capable of reproducing smiles, but to do so in a financially realistic manner, and by retaining as much as possible the computational techniques that had made the LIBOR market model possible and successful. In particular, ease of calibration had to be retained if the new approaches were to win any degree of market acceptance. By 1998 the pricing community had almost completely converted to the way of thinking associated with the LMM, and it was therefore natural to look for extensions that could be naturally grafted onto the existing structure.

10.1 First-Generation Smiles

In the first phase (pre-1998), the smiles were typically monotonically decreasing from the high- to the low-strike. The fact that they first appeared in JPY,

³³The name 'smile' originates from the equity-index option world, where deviations from the horizontal line were originally observed to appear after the 1987 market crash (Rubinstein and Jackwerth). The shape of the implied volatility curve as a function of strike was observed to be an increase in implied volatilities both for out-of-the-money calls and puts (with respect to the at-the-money volatility). The resulting upward-pointing half moon gave rise to the name 'smile'. Despite the fact that more complex shapes have appeared, I use the term smile to denote any non-falt implied volatility curve.

where rates were heading towards very low levels as early as 1995, suggested that the implicit assumption of proportionality of the random shock to the prevailing level of rates was being questioned. However, with rates in USD and in the major European currencies in the 6-to-9% range, the case of the JPY was originally regarded as an aberration, and the USD and European-currency implied volatility surfaces displayed little or no smile.

However, in the second half of the 1990s (especially with the convergence of many continental currencies to the Euro) rates began to decline significantly in Europe as well. Traders were still reluctant to abandon the Black pricing paradigm, but it began to be felt that the out-of-the-money floor (receiver-swaption) area was a 'special case', to be handled as a sort of 'exception' to the log-normal rule. So little was the initial consensus as to how to handle these 'special situations' that one of the most notorious cases of option mispricing of that period was associated with a UK bank failing to take into account smiles for deeply out-of-the-money floors.

By, approximately, the summer of 1998 a monotonically decreasing smile was clearly observable for, and taken into account into the pricing of, plain-vanilla caps and swaptions. Virtually no trading houses, however, were incorporating smiles in the pricing of complex interest-rate derivative products in a systematic way.

10.2 Second-Generation Smiles

The events that followed the Asian South-East Asia currency crisis of October 1997, the Russia default and the near-collapse of LTCM in 1998 brought about unprecedented dislocations in several markets. The interest rate volatility surfaces were not immune from these upheavals, and many exotic traders found themselves experiencing heavy losses (IFR). One of the reasons for these losses was that the deterministic-volatility LIBOR market model implied one and only one set of future caplet and swaption surfaces. In 'normal' periods, implied volatilities 'vibrate' in such a way as to retain, at least approximately, the shape of the caplet and swaption matrix³⁴. This situation can be adequately handled by a deterministic-volatility model as long as the model-implied future smile surface is not too dissimilar from the average of the realized implied volatilities. However, as the volatility surface underwent dramatic and sudden shape changes in 1998, the hedges suggested by the LIBOR market models calibrated to the normal volatility regimes often proved dramatically wrong. (see [Rebonato and Joshi (2002)] for an empirical discussion of the modes of deformation of market swaption matrices during this period).

Finally, the situation was made worse by the realization that certain products could display pronounced negative 'gamma-vega', which would be systematically

³⁴see, eg, [Rebonato (1999)] or [Rebonato (2002)] for an illustration of the persistence of the qualitative shape of caplet surfaces during 'normal' market conditions

ignored by deterministic-volatility models³⁵.

10.3 The Poverty of Pure Price Fitting

Because of this state of affairs the LMM had to be radically altered if it was to retain mainstream acceptance. By drawing on the experience accumulated in the equity and FX area, where smiles had been observed for a much longer period, it was not difficult to think of ways (eg, stochastic volatility, jump-diffusions, local volatilities, etc) to produce smiley volatility surfaces. The question immediately arose, however, as to how one should choose between these different possible mechanisms.

As in the areas of derivatives pricing that had experienced smiles for a much longer period of time, at the beginning price fitting across strikes was perceived to be the most important requirement. It was soon realized, however, that many very different modelling approaches could provide fits of very similar quality. Indeed, in a very insightful paper, [Britten-Jones and Neuberger (1998)] show that, given *any* stochastic volatility process, it is always possible to add a volatility component functionally dependent on the underlying, such that an arbitrary (admissible) exogenous smile surface is exactly recovered³⁶.

If a good fit to market option prices provides a necessary but not sufficient condition for the acceptance of a model, what criteria can be used to choose a satisfactory modelling approach? The answer to this question must take into account that, as we allow either for stochastic volatility or for discontinuous processes, markets are now certainly incomplete³⁷. Therefore, 'locking in' market-implied quantities via dynamic trading is no longer an option, and it becomes necessary to have an adequate actuarial description of what cannot be exactly hedged. In particular, of paramount importance will be the ability of

³⁵The term gamma-vega denotes the change in vega as the level of rates changes:


$$\frac{\partial^2 V}{\partial \sigma^2}$$

Traders fear products such that require the trader, in order to remain vega neutral, to buy volatility as it increases and to sell it when it falls. This situation (reminiscent of being short gamma in the underlying, whence the name) would produce systematic reheding costs even in the absence of bid-offer spreads. Deterministic-volatility models do not allow for stochastic movements in the volatility, and therefore automatically assign zero cost to this re-hedging strategy. A stochastic-volatility model 'knows' about (some) possible modes of vibration of the smile surface, and can therefore incorporate information about this price bias. From a trader's perspective this is one of the most desirable features of stochastic-volatility models.

³⁶This result is conceptually similar to the approach by Hull and White, who showed how, given an arbitrary market yield curve, it is always possible to recover it exactly by superimposing a deterministic function of time to the drift of the short rate. See Equation 54.

³⁷It is sometimes said (see, eg, [Duffie (1996)]) that, if a finite number of plain-vanilla options are added to the set of hedging instruments, the market can be completed even in the presence of stochastic volatility. Similarly, if only a finite number of discontinuous jumps can occur, market completion can in theory be achieved by suing in the replicating portfolio as many plain-vanilla options as possible jump amplitudes. This is correct, however, only if the full process (and not just the price today, of these hedging plain-vanilla options were known exactly. To ensure that these price processes allow no arbitrage with the process for the underlying(s), exogenous risk vectors have to be specified, *and assumed to be constant*.

Fig 1: US\$ data (April 1998- November 2002). Scatter plot of the 1 x 1 implied volatility (y axis) versus swap rate (x axis)



the model to recover, at least in a statistical sense, the relevant features of the future smile surface encountered during the life of the option (see Section 2.2). I shall discuss in Section 11 whether these real-world-measure estimates can be obtained from traded option prices. Irrespective of how these estimates should be arrived at, a 'good' model will be one that incorporates in a satisfactory manner the relevant elements of the smile dynamics³⁸. This brings about a new emphasis on econometric information, and therefore a brief analysis of stylized empirical facts about interest rate smiles is presented below.


10.4 Decomposition of the Volatility Drivers

An important insight into the best way to account for stochastic volatility can be obtained by looking at the joint behaviour of swap rates and of the associated implied volatilities. Fig. 1 shows a scatter plot of the changes in implied volatilities of the USD 1 x 1 swaption rate over the period X to Y against changes in the associated forward swap rate. Fig. 2 then shows a scatter plot of the changes in the swap rate against the quantity y defined as

$$y = \sigma_{impl} f^\beta \tag{91}$$

³⁸The expression 'relative elements' refers to the fact that if, say, a stochastic-volatility approach is chosen, the drift of the volatility will undergo a Girsanov transformation from the objective to the pricing measure (see, eg, [Lewis (2000)]), and therefore drift-related statistics in the real and model world cannot be naively compared. [Rebonato and Joshi (2002)] show how to circumvent this problem.

Fig 2: US\$ data (April 1998- November 2002). Scatter plot of the 1 x 1 transformed implied volatility (y axis) versus the swap rate (x axis)



(where σ_{impl} is the implied volatility, f is the forward swap rate, and β a constant). It is clear that this simple change of variables is sufficient to account for a large part of the variability in the implied volatility. See also Fig. 3, where both the percentage and the 'normalized' volatility (Equation 91) were arbitrarily rescaled to 1 at the beginning of April 2002: if one looked at the top line, which describes the percentage volatility usually quoted in the market, one would conclude that it is necessary to use a model where the volatility is highly variable. Indeed the levels reached by the *percentage* volatility in the second half of 2002 are exceptionally high by historical standards, and one would guess that it is desirable to use a stochastic-volatility model capable of doubling the implied volatility over a few months' period. However, the picture changes dramatically if one looks at the rescaled volatility, y , whose time series is depicted in the bottom line. The variability of this quantity is much more limited, and, if anything, this rescaled volatility is seen to be *decreasing* over the same period. The message is reinforced by Fig. 4, which shows the 1 x 1 forward swap rate and implied volatility over a two-year period. The very high degree of dependence between the two stochastic quantities is apparent, and clearly indicates that prescribing a stochastic behaviour for the *percentage* volatility is a very inefficient way to account for the observed data. The conclusion from this discussion is that the choice of log-normal 'co-ordinates' used in the standard LIBOR market model becomes inappropriate when dealing with stochastic volatility, and that it is advisable to carry out a transformation of variables first.

Fig 3: Rescaled (lower curve) and percentage (upper curve) volatilities (1 x 1 series) rebased at 1 on 1-Apr-2002

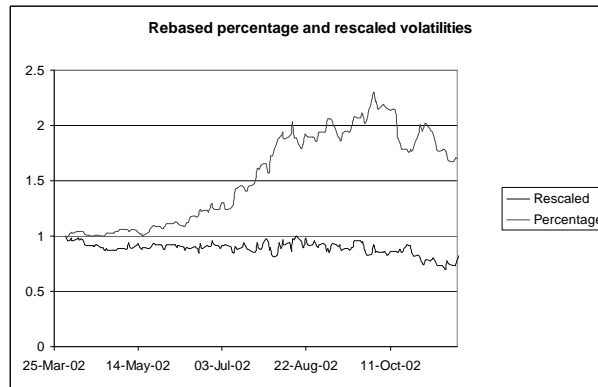



Fig 4: Swap Rate multiplied by 10 (upper curve) and swaption implied volatilities (lower curve) for the 1 x 1 series. US\$ data for the period 5-Jun-98 to 22-Nov-2002



10.5 The Proposed Approaches

The implications of these empirical observations for the lognormal or otherwise nature of the forward rate process were quickly taken on board by most traders and researchers. Therefore, early attempts to account for smiles by introducing a discontinuous (jump) component to the log-normal diffusive process for the forward rates (see, eg, [Glasserman and Kou (2000)], [Glasserman and Merener (2001)], [Jamshidian (1999)]) met with little acceptance. More popular were early extension of the LIBOR market model, such as [Andersen and Andreasen (1997)], [Andersen and Andreasen (2000)], and [Zuehlsdorff (2001)], who advocated a CEV ([Beckers (1980)]) process (see [Andersen and Andreasen (2002)] for a more general class of models in the same spirit). [Marris (1999)] then pointed out that, over a wide range of interest rate levels, the much more tractable displaced-diffusion process [Rubinstein (1983)] produces risk-neutral densities (and hence prices) extremely similar to the CEV approach, and can therefore be used as an effective and computationally simpler alternative.

CEV (or displaced-diffusion) processes constitute a more appealing set of 'coordinates', but still imply a perfect functional dependence of the volatility on the forward rates (and a monotonically decreasing smile). The next modelling step was therefore to require that one or more independent Brownian shock affect the volatility process. This is the route followed by [Joshi and Rebonato (2001)], Andersen (Barcelona). In particular, [Joshi and Rebonato (2001)] show how to extend most of the implementation techniques developed for the log-normal market model (eg, ease of calibration to caplets and swaption, optimal recovery of co-terminal option prices, etc). The quality of their fits to market data with time-homogeneous parameters is good. They also show that the (complex) shape of the model-produced eigenvectors obtained by orthogonalizing the matrix of the changes in implied volatilities bears a strong resemblance with the empirical data [Rebonato and Joshi (2002)]. Their use of the simpler displaced-diffusion approach, which does not guarantee positive interest rates, can, however, give rise to some concerns in the very-low-interest-rate environments experienced in most currencies at the time of writing.

11 The Calibration Debate

The treatment so far has shown that, in order to price interest-rate derivatives in the complete-market framework, one could either prescribe the whole real-world dynamics for the driving factor(s) (eg, the short rate) and for the associated risk premia (absolute pricing), or assign the volatility and correlation functions (the covariance structure, for short) of the stochastic state variables (relative pricing).

While both routes are in principle possible, for practical trading purposes the relative-pricing route has been almost universally adopted. Therefore, specification of a relative-pricing, arbitrage-free model in a complete-market setting has in trading practice become tantamount to assigning the covariance structure.

This is particularly transparent in the HJM approach (see Equation 21), but, given the equivalence among the different formulations, obviously also applies to all the other models.

After HJM, the statement that volatilities and correlations are 'all that matters' in derivatives pricing has become rather common-place. When this claim is combined with the market practices described in Section 2.4 (out-of-model hedging and model re-calibration), however, it produces some important corollaries, which have a direct bearing on calibration. This is because, ultimately, any choice of volatilities and correlations will determine the model-implied future conditional prices of caplets and swaptions where the future re-hedging trades will be transacted. But, as pointed out in section 2.4, the universal practice of re-calibrating the model and of rebalancing the vega hedges during the life of the complex trade requires that the model should recover to a satisfactory degree the future prices of the hedging instruments. The fundamental calibration question therefore becomes: 'What sources of information can most reliably provide an estimate of the covariance structure capable of producing these desirable future prices?' The answer to this question is as, if not more, important than choosing the 'best' model.

11.1 Historical versus Implied Calibration

When the problem is looked at in this light, the estimation of the volatilities and correlations can be arrived at either using historical estimation or via the implied route. When models have been regarded as pricing tools rather than general descriptions of the yield curve dynamics, both academics and practitioners have tended to embrace the implied route with far more enthusiasm than the statistical approach. Furthermore, the ability of a model to recover simultaneously as many 'market-implied' features as possible (eg, implied instantaneous volatilities *and* correlations from the prices of caplets *and* swaptions), has generally been regarded as highly desirable. See, for instance, [Schoenmakers and Coffey (2000)], [Brace and Womersley (2000)], [De Jong, Driessen and Pelsser (1999)], [Lane and Marris (2002)], Andersen Barcelona³⁹. Much as the implied route might appear 'natural' to a trading community trained in footsteps of the BS approach, it should not be taken as self-evident. Indeed, if in the hedging of a product (say, a Bermudan swaption) one set of plain-vanilla instruments (say, the associated co-terminal swaptions) will mainly be used, it should be at least open to debate whether it is more important for the chosen model to reproduce accurately the *current and future* prices of the required hedging swaptions, or to recover the *current* prices of as many plain-vanilla options as possible. See again the discussion in Section 1.6 and in the last sections of [Jamshidian (1997)].⁴⁰

³⁹Incidentally, this practice might have been partly motivated or encouraged by the fact that the LIBOR market model has just enough degrees of freedom to fit exactly, if one so wanted, all the caplet and European swaption prices.

⁴⁰The question is far from academic, since one of the most heated debates in interest-rate derivatives pricing has been brought about by the a recent paper suggestively titled '*Throwing Away a Billion Dollars: The Cost of Sub-Optimal Exercise*

11.2 The Logical Underpinning of the Implied Approach

Complex traders will typically want to remain vega neutral during the life of a trade. I have highlighted in Section 1.6 that this practice should be analyzed together with the requirement to generate as stable a calibration as possible. What estimation methodology will give rise to the most stable calibration? If a trader believes the 'implied' route to be the best way to estimate the inputs to a model, this must be justified on the basis of one of these possible alternatives:

1. The input functions (eg, volatilities or correlations of forward rates) are deterministic and perfectly known by the market. In this case the covariance matrix (that fully determines pricing) is perfectly deterministic. We are in a pure relative-pricing/ perfect-replication setting, accurate simultaneous calibration of the model to plain-vanilla prices (caplets and swaptions) is a must, and the resulting implied deterministic volatilities and correlations are all that matters for pricing. This is the approach implicitly or explicitly taken by [Schoenmakers and Coffey (2000)], among others.
2. The input functions are deterministic, but they are not perfectly known by the market. There exist additional liquid benchmark instrument, however, in addition to caplets and swaptions, that allow to complete the market, and to 'lock in' whatever values for the covariance elements are implied by the market. These instruments would be serial options (see [Rebonato (2002)]). Taken together, the prices of caplets, swaptions and serial options allow a unique determination of the covariance structure that determines the prices, not only of the plain-vanilla, but also of the exotic LIBOR products([Rebonato (2002)]). Since we are still in a pure relative-pricing/ perfect-replication setting, the trader need not worry whether these implied values are econometrically plausible, (ie, whether the market misestimated the covariance structure), because she can 'lock in' the implied prices via dynamic trading⁴¹.
3. *Either* [the input functions are deterministic and not perfectly known by the market, and there exist no additional liquid instrument besides caplets and swaptions], *or* [the input functions are stochastic], *but*, in either case, there is no systematic imbalance of supply or demand for plain-vanilla options. The market is now incomplete, but the market-implied values for

Strategies in the Swaptions Markets'. See [Longstaff, Santa Clara and Schwartz (2000a)], [Longstaff, Santa Clara and Schwartz (2000b)], who argued that using low-dimensional implementations of the LIBOR market model does not allow to capture adequately the value embedded in Bermudan swaptions, and the response by [Andersen and Andreasen (2001)]. A careful analysis of the arguments indeed reveals that either conclusion can be justified, depending on what else the model is required to fit to, in addition to the co-terminal swaption prices. See also [Joshi and Theis (2002)].

⁴¹The situation is similar to the pricing of forward-rate-agreements (FRAs), *given the prices of discount bonds*. As long as a trader can freely take long and short positions in the bonds, the prices of the FRAs are arrived purely by a no-arbitrage arguments, and the 'real-world plausibility' of the implied forward rates is totally irrelevant.

volatilities and correlations are an unbiased estimate of their assessment by the market without any correction for risk (even if the individual players are risk-averse). Hedging is now in general imperfect, and a mixture of relative and absolute pricing is appropriate. For the pricing and hedging of the undiversifiable risk the real-world realization of the volatilities is relevant, and, given the assumption of balanced supply and demand, the market still provides an unbiased estimates. Calibration of the model to the market-implied volatilities and correlations should be carried out.

4. We are in the same situation as in 3., but there now is a systematic imbalance of supply or demand for some plain-vanilla options. However pseudo-arbitrageurs are active with no size constraints. So, for instance, if the implied correlations were deemed on the basis of statistical analysis to be, say, too high, the pseudo-arbitrageurs would enter trades whereby they would be 'long correlation' (perhaps by trading in swaptions), and hedge the volatility exposure (perhaps by trading in correlation-independent caplets). As a result, this would bring the implied correlation in line with fundamentals. Effectively we are in the same situation as in case 3., and calibration of the model to all the plain-vanilla prices should still be attempted.
5. We are in the same situation as in 4., *but pseudo-arbitrageurs are not able to carry out their trades*. If this is the case, as in case 4., we are no longer in a pure relative-pricing/perfect-replication setting, realizations in the objective measure do matter, but, unlike case 4, the trader can no longer rely on the market prices to provide an unbiased estimate of volatilities and correlations. Implying the values of volatilities and correlations from the simultaneous calibration to all the plain-vanilla prices should not necessarily be attempted. Whenever plain-vanilla prices are recovered by a calibration strategy, care should be given to ensuring that some real-world quantities are also plausibly reproduced⁴².

Very few traders or researchers would subscribe to 1 (if volatilities were known and deterministic, for instance, it would be difficult to justify the universal practice of vega hedging). Therefore advocating an 'implied' estimation relies either on perfect replication (possibility 2) or on the informational efficiency of prices (options 3 and 4).

Given that serial options are both illiquid and only available for a short range of maturities/expiries, perfect replication of plain-vanilla payoffs is not an option even if volatilities were truly perfectly deterministic (but imperfectly

⁴²Recall, however, that, even for Girsanov-transformed quantities, the estimate of their risk-adjusted values from the market prices contains some useful information. For instance, if a residual imbalance of, say, *demand* is assumed to be at play, this would at least determine the sign of the deviation from the expectation in the objective probability measure: if equity smiles are due to portfolio managers seeking insurance against undiversifiable jumps, the prices of out-of-the-money puts should be obtained with a risk-adjusted jump frequency *higher*, and not just *different*, from the objective one.

known). Therefore possibility 2 does not appear convincing. Fitting a model simultaneously to all the available market information must therefore rely on the assumed informational efficiency of the plain-vanilla market prices. A failure of efficiency however requires two simultaneous conditions: i) that a systematic imbalance in supply and demand for a derivative exist and ii) that limitations be in place to the actions of pseudo-arbitrageurs. I tackle these points in the following sections.

11.3 Are Interest-Rate Derivatives Market Informationally Efficient?

A large body of literature has appeared in the last ten years or so, which challenges one of the pillars of the classical financial asset pricing, namely the Efficient Market Hypothesis (EMH). The name generically applied to these rather disparate studies is that of 'behavioural finance'. In short, two joint claims are implicitly made (although not always explicitly articulated) by the proponents of this school, namely that i) at least some investors arrive at decisions that are not informationally efficient and ii) that mechanisms that would allow better-informed traders to exploit and eliminate the results of these 'irrationalities' are not always effective. Since the prime mechanism to enforce efficiency is the ability to carry out (pseudo-)arbitrage, an important line of critique of the EMH has been developed (see, eg, [Shleifer and Vishny (1997)]) which shows that pseudo-arbitrage can in reality be very risky, and that, therefore, the pricing results of irrational decisions made on the basis of psychological features such as, say, overconfidence might persist over long periods of time.

In order to account for the *origin* of the pricing inefficiencies, the original emphasis was put on psychological features, such as, for instance, overconfidence, anchoring, framing effects, etc (see, eg, [Shefrin (2000)], [Shiller (2000)]), whence the name *behavioural* finance. The argument based on the difficulty and riskiness of pseudo-arbitrage can, however, still be applied if the price of an asset (and, in our case, of a derivative) is disconnected from 'fundamentals' for any (ie, not necessarily for psychological) reasons: agency conditions, for instance, can give rise to failure of the EMH even if all the players are fully rational. This is important, because in the interest-rate derivatives area, (largely the arena of professional traders), it is more likely that institutional constraints, rather than psychological biases, might be at the root of possible price deviations from fundamentals.

The relevance of these possible informational inefficiencies for derivatives pricing can be seen as follows. First, according to the EMH, prices are arrived at by discounting future expected payoffs using an appropriate discount factor⁴³. The second step in the argument is that new information (a change in

⁴³'Appropriate' in this context means, on the one hand that it takes the riskiness of the cashflows into account, but, on the other, that it is only affected by non-diversifiable risk. So, if a security (an option) can be replicated by another (the hedging portfolio) no idiosyncratic risk will be left and the appropriate discount factor is derived from risk-less bonds.

'fundamentals') can lead the trader to reassess the current price for a derivative (a new expectation is produced by the augmented filtration), but supply and demand pressures *per se* cannot: if the 'fundamentals' have not changed, a demand-driven increase in the price of substitutable security will immediately entice pseudo-arbitrageurs to short the irrationally expensive security, and bring it back in line with fundamentals. The more two securities (or bundles of securities) are similar, the less undiversifiable risk will remain, and the more pseudo-arbitrageurs will be enticed to enter 'correcting' trades.

So, answering the question, 'To what extent should one make use of market-implied quantities as input to a model?' means addressing the joint two questions: 'Are there reasons to believe that a systematic imbalance of supply or demand might be present in the interest-rate plain-vanilla market?' and 'Are there reasons to believe that the activity of pseudo-arbitrageurs might entail substantial risks?'

11.3.1 Possible Mechanisms to Create a Supply/Demand Imbalance

The dynamics of supply of and demand for interest-rate derivatives products are very complex, especially in US\$, where the mortgage-backed securities market creates a large demand for a variety of derivatives products. In broad terms, however, some relatively simple patterns can be identified: on the one hand there are investors looking for 'yield enhancement' and issuers in search of 'advantageous' funding rates; on the other hand there are floating-rate borrowers who want to reduce their risk by purchasing interest rate protection. In order to obtain the advantageous yields or funding rates investors or issuers, respectively, tend to sell the right to call or put a bond, ie, swaption-type optionality, which is typically 'sold-on' to investment houses. See [Rebonato (1998)] for a description of these structures. These will therefore find themselves systematically *long* swaption optionality.

At the same time, floating-rate corporate borrowers will seek to limit their exposure to rising rates by purchasing caps from the same trading desks. The latter will therefore find themselves systematically long swaption optionality and short caplet optionality.

A case has therefore been made as to why there might be a systematic excess demand of cap volatility and excess supply of swaption volatility. Clearly, however, both caps and swaptions share the same underlyings (ultimately, forward rates). Therefore pseudo-arbitrageurs should be enticed to take advantage of whatever move away from 'fundamentals' the supply/demand imbalance might create in the relative prices of caps and swaptions. Are there reasons to believe that their effectiveness might in practice be hampered?

11.3.2 Possible Limitations to Pseudo-Arbitrageur Activity

What can prevent pseudo-arbitrageurs from carrying out their task of bringing prices in line with fundamentals? To begin with, (see [Shleifer (2000)] or [Shleifer and Vishny (1997)] for a fuller discussion) these pseudo-arbitrageurs

(hedge funds, relative-value traders, etc) often take positions not with their own money, but as agents of investors or shareholders. If the product is complex, and so is the model necessary to arrive at its price, the ultimate owners of the funds at risk might lack the knowledge, expertise or inclination to assess the fair value, and will have to rely on their agent's judgement. This trust, however, will not be extended for too long a period of time, and certainly not for many years. Therefore, the time span over which securities are to revert to their fundamental value must be relatively short. If the supply-and-demand dynamics were such that the mispriced instrument might move even more violently out of line with fundamentals, the position of the pseudo-arbitrageur could swing further into the red, and the 'trust-me-I-am-a-pseudo-arbitrageur' argument might rapidly lose its appeal with the investors and shareholders.

Another source of danger for relative-value traders is the existence of institutional and regulatory constraints that might force the liquidation of positions before they can be shown to be 'right': the EMH does not know about the existence of stop-loss limits, VaR limits, size constraints, concentration limits etc.

Similarly, poor liquidity, often compounded with the ability of the market to guess the position of a large relative-value player, also contributes to the difficulties of pseudo-arbitrageurs. In this context, the role played by pseudo-arbitrageurs as ultimate providers of liquidity has been discussed by [Scholes (2000)].

Finally, very high information costs might act as a barrier to entry, or limit the number, of pseudo-arbitrageurs. Reliable models require teams of quants to devise them, scores of programmers to implement them, powerful computers to run them and expensive data sources to validate them. The perceived market inefficiency must therefore be sufficiently large not only to allow risk-adjusted exceptional profits after bid-offer spreads, but also to justify the initial investment.

In short, because of all of the above, even in the presence of a significant imbalance of supply or demand, relative-value traders might be more reluctant to step in and bring prices in line with fundamentals than the EMH assumes.

11.3.3 Empirical Evidence

The literature covering empirical tests of market efficiency is far too large to survey even in a cursory manner (a recent count of papers in behavioural finance aimed at displaying failures of the EMH exceeded the 2,000 entries), and, beyond the statement that markets are 'by and large' efficient, there has been a strong, and unfortunate, polarization of academic opinion. However, questions of more limited scope can be posed, such as: 'Is there any evidence about that the mechanisms discussed above do hamper the activity of pseudo-arbitrageurs?', or 'Does the relative price of caplets and swaptions provide any indication of price deviations from fundamentals compatible with the supply/demand dynamics discussed in Section 11.3.1?'. Both questions are difficult to answer, the first because of the clearly secretive nature of the activities of pseudo-arbitrageurs (hedge funds and proprietary traders); the second, because showing that some

prices are 'incorrect' always requires working under a joint hypothesis: what it tested is the deviation from a pricing model, given that the pricing model itself is correct. [Reference here]. Nonetheless, some pertinent observations *can* be made.

Starting from the question regarding the actual impact of the factors discussed above on the effectiveness of the pseudo-arbitrageurs, some indirect evidence can be obtained from reconstructions of the market events that surrounded the near-collapse of the LTCM hedge fund in 1998. (See [Jorion (2000)] and [Das (2002)] for a description of the events from a risk management perspective). Both [Dunbar (2000)] and [Scholes (2000)], although from different perspectives and drawing different conclusions, describe a situation when the long-dated (5-year) equity implied volatility had reached in the autumn of 1998 levels that would imply for the next several years a realized volatility much higher than what ever observed over similar periods. Yet traders (LTCM *in primis*) who attempted to short volatility found their positions moving further into the red before the volatility finally declined. Many traders/arbitrageurs, including LTCM, faced with margin calls and with their request for additional 'arbitrage' capital from the (technically uninformed) investors turned down, had to cut the positions at a loss before the 'correctness' of their views could be proven. Similarly, swap spreads which are but tenuously linked to the default risk of banks, reached during the same period levels difficult to reconcile with any plausibly-risk-adjusted probability of bank default⁴⁴. Finally, and most relevantly for the topic of this survey, a reconstruction of events from market participants who prefer to retain anonymity suggests that a major international investment house observed during the same period swaptions and caplet volatilities to move away from levels that most plausible models could explain. In particular, for most plausible instantaneous volatility functions that recovered caplet prices, the 'implied correlation' was very different from correlations estimated statistically. The same house is widely thought to have put large, (and 'correct'), swaption-caplet 'arbitrage' trades in place, only to have to unwind them at a loss as the positions temporarily moved even more strongly away from 'fundamentals'.

The possible impact of the systematic supply/demand imbalances on the relative prices of caplets and swaptions is particularly relevant to the topic of this survey. Since a swap rate is a combination of forward rates, both caplets and swaptions can be expressed as a function of the dynamics of either forward rates or swap rates alone (if the volatility and correlation functions are assigned). If the account given above is correct, and if therefore supply and demand cannot be arbitrated away efficiently by proprietary traders, one would expect to see

⁴⁴Given the role of pseudo arbitrageurs as providers of liquidity alluded to above, [Scholes (2000)] argues that, in order to assess the 'fair' level of the swap spread, one should factor in the (time-varying) price for liquidity. If this view is correct, there would be no market inefficiency at play. From the point of view of the calibration of a model, however, the existence of an important, and unhedgeable, risk factor (in this case liquidity) neglected in the pricing would raise similar concerns as to the appropriateness of using the market-implied estimate.

the market-implied forward-rate instantaneous volatilities estimated from swaption prices systematically lower than the same quantities estimated from caplet prices. This is a testable hypothesis. Is it borne out by empirical evidence? [Rebonato (2002)] displays graphs of the instantaneous volatilities of forward rates for several currencies estimated from the caplet and from the swaption markets⁴⁵. The instantaneous volatility functions turned out to have a very similar qualitative shape irrespective of the instruments used for their estimation, but to be systematically lower in all currencies when estimated from swaption data. Similarly, using different market data, [Rebonato (2000)] finds that the implied correlation required to price a set of co-terminal swaptions given the market prices of the caplets is much lower than what historically observed.

It must be stressed that these results must be interpreted with care, because what is tested is the joint assumption that supply and demand skew the prices of a theoretically replicable set of securities *and* that the model used for the pricing (with the chosen parametric forms for the volatility and the correlation) is correct. Even with these caveats, the results appear to provide corroboration for the hypothesis that supply/demand imbalances do affect the relative prices of caplets and swaptions.

11.4 Conclusions

This discussion brings us back to the calibration issue. The prevalent market practice, as evidenced by the references quoted above, seems to favour the 'implied' estimation approach. There appear to be sufficient reasons, however, to doubt the informational efficiency of the plain-vanilla instruments used for model calibration. Therefore the generally accepted practice to fit the free parameters of a model so as to recover the prices of as many plain-vanilla instruments as possible should be strongly questioned. I have argued that a more relevant criterion for choosing these input functions should be their ability to recover plausible future prices of the re-hedging instruments.

12 Summary and Further Perspectives

I have reviewed in this survey the interaction between the theoretical developments and the market practice in the context of interest-rate derivatives pricing. I have tried to show that this interaction is rather complex, and a 'linear' evolutionary account would not only be historically incorrect, but probably suggest inaccurate pointers as to where the theory might evolve in the near future. In particular, I have highlighted that models are not abandoned unless a *computationally feasible* better alternative is found, and that new directions of research become established only once new technological developments (more powerful

⁴⁵A variety of correlation functions were used in the calibration to swaption prices. For econometrically plausible correlations the results displayed weak dependence on the level and details of the correlation functions. See also [De Jong, Driessen and Pelsser (1999)] on this point.

computers, low-discrepancy series, etc) become available. In this 'punctuated evolutionary' account, familiarity with a set of modelling co-ordinates (eg, log-normal rates or prices) plays a strong role in establishing the modelling direction. Statistical information has obviously played a role in this process, but, again, a simple picture of 'falsification by experiment' appears to be fundamentally inadequate to account for the observed evolution of the modelling theory.

Moving to more practical aspects, I have stressed the importance of the calibration process, and questioned the soundness of what appears to be prevalent market practice. I have also argued that recent market developments (towards slimmer margins and more complex products) are putting under strain the pricing approach (based on the assumption of perfect payoff replicability) that was appropriate for the first- (and second-) generation derivatives products. A comparison with pricing practice in a germane area, ie the mortgage-backed-securities market, is in this respect very illuminating.

12.1 Comparison of Pricing and Modelling Practices

In the mortgage-backed-securities (MBS) area pre-payment models are coupled with interest-rate models in order to produce the present value of the expected cashflows arising from a pool of mortgages (see, eg, [?] or [Hayre (2001)] for a description of the market and of the prevalent pricing approaches). As a first stage in arriving at a price for, say, a pass-through the cash-flows (including pre-payments) are discounted at the riskless (LIBOR) rate. The implicit assumption in doing so is that not only the interest-rate, but also the pre-payment model should provide a perfect hedge for the cash-flow uncertainty. From the model point of view, this is perfectly justifiable, because the vast majority of pre-payment models use interest rates as state variables, and therefore allow for theoretically perfect hedging the pre-payment risk by trading in interest-rate sensitive underlying products (swaps, caps, swaptions, indexed-principal swaps, spread locks, etc). However, it has always been recognized in the MBS market that other variables (such as unemployment, GDP growth, etc) strongly affect pre-payments, and that these variables are very imperfectly correlated with the interest-rate state variables.

Because of this, the concept of the option-adjusted spread (OAS) has been introduced. The OAS is defined to be the spread to be added to the LIBOR-derived forward rates in order to obtain the discount factors to present-value the expected cashflows. A non-zero OAS therefore explicitly adjusts the price for all the undiversifiable (unhedgeable) sources of risk, for model uncertainty, for liquidity effects, etc. It is by no means a second-order effect, since, especially in periods of great pre-payment uncertainty (eg, during the unprecedented wave of mortgage refinancing of 2002, when pre-payment models were constantly 're-calibrated' by trading houses, and the coupon on every outstanding issue was higher than the current par coupon) it reached values well over 100 basis points.

Why has the equivalent of an OAS not developed in the interest-rate (LIBOR) derivatives area? Apart from issues of product liquidity and standardization, I believe that an important reason has been the different 'starting points'

for the two markets. Even the first mortgage-backed-securities (pass-throughs) have always been complex, because of the inherent difficulty in hedging the non-interest-rate-related risk factors. The appearance of more complex products (IOs, POs, sequentials, PACs, etc) simply added to an existing substantial modelling complexity, and to the relatively poor ability to hedge. Assuming perfect replication, in other terms, was never a realistic working hypothesis.

First-generation derivatives products, on the other hand, were relatively simple, and, given the well-known robustness of the BS model to reasonable misspecification of the input volatility, payoff replicability (with the corollary of risk-less discounting) was a very reasonable working assumption. As new products have been introduced, each incremental increase in complexity has not been so big as to require a totally new and fresh pricing approach. The cumulative effect of this ever-increasing complexity, however, has produced products whose pricing requires the simultaneous modelling of compound optionality arising from the evolution over thirty years or more of two yield curves, of their volatilities and correlations, of the correlations among the forward rates of the two currencies, of the spot FX rate, and of its correlation with the interest forward rates. Pricing developments in related areas (credit derivatives, and n -th-to-default swaps in particular) bring about even greater modelling challenges (and implicitly make liquidity assumptions for the hedging instruments which are even more difficult to justify). One can therefore argue that these products have become no simpler, and their payoff replication not any easier, than the first pass-throughs. Nonetheless, no equivalent of the OAS has been introduced in pricing of these assets, and the paradigm of risk-neutral valuation still reigns supreme. Model reserves are sometimes applied when recognizing the book value of these products, but this has not affected the 'mid' marking to model. The reasons for this, I believe, can be traced to the power of a robust and elegant conceptual framework (the BS replication insight) and the self-sustaining nature of the 'inertial momentum' that a successful modelling framework generates (see Section 2.3).

If this analysis is correct the implications for interest-rate derivatives pricing are not that the approaches described in the preceding sections are of little use: even in the MBS arena state-of-the-art models are continuously refined and developed for the diversifiable risk factors, and the interest-rate models of choice have closely followed the evolution of the LIBOR market. What is required, I believe, is a re-assessment of the limitations of the pure-replication-based pricing philosophy, and the introduction in the price-making process of explicit recognition of the existence of substantial unhedgeable components. Perhaps this 'LIBOR-OAS' could be arrived at in a coherent and theoretically robust manner by following one of the approaches (see, eg, [Cochrane and Saa-Requejo (2000)]) recently introduced in the literature to account for this very state of affairs. I can appropriately close by quoting [Cochrane (2000)]:

Holding [an] option entails some risk, and the value of that option depends on the 'market price' of that risk - the covariance of the risk with an appropriate discount factor. Nonetheless we would like not

to [...] go back to 'absolute' methods that try to price all assets. We can [...] still form an approximate hedge based on [...] a portfolio of basis assets 'closest to' the focus payoff. [...]. Then the uncertainty about the option value is reduced only to figuring out the price of the residual.

Acknowledgement 1 *It is a pleasure to acknowledge many profitable discussions with Dr Mark Joshi, and useful suggestions from Dr Thomas Gustafsson. I remain solely responsible for all remaining errors.*

References

- [Ahn et al. (2002)] Ahn, D-H, Dittmar R F, Galant A R, (2002), 'Quadratic Term Structure Models: Theory and Evidence', Review of Financial Studies, 15, 243-288
- [Ait-Sahalia (1996)] Ait-Sahalia Y, 1996, 'Testing Continuous-Time Models of the Spot Interest Rate', Review of Financial Studies, 9, 385-426
- [Alexander (2000)] Alexander C, (2000), 'Principals Component Analysis of Implied Volatility Smiles and Skews', ISMA Centre discussion paper in Finance, December
- [Alexander (2001)] Alexander C, (2001), 'Market Models', John Wiley, Chichester
- [Andersen and Andreasen (1997)] Andersen L, Andreasen J, 'Volatility Skews and Extensions of the LIBOR Market Model', working paper, GenRe Financial Products
- [Andersen and Andreasen (2000)] Andersen L, Andreasen J, (2000), 'Volatility Skews and Extensions of the LIBOR Market Model', Applied Mathematical Finance, 7, March, 1-32
- [Andersen and Andreasen (2001)] Andersen L, Andreasen J, (2001), 'Factor Dependence of

- Bermudan Swaptions: Fact or Fiction?', *Journal of Financial Economics*, 62, 3-37
- [Andersen and Andreasen (2002)] Andersen L, Andreasen J, (2002), 'Volatile Volatilities', *Risk*, Vol 15, n 12, (December), 163-168
- [Andersen and Broadie (2001)] Andersen, L., and M.Broadie, (2001), "A Primal-Dual Simulation Algorithm for Pricing Multi-Dimensional American Options," working paper, Columbia University, New York.
- [And and Bekaert (1998)] Ang A, Bekaert G (1998), 'Regime Swirches in Interest Rates', NBER working paper 6508, 91, 61-87
- [Babbs (1997)] Babbs, S, H, (1997), 'Rational Bounds', working paper, First national Bank of Chicago
- [Bansal and Zhou (2002)] Bansal R, Zhou H, (2002), 'Term Structure of Interest Rates with Regime Shifts', *Journal of Finance*, LVII, 5, October, 1997-2043
- [Beckers (1980)] Beckers S (1980), 'The Constant Elasticity of Variance Model and Its Implications for Option Pricing', *Journal of Finance*, 35, (3), June, 661-673
- [Bjork (1998)] Bjork, T, (1998), 'Arbitrage Theory in Continuous Time', Oxford Univeristy Press, Oxford
- [Black (1976)] Black F, (1976), 'The Pricing of Commodity Contracts', *Journal of Financial Economics*, 3, 167-179

- [Black, Derman and Toy (1990)] Black F, Derman E, Toy W, (1990), "A one-factor model of interest rates and its application to Treasury bond options", Financial Analyst Journal, 1990, 33-339
- [Black and Karasinski (1991)] Black F, Karasinski P, (1991) "Bond and option pricing when short rates are lognormal", Financial Analyst Journal, July-August 1991
- [Black and Scholes (1973)] Black, F, Scholes, M, (1973) 'The Pricing of Options on Corporate Liabilities' Journal of Political Economy, 81, 637-59
- [Bouchaud et al. (1998)] Bouchaud J-P, Sagna N, Rama Cont, El-Karoui N, Potters M, (1998) 'Strings Attached', Risk, 11 (7), 56-9
- [Boyle (1977)] Boyle P P, (1977) 'Options: a Monte Carlo Approach', Journal of Financial Economics, 4, 323-338
- [Brace, Gatarek and Musiela (1995)] Brace A, Gatarek D, Musiela M, (1995) 'The Market Model of Interest Rate Dynamics', working paper, School of Mathematics, University of New South Wales, Australia
- [Brace, Gatarek and Musiela (1996)] Brace A, Gatarek D, Musiela M, (1996) 'The Market Model of Interest Rate Dynamics', Mathematical Finance, 7, 127-154
- [Brace and Womersley (2000)] Brace A, Womersley S, (2000) 'Exact Fit to Swaption Volatility Using Semidefinite Programming', working paper presented at the ICBI Global Derivatives Conference, Paris, April 2000
- [Brennan and Schwartz (1978)] Brennan M J, Schwartz E S (1982) 'Finite Difference

- Method and Jump Processes Arising in the Pricing of Contingent Claims', *Journal of Financial and Quantitative Analysis*, 13, September, 461-474
- [Brennan and Schwartz (1982)] Brennan M J, Schwartz E S (1982) 'An Equilibrium Model of Bond Pricing and a Test of Market Efficiency', *Journal of Financial Quantitative Analysis*, 17,, 301-329
- [Brennan and Schwartz (1983)] Brennan M J, Schwartz E S (1982) 'Alternative Methods for Valuing Debt Options', *Finance*, 4, 119-138
- [Brigo and Mercurio (2001)] Brigo D, Mercurio F (2001) 'Interest Rate Models - Theory and Practice', Springer Verlag, Berlin
- [Britten-Jones and Neuberger (1998)] Britten-Jones M, Neuberger A, (1998) 'Option Prices, Implied Price Processes and Stochastic Volatility', London Business School working paper, available at www.london.edu/ifa
- [Broadie and Glasserman (1997)] Broadie M, Glasserman P, (1997) 'A Stochastic Mesh Method for Pricing High-Dimension American Options', Working Paper, Columbia University, New York, to appear in *Computational Finance*
- [Broadie and Glasserman (1997b)] Broadie, M., and P.Glasserman, (1997b), "Pricing American-Style Securities Using Simulation," *Journal of Economic Dynamics and Control*, Vol.21, Nos.8-9, 1323-1352.
- [Broadie, Glasserman and Jain (1997)] Broadie, M., P.Glasserman, and G.Jain, 1997, "Enhanced Monte

- Carlo Estimates for American Option Prices," *Journal of Derivatives*, Vol.5, No.1 (Fall), 25-44.
- [Brown and Dybvig (1986)] Brown S, Dybvig P, (1986) 'The Empirical Implications of the CIR Theory of the Term Structure of Interest Rates', *Journal of Finance*, 41, 617-632
- [Broadie and Cao (2003)] Broadie, M, and M.Cao, 2003, "Improving the Efficiency of the Primal-Dual Simulation Algorithm for Pricing American-Style Options," workingpaper, Columbia University, New York
- [Brown and Schaefer (1991)] Brown S, Scafefer S, (1991) 'Interest Rate Volatility and the Term Structure', working paper, London Business School
- [Brown and Schafer (1994)] Brown S, Scafefer S, (1994) 'The Term Structure of Real Interest Rates and the CIR Model', *Journal of Financial Economics*, 35, 3-42
- [Cambell, Lo and MacKinley (1996)] Campbell, J Y, Lo A W, MacKinlay A C, (1996) 'The Econometrics of Financial Markets', Princeton University Press, New Jersey
- [Carverhill A (1992)] Carverhill A (1992) 'A Binomial Procedure for Term Structure Options: When Is the Short Rate Markovian?', working paper, Hong Kong University of Science and Technology, Clearwater Bay, Hong Kong, January
- [Chan et al. (1992)] Chan KC, Karolyi GA, Longstaff FA, Sanders AB, (1992) 'An Empirical Comparison of Alternative Models of the Short-Term Interest Rate',

- Journal of Finance, 57, May 1992
- [Chen and Scott (1993)] Chen R-R, Scott, L, (1993). 'Maximum-Likelihood Estimation for a Multifactor Equilibrium Model of the Term Structure of Interest Rates', Journal of Fixed Income, 3, 13-41
- [Clewlow and Strickland (1994)] Clewlow L, Strickland C, (1994) 'A Note on the Parameter Estimation in the Two-Factor Longstaff and Schwartz Interest Rate Model', Journal of Fixed Income, March, 95-100
- [Cochrane (2000)] Cochrane, J. H. (2001) 'Asset Pricing', Princeton University Press, Princeton (New Jersey) and Oxford
- [Cochrane and Saa-Requejo (2000)] Cochrane, J. H., Jesus Saa-Requejo, (2000), 'Beyond Arbitrage: Good Deal Asset price Bounds in Incomplete Markets', Journal of Political Economy, 108, 79-119
- [Cont (2001)] Cont Rama (2001) 'Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues', Quantitative Finance, Vol. 1, No 2, (March 2001) 223-236
- [Constantinides (1992)] Constantinides, G M, (1992) 'A Theory of the Nominal Term Structure of Interest Rates', Review of Financial Studies, 5 (4), 531-552
- [Cox, Ingersoll and Ross (1985)] Cox J., Ingersoll J.E., Ross S.A., (1985) 'A Theory of the Term Structure of Interest Rates' Econometrica, 53, 385-407

- [Dai and Singleton (2000)] Dai Q, Singleton K J, (2000), 'Expectation Puzzles, Time-Varying Risk Premia and Dynamic Models of the Term Structure', accepted for publication in *Journal of Financial Economics*
- [Das (2002)] Das S, (2002), 'Liquidity Risk Part 3 - Long-Term Capital Management', *Futures and Options World (FOW)*, February, 55-62
- [De Jong, Driessen and Pelsser (1999)] De Jong, F, Driessen L, Pelsser, A, (1999) 'LIBOR and Swap Market Models for the Pricing of Interest-Rate Derivatives: an Empirical Comparison', working paper
- [Del Bano (2002)] Del Bano S R, (2002), 'Long-Dated FX Products', working paper, Royal Bank of Scotland
- [Derman (1999)] Derman E, (1999), 'Volatility regimes', *Risk*, 12, 4, 55-59
- [Derman and Kamal (1997)] Derman E, Kamal M (1997) 'The Patterns of Change in Implied Index Volatilities' *Quantitative Strategies research Notes*, Goldman Sachs
- [Dodds (1998)] Dodds S. (1998), personal communication.
- [Dodds(2002)] Dodds S, (2002) 'Pricing and Hedging Bermudan Swaptions with the BGM Model', Barcalys Capital working paper, presented at the Global Derivatives Conference in Barcelona, May 2002.
- [Dothan (1990)] Dothan M U, (1990), 'Prices in Financial Markets', Oxford University Press, Oxford

- [Doust (1995)] Doust, P., (1995) 'Relative Pricing Techniques in the Swaps and Options Markets', Journal of Financial Engineering, March 1995, 45-71
- [Duffee (2002)] Duffee G, (2002), 'Term Premia and Interest-Rate Forecasts in Affine Models', Journal of Finance, 57, 405-443
- [Duffie (1996)] Duffie, D., (1996) 'Dynamic Asset Pricing Theory' , Princeton University Press, Princeton, New Jersey, Second Edition
- [Duffie, Ma and Yong (1995)] Duffie D, Ma J, Yong J, 'Black's Consol Rate Conjecture', working paper, Graduate School of Business, Stanford University, Stanford, CA
- [Duffie and Kan (1996)] Duffie, D., Kan, R, (1996) 'A Yield-Factor Model of Interest Rates', Mathematical Finance, 6, 379-406
- [Dunbar (2000)] Dunbar N (2000), Inventing Money, Wiley, Chichester
- [Dybvig (1998)] Dybvig, P, H (1998), 'Bond and Bond Option Pricing Based on the Current Term Structure', working paper, Washington University in St Louis, Missouri
- [Eberlein and Raible (1999)] Eberlein, E, Raible, S, (1999), 'Term Structure Models Driven by General Levy Processes', Mathematical Finance, 9, 31-54
- [1] El Karoui N, Geman H, Rochet J-C, (1995) Journal of Applied Probability
- [Engl (1993)] Engl H E (1993) 'Regularization Methods for the Stable Solution of Inverse Problems', Sur-

- veys on Mathematics for Industry, Springer Verlag, 3, 71-143
- [Fabozzi (2001)] Fabozzi F J, (2001), 'The Handbook of Mortgage-backed Securities', Fabozzi, F, J editor, MacGraw-Hill, New York
- [Flesaker and Hughston (1996)] Flesaker B, Hughston L, (1996), 'Positive Interest', Risk, 9, 46-49
- [Garcia and Perron (1996)] Garcia P, Perron P, (1996) 'An Analysis of the Real Inerest Rate Under Regime Shifts', Review of Economics and Statistics, 78, 111-125
- [Gatarek (2001)] Gatarek D (2001) 'LIBOR Market Model with Stochastic Volatility' BRE Bank and Systems Research Institute Working Paper presented at the Maths Week Risk Conference – London 28th November 2001
- [Glasserman and Kou (2000)] Glasserman P, Kou, S.G (2000) 'The Term Structure of Simple Forward Rates with Jump Risk' – Working paper - Columbia University
- [Glasserman and Merener (2001)] Glasserman P, Merener N, (2001) 'Numerical Solutions of Jump-Diffusion LIBOR Market Models' (2001), Working paper, Columbia University
- [Gray (1996)] Gray, S F, (1996) 'Modeling the Distribution of Interest Rates as a Regime-Switching Process', Journal of Financial Economics, 42, 27-62
- [Gustafsson (1992)] Gustafsson T, (1992) 'No-Arbitrage Pricing and the Term Structure of Interest Rates' Working Paper, Department

- of Economics, University of Uppsala, Sweden, Economic Studies, 2
- [Gustafsson (1997)] Gustafsson T, 'On the Pricing of European Swaptions', Working Paper, Department of Economics, University of Uppsala, Sweden, March
- [Hamilton (1989)] Hamilton J (1989) 'A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle', *Econometrica*, 57, 357-384
- [Harrison and Kreps (1979)] Harrison J M, Kreps D (1979) 'Martingales and Arbitrage in Multi-period Securities Markets', *Journal of Economic Theory*, 20, 381-408
- [Harrison and Pliska (1981)] Harrison J M, Pliska S (1981) 'Martingales and Stochastic Integrals in the Theory of Continuous trading', *Stochastic Processes and their Applications*, 11, 215-260
- [Haug and Kogan (2000)] Haug M. B., Kogan L (2000) 'Pricing American Options: a Duality Approach' working paper, MIT and the Wharton School
- [Hayre (2001)] Hayre L (2001) 'A Guide to Mortgage-Backed and Asset-Backed Securities', Lakhbir Haire editor, John Wiley, Chichester
- [Heath et al (1989)] Heath D, Jarrow R A, Morton A, (1989) "Bond Pricing and the Term Structure of Interest Rates: A New Methodology", working paper, (revised edition) , Cornell University

- [Heath et al (1992)] Heath D, Jarrow R A, Morton A, (1992) "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claim Valuation", *Econometrica*, 60, 77-105
- [Hogan (1993)] Hogan M (1993) 'Problems in Certain Two-Factor Term Structure Models', *Annals of Applied Probability*, 3, 576-591
- [Hughston (2000)] Hughston L (2000) 'The New Interest Rate Models', Edited by Lane Hughston, Risk Books, London
- [Hughston and Brody (2000)] Hughston, L, Brody, DC, (2000) 'Modern Theory of Interest Rates - With Extensions to Foreign Exchange, Inflation and Credit', working paper, King's College
- [Hull (1993)] Hull, J, (1993), 'Options, Futures and Other Derivative Securities', 2nd Edition, Prentice-Hall International Editions
- [Hull and White (1987)] Hull J, White A, (1987) 'The Pricing of Options on Assets with Stochastic Volatilities' *The Journal of Finance*, Vol XLII, No 2
- [Hull and White (1988)] Hull J, White A, (1993) 'The Use of Control Variates in Option Pricing', *Journal of Financial and Quantitative Analysis*, 23, 237-251
- [Hull and White (1990)] Hull J, White A, (1993) 'Valuing Derivatives Securities Using the Explicit Finite Differences Method', *Journal of Financial and Quantitative Analysis*, 25, 87-100

- [Hull and White (1993)] Hull J, White A, (1993) 'Bond Option Pricing Based on a Model for the Evolution of Bond Prices', *Advances in Futures and Option Research*, 6
- [Hull and White (1994a)] Hull J, White A, (1994a) 'Numerical Procedures for Implementing Term Structure Models I: single-factor models', *Journal of Derivatives*, Autumn, 7-16
- [Hull and White (1994b)] Hull J, White A, (1994b) 'Numerical Procedures for Implementing Term Structure Models II: two-factor models', *Journal of Derivatives*, Winter, 37-49
- [Hull and White (2000a)] Hull J, White A, (2000a) 'The Essentials of LMM', *Risk Magazine*, December 2000
- [Hull and White (2000b)] Hull J, White A, (2000b) 'Forward Rate Volatilities, Swap Rate Volatilities and the Implementation of the LIBOR Market Model', *Journal of Fixed Income*, 10(2), September, 46-62,
- [Hunt and Kennedy (2000)] Hunt, P J, Kennedy J E, (2000) 'Financial Derivatives in Theory and Practice' John Wiley, Chichester
- [Hunt, Kennedy and Pelsler (2000)] Hunt, P J, Kennedy J E, Pelsler A, (2000) 'Markov-Functional Interest-Rate Models' in 'The New Interest Rate Models', Hughston L ed., Risk Books, London, also forthcoming in *Finance and Stochastics*
- [Hunter, Jaeckel and Joshi (2001)] Hunter C, Jaeckel P and Joshi M (2001) 'Drift approximations in a LIBOR market model' accepted for publication in *Risk Magazine*, also QUARC (Quantitative Research Centre)

- Working Paper available at www.Rebonato.com
- [Jagannathan et al (2000)] Jagannathan R, Kaplin A, Sun G, (2000), 'An Evaluation of Multi-Factor CIR Models using LIBOR, Swap Rates and Cap and Swaption Prices', working paper, Kellogg Graduate School of Management, Northwestern University
- [Jaeckel (2000)] Jaeckel P (2000) 'A Simple Method to Evaluate Bermudan Swaptions in the LIBOR Market Model Framework', QUARC working paper available at www.rebonato.com
- [Jaeckel (2002)] Jaeckel P, (2002), 'Monte Carlo Methods in Finance', John Wiley, Chichester
- [Jaeckel and Rebonato (2000)] Jaeckel P, Rebonato R, (2001) 'Linking caplet and swaption volatilities in a LIBOR market model setting', accepted for publication in Journal of Computational Finance and RBoS QUARC (Quantitative Research Centre) working paper, available at www.Rebonato.com
- [James and Webber (2000)] James J, Webber N, (2000), 'Interest Rate Modelling', John Wiley, Chichester
- [Jamshidian (1991)] Jamshidian F, (1991), 'Forward Induction and Construction of Yield Curve Diffusion Models', working paper, Financial Strategies Group, Merrill Lynch Capital Markets, New York, USA, and then Journal of Fixed Income, 1, 62-74
- [Jamshidian (1997)] Jamshidian F, (1997) 'LIBOR and Swap Market Models and

- Measures' Finance and Stochastic, Vol. 1, 4, September 1997, 293-330
- [Jamshidian (1999)] Jamshidian F, (1999) 'LIBOR Market Model with Semimartingales', working paper, NetAnalytic Ltd., London
- [Jarrow (1996)] Jarrow R A, (1996), 'Modelling Fixed Income Securities and Interest Rate Options', Mac Graw-Hill, New York
- [Joshi and Rebonato (2001)] Joshi M, Rebonato R, (2001) 'A Stochastic-Volatility, Displaced-Diffusion Extension of the LIBOR Market Model', QUARC (Quantitative Research Centre) working paper available at www.Rebonato.com, and submitted to Quantitative Finance
- [Joshi and Theis (2002)] Joshi M, Theis J, (2002) 'Some Title', Quantitative Finance
- [Jorion (2000)] Jorion, P, (2000) 'How Long Term Lost Its Capital'
- [Kennedy (1997)] Kennedy D P, (1997), 'Characterising Gaussian Models of the Term Structure of Interest Rates', Mathematical Finance, 7, (2), 107-118
- [Kloeden and Platen (1992)] Kloeden, P.E., Platen E., (1992) 'Numerical Solutions of Stochastic Differential Equations', Springer-Verlag, Berlin, Heidelberg, New York
- [Kurbanmuradov et al. (2002)] Kurbanmuradov, O, Sabefeld, K, Schoenmakers J (2002), 'Lognormal Approximations to LIBOR Market Models', Journal of Computational Finance, 6 (1) 356-72

- [Lane and Marris (2002)] Lane, A, Marris, D, (2002), 'Fast calibration of LIBOR Market Models', working paper, Royal Bank of Scotland and Centre for Quantitative Finance, Imperial College
- [Lewis (2000)] Lewis A, (2000), 'Option Valuation Under Stochastic Volatility', Finance Press, Newport Beach, California
- [Li et al (1995)] Li, A, Ritcher, P, Sanskarasubramanian L, (1995), 'Lattice Models for Pricing American Interest Rate Claims', Journal of Finance, 50, 719-737
- [Long (1990)] Long, J L, (1990), 'The Numeraire Portfolio', Journal of Financial Economics, 26, 29-69
- [Longstaff and Schwartz (1992)] Longstaff, F A, Schwartz, E S (1992) 'Interest Rate Volatility and the Term Structure: a Two-Factor General Equilibrium Model', Journal of Finance, 47, 1259-1282
- [Longstaff and Schwarz (1994)] Longstaff, F A, Schwartz, E S (1992) 'Comments on: A Note on the Parameter Estimation in the Two-Factor Longstaff and Schwartz Interest Rate Model', Journal of Fixed Income, March, 101-102
- [Longstaff, Santa Clara and Schwartz (2000a)] Longstaff F.A., Santa-Clara P, Schwartz, ES, (2000) 'Throwing Away a Billion Dollars: The Cost of Sub-Optimal Exercise Strategies in the Swaptions Markets', Working Paper, UCLA, presented at the ICBI Risk Conference – Geneva (2000)

- [Longstaff, Santa Clara and Schwartz (2000b)] Longstaff FA, Santa-Clara P, Schwartz, ES, (2000) 'The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence', Working Paper, UCLA, presented at the ICBI Risk Conference – Geneva (2000)
- [Madan, Carr and Chang (1998)] Madan D.B., Carr P, Chang E.C. (1998) 'The Variance Gamma Process and Option Pricing', University of Maryland working paper
- [Marris (1999)] Marris D, (1999) 'Financial Option Pricing and Skewed Volatility', M. Phil thesis, Statistical Laboratory, University of Cambridge
- [Martellini and Priaulet (2001)] Martellini, L, Priaulet, P (2001) 'Fixed-Income Securities', John Wiley, Chichester
- [Mikosh (1998)] Mikosh, T, (1998) 'Stochastic Calculus with Finance in View', Advanced Series on Statistical Science and Applied Probability, Vol. 6, World Scientific Publishers, Singapore, New Jersey, London, Hong Kong
- [Musiela and Rutkowski (1997)] Musiela, M, Rutkowski M (1997): 'Continuous-Time Term Structure Models: Forward-Measure Approach' Finance and Stochastic, Vol. 1, 4, September 1997, 261-292
- [Musiela and Rutkowski (1997b)] Musiela, M, Rutkowski M (1997b): 'Martingale Methods In Financial Modelling' Springer Verlag, Berlin
- [Naik and Lee (1997)] Naik V, Lee M H, (1997), 'Yield Curve Dynamics with Discrete Shifts in Economic Regimes:

- Theory and Estimation', working paper, Faculty of Commerce, University of British Columbia
- [Nelson and Ramaswamy (1990)] Nelson, D, B, Ramaswamy K, (1990) 'Simple Binomial Approximations in Financial Models', *Review of Financial Studies*, 3, 393-430
- [Neuberger (1990)] Neuberger A (1990) 'Pricing Swap Options Using the Forward Swap Market', IFA working paper, London Business School
- [Nielsen (1999)] Nielsen L, T, (1999), 'Pricing and Hedging of Derivatives Securities', Oxford University Press, Oxford
- [Pelsser (2000)] Pelsser A, (2000), 'Efficient Methods for Valuing Interest Rate Derivatives', Springer Verlag, London, Berlin, Heidelberg
- [Pietersz et al. (2002)] Pietersz R, Pelsser A, van Regenmortel M, (2002), 'Fast Drift Approximated Pricing in the BGM Model', working paper, Erasmus University, Erasmus Institute of Management
- [Rebonato (1997)] Rebonato R, (1997) 'A Class of Arbitrage-Free, Short-Rate, Two-Factor Models', *Applied Mathematical Finance*, 4, 1-14
- [Rebonato (1998)] Rebonato R, (1998) 'Interest-Rate Option Models', Second Edition, John Wiley, Chichester
- [Rebonato (1999)] Rebonato R, (1999) 'On the pricing implications of the joint log-normality assumption for the cap and swaption markets', *Journal of Computational Finance*, Vol. 2, 3, 30-52 and

- QUARC Working Paper available at www.rebonato.com
- [Rebonato (1999a)] Rebonato R (1999a) "On the Simultaneous Calibration of Multi-Factor Log-Normal Interest-Rate Models to Black Volatilities and to the Correlation Matrix", *Journal of Computational Finance*, Vol. 2, 4, 5-27 and QUARC (Quantitative Research Centre) Working Paper available at www.Rebonato.com
- [Rebonato (1999b)] (1999) 'Volatility and Correlation', John Wiley, Chicester
- [Rebonato (2000)] Rebonato R (2000) 'Calibration to Co-Terminal European Swaptions in a BGM Setting' QUARC (Quantitative Research Centre) working paper, available at www.Rebonato.com
- [Rebonato (2002)] Rebonato R (2002) 'Modern Pricing of Interest-Rate Derivatives: the LIBOR Market Model and beyond', Princeton University Press, Princeton, New Jersey
- [Rebonato (2003)] Rebonato, R, (2003) 'Which Process Gives Rise to the Observed Dependence of Swaption Implied Volatilities on the Underlying?', working paper, QUARC (Quantitative Research Centre), available at www.rebonato.com, and submitted to *International Journal of Theoretical and Applied Finance*
- [Rebonato and Cooper (1995)] Rebonato R, Cooper I, (1995) 'Limitations of Simple Two-Factor Interest Rate Models',

- Journal of Financial Engineering, 5, 1-16
- [Rebonato and Joshi (2002)] Rebonato R, Joshi M (2002) 'A Joint Empirical/Theoretical Investigation of the Modes of Deformation of Swaption Matrices: Implications for the Stochastic-Volatility LIBOR Market Model', International Journal of Theoretical and Applied Finance, 5, (7) 667-694 and Working Paper, QUARC (Quantitative Research Centre), available at www.rebonato.com
- [Rebonato and Kazziha (1997)] Rebonato R, Kazziha S, (1997) 'Unconditional Variance, Mean Reversion and Short-Rate Volatility in the Calibration of the BDT Model and of Tree-Based Models in General', Net Exposure, 2, November (no page number, electronic journal)
- [Ritchen et al. (1995)] Ritchen P, Sanskarasubramanyan L, (1995), 'Volatility Structures of the Forward Rates and the Dynamcis of the Term Structure', Mathematical Finance, 5, 55-72
- [Rogers (1997)] Rogers L C G, (1997), 'The Potential Approach to the Term Structure if Interest Rates and Foreign Exchange Rates', Mathematical Finance, 7, 157-176
- [Rubinstein (1983)] Rubinstein M, (1983), 'Displaced Diffusion Option Pricing', Journal of Finance, 38, March 1983, 213-217
- [Rutkowski (1997)] Rutkowski M, (1997), 'Models of Forward LIBOR and Swap

- Rates' Working Paper, University of New South Wales
- [Rutkowski (1997b)] Rutkowski M, (1997b) 'A Note on the Flesaker-Hughston Model of the Term Structure of Interest Rates', Applied mathematical Finance, 4, 151-163
- [Rutkowski (1998)] Rutkowski M, (1998), 'Dynamics of Spot, Forward and Futures LIBOR Rates' International Journal of Theoretical and Applied Finance, Vol 1, No. 3, July 1998, 425-445
- [Schaefer (1977)] Schaefer, S. M., (1977) 'The Problem with Redemption Yields' Financial Analyst Journal, July-August 1977, 59-67
- [Schoenmakers and Coffey (2000)] Schoenmakers J, Coffey B, (2000), 'Stable Implied Calibration of a Multi-Factor Libor Model Via a Semi-Parametric Correlation Structure' Weierstrass-Institut fuer Angewandte Analysis und Stochastik, Working Paper N. 611, Berlin 2000
- [Scholes (2000)] (2000) 'Crisis and Risk', Risk, May, 50-54
- [Schwartz (1977)] Schawartz, E S, (1977), 'The Valuation of Warrants: Implementing a New Approach', Journal of Financial Economics, 4, January, 79-93
- [Shefrin (2000)] Shefrin, H, (2000), 'Beyond Greed and Fear', Harvard Business School Press, Boston, Massachusetts

- [Shiller (2000)] Shiller, R J, (2000) 'Irrational Exuberance', Princeton University Press, Princeton, New Jersey
- [Shleifer (2000)] Shleifer A (2000) 'Inefficient Markets – An Introduction to Behavioural Finance', Clarendon Lectures in Economics, Oxford University Press, Oxford
- [Shleifer and Vishny (1997)] Shleifer A and Vishny R (1997) 'The Limits of Arbitrage', Journal of Finance, 52, 35-55
- [Sidenius (2000)] Sidenius J, 'LIBOR Market Models in Practice', Journal of Computational Finance, 3, Spring, 75-99
- [Vasicek (1977)] Vasicek O, (1977) 'An Equilibrium Characterization of the Term Structure' Journal of Financial Economics, 5, 177-188
- [Zuehlsdorff (2001)] Zuehlsdorff, C, (2001), 'Extended LIBOR Market Models with Affine and Quadratic Volatility', Working Paper, Department of Statistics, Rheinische Friederich-Wilhelms-Universitaet, Bonn, Germany