A stochastic-volatility, displaced-diffusion extension of the LIBOR Market Model

Mark Joshi
QUARC (Quantitative Research Centre of the Royal Bank of Scotland Group)
Riccardo Rebonato
QUARC (Quantitative Research Centre of the Royal Bank of Scotland Group)
University of Oxford – Institute of Applied Mathematics

Abstract

We present an extension of the LIBOR market model which allows for stochastic instantaneous volatilities of the forward rates in a displaced diffusion setting. We show that virtually all the powerful and important approximations that apply in the deterministic setting can be successfully and naturally extended to the stochastic volatility case. In particular we show that i) the caplet market can still be efficiently and accurately fit; ii) that the drift approximations that allow the evolution of the forward rates over time steps as long as several years are still valid; iii) that in the new setting the European swaption matrix implied by a given choice of volatility parameters can be efficiently approximated with a closed-form expression without having to carry out a Monte Carlo simulation for the forward-rate process; and iv) that it is still possible to calibrate the model virtually perfectly via simply matrix manipulations so that the prices of the co-terminal swaptions underlying a given Bermudan swaption will be exactly recovered, while retaining a desirable behaviour for the evolution of the term structure of volatilities. We also show that, even after reducing the number of the possible fitting parameters, the market caplet surface across strikes and maturities can be well recovered. We notice the existence of what appears to be a systematic discrepancy for very low strikes, but we have refrained from attempting to recover this feature.

1 – Introduction and motivation for the present study

In the last few years the so-called 'LIBOR market model' (see, e.g. Musiela and Rutkowski (1997), Jamshidian (1997), Brace, Gatarek and Musiela (1996)), based on the assumption of deterministic volatility and log-normal distribution (in the appropriate measure) for forward or swap rates, has won very wide market acceptance. One of the many virtues of this approach is that it allows exact recovery of the Black prices of the plain-vanilla options associated with at least one set of state variables (i.e. caplets or European swaptions). To the extent that the Black formula (with the same implied volatility across strikes) for a fixed maturity held true, this important feature earned the approach the label of 'the market model'. Unfortunately, in the last few years the market in plain-vanilla interest-rate options has ceased to agree with the market model to an increasing extent. As early as, approximately, 1996 the implied volatility for caplets of a given maturity began to assume a marked monotonically decreasing behaviour as a function of strikes. (The case of JPY, where extremely low levels of rates were prevalent before 1996, is rather unique and has never been satisfactorily dealt with in the context of the classic market model approach). This discrepancy from log-normal behaviour can be accounted for in a relatively straightforward way, and is amenable to a simple and convincing financial explanation. From the perspective of the market model, not only is accounting for this feature relatively straightforward, but the powerful calibration procedures described in the log-normal setting can be ‘rescued’ and adapted in a simple way. The degree of ‘surgery’ needed in order to bring about a satisfactory agreement with the market prices for the plain-vanilla instruments, and to price exotic instruments consistently, in the presence of a monotonically decaying implied volatility curve is therefore rather limited. The most common modelling approaches to incorporate a monotonically decreasing smile (as a function of strike) have been Constant Elasticity of Variance (CEV) and Displaced Diffusions (DD) models. Both approaches have a long and illustrious pedigree, and, despite the fact that no paper seems to have appeared dealing with issues of calibration in the market model context, the modifications for the latter are both sufficiently obvious and well known by practitioners that they would hardly warrant *per se* a detailed explicit treatment.

During the second half of 1998, however, the dramatic and unprecedented changes in cap and swaption implied volatilities that occurred in the aftermath of the Russia crisis highlighted the fact that a non-flat but still deterministic smile surface (as implied by CEV and DD) was intrinsically unsuited to capture the reality of market dynamics. Not surprisingly, the implied volatility ‘smile’[1] began to assume more complex shapes (‘hockey stick’), (see Figs. 1, 2 and 3) which cannot be reproduced by simple ‘tinkering at the edges’ of the market-model approach.

FIGURES 1 2 AND 3 APROXIMATELY HERE

[1] In this paper, we define ‘smile’ any non-flat function expressing the implied volatility (discussed and defined precisely in the following) as a function of strike, for a fixed maturity, irrespective of the shape of the function itself.

In view of this historical development, we propose in this paper that that this more recent change in the shape of the implied volatility smile should be explained in terms of a new source of discrepancy from the log-normal behaviour *added to*, but not substituting, whatever cause was responsible for the monotonically decreasing smile which had been visible since the second half of the 1990s.

In other words, since the original reasons for deviation from the log-normal behaviour have not disappeared after the autumn of 1998, we argue that a simultaneous description of the decreasing behaviour of the implied volatility as a function of strike and of the more recent hockey stick behaviour by means of a single mechanism could be a tempting but misguided effort. We therefore propose in this work to explain the observed implied volatility smiles (at least for the EUR and GBP currencies) by means of a combined displaced-diffusion stochastic-instantaneous-volatility approach.

We want to stress that we find the CEV approach conceptually more appealing. Unfortunately, for arbitrary values of the exponent it does not readily lend itself to simple closed-form solutions. Marris (1999) has recently shown, however, that a close, albeit approximate, correspondence between the two can be established (see also Section 8). In particular, there exists a simple relationship between the exponent of the CEV model and the displacement factor of the displaced diffusion. In this paper we therefore choose DD over CEV simply for computational simplicity, and keep in the back of our minds the financially more appealing CEV setting.

Another casualty of the more complex world brought about by the existence of smiles is the ability to identify the implied volatility with the root-mean-square of the instantaneous volatility: as soon as implied volatility surfaces are non-flat as a function of strike it is no longer true that

$$\sigma_{Black}^2 T = \int_0^T \sigma_{inst}(u,T)^2 \, du$$

and the implied volatility simply becomes the ‘wrong number to put in the wrong (Black) formula to obtain the right price’ (see Rebonato (1999)). The conceptual changes brought about by this state of affairs are not insurmountable, especially if one regards, in the presence of smiles, the identification of the root-mean-square instantaneous volatility with the implied volatility as a useful, albeit imprecise, first-order approximation. What becomes much more delicate, however, is the calibration procedure to caplet prices. This observation is important, because the ease and flexibility of calibration to caplet prices constituted one the strongest points in favour of the LIBOR market models. We shall show in the following that efficient and accurate calibration is still possible, but now one requires a more careful treatment where it was previously possible, for instance, to achieve perfect agreement between model and market caplet prices for any choice of instantaneous volatility by taking a simple ratio (see Section 4).

In concluding this introductory section, we want to emphasize that we consider the main contribution of this paper not so much the formulation of the stochastic volatility version of the LIBOR market model – which is, after all, rather straightforward; instead, our aim is to present a new set of precise and numerically efficient calibration and implementation procedures, which translate to the new setting the results by now well established for the deterministic-volatility case, that have played such an important role in the acceptance of the LIBOR market model as a market standard. More precisely we demonstrate that the following are still possible:

- the efficient recovery of at-the-money caplet prices;
- the accurate semi-analytic evaluation of European swaption prices given the parameters of the instantaneous volatility and correlation functions;
- the accurate recovery of the prices of the co-terminal European swaptions underlying a given Bermudan swaption;
- the rapid pricing of path-dependent exotic options.

Finally, we do not regard the main purpose of this paper to fit the caplet/swaption volatility surface as closely as possible. There would indeed be ways, if one so wanted, to produce a closer fit. We have attempted, however, to retain a transparent financial justification for the smile-inducing mechanisms proposed. We are not aware of approaches capable of producing a tighter fit to the smile surface across markets, without, at the same time, introducing opaque, financially implausible, *ad hoc* or outright undesirable features (such as strongly time-inhomogeneous term structure of volatilities). In other terms, we are ready to pay what turns out to be a small price in terms of closeness of fit (see Section 8.6) in order to retain financial plausibility and explanatory power for the model.

2 –The modelling framework

One of the central building blocks of the standard LIBOR market model is the specification of the instantaneous volatility function. Indeed, in the classic treatment (see, e.g., Musiela and Rutkowski (1997) or Jamshidian (1997), the LIBOR market model is characterized by imposing that the volatility of forward rates should be deterministic: $\sigma_{inst}=\sigma_{inst}(T,t)$, where t is calendar time and T denotes the maturity of a given forward rate. Several functional forms for this function have been suggested in the literature. In the following we shall refer to the particular specification

$$\sigma_{inst}(t,T) = k(T)\, g(T-t) \qquad (1)$$

whereby the total instantaneous volatility is expressed in separable form as the product of a time-homogeneous component, g(T-t), times an idiosyncratic (forward-rate-specific) part, k(T). In the standard deterministic LIBOR market mode setting and in the absence of smiles, by means of this separation one can ensure correct pricing of caplets for any choice for g ( ) simply by imposing that

$$k(T)^2 = \sigma_{Black}^2\, T / \int_0^T g(T-u)^2\, du \qquad (2)$$

Furthermore, as discussed as length in Rebonato (1999) and Rebonato (2001), if one chooses the time-homogeneous function g( ) in such a way that the function k( ), as implied by (2), varies as little as possible as a function of the forward rates, one can rest assured that, for a given g( ), the evolution will be as time-homogeneous as possible. See the discussion in Section 4. We have found that the functional form (Rebonato (1999))

$$g(T\text{-}t) = [a + b\,(T\text{-}t)]\,\exp(-c\,(T\text{-}t)) + d \qquad (3)$$

provides a sufficiently flexible, easy-to-interpret, and analytically tractable choice. We shall therefore use Equation (3) as our reference deterministic instantaneous volatility, and report our results based on the stochastic extensions of this function indicated below. The treatment we propose is however more general, and can be easily adapted to different parametric functional choices. Our proposed stochastic extension of the LIBOR market model is therefore the following:

$$d(f_i(t)+\alpha)/(f_i(t)+\alpha)= \mu_i^{\alpha}\,(\{\mathbf{f}\}, t)\,dt + \sigma_i(t,T_i)\,dz_i \qquad (4)$$

$$\sigma_i^{\alpha}(t,T_i) = [a_t + b_t\,(T_i\text{-}t)]\,\exp(-c_t\,(T_i\text{-}t)) + d_t \qquad (5)$$

$$da_t = RS_a\,(a - RL_a)\,dt + \sigma_a\,dz_a \qquad (6)$$

$$db_t = RS_b\,(b - RL_b)\,dt + \sigma_b\,dz_b \qquad (7)$$

$$d\,[\ln(c_t)] = RS_c\,(\ln(c) - RL_c)\,dt + \sigma_c\,dz_c \qquad (8)$$

$$d\,[\ln(d_t)] = RS_d\,(\ln(d) - RL_d)\,dt + \sigma_d\,dz_d \qquad (9)$$

$$E[dz_i\,dz_a] = 0$$
$$E[dz_i\,dz_b] = 0$$
$$E[dz_i\,dz_c] = 0$$
$$E[dz_i\,dz_d] = 0 \qquad (10)$$

$$E[dz_b\,dz_a] = 0 \qquad E[dz_b\,dz_c] = 0 \qquad E[dz_d\,dz_c] = 0$$
$$E[dz_c\,dz_a] = 0 \qquad E[dz_b\,dz_d] = 0$$
$$E[dz_d\,dz_a] = 0 \qquad (10')$$

with Rs_a, RS_b, RS_c, RS_d, RL_a, RL_b, RL_c and RL_d the reversion speeds and reversion levels of a, b, ln(c) or ln(d), respectively, and σ_a, σ_b, σ_c and σ_d their volatilities. We take the instantaneous correlation of the forward rates to be of the form $e^{-\beta|ti-tj|}$.

Before proceeding it is worth commenting briefly on the number of parameters that the approach requires. If the reversion speeds, reversion levels and the volatilities σ_a, σ_b, σ_c and σ_d were all be to used as fitting parameters, the approach could be criticized on the ground of its non-parsimonious nature. The dangers of working with over-parametrized models are particularly acute when the financial interpretation of the parameters is opaque, and their main justification is the achievement of a tighter fit to a set of market data. As discussed in the previous section, however, obtaining a very close match to the smile surface is not the sole, or even the main, purpose of the approach, and the stochastic features introduced by Equations (5) to (10) are meant to provide sufficient flexibility to describe in a financially realistic way not just today's smile surface, but also the observed changes in the market term structure of volatilities. See Figures 10 to 13 below, and the accompanying discussion. This qualitative but important inter-temporal check can therefore afford a useful independent 'corroboration' of the parametrization obtainable from a single-time cross sectional fit. Furthermore, one can empirically observe that the caplet prices are mainly affected by the stochastic behaviour of the d coefficient, and that they are almost independent of the stochasticity of a and b (needless to say, this is not the case for more complex products). Therefore the quality of the fit to caplet prices, displayed in Section 8.6, is not much worse if only the d parameters were allowed to be stochastic. If the trader felt more comfortable with a small number of parameters, and a more transparent modelling approach, she could therefore work with a, b and c deterministic, thereby gaining in 'control' over the model and in computational efficiency, while losing relatively little in terms of quality of fit. Finally, despite the fact that we present the general framework allowing, for maximum generality, for the ability to optimize over the sixteen parameters, when it comes to fitting market data we have chosen to restrict this flexibility by requiring that the reversion levels should be equal to today's value of a, b ln[c] and ln[d], as appropriate. Since, as we show in detail in Section 8.4, even if only the d parameter is allowed to display a stochastic behaviour, the quality of the resulting fit to the whole caplet smile surface is not substantially impaired, the reader who felt uncomfortable with the possibility of working with an over-parametrized model could therefore follow this simpler version (a, b and c deterministic and RL_d =

ln[d(0)]) with relatively little loss of fitting accuracy. Note that this would involve only two new parameters: the volatility of d and the reversion speed of d. Finally, we shall show later on that, in order to obtain a perfect fit to the caplet prices, as many re-scaling constants will be introduced as forward rates. This would appear to increase dramatically the number of fitting parameters. As shown in the result section, however, these extra multiplicative parameters are constructed to be as close to unity as possible, and only provide a very modest last-step fine-tuning, without altering in any material way the dynamics described by Equations (4) to (10).

The precise expression for the no-arbitrage drifts of the forward rates, symbolically shown as $\mu_i(\{\mathbf{f}\},t)$ in Equation (4), depends on the specific choice of numeraire (in the following this will always chosen to be on of the n+1 pure discount bonds that define the n spanning forward rates). In particular, there always is a particular discount bond, $B(t,T_{j+1})$, that can be chosen as the 'natural' numeraire, N_j, for forward rate j, f_j, such that the forward rate itself is a martingale. This 'natural' numeraire coincides with the discount bond maturing at the payoff time of the j-th forward rate. For any other numeraire the forward rate drift will be non-zero. In order to clarify the terminology and notation, the precise numbering of the indices is shown in Tab. I for the case of three forward rates.

FIGURE 4 APPROXIMATELY HERE

There exist several 'versions' of the standard (i.e., deterministic-volatility) LIBOR market model. We shall broadly follow in our treatment the conceptual path laid out by Musiela and Rutkowski (1997) and Jamshidian (1997). In particular, following Musiela and Rutkowski, we avoid the need to specify the volatility structure for the spot bond processes, and we directly work in forward space. Furthermore, we have chosen to extend and generalize in the two directions mentioned above one particular version of LIBOR-market-model approach, namely the one which is characterized by requiring that the volatilities of the *forward-rate* processes, as opposed to the volatilities of the (forward) *bond* processes, should be deterministic. We part company, by so doing, with

Rutkowski (1998) and the related literature. In addition, as explained at greater length below, we shall restrict our attention to payoffs that satisfy the measurability conditions (see Jamshidian (1997)) necessary to ensure that the instantaneously-compounded money-market account and the short rate do not have to be introduced. Also, given the nature of the financial instruments the LIBOR market model is employed to price, we shall only consider the finite-tenor case. Finally, we shall assume in our treatment that all the necessary regularity and technical conditions are met, the most important of which is probably the square integrability of the volatility function.

The derivation of the no-arbitrage drifts for the forward rates in the standard LIBOR market model presented by Musiela and Rutkowski (1997) and Jamshidian (1997) hinges i) on the definition of bond processes as strictly positive continuous semi-martingales with finite second (co)-variation defined on a probability space ($\mathbf{Q}$, Ω, $\Im$); ii) on the construction of forward price and forward rate processes, also described by similar strictly positive semi-martingales; iii) on the requirement that the volatilities of the forward rates (and of no other quantities) should be deterministic functions of time; iv) on the application of Ito's lemma and Girsanov's theorem coupled with a well-known backward induction procedure to obtain the no-arbitrage drifts of the forward rates in a particular measure. Furthermore, Jamshidian (1997) shows that, as long as the payoff of the derivatives to be priced satisfy the conditions of i) being a function homogeneous of degree one in the bond prices and ii) being measurable with respect to the filtration generated by the forward rates on their reset times, the instantaneously compounded money-market account, and the short rate, need not be introduced. Notice that condition i) is equivalent to requiring that that the payoff function should be expressible as the product of an arbitrary function of the forward rates (which are a function of degree zero in the bond prices), times a discount bond, and therefore affords all the flexibility we need for the common LIBOR products encountered in market practice.

In dealing with the case of stochastic instantaneous volatilities, we place ourselves exactly in this joint framework, and our treatment carries through virtually unaltered with two important exceptions: at point iii) above we require that the volatility of the forward rates should be stochastic, as per Equations (4) to (10); furthermore, we require our underlying probability space to be endowed with a finer filtration, generated not only by the forward rates at their reset times, but also by their stochastic variances on the same reset dates. In practice, since the Ornstein-Uhlenbeck behaviour postulated for the coefficients of the instantaneous volatility (or their logarithms) does not yield a closed-form solution for the variance, we shall have to numerically sample their processes with finer frequency in order to estimate the required variances. It should be kept in mind, however, that this is a technical (numerical), rather than a conceptual (financial) constraint. As for the payoffs that we deal with, we retain the same above-mentioned measurability and homogeneity conditions imposed by Jamshidian (1997).

Since point iv) in Musiela and Rutkowski's derivation (and, in particular, the application of Ito's lemma and Girsanov's theorem) holds true irrespective of whether the volatilities are deterministic or stochastic, the no-arbitrage drifts retain the same formal appearance, with the volatilities now a stochastic quantity. With the notation introduced in Fig. 4, the drifts μ_i, symbolically denoted in Equation (4) by the expression μ_i ({**f**}, t), therefore become

$$\mu_i = \sigma_i(t) \sum_{k=j,i} \sigma_k(t)\, \rho_{ik}(t)\, f_k(t)\tau / [1 + f_k(t)\tau] \qquad \text{if } i > j \qquad (11)$$

$$\mu_i = - \sigma_i(t) \sum_{k=i+1,j} \sigma_k(t)\, \rho_{ik}(t)\, f_k(t)\tau / [1 + f_k(t)\tau] \qquad \text{if } i < j \qquad (12)$$

$$\mu_i = 0 \qquad \text{if } i = j \qquad (13)$$

A simple calculation shows that, if the process is turned into a displaced diffusion, the equations above simply read

$$\mu^\alpha_i = \sigma^\alpha_i(t) \sum_{k=j,i} \sigma^\alpha_k(t)\, \rho_{ik}(t)\, (f_k(t)+\alpha)\tau / [1 + f_k(t)\tau] \qquad \text{if } i > j \qquad (11')$$

$$\mu^\alpha_i = - \sigma^\alpha_i(t) \sum_{k=i+1,j} \sigma^\alpha_k(t)\, \rho_{ik}(t)\, (f_k(t)+\alpha)\tau / [1 + f_k(t)\tau] \qquad \text{if } i < j \qquad (12')$$

$$\mu^\alpha_i = 0 \qquad \text{if } i = j \qquad (13')$$

In Equations (11'), (12') and (13'), the symbols μ^α_i and σ^α_i indicate the percentage drift and percentage volatility of the quantity ($f_i + \alpha$). These equations, together with Equations (5) to (10), therefore fully describe the no-arbitrage evolution of the spanning forward rates in the problem. Note that whilst we have a chosen a particular form for the drifts of the volatility parameters, this constraint is not a no-arbitrage condition as any drift for the volatility will result in a non-arbitrageable price provided the volatility of the parameter is non-zero. Our choice is instead guided by the desire to provide a realistic evolution of the term structure of volatilities.

The financial and mathematical motivations for the choice of processes for the volatility coefficients are as follows. As for the conditions of linear independence (10), we assume that the instantaneous volatility function has a stochastic behaviour that can be separated into a component that is perfectly correlated with the level of the forward rate in a CEV fashion; and further components that are totally uncorrelated with the forward rates (and, for simplicity, among themselves). In other terms, the use of the CEV, or displaced-diffusion dynamics, produces an effective correlation between the percentage volatility and the level of the forward rates, and we therefore make the assumption (10) that there are no other sources of correlation. We implicitly endorse, by this modelling choice, the explanation of the monotonically decreasing part of the smile in terms of traders' beliefs that the absolute (basis point) moves of forward rates are less-than-proportional to the level of the forward rates themselves. We then use, *purely for computational reasons*, a displaced diffusion as a proxy for the CEV dynamics following the equivalence and transformation rules presented by Marris (1999). The accuarcy and validity range of this transformation is illustrated in Fig. 13bis in Section 8.

Having accounted for this component of the smile, we then attempt to capture the remaining, and recently appeared, component of the smile surface by means of a diffusive stochastic instantaneous volatility driven by shocks that have no residual correlation with the level of the forward rate. We do so by imposing an Ornstein-Uhlenbeck mean-reverting behaviour to the volatility parameters, thereby ensuring that they not to stray too much from 'reasonable' values; furthermore, closed-form solutions for their realization after a finite time interval are available. The various coefficients in the functional form (3) can be given transparent financial interpretations (by noticing, for instance, that d is equal to the asymptotic long-residual-maturity level of the volatility; that the quantity 1/c – a/b gives the location of the maximum; that b controls whether the initial slope of the instantaneous volatility will be negative or positive; etc.). Introducing a diffusive behaviour for these quantities can therefore be interpreted as allowing the level of the volatility to vary (via d); the location of the maximum of the hump (when present) to change; the hump to become more or less pronounced (via b and c); etc.

FIGURES 5 TO 9 APPROXIMATELY HERE

TABLE I APPROXIMATELY HERE

Figures 5 to 8 display the qualitative features of the stochastic volatility functions when one coefficient after another is allowed to assume a stochastic mean-reverting behaviour. Figure 9 then shows the combined effect of four stochastic coefficients. The sampled paths displayed in these figures together with the deterministic volatility curve were obtained assuming the initial values of the processes for the coefficients, a(0), b(0), c(0) and d(0), to be known exactly (as they could be, for instance, after a fitting procedure has been carried out). Given the perfect knowledge of the parameters a(0), b(0), c(0) and d(0) 'today', the current instantaneous volatility of every forward rate is perfectly known. In particular, for the seven-year-residual-maturity forward rate whose volatility is examined in Figs. 5 to 9, its instantaneous volatility is known exactly from a(0), b(0), c(0) and d(0) to have the value 14.00%. This value corresponds to the left-most point on the curves. The volatility of this forward rate is then evolved through the seven years of the life of the forward rate up to its expiry (the right-most points in each graph) using the coefficients in Tab. I.

Since, in the calibration procedure proposed in the following, the values for the reversion speeds and levels will be obtained from quoted option prices, they will be risk-adjusted, rather than real-world, quantities. Even with this *caveat*, however, the qualitative behaviour appears reassuringly plausible (see below). In particular, the stochastic behaviour of a and b appear to affect mainly the portion of the instantaneous volatility curve immediately prior to the reset of the forward rate, and, in particular to influence the magnitude of the volatility hump. This feature, in turn, has been shown (see, e.g. Rebonato (2001)) to have a noticeable influence on the terminal de-correlation between forward rates, and can therefore be expected to have a significant influence on the covariance elements that enter the pricing of the complex products. Given the damping effect of the decaying exponential, one can reasonably expect the stochastic behaviour of a and b to have a relatively small impact on the prices of caplets, but a stronger influence on the pricing of those securities that strongly depend on the decorrelation among forward rates brought about by non-flat volatility functions. The volatility of the coefficients c and d, on the other hand, have a more pronounced effect throughout the life of the forward rate, and are expected to produce an effect on caplet prices similar to the 'classic' (Hull and White (1987)) stochastic volatility model. In particular, they are expected to produce shallow smiles for short maturities, and more pronounced ones as the expiries increase. Furthermore, given assumptions (10), the smiles produced by the stochastic behaviour of c and d will display little or no skew (this latter feature, of course, is introduced by the CEV-DD behaviour).

Since it might be difficult to appreciate the degree of plausibility of a non-directly observable quantity such as the instantaneous volatility, we have integrated these functions in order to obtain the corresponding stochastic term structure of (at-the-money) volatilities implied by our model. Figures 10 and 11 display two typical results. These qualitative patterns can then be more transparently compared with observed time series of market term structures of volatilities. See Figures 12 and 13.

FIGURE 10 APPROXIMATELY HERE

Clearly, we are not in a position to make any statistically testable statement, but we find that the similarity between the observed and model changes in the term structure of volatilities (which is indeed the feature the trader using a deterministic volatility LIBOR market model is exposed to) is encouraging. (Note that, in producing Figs. 11 and 12, the same parameters in Tab. I were used, and these were not obtained by fitting the historical term structure of volatilities shown in Fig. 13).

FIGURES 11 TO 13 APPROXIMATELY HERE

In order to fulfil the programme highlighted in Section 1, we now proceed to examine if, and to what extent, the powerful approximation procedures which have been introduced over the last few years, and that make the LIBOR market model a powerful and practical pricing tool, retain their validity in the stochastic volatility setting we have presented above.

3 – Approximations for the drifts and evolution of the forward rates

We propose an efficient general computational procedure in order to sample the joint distribution of the forward rates given the stochastic volatility structure postulated in Section 2. More precisely, we recommend performing a simple short-stepped, first-order Euler-scheme Monte Carlo evolution for the volatility process (given our specification of the processes for the coefficients we can use the Ornstein-Uhlenbeck solutions – this might not be the case for a more general choice of processes for the volatility parameters); once the appropriate covariance matrix for the forward rates for the relevant portion of the path has been obtained, we then suggest a 'long-jump' evolution of the forward rates form one price-sensitive event to the next – typically, from one forward-rate reset to the next. In other terms, once the sampling of the volatility process by a short-stepped Monte Carlo evolution has been carried out, the problem is conditionally (i.e. for a given realized covariance matrix) reduced to the deterministic case. The procedure is rendered particularly simple by the assumed independence between the forward rates and the volatility process.

The evolution of the forward rates conditional on a particular covariance matrix having been realized, however, is not totally straightforward. For a given numeraire, in fact, at most one of the forward rates will be drift-less, and therefore truly log-normally distributed, and the drifts of all the others will contain the forward rates (see Equations (11) and (12)). When the stochastic forward rates enter the drifts, there is no closed-form solution for the evolution of a forward rate over a finite time interval. In order to avoid a short-stepped Monte Carlo sampling of the path, it is customary (Hull and White (2000)) to advocate an approximation whereby the quantities $f_j(t)/(1+f_j(t)\tau)$, $T_i \leq t \leq T_{i+1}$, in the drift terms (11) and (12) are assumed to be constant and approximated by their values before the evolution, $f_j(T_i)/(1+f_j(T_i)\tau)$. This approximation can be acceptable for price-sensitive events separated by time intervals of the order of three or six months. In the presence of non-call features, however, the interval between price-sensitive events can be as long as a few years, and the approximation above soon ceases to be, in Hull and White's terminology, 'innocuous'. More importantly and fundamentally, there exist an important class of securities (multi-look trigger swaps are an example) characterized by the fact that their payoff can be expressed as a function of the various forward rates *on their own reset dates*. When this is the case, by far the most efficient Monte Carlo evaluation is achieved by constructing the terminal covariance matrix among the forward rates; orthogonalizing it; and evolving using the resulting eigenvectors and eigenvalues simultaneously all the forward rates *to the last maturity date*, which could be ten years or more (see Rebonato (2001) for a discussion). In this case the simple approximation mentioned above certainly ceases to provide a adequate accuracy for pricing.

Hunter, Jaeckel and Joshi (2001), however, have recently presented a simple but powerful extension of the 'naïve' approximation by applying to the financial area a Runge-Kutta (predictor/corrector) scheme first introduced by Platen (1992). Simply put, this approximation consists of estimating the drift as the average of the terminal drift, predicted by the naïve approximation, and the initial drift. In the deterministic-volatility case the approach recommended by Hunter, Jaeckel and Joshi has been shown to yield sufficient accuracy for steps as long as ten or twenty years (for typical market values of the volatilities). We shall show in Section 7 by numerical experiments that it retains its validity also for displaced diffusions and in the presence of stochastic volatilities.

As for the evolution of the forward rates, once the (conditionally deterministic) drift has been obtained we proceed as follows: as in the rest of this paper, we recall that the evolution of the volatility parameters is independent of the movements of the forward rates (see Equations (10)). We then divide the time intervals into small steps, Δs_l, $l = 1, \ldots, m$, and we evolve a, b, c and d over these small steps. For each small step we create a forward-rate marginal covariance matrix for that step C^f_l. More precisely, the (j,k)-th entry of C^f_l will be

$$C^f_l(j,k) = ((a_{s(l)} + b_{s(l)} (t_j - s_l))\exp(-c_{s(l)} (t_j - s_l) + d_{s(l)}) \\ ((a_{s(l)} + b_{s(l)} (t_k - s_l))\exp(-c_{s(l)} (t_k - s_l) + d_{s(l)}) \rho_{jk} \Delta s_l$$

After forming a covariance matrix C^f by summing C^f_l for l=1,…m, we proceed to calculating a pseudo-square root, A^f, of C^f by Cholesky decomposition, i.e. we determine the lower-triangular matrix A^f such that

$$C^f=A^f(A^f)^T.$$

We can now evolve the forward rates across the interval $[T_r, T_{r+1}]$ according to

$$\log f_j(T_{r+1} + \alpha) = \log(f_j(T_r + \alpha)) + \mu_j(T_r) (T_{r+1} - T_r) + \Sigma A_{ij} Z_j$$

(with {Z} i.i.d. standard normal variates) applying the same predictor-corrector approximation presented by Hunter Jaeckel and Joshi (2001) on a path-by-path basis. The advantage of the technique is that only the volatility process need be short stepped, and the only loss of accuracy in the evolution of the forward rates is in the approximation to the drifts, which already exists in the non-stochastic volatility case. Furthermore, since the volatility process is only four-dimensional, while the forward rate process can easily be twenty- or forty-dimensional, the procedure affords significant computational savings. Note that we use Cholesky decomposition in the stochastic volatility case instead of orthogonalization because the procedure has to be carried out on a path by path basis, and the time required to carry out the orthogonalization would therefore be prohibitive.

4 – Exact calibration to caplet prices

The next task is the fast and numerically efficient calibration of the model to the caplet prices. We first briefly recall the strategy one of us (RR) has suggested elsewhere (Rebonato (1999)) in order to determine the coefficients of the instantaneous volatility function g( ) in Equation (3) in the deterministic case. For an arbitrary set of coefficients a, b, c and d, we calculate analytically for all the caplets in the term structure of volatilities the integral

$$\int_0^{T(i)} \sigma_{inst}(T_i - u)^2 \, du = \int_0^{T(i)} \{[a + b\,(T-u)] \exp(-c\,(T-u)) + d\}^2 \, du \equiv {\sigma^i_{BlackModel}}^2 \, T_i \qquad (14)$$

One can then vary the coefficients in such a way as to minimize

$$\chi^2 = \sum_{i=1,n} [\sigma^i_{BlackModel} - \sigma^i_{Black}]^2. \qquad (15)$$

By so doing one can rest assured that, given the choice for the function g( ), as much of today's term structure of volatilities will be explained by a time-invariant component. This in turn, will imply that today's term structure of volatilities will also be as time-homogenous (i.e. evolve in a self-similar way) as possible, given today's market prices and the function g( ). Empirical observation of the time evolution of term structure of volatilites, and the arguments presented, e.g., in Longstaff, Santa Clara and Schwartz (2000a) and Longstaff Santa Clara and Schwartz (2000b) suggest that, *for a deterministic volatility model*, this feature is both a desirable and an important feature. In general, however, the minimization (15) will not yield the caplet market prices precisely. The instantaneous volatility function is therefore separated into a time-homogeneous component, and the idiosyncratic (i.e. forward-rate-specific) multiplicative correction factor, $k(T_i)$, is determined by imposing that

$$k(T_i)^2 = {\sigma^i_{Black}}^2 \, T_i / \int_0^{T(i)} g(T_i - u)^2 \, du \qquad (16)$$

The smaller the variation of the function k( ) (i.e., without loss of generality, the closer the various $k(T_i)$ are to unity), the more time homogeneous the evolution of the term structure of volatilities will be.

We modify this calibration approach to deal with stochastic instantaneous volatilities in the following way. First, we subdivide the time interval from today to the expiry of the i-th caplet, $[0\ T_i]$, i = 0,1, …, m, in steps, Δs_r, r = 1, …, m, sufficiently small to allow accurate sampling of the volatility path. We evolve a, b, c and d over these small steps. For each small step we create one element of a marginal variance matrix, V_{ri}:

$$V_{ri} = \{[a(s_r) + b(s_r)\,(T_i-s_r)]\,\exp(-c(s_r)\,(T_i-s_r)) + d(s_r)\}^2\,\Delta s_r \qquad (17)$$

We then form a forward-rate-specific variance vector $\mathbf{V_i}$ by summing the elements V_{ri}, r=1, 2,…, m:

$$\mathbf{V_i} = \sum_{r=1,\,m} V_{ri} \qquad (18)$$

To price the i-th caplet, for this volatility path we compute the root-mean-squared volatility out to its expiry and use it in the displaced-diffusion Black formula, Black(T,V). Given the assumption of independence between the Brownian increments of the volatility coefficients on the one hand and the forward rates on the other, the price, P_i, of the i-th caplet will be given by

$$P_i = \int \text{Black}(T_i,V_i)\,\phi(V_i)\,dV_i. \qquad (19)$$

This is an adaptation of a result of Hull and White, (1987). The density $\phi(\,)$ is not known analytically. Expression (19) can however be sampled very efficiently using low-discrepancy numbers – we found that, in practice, as few as 64 volatility paths were sufficient to ensure convergence. We then mirror the procedure used in the deterministic-volatility case by minimizing over all the parameters using the market prices, P_i^{Market}, for all available maturities *and strikes*

$$\chi^2 = \sum_{i=1,N} [P_i - P_i^{Market}]^2 \qquad (20)$$

As in the deterministic case, the agreement between model and market prices will not be perfect. In order to match these two quantities exactly at the money, we now invoke the well-known result that the Black function is almost exactly linear in volatility for the at-the-money strike (see Wilmott (1998), or Brenner and Subrahmanyam (1994)). (This property also holds in the displaced diffusion setting.) As above, let $k(T_i)$ be the ratio between the market and the model price. Given the linearity property, we then set up an alternative volatility process for the i-th forward rate given by

$$\sigma'_{inst}(t) = k(T_i)\, \sigma_{inst}(t). \tag{21}$$

For any given path, the root-mean-squared volatility for σ' will simply be equal to $k(T_i)$ times the root-mean-squared volatility for σ. Given the approximate, but very accurate, linearity, the price implied by the Monte Carlo simulation will also be multiplied by the same scaling factor, thereby ensuring correct pricing of each at-the-money caplet.

The effectiveness of the procedure is shown in the numerical experiments reported in the result section, but it is important to stress already at this stage that we run one simulation to price all the caplets at once: for a given volatility path we calculate all the strikes, and we use different stopping points along a path for different maturities. An entire caplet surface can therefore be evaluated in 0.5 seconds using a Pentium II, 200-MHz computer.

5 – Fast pricing of European swaptions

Another important feature of the deterministic-volatility LIBOR market model is the ability to approximate accurately and simply the at-the-money price of a European swaption, given the instantaneous volatility of, and the correlation among, the forward rates. See, e.g., Hull and White (2000) and Rebonato and Jaeckel (2001). This is of great importance in order to explore the congruence and consistency of the caplet and swaption markets, or to ensure, if the trader so desires, that a given choice of forward-rate volatilities should reproduce as well as possible a given swaption matrix.

If we denote by $\sigma^{\alpha}_{SR(k)}(t)$ the instantaneous volatility of the k-th swaption when the displacement of the diffusion has value α, the approximations in the references above clearly generalize to

$$\sigma^{\alpha}_{SR(k)}(T)^2\, T \cong$$

$$\sum_{i,j=1,n(k)}(f_j(0)+\alpha)\,(f_i(0)+\alpha)\,\partial SR_k/\partial f_i\,\partial SR_k/\partial f_j \int_0^T \sigma^{\alpha}{}_i(s)\sigma^{\alpha}{}_j(s)\rho_{ij}\,ds/[SR_k+\alpha]^2 \equiv$$
$$\sum_{i,j=1,n(k)}Z_{ik}\,Z_{jk}\int_0^T \sigma^{\alpha}{}_i(s)\sigma^{\alpha}{}_j(s)\rho_{ij}\,ds/[SR_k+\alpha]^2 \qquad (21')$$

where n(k) is the number of forward rates in the k-th swap, and Z_{ik} is given by

$$Z_{ik} = [\partial SR_k/\partial f_i][(f_i+\alpha)/(SR_k+\alpha)] \qquad (22)$$

which will be zero when k < i. Note in passing that the matrix $Z = (Z_{ik})$ is upper triangular and non-zero on the diagonal; it will therefore have non-zero determinant and an upper-triangular inverse. We shall use this fact in the next section.

The approach we take to extend this approximation to the stochastic setting is very similar to the one presented for the calibration to caplets. Let T be the expiry of the swaption. We divide the interval [0, T] into m small steps, Δs_r, r=1,2,…,m, and evolve a, b, c and d over these small steps. As in Section 3 we form the forward-rate covariance matrix of elements

$$C^f{}_l(j,k) = ((a_{s(l)} + b_{s(l)}(t_j - s_l))\exp(-c_{s(l)}(t_j - s_l) + d_{s(l)}) \\ ((a_{s(l)} + b_{s(l)}(t_k - s_l))\exp(-c_{s(l)}(t_k - s_l) + ds_{(l)})\,\rho_{jk}\,\Delta s_l$$

This covariance matrix can now be used to price the swaption using the Black formula with the input volatility given by Equations (21’) and (22). The resulting price is therefore associated to a particular volatility path and the swaption price in the presence of stochastic volatility is simply obtained by Monte Carlo averaging, as shown in the previous section in the case of caplets. Detailed results are presented in Section 8.4, but, given that the approximation in (21’) and (22) is highly accurate in the deterministic case, and if we recall that once a stochastic volatility path has been drawn we are effectively in the deterministic case, we can expect the approximation to work equally well also in the present setting.

6 - Calibration to co-terminal swaptions

In this section, we look at the problem of calibrating a forward-rate-based displaced-diffusion stochastic-volatility LIBOR market model to the prices of a set of co-terminal swaptions. This problem is of particular relevance in the pricing of those complex derivatives instruments which are 'naturally' hedged using plain-vanilla European swaptions. CMS-based products and Bermudan swaptions are the most obvious examples. The deterministic case has been treated in Rebonato (2000). One advantage of the procedure proposed there is that virtually exact recovery of the prices of the desired set of co-terminal swaptions can be achieved even when working with a FRA-based LIBOR market model. The no-arbitrage drifts for the latter are, in turn, considerably easier to implement than the corresponding quantities in the swap-rate-based version of the market model (Jamshidian (1997)). Furthermore, Rebonato (2000) shows how, in the deterministic-volatility case, this recovery of the co-terminal swaption prices can be obtained in an infinity of ways. Making use of this feature, the proposed method indicates how one can simultaneously obtain the best fit to a forward-rate covariance matrix.

The conceptual path followed in the deterministic-volatility setting is the following. Let $0 < t_0 < t_1 < \ldots < t_n$ be the starting times of a set of swaptions, let $SR_j(t)$ denote the swap rate at time t of the swap associated to the times $t_j < t_{j+1} < \ldots < t_n$ and let f_j denote the forward rate from t_j to t_{j+1}, (as seen from time t). Applying Ito's Lemma and ignoring drifts (irrelevant for the calculation of the volatilities), one can write

$$d(SR_k+\alpha)/(SR_k+\alpha) = \Sigma_{i=1,n(k)}[\partial SR_k/\partial f_i]\ (f_i +\alpha)/(SR_k+\alpha)\ \sigma^{\alpha}_{f(i)}\ dw_i,$$

$$d(f_i+\alpha)/(f_i+\alpha) = \sigma^{\alpha}_{f(i)}\ dw_i = \sigma^{\alpha}_{f(i)}\ [\Sigma_m\ b^f_{im}\ dz_m]$$

$$d(SR_k+\alpha)/(SR_k+\alpha) = \Sigma_{i=1,n(k)} Z_{ik}\ \sigma^{\alpha}_{f(i)}\ dw_i = \Sigma_{i=1,n(k)} Z_{ik}\ \sigma^{\alpha}_{f(i)}\ [\Sigma_m\ b^f_{im}\ dz_m]$$

and deduce that

$$\sigma^{\alpha}_{SR(k)}{}^2 = \Sigma_{i,j=1,n(k)} \sigma^{\alpha}_{f(i)}\ \sigma^{\alpha}_{f(j)} \rho_{ij} Z_{ik} Z_{jk} \equiv \Sigma_{i,j=1,n(k)}\ C^f_{ij}\ Z_{ik} Z_{jk}, \qquad (23)$$

In the expressions above $E[dw_i \ dw_j]= \rho^f_{ij}$, {dz} are the increments of orthogonal Brownian motions, n(k) is the number of forward rates in the k-th swap rate, C^f_{ij} is the (i,j)-th element of the forward-rate covariance matrix, and the weights $\{Z_{jk}\}$ are given by Equation (22), and the constraint $[\sum_m b^f_{im}{}^2] = 1$ on the loadings $\{\mathbf{b^f}\}$ is introduced to preserve the correct recovery of the desired forward-rate instantaneous volatility. The covariance element between the swap rates r and s, C^{SR}_{rs}, is then given by

$$E[d(SR_r+\alpha)/(SR_r+\alpha) \ d(SR_s+\alpha)/(SR_s+\alpha)] =$$
$$E\{\sum_{i=1,n(r)}Z_{ir} \ \sigma^{\alpha}_{f(i)} \ [\sum_m b^f_{im} \ dz_m] \ \sum_{j=1,n(s)}Z_{js} \ \sigma^{\alpha}_{f(j)} \ [\sum_q b^f_{jq} \ dz_q] \ \} \qquad (23')$$

In matrix form (23) can be written as

$$C^{SR}= Z \ \sigma^f \beta \ dz \ dz^T \beta^T(\sigma^f)^T Z^T = Z \ \sigma^f \beta \ \beta^T(\sigma^f)^T Z^T= Z \ \sigma^f \rho^f \ (\sigma^f)^T Z^T= Z \ C^f \ Z^T \qquad (24)$$

where C^{SR} (C^f) is the covariance matrix among the swap (forward) rates, β is the matrix containing the elements $\{b^f_{ij}\}$ and ρ^f is the correlation matrix among forward rates. If one then defines the matrix A to be given by $A= \sigma^f \beta$, in matrix form, and neglecting drifts again, it therefore follows that

$$d(f+\alpha)/(f+\alpha) = \quad A \ dW \qquad (25)$$
$$d(SR+\alpha)/(SR+\alpha) = \quad Z \ A \ dW \qquad (26)$$

Expressions (25) and (26) are instantaneously correct. The value of the swap rate vector after a finite time will be given by

$$\int d(SR_u +\alpha)/(SR_u+\alpha) \ du = \quad \int Z_u \ A_u \ dW_u$$

The matrix Z_u is stochastic, and in general no closed-form solution for the equation above exists. In the deterministic case the following approximation can therefore be made with little loss of accuracy (its validity is justified and analyzed in Rebonato and Jaeckel (2000)):

$$d(SR+\alpha)/(SR+\alpha) = \quad Z(0)\, A\, dW$$

where Z(0) is the value of the (constant) matrix Z evaluated *given the knowledge of today's yield curve*. With this approximation one can write

$$\int d(SR_u + \alpha)/(SR_u + \alpha)\, du = Z(0) \int A_u\, dW_u$$

and all the quantities under the integral sign are now deterministic and the integral can therefore be easily evaluated analytically. Almost-exact pricing of the co-terminal European swaptions can therefore be achieved by requiring that

$$C^f = Z^{-1} C^{SR} (Z^T)^{-1} \qquad (27)$$

where C^{SR} is a swap rate covariance matrix such that its diagonal elements integrate correctly to produce the desired root-mean-squared volatility for each swap. Given the existence of an infinity of such matrices C^{SR}, in Rebonato (2001) a procedure based on a trigonometric decomposition is provided in the deterministic-volatility setting to choose the matrix that simultaneously produces the closest match to an exogenous forward rate covariance matrix.

With the procedure for the deterministic-volatility case clearly in mind, we can now move to the stochastic-volatility setting. We shall mirror the treatment presented above up to the very last step, at which point a different route to ensure reasonable (although no longer optimal) simultaneous fit to an exogenous forward rate covariance matrix must be taken.

As in Section 4, we divide the time intervals into small steps, $\Delta_{s(l)}$, l =0, ..., m. We evolve a, b, c and d over these small steps. For each small step we create a marginal covariance matrix for that step C^f_l . The jj entry of C^f_l will be

$$C^f_l(j,j) = ((a_{s(l)} + b_{s(l)} (t_j - s(l)))\exp(-c_{s(l)} (t_j - s_l)) + d_{s(l)})^2 (s_{l+1}-s_l)$$

The off-diagonal element in position jk is

$$C^f_l(j,k) = \quad ((a_{s(l)} + b_{s(l)} (t_j - s_l))\exp(-c_{s(l)} (t_j - s_l)) + d_{s(l)}) \\ ((a_{s(l)} + b_{s(l)} (t_k - s_l))\exp(-c_{s(l)} (t_k - s_l)) + d_{s(l)}) \rho_{jk} \Delta s_l$$

We then form a forward-rate covariance matrix C^f by summing C^f_l for l =0, ... , m-1. From this we can obtain an implied swap-rate covariance matrix *for this particular volatility path* via $C^{SR} = Z\ C^f\ Z^T$. This covariance matrix can now be used to price the swaptions using the Black formula with volatilities equal to the root-mean-squared volatilities implied by the diagonal elements. These are the prices of the co-terminal European swaptions implied by this particular volatility path. By averaging these prices using a Monte Carlo sampling of the covariances implied by the chosen coefficient dynamics as outlined in Section 8.5 we obtain the 'first-pass' model swaption prices.

For a generic set of parameters a, b, c and d the co-terminal swaptions will, in general, not be correctly priced. As shown below, this is easy to remedy, but, as in the deterministic-volatility case, the challenging task is not the recovery of the market prices, but the ability to reproduce at the same time a believable and desirable evolution for such observables as the term structure of volatilities. In this respect, we fully concur with the analysis and conclusions in Longstaff, Santa Clara and Schwartz (2000a) and Longstaff, Santa Clara and Schwartz (2000b) in their study of Bermudan swaptions, who emphasize the desirability for pricing purposes of producing a realistic evolution of the term structure of volatilities. We therefore propose the following procedure. An initial set of forward-rate-volatility parameters are first chosen so as to produce a desirable evolution of the term structure of volatilities. The procedure outlined in Section 4 can be used for this purpose. With these forward-rate-market-calibrated volatility parameters, we can expect to obtain reasonable, but not exact, prices for the co-terminal swaptions we are interested in. We cannot apply the strategy used in the deterministic-volatility case, because we can no longer specify the forward-rate covariance matrix simply to be $Z^{-1}C^{SR}(Z^T)^{-1}$. This is due to the fact that the covariance matrix is specified on a path-by-path basis by the volatility process. However, if we multiply the prices of each swaption on each path by a particular set of constants that are specific to the swaption but independent of the path, then by the linearity of expectation the final prices of each of the swaptions will be multiplied by the same set of constants.

We then make use again of the result that the at-the-money (displaced) Black formula is almost exactly linear in the volatility, or, equivalently, linear in the square root of the variance. This means that if we adjust the variance of the i-th swaption by a constant K_i^2 on every path, then the final at-the-money price will be multiplied by K_i. The linearity of the at-the-money Black formula means that the implied volatility will also be multiplied by K_i. Let λ be the matrix with the numbers K_i on the diagonal and zero elsewhere. If C^f is the drawn forward-rate covariance matrix for a given path, then the corresponding swap-rate covariance matrix for that path is $C^{SR} = Z\ C^f\ Z^T$, whereas, to price the co-terminal swaptions correctly (at the money), we would like it to be

$$C^{SR} = \lambda\ Z\ C^f\ Z^T \lambda$$

(Note that $\lambda=\lambda^T$). If A is a pseudo-square-root of C then, if we replace A by

$$B = Z^{-1}\ \lambda\ Z\ A$$

then we have the effective swap-rate covariance matrix

$$Z\ B\ B^T\ Z^T = \lambda\ Z\ C\ Z^T\ \lambda^T$$

and the diagonal elements have been multiplied by K_i^2, as required. This means that the average price, and therefore the implied volatility of the i-th swap rate, is also multiplied by K_i, and the at-the-money price is therefore (almost) exactly recovered. This procedure has the desirable feature that it can give an idea of the extent by which the forward-rate volatilities have to be altered in order to obtain an exact fit to the relevant swaption prices: the closer the matrix λ is to (a multiple of) the identity matrix, the more congruent the observed market prices of the co-terminal swaptions and of the underlying forward rates will be.

As for the accuracy of the procedure, any errors will come either from the inaccuracy of the equivalent swaption formula (see Section 5) or from failure of linearity of the at-the-money Black price as a function of volatility. Any errors stemming from the first possible source of inaccuracy would be equally prevalent in the deterministic volatility case. Hull and White (2000) and Rebonato and Jaeckel (2000), however, show that these pricing errors are hardly noticeable on the scale of a typical bid-offer spread. As for the second possible source of inaccuracy, the approximate linearity is (at-the-money) extremely accurate (see again Wilmott (1998), or Brenner and Subrahmanyam (1994)). Detailed results of the relevant numerical experiments are presented in Section 8.5.

7 – Rapid pricing of complex derivatives

Once the calibration of the model has been carried out using the techniques presented in the paper so far, the scene is set for the actual pricing of a complex derivatives product. As mentioned in Section 2, we always assume that the payoff enjoys the measurability and homogeneity conditions highlighted in Jamshidian (1997). As for the evolution of the forward rates, Monte Carlo is again the obvious numerical technique of choice, at least for path-dependent derivatives. We would like to stress, however, that, given the results recently presented by Andersen (1997) and Broadie (1997) amongst others, more complex free-boundary problems can be reduced to the path-dependent case after estimation of the exercise boundary. Needless to say, the Monte Carlo evolution of the forward rates should be carried out using the methodology presented in the section devoted to the drift approximation, which can obviously be naturally applied to the case of path-dependent securities.

The main obstacle would appear to be the sheer computational burden that would be encountered if one carried out the evolution of the forward rates using a short-stepped Monte Carlo simulation. This route would *prima facie* seem necessary in order to accommodate the stochastic nature of the volatility process. As discussed in Section 3, however, in the deterministic-volatility case the ability to carry out moves of the yield curve over periods much longer than the typical three- or six-month tenors becomes of crucial importance not only in the presence of non-call features, but for all those instruments where the evolution of all the forward rates can be carried out in a single very long jump all the way to the final expiry (e.g. ratchet caps or multi-look triggers). Fortunately, by virtue of Equations (10) and (10'), the four-dimensional stochastic integration of the volatility process can be de-coupled from the evolution of the forward rates. As a consequence, conditionally on a given volatility path having been achieved, the evolution of the forward rates over very long steps (perhaps of ten years or more) remains possible also in the stochastic-volatility setting. Therefore, by combining the fast calibration procedures presented in Sections 4, 5 and 6 with the efficient evolution of the yield curve using the predictor-corrector approximation discussed in Section 3, the pricing of complex derivatives in a trading environment becomes a demanding but feasible computational task.

8 – Results

In this section we substantiate with numerical experiments the claims about the accuracy and effectiveness of the approximations presented in the paper so far. Since we mention at several points in the paper the approximate but very accurate equivalence between the caplet prices obtained from the CEV and the displaced-diffusion approaches once the at-the-money prices are matched, we preface our results with a figure from Marris (1999) that shows the quality of this approximation.

FIGURE 13BIS APPROXIMATELY HERE

8.1 Convergence of the sampling of the volatility distribution

First, we want to substantiate with numerical experiments the claims about the accuracy and effectiveness of the approximations presented above. To begin with, we stated in Sections 5 and 6 that, by using low-discrepancy sequences, sufficient convergence to the desired price can be obtained with a relatively small number of volatility paths (very often, as few as 64). Furthermore, as mentioned above, the procedure we recommend is based on long-jumps for the forward rates once the covariance matrix along one particular path has been computed, and a short-stepped Monte Carlo for sampling the distribution of the volatility. We must therefore also explore the issue of price convergence as a function of the step size. We address the two issues in turn.

We begin with a fixed step size (chosen to be equal to 0.08 years – see the discussion below) and we explore the convergence of the model price of European swaptions evaluated using the approximation in Section 5 as a function of the number of sampled volatility paths.

TABLE III APPROXIMATELY HERE

More precisely, fifteen European swaptions of varying expiries and lengths were valued using the approximation in Section 5 in a stochastic-a,-b,-c,-d model. The parameters for the stochastic processes of the coefficients are shown in Tab. III. The swaption prices were evaluated for a fixed step size and for a number of paths increasing as powers of two. Graphs are shown in Figures 14 and 15 for the 1 x 1 and 10 x 10 swaptions. One can notice that, after as few as 64 paths, the 1 x 1 swaption is converged to within a small fraction of a basis point, and the 10 x 10 swaption is converged to within approximately 0.4 basis points. In both cases these price inaccuracies correspond to approximately a twentieth of a vega. Over a variety of tests (see Tab. V), 64 paths always produced pricing inaccuracies substantially smaller than a tenth of a vega (taken as a proxy of a typical bid/offer spread). Table IV reports, for the purpose of providing a comparison yardstick, the approximate vega for the swaptions analyzed in the tests.

TABLE IV APPROXIMATELY HERE

As one can appreciate from these tables and figures the convergence as a function of number of paths is indeed very rapid, and the claim made in the previous sections that as few as 64 (well chosen) paths can produce a convergence well within market bid-offer spreads even for a stress case such as a 10 x 10-year swaption has been substantiated.

The second convergence issue to explore is the behaviour of the estimated price as a function of step size. This aspect of the convergence is addressed in Fig. 16 for the 'stress' case of 10 x 10 swaption, for step sizes ranging from 1.25 years (8 steps over ten years) to 0.08 years (128 steps over ten years). One can see from this figure that, despite the fact that 8 steps over ten years are clearly not sufficient to produce a fully converged solution, even in this case the price error corresponds to only 8 basis points, i.e. a quarter of the typical bid-offer spread for the 10 x 10-year swaption under study. On the basis of these results one can therefore conclude that 64 paths with a step size of 1.25 or 0.625 years should produce results well within the accuracy required for practical applications.

TABLE V APPROXIMATELY HERE

FIGURES 14 TO 16 APPROXIMATELY HERE

8.2 Accuracy of the drift approximation

The next issue we want to explore is the accuracy of the drift approximation in the stochastic volatility case. In order to carry out this test, we chose a ten-year caplet, whose natural numeraire, using the terminology and notation introduced in Fig. 4, would be called N(10). We recall that, under the measure associated with the natural numeraire, the underlying forward rate would be drift-less. For any other numeraire the drift would be non-zero, and is given by expressions (11') or (12'), as appropriate. The greater the mismatch between the expiry of the natural numeraire and any other numeraire (chosen among the universe of traded assets, i.e. pure discount bonds), the larger the drift term, and the more demanding the task for the approximation. Similarly, the longer the expiry, T, of the forward rate, the larger the exponent (because of its linear dependence on T), and the more severe the 'stress case' for the drift approximation. As a test we therefore chose to select an approximately at-the-money (5%), 10-year-expiry caplet, to evolve with a single 10-year jump the 10-year-expiry forward rate under a variety of numeraires ranging from N(1) (largest drift) to N(10) (no drift), to evaluate by Monte Carlo averaging as indicated in Section 3 the value of both a caplet and a floorlet, and to obtain the value of the forward-rate agreement via call-put parity. This model value is then compared with the exact theoretical value. The results obtained for this case study and the parameters used in the simulations are displayed in the tables below.

TABLES VI AND VII APPROXIMATELY HERE

Since we are focussing the attention on the quality of the drift approximation, for the simulations 32 steps and 16,384 (2^{14}) paths were used to ensure that numerical convergence to within one part in 10^6 in the Monte Carlo sampling had been achieved. It is clear from the results in Tab. VII that, even for this very severe test, the drift approximation produces results of excellent numerical quality. It can therefore be reliably used in order to carry out very-long-stepped evolutions of the forward rates either across non-call features, or, when appropriate, to the final expiry of a complex product.

8.3 Accuracy of the fitting to caplet prices

The next investigation has focussed on the accuracy of the procedure suggested in Section 4 order to produce fast and accurate fitting to the at-the money caplet prices by a simple extension of the fitting procedure suggested by Rebonato (1999) in the deterministic case, and briefly explained in Section 4. A global (all-strikes, all-maturities) fit to the GBP caplet market (February 2001) was first of all carried out, and this gave rise to the coefficients, initial values and displacement coefficient displayed in Tab. VIII. The issue of the quality of the market fit is explored separately (see Section 8.6), but these values could not reproduce the market prices exactly (as, in general, will always be the case).

TABLE VIII APPROXIMATELY HERE

The global $k(T_i)$ were chosen to given the best overall fit across all strikes. The re-scaling factors $k(T_i)$ were determined as per Equation (21) in order to fit the at-the-money prices precisely, and the results are displayed in Tab IX. This table shows that the agreement between the model and the target prices (whose accuracy basically hinges on the linearity of the Black formula in the at-the-money volatility) is virtually perfect. We would like to stress that, in carrying out this numerical test, our goal was to show the effectiveness of the procedure even when the required re-scaling is substantial (see column 'Rescaled Ks' in Tab. IX). Therefore, we did not seek, in this particular test, to ensure that the factors $k(T_i)$ should be as close as possible, thereby ensuring as time-homogenous an evolution of the term structure of volatilities as possible.

TABLE IX APPROXIMATELY HERE

8.4 Accuracy of the European swaption approximation

We now examine the accuracy of the extension of the European swaption approximation formula to the displaced diffusion/stochastic volatility setting, as discussed in Section 5. We used again for the purpose the same coefficients shown in Tab VIII. We first valued a series of swaptions using a full stochastic-volatility Monte-Carlo-simulation approach (results labelled 'Full Monte Carlo') in which the forward rates were evolved to the swaption maturity using as many factors as the number of forward rates themselves, and the volatility paths were sampled using for accuracy as many as 512 paths with sufficiently short steps to ensure convergence in price to within a basis point. The only approximation used in this calculation was the use of the extension Hunter-Jaeckel-Joshi (2001) drift correction to the stochastic setting, discussed in Section 3. Given the empirical results presented in Section 8.2, we believe the errors introduced by using this approximation were totally immaterial. We then used the approximate formulae given in Section 5 to obtain the 'quick' prices for the same swaptions, and compared the results. We display below the results for a handful of typical cases, including, as usual, the 'stress' case of a 10 x 10-year semi-annual swaption:

TABLE XI APPROXIMATELY HERE

The vegas for the swaptions (i.e. a reasonable proxy for the bid-offer spread) can be read in Tab. IV. For the most severe 'stress case', i.e. the 10 x 10-year swaption, the vega is approximately 34 basis points, while the approximation error is just two basis points. Also this extension of the procedure known to work well in the deterministic-volatility case has therefore been shown to be very effective in the new setting.

8.5 Accuracy of the simultaneous calibration to a set of co-terminal swaptions

In order to investigate the accuracy of the procedure suggested in Section 6 to produce a virtually exact fit to a set of co-terminal European swaptions (as would typically be found in a Bermudan swaption problem) we followed a similar procedure to the one presented above. We calibrated the model to a set of co-terminal swaptions as described above. A Monte Carlo simulation was then carried out, and the resultant swaption prices were then used to compute implied volatilities. The results are shown in Tab. XII and show again the excellent quality of the approximation.

TAB XII APPROXIMATELY HERE

For the sake of brevity we do not display information about tests run for other currencies and/or trading dates. The results we obtained were however always of the same numerical quality as the ones displayed for this case study. Once again we can therefore conclude that the quality of the numerical approximation is excellent.

8.6 Global fitting to the caplet market smile surface

Despite the fact that, as discussed in the introductory section, our main purpose in this work has not been to produce as close a fit to the caplet smile surface *per se*, it is obviously important to check the quality of the match between model and market prices. We present the results obtained for the global fit to the market smile caplet surface (GBP

1-Sept-2000), and we emphasize i) that we used the same parameters for the whole surface (i.e. all strikes and all maturities) and ii) that the reversion levels were not allowed to be fitting parameters, but were constrained to be equal to today's values of the appropriate quantity (thereby enhancing time-homogeneity). Similar results were obtained in the case of the EUR caplet surface but, again in the interest of brevity, they have not been presented. We ran two distinct optimizations, one allowing for a stochastic behaviour for a, b, c and d (curves labelled 'all stochastic'), and the other by keeping a, b and c deterministic (curves labelled 'Fit'). More precisely, in the 'all-stochastic' fit we optimized the single displaced-diffusion coefficient α, the initial values a(0), b(0), c(0) and d(0), and the reversion speed and volatility of a, b, ln[c] and/or ln[c] to obtain a best global fit to the whole surface. We then imposed perfect fit of the at-the-money prices by using the re-scaling factors, $k(T_i)$, as described in Section 4. Cross-sections of the resulting fitted smile surface are shown in Figures 17 to 21.

TAB. XIII APPROXIMATELY HERE

The following features in these graphs above are worthwhile mentioning. First, with the exception of the shortest-maturity, the agreement between the market and the model smile surface remains very good in a wide region around the at-the-money strike. Allowing for a stochastic behaviour for all the coefficients obviously provides a better overall fit; the only-d-stochastic fit, however, is of comparable quality, and only appears to fail noticeably for extremely out-of-the-money strikes, where the reliability of market quotes is very questionable in any case. Furthermore, (see Table XIII) the optimal solutions obtained for the parameters of the d process and for the initial values a(0), b(0), c(0) and d(0) turned out to be very similar irrespective of whether an optimization over all the reversion speeds, volatilities (and α) was carried out, or simply over the corresponding quantities for d (and α). The similarity of the solutions (as found in this example and in other cases not reported for the sake of brevity) lends some reassurance as to the robustness of the procedure. It is also reassuring that the parameters obtained in the course of the optimization to the caplet surface turned out to be econometrically plausible, without our imposing any constraints to this effect. More precisely, the long-term reversion level, d, was naturally found by the optimization procedure to have a plausible magnitude of 11.4%; the single displacement coefficient α, which turned out to be approximately equal to 2%, would correspond, after carrying out Marris (1999) transformation to a CEV model, to an exponent β approximately equal to 0.60 for an initial forward of 6.00%. This compares well with 'market lore'[2], that favours a value around 0.5; the initial values a(0), b(0), c(0) and d(0) turned out to be very similar to the corresponding values obtained in the deterministic case, and gave rise to a maximum in the hump for the instantaneous volatility around the one-year region, which agrees well with the maturity of the most volatile futures contracts, as discussed in Rebonato (1999); finally, the values for the mean-reversion coefficient turned out to be sufficiently high to ensure that, as desired, the dispersion around the initial values should not be too high, but not so large as to 'kill' the stochastic behaviour of any of the coefficients.

It is perhaps even more important and encouraging to draw attention to the fact that the re-scaling coefficients, $k(T_i)$, necessary to ensure perfect fit to the at-the-money prices, turned out to be extremely close to unity, at least after a brief initial period of less than one year. For the (typical) example reported below, all but four out of sixty caplets perfect fit could be obtained with re-scaling coefficients between 1.05 and 0.95. This is important, because a set of perfectly constant coefficients $k(T_i)$ would imply an exactly time-homogeneous behaviour for the stochastic evolution of the term structure of volatilities (see, e.g. the discussion in Rebonato (1998) and Rebonato (2001)). We have argued above that, in agreement with Longstaff, Santa Clara and Schwartz (2000a) and Longstaff, Santa Clara and Schwartz (2000b), this is a highly desirable feature. We cannot stress strongly enough the importance of this result. By virtue of the fact that the optimized coefficients allow for an approximate forward-rate-independence of the scaling factors, in fact, we have achieved not only a perfect fit to *today's* at-the-money prices and a very good fit to *today's* smile surface, but we have also simultaneously ensured a time-homogeneous *future* evolution of the latter. We consider this ability to account simultaneously and convincingly not only for today's caplet surface but also for its evolution possibly the most important feature of the approach we propose.

---

[2] There is surprisingly little 'hard' statistical evidence in this area: Marsh and Rosenfeld (1983), for instance, find no evidence to reject the log-normal model, while Chan, Karolyi, Longstaff and Sanders

We stress that it would indeed be possible to obtain a closer fit by allowing, for example, the displacement-diffusion coefficient α to become forward-rate specific; or by allowing for the reversion levels of the coefficients a, b, c and d to be different than today's values. By so doing, however, we feel that these modifications of the proposed approach would be very *ad hoc*, and that little would be added in terms of financial modelling. Furthermore, we believe that the trader will be better served by using an imperfectly-fitting, but transparent and intuitively understandable model, than by employing an over-parametrized approach.

FIGURES 17 TO 22 APROXIMATELY HERE

9 – Conclusions and suggestion for future work

We have presented a displaced-diffusion extension of the LIBOR market model which allows for stochastic instantaneous volatilities of the forward rates. We have shown that virtually all the powerful approximations which have been devised in the last few years in the deterministic-volatility setting can be successfully and naturally extended to the stochastic volatility case. In particular we have shown that the caplet market can be efficiently and accurately fit; that the drift approximations that allow the evolution of the forward rates over very large steps can be recovered; that the European swaption matrix implied by a given choice of volatility parameters can be efficiently approximated with a closed-form expression without having to carry out a Monte Carlo simulation of the forward rates; that, just as in the deterministic-volatility case, it is still possible to calibrate the model virtually perfectly via simply matrix manipulations so that the prices of the co-terminal swaptions underlying a given Bermudan swaption can be exactly recovered, while retaining a desirable behaviour for the evolution of the term structure of volatilities.

We have also shown that the proposed approach is capable not only of pricing perfectly today's at-the-money caplets, and of obtaining a very good fit to today's smile surface, but also of producing a financially desirable time-homogeneous evolution for the

(1992) find evidence for β *greater* than 1.

term structure of volatilities. In other terms, a cross-sectional, single-time estimation has naturally and automatically produced in a very satisfactory way a dynamic, inter-temporal feature, without any constraint that this should be the case.

There are several natural extensions to this work: first, one can explore the pricing impact of the approach we have proposed for different types of complex derivatives products. We have already observed, for instance, that, for the simplest instrument of all, i.e. for at-the-money caplets, allowing a stochastic behaviour for the d coefficient has a strong pricing impact, but the effect of the stochasticity of a and c is very limited. In general, we would expect that those complex instruments whose price strongly depends on the terminal de-correlation brought about by non-constant volatility, and, in particular, by the presence of a hump in the instantaneous volatility curve, should be noticeably affected by the stochastic behaviour of the volatility. Trigger swaps, flexi caps, and Bermudan swaptions would clearly warrant close scrutiny.

The second obvious avenue would be the analysis of the hedging performance of the model using real-world data for the analysis. Also very interesting would be the exploration of the ability to obtain, at the same time, a good cross-sectional fit to a single-time smile surface and a convincing recovery of the most salient features of the inter-temporal evolution of the at-the-money implied term structure of volatilities. Work is currently afoot in these three directions.

Finally, comparison, and, possibly, integration with work which has recently been done (Glasserman and Kou, (2000), Jamshidian (2000)) in extending the LIBOR market model to a jump-diffusion setting would be of great interest.

References

Andersen, L, (1999) ‘A simple approach to the pricing of Bermudan swaptions in the multi-factor LIBOR market model’ Journal of Computational Finance, (3) 2, Winter 1999, 5-32.

Andersen L and Andreasen J (1997) ‘Volatility Skews and extension of the LIBOR Market Model’, Gen Re Financial Products working paper, to appear in Applied Mathematical Finance

Beckers S, (1980) ‘The Constant Elasticity of Variance Model and its Implications for Option Pricing’, Journal of Finance, 35, June 1980.

Brace A, Gatarek D, Musiela M, (1996) ‘The Market Model of Interest Rate Dynamics’, Mathematical Finance, 7, 127-154

Brenner M, Subrahmanyam MG, (1994) A simple approach to option valuation in the Black-Scholes model. Financial Analysts Journal, 25-28

Broadie M, Glasserman P, (1997) ‘A stochastic mesh method for pricing high-dimension American options’, Working Paper, Columbia University, New York

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.

Glasserman P, Kou, S.G (2000) ‘The term structure of simple forward rates with jump risk’ **–** Working paper - Columbia University

Hull J, White A, (1987) ‘The Pricing of Options on Assets with Stochastic Volatilities’ The Journal of Finance Vol XLII No 2

Hull J, White A, (2000) 'The Essentials of LMM', Risk Magazine, December 2000

Hunter C, Jaeckel P and Joshi M (2001) 'Drift approximations in a LIBOR market model' accepted for publication in *Risk* magazine and available as RboS QUARC (Quantitative Research Centre) working paper at www.rebonato.com

Jamshidian F (1997) 'LIBOR and Swap Market Models and Measures' Finance and Stochastic, Vol. 1, 4, September 1997, 293-330

Jamshidian F (2000) 'Jump LMM' , – Working paper – presented at the ICBI conference, Paris 2000.

Longstaff FA, 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 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)

Marris D., (1999) 'Financial Option Pricing and Skewed Volatility', M. Phil thesis, Statistical Laboratory, University of Cambridge, 1999

Marsh, T.A., Rosenfeld, E.R. (1983) 'Stochastic Processes for Interest Rates and Equilibrium Bond Prices', Journal of Finance, 38, 635-646

Musiela, M, Rutkowski M (1997): 'Continuous-Time Term Structure Models: Forward-Measure Approach' Finance and Stochastic, Vol. 1, 4, September 1997, 261-292

Rebonato (1999) 'Volatility and Correlation', John Wiley, Chichester, New York, Weinheim, Brisbane, Singapore, Toronto

Rebonato R (2000) 'Calibration to co-terminal European swaptions in a BGM setting' submitted to Journal of Computational Finance and RBoS QUARC (Quantitative Research Centre) working paper, available at www.rebonato.com

Rebonato and Jaeckel (2001) 'Linking caplet and swaption volatilities', accepted for publication in Journal of Computational Finance and RBoS QUARC (Quantitative Research Centre) working paper, available at www.rebonato.com

Rebonato (2001) 'Modern Pricing of Interest Rate Derivatives: the LIBOR market models and beyond', Princeton University Press

Rebonato R, Jaeckel P, (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

Rubinstein M, (1983), 'Displaced Diffusion Option Pricing', Journal of Finance, 38, March 1983, 213-217

Rutkowski, M, (1997), 'Models of Forward LIBOR and Swap Rates' Working Paper, University of New South Wales.

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

Wilmott P. (1998), 'Derivatives: the theory and practice of financial engineering' John Wiley, Chichester, New York, Weinheim, Brisbane, Singapore, Toronto

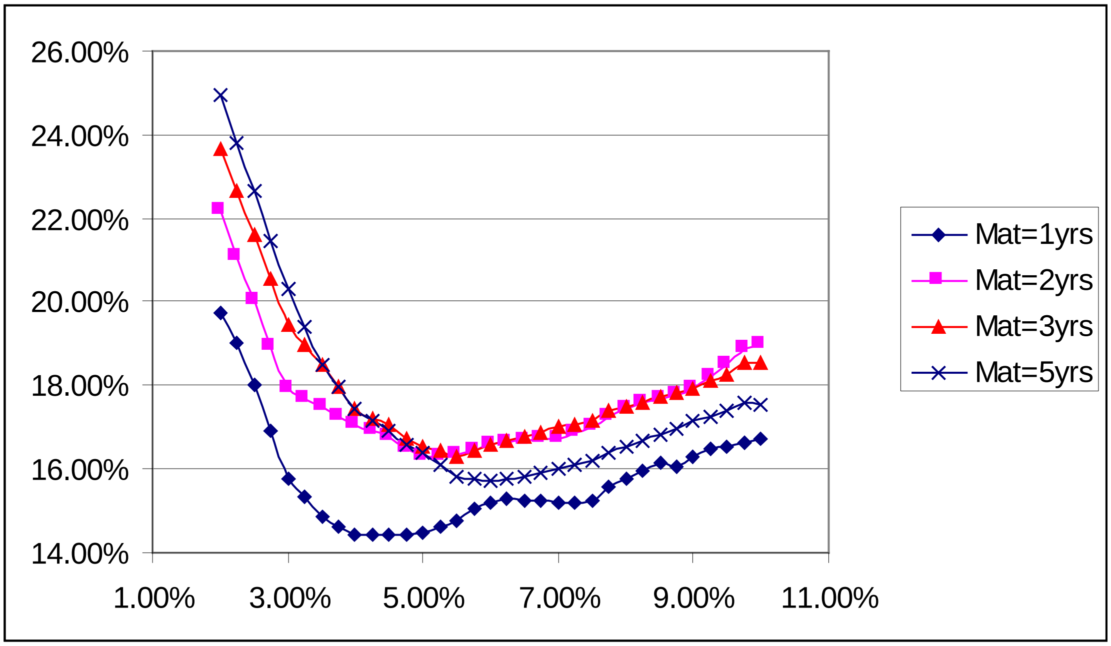

Fig. 1: The implied volatility as a function of strike for selected caplet expiries (USD curve 28-Jul-2000), displaying a very pronounced hockey stick shape at short maturities, and a less strong, but still visible, minimum for maturities out to five years.

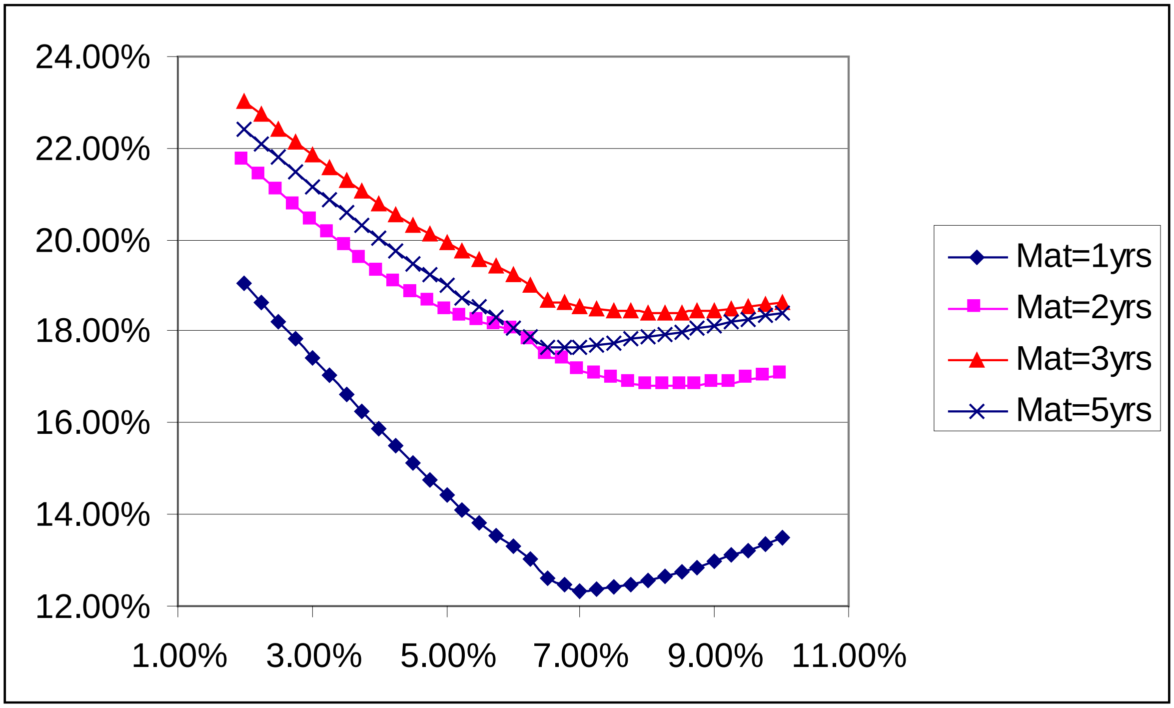

Fig. 2: The implied volatility as a function of strike for selected caplet expiries (EUR curve 4-Aug-2000), displaying a noticeable pronounced hockey stick shape at all maturities.

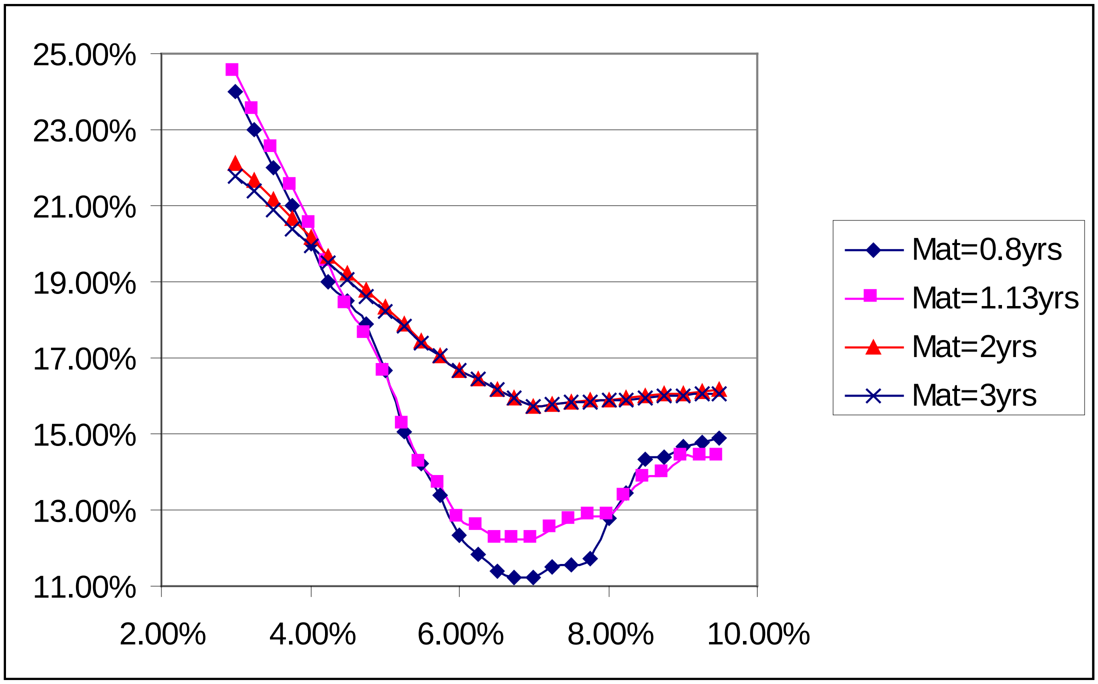

Fig. 3: The implied volatility as a function of strike for selected caplet expiries (GBP curve 1-Aug-2000), displaying a noticeable pronounced hockey stick shape at all maturities.

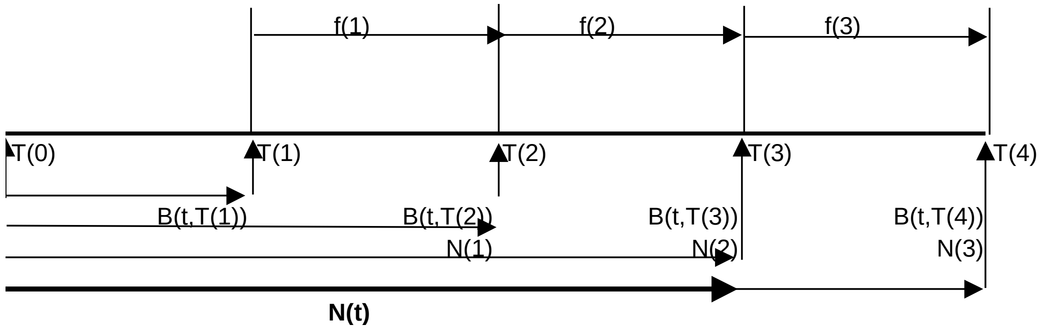

Fig. 4: The securities and timings of cashflows associated with a set of three spanning forward rates. The instruments traded in the economy are the four discount bonds B(t,T(i)), 1≤i≤4. N(1), N(2) and N(3) indicate the natural payoffs of f(1), f(2) and f(3). Note that the (i+1)-th bond is the i-th natural payoff associated with the i-th forward rate. In this example a common numeraire, N(t), was chosen for all the forward rates, namely the natural numraire of the second forward rate. With this numeraire, forward rate f(1) will have a drift given by Equation (12), forward rate f(2) will have no drift (Equation (13)), and forward rate f(3) will have a drift given by expression (11).

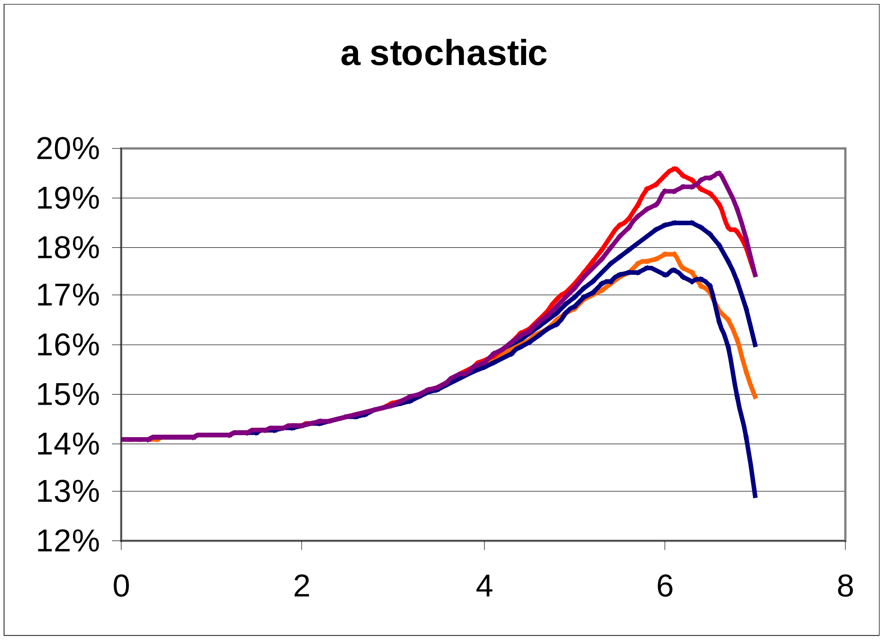

Fig. 5: A few sample instantaneous volatility curves for a 7-year expiry forward rate when only the a coefficient is stochastic according to Equations (6) and (10). The deterministic curve is also given as a reference. The parameters used are shown in Tab. I. See the text for more details.

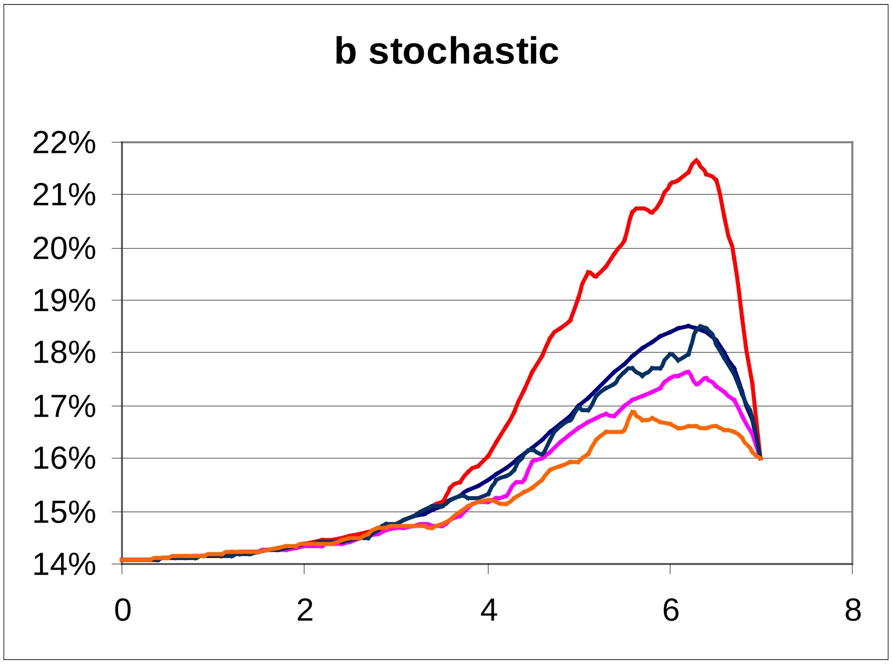

Fig. 6: Same as Fig. 5 when only the b coefficient is stochastic according to Equations (7) and (10). The deterministic curve is also given as a reference. The parameters used are shown in Tab. I. See the text for more details.

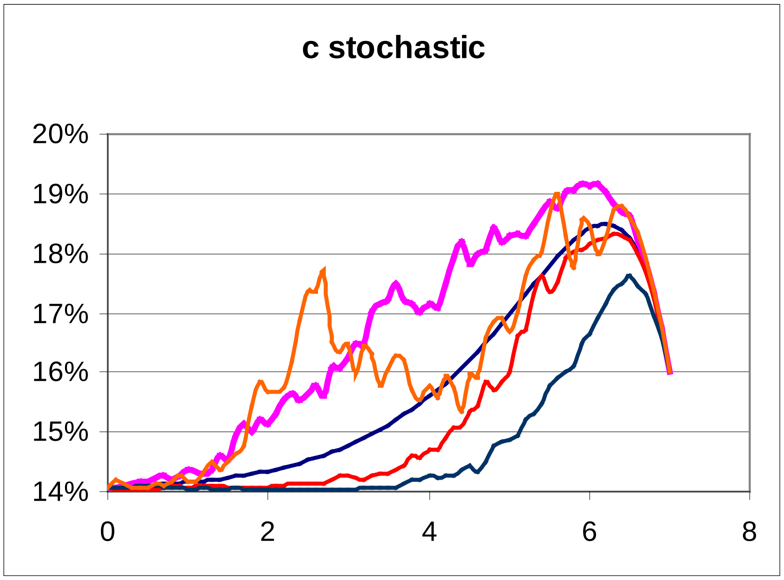

Fig. 7: Same as Fig. 5 when only the b coefficient is stochastic according to Equations (8) and (10). The deterministic curve is also given as a reference. The parameters used are shown in Tab. I. See the text for more details.

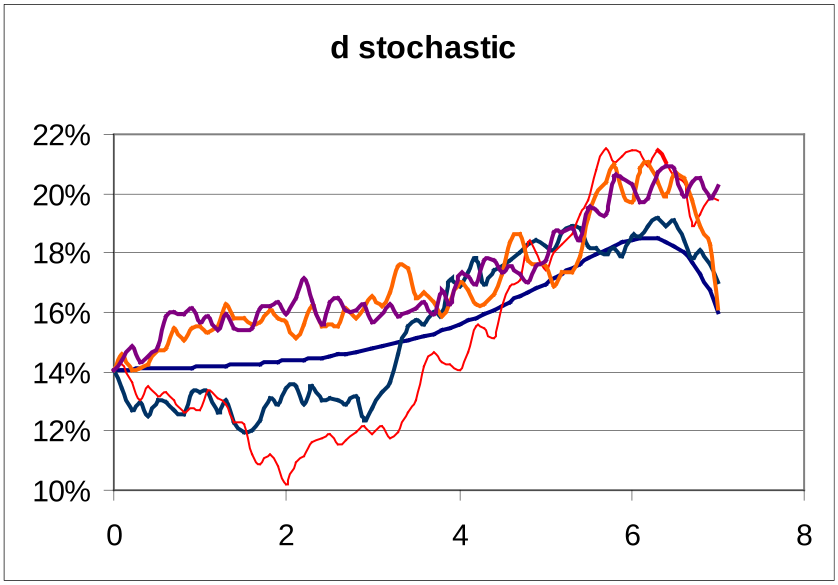

Fig. 8: Same as Fig. 5 when only the b coefficient is stochastic according to Equations (9) and (10). The deterministic curve is also given as a reference. The parameters used are shown in Tab. I. See the text for more details.

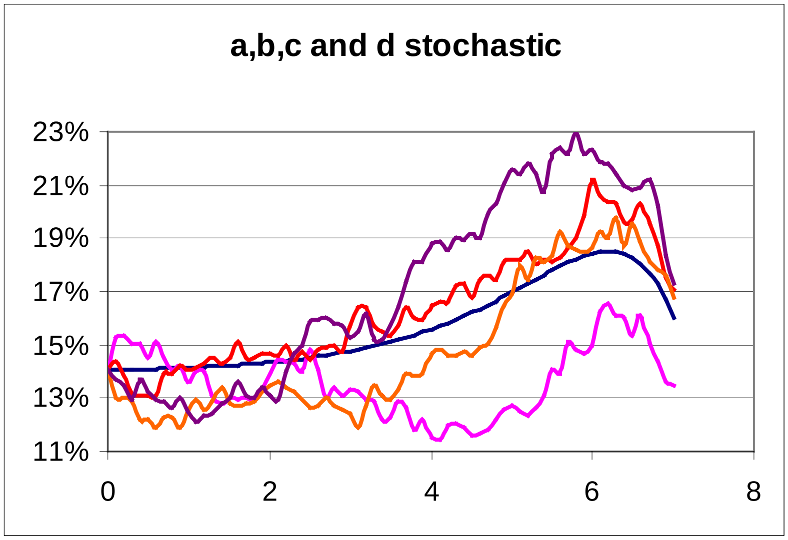

Fig. 9: Same as Fig. 5 when only the all the coefficient are allowed to be stochastic according to Equations (5) to (10). The deterministic curve is also given as a reference. The parameters used are shown in Tab. I. See the text for more details.

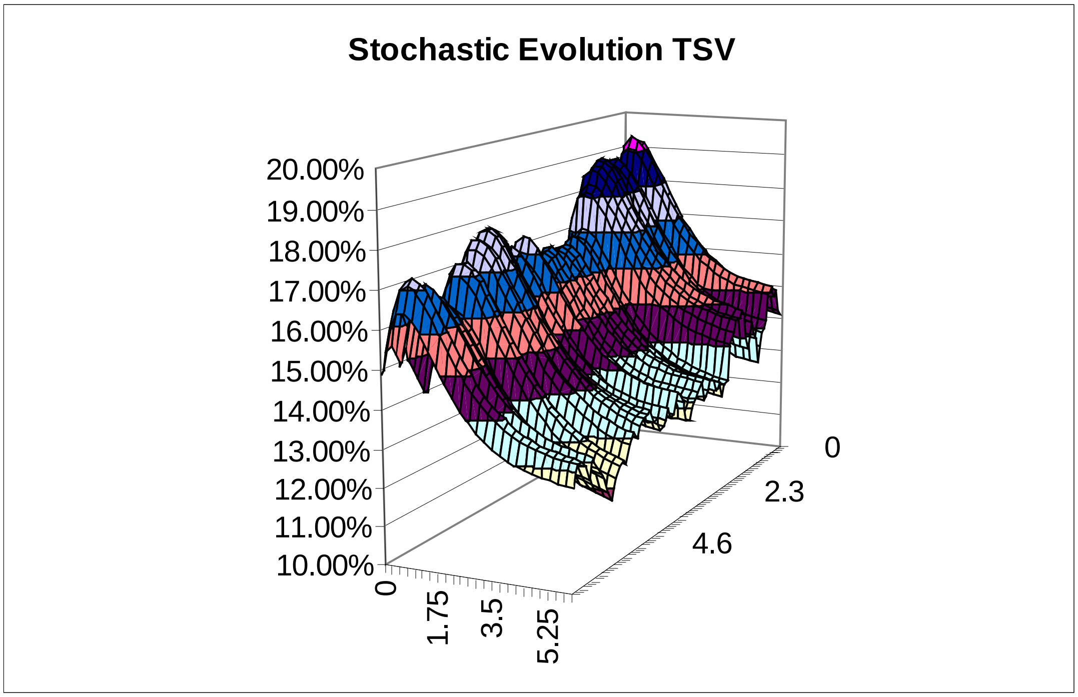

Fig. 10: A model term structure of volatilities (at-the-money) obtained with the parameters in Tab. I

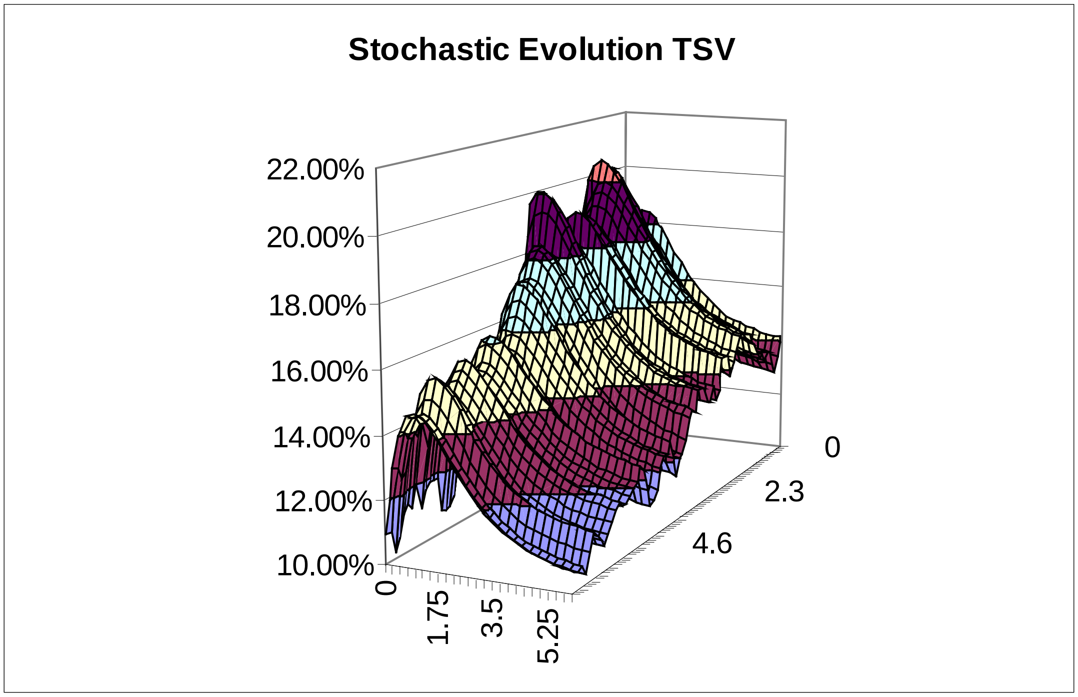

Fig. 11: Another model term structure of volatilities (at-the-money) obtained with the parameters in Tab. I

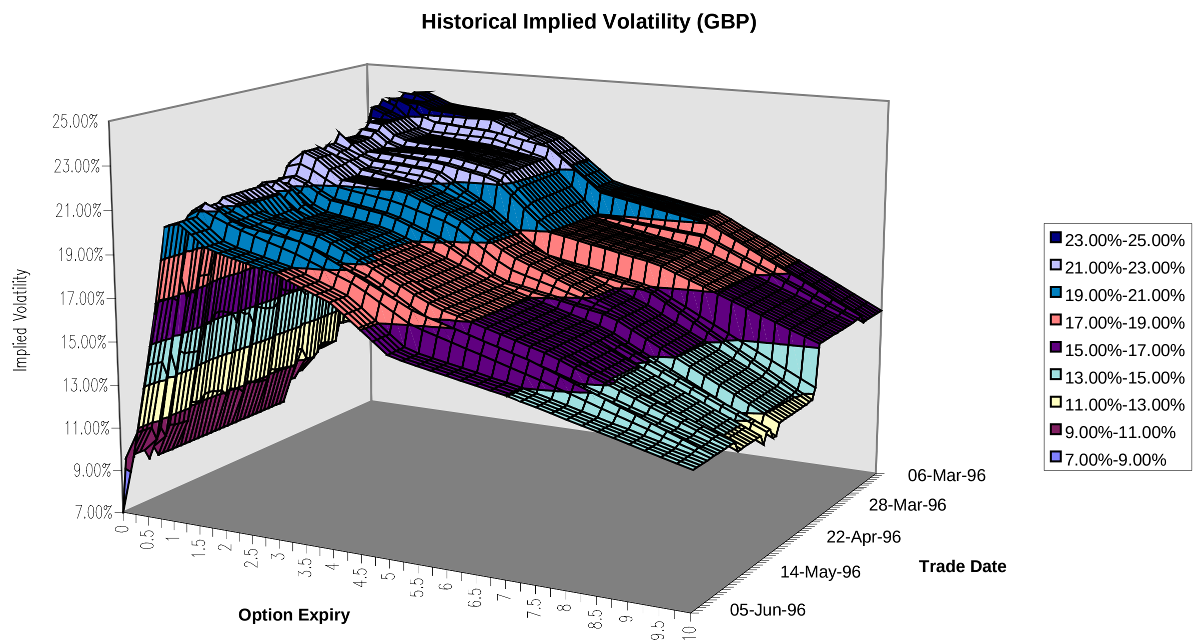

Fig 12: Implied volatilities for different maturities collected over a period of approximately 3 months in 1996 for GBP

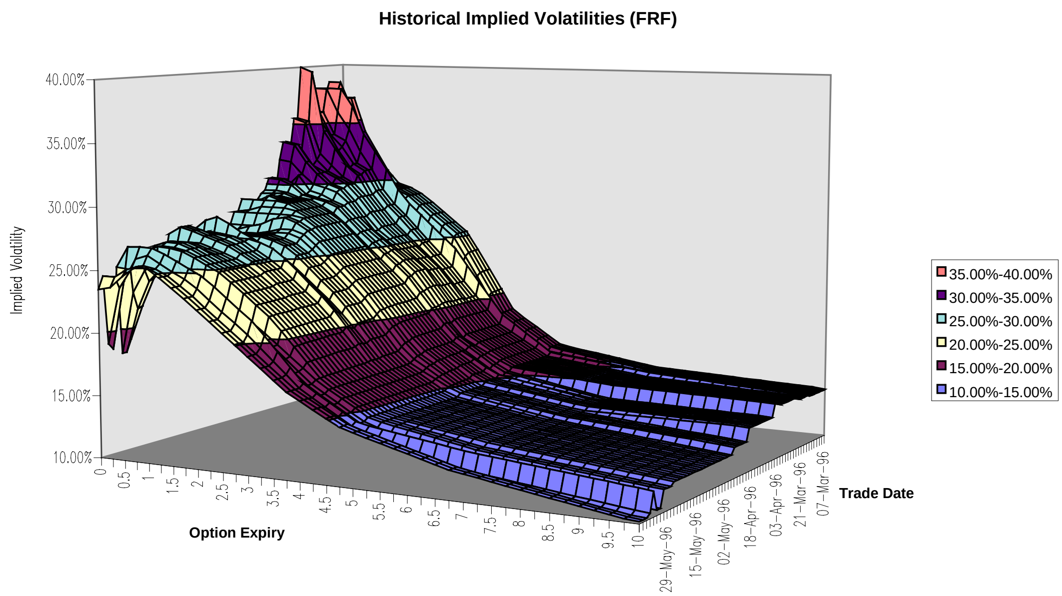

Fig 13: Implied volatilities for different maturities collected over a period of approximately 3 months in 1996 for FRF

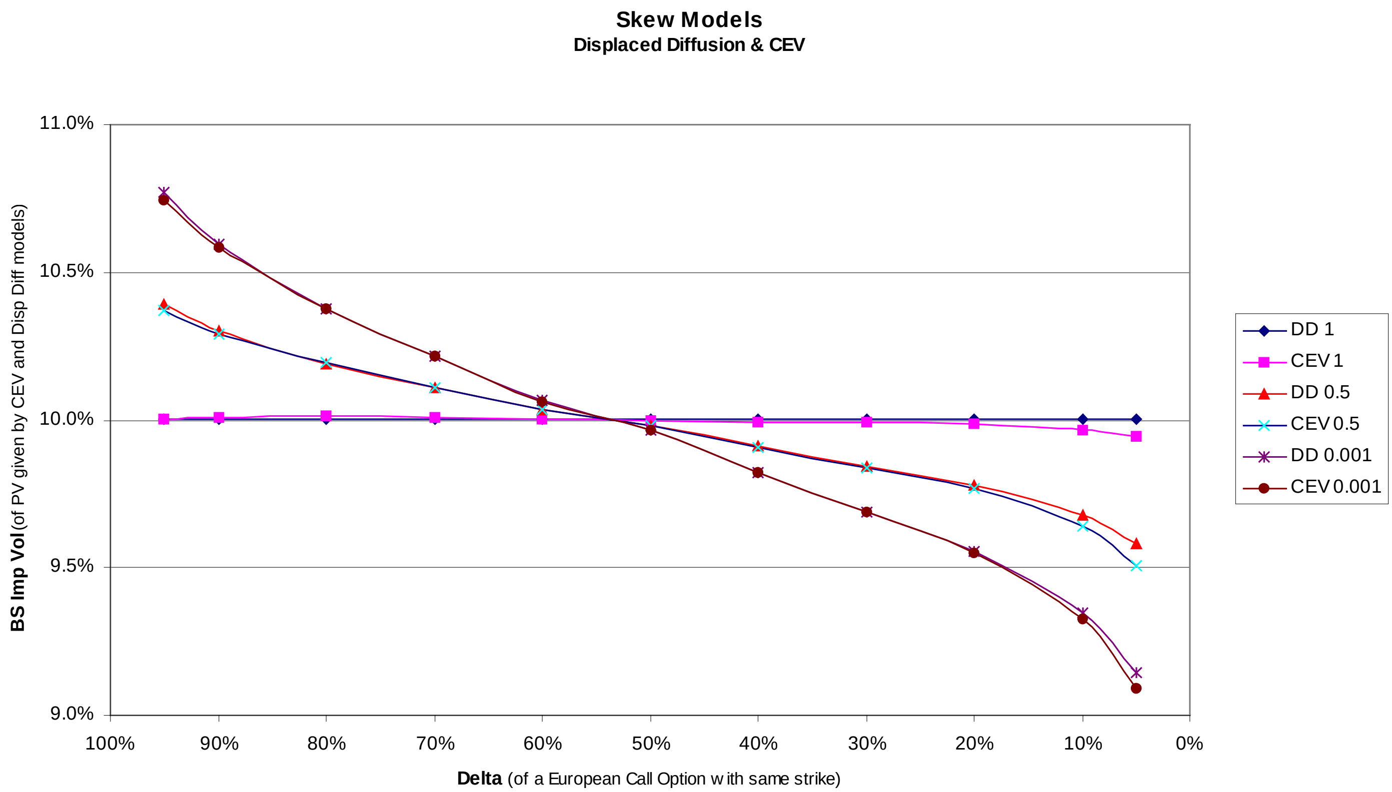

Fig. 13bis: The implied volatility obtained for call option prices with strikes corresponding to different (Balck)-delta values from 95% to 5%, for the log-normal case (DD 1 CEV 1), the square-root process (DD 0.5, CEV 0.5) and the (almost) normal case (DD 0.001, CEV 0.001), from work by Marris (1999), whose help is gratefully acknowledged.

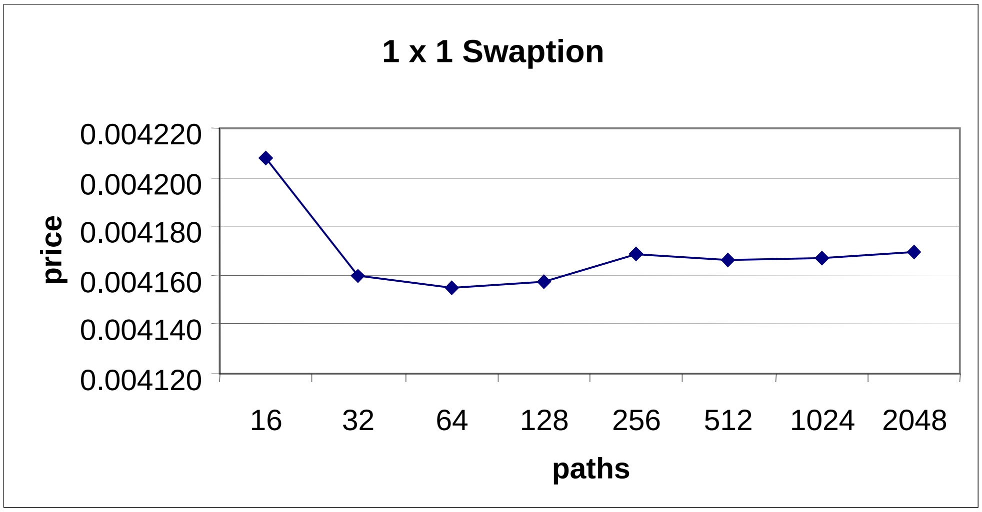

Fig 14: Convergence as a function of number of paths for the 1 x 1 swaption.

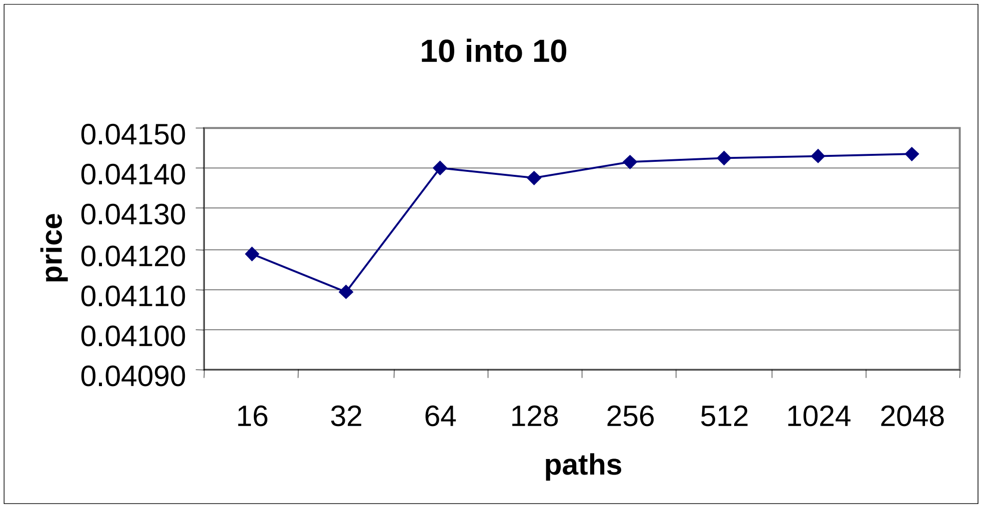

Fig 15: Same as figure 14 for a 10 x 10 swaption.

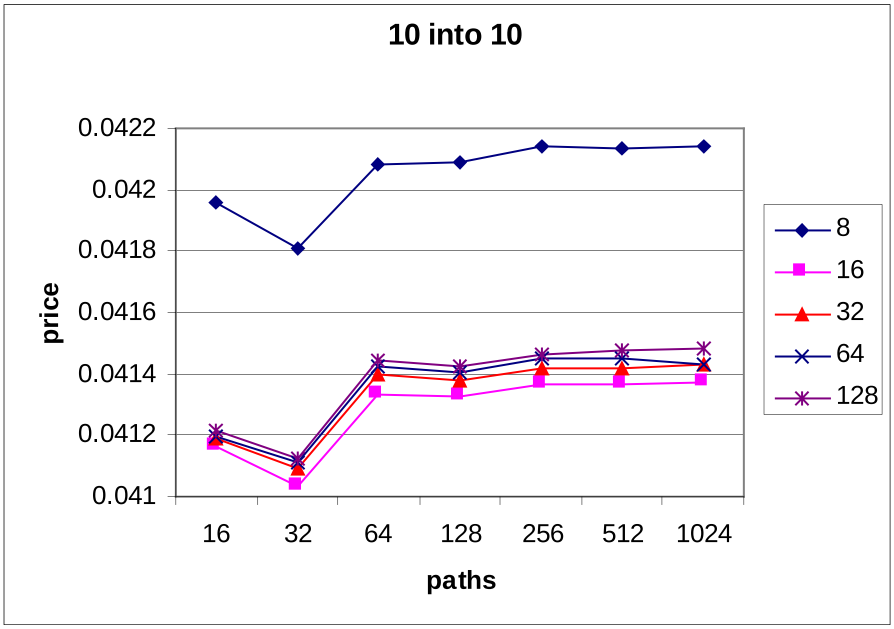

Fig. 16: Convergence as a function of number of paths and step size.

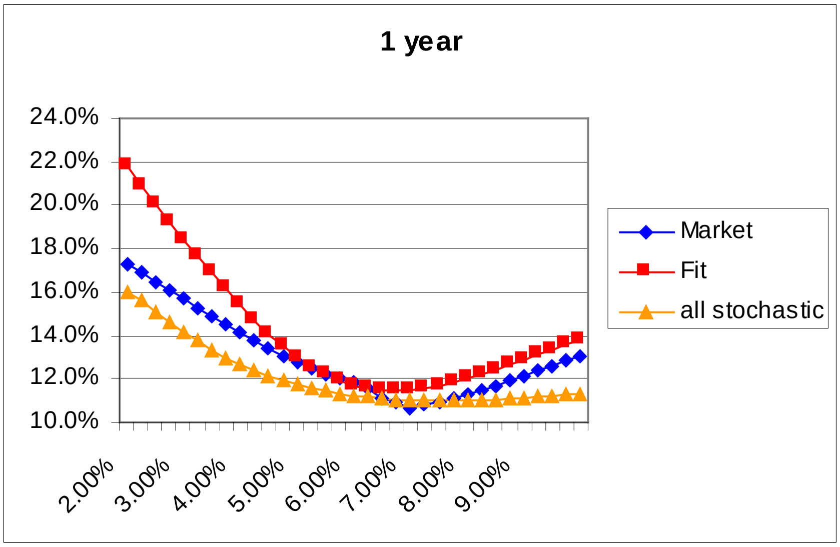

Fig. 17: See below common captions for Figs 17 to 21.

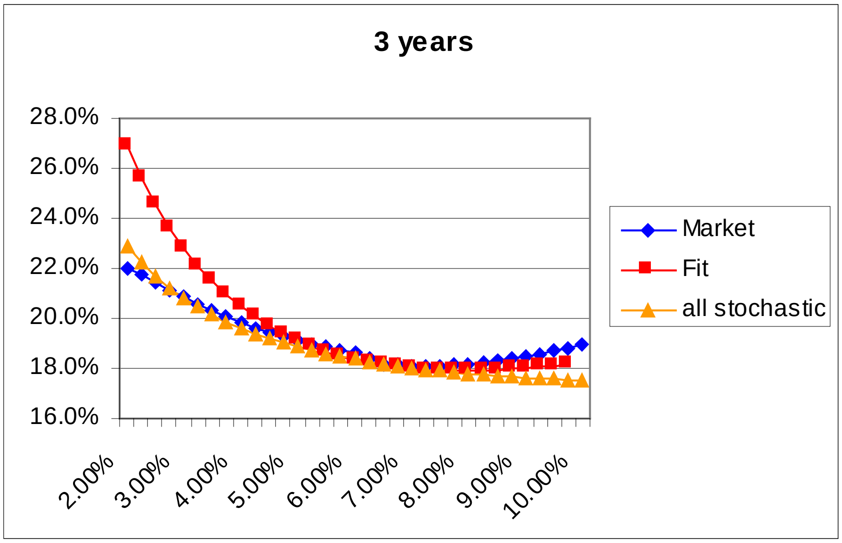
3 years
28.0%
26.0%
24.0%
22.0%
20.0%
18.0%
16.0%
2.00%
3.00%
4.00%
5.00%
6.00%
7.00%
8.00%
9.00%
10.00%
Market
Fit
all stochastic

Fig 18

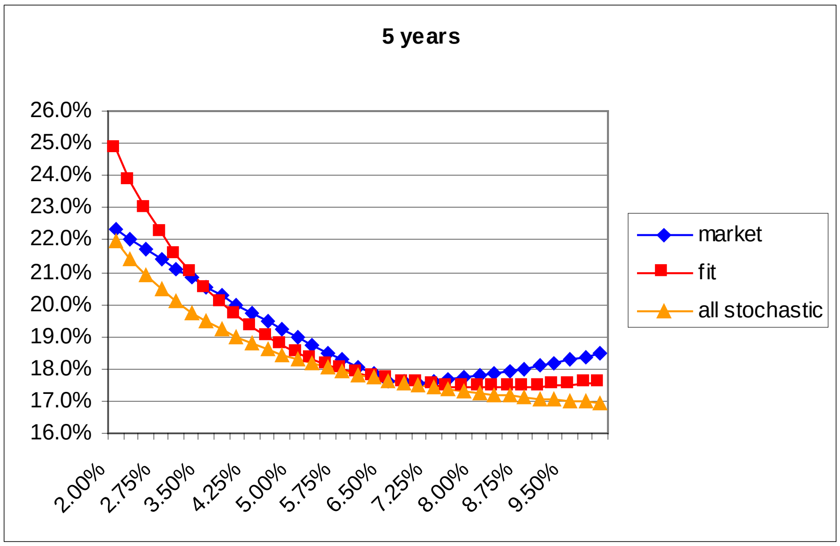
5 years
26.0%
25.0%
24.0%
23.0%
22.0%
21.0%
20.0%
19.0%
18.0%
17.0%
16.0%
2.00%
2.75%
3.50%
4.25%
5.00%
5.75%
6.50%
7.25%
8.00%
8.75%
9.50%
market
fit
all stochastic

Fig 19

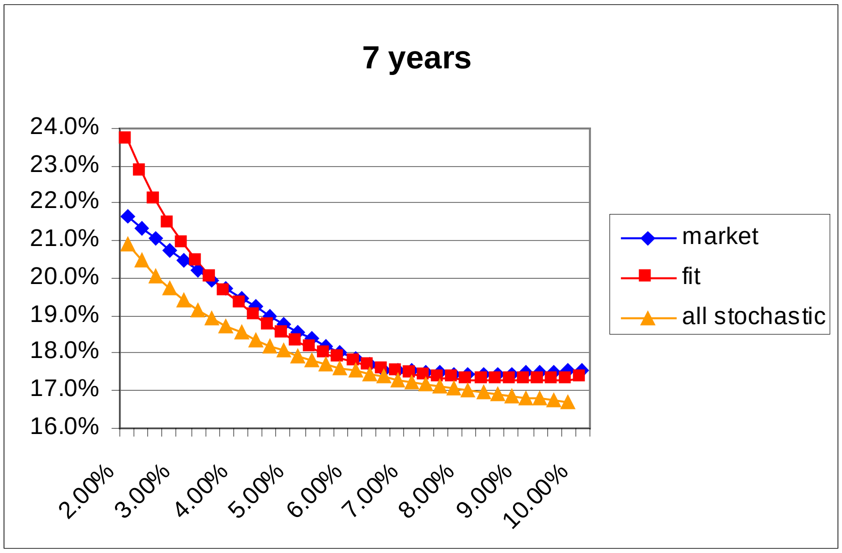
7 years
24.0%
23.0%
22.0%
21.0%
20.0%
19.0%
18.0%
17.0%
16.0%
2.00%
3.00%
4.00%
5.00%
6.00%
7.00%
8.00%
9.00%
10.00%
market
fit
all stochastic

Fig. 20

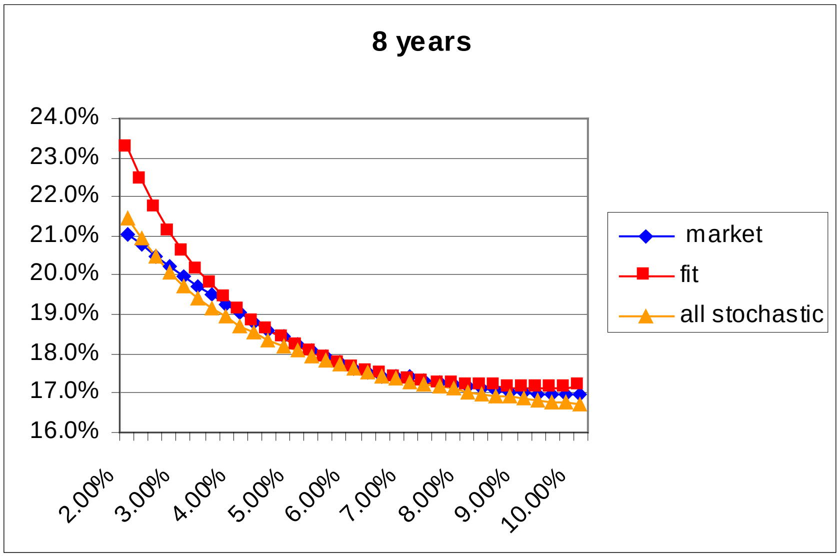

Fig. 21

Figs. 17 to 21: The fit across strikes to the market implied volatilities (curve labelled 'market') obtained with a, b, c and d stochastic (curve labelled 'all stochastic'), or just d stochastic (curve labelled 'Fit'), for different maturities ranging from 1 (Fig. 17) to 8 years (Fig.21) for the GBP surface discussed in the text.

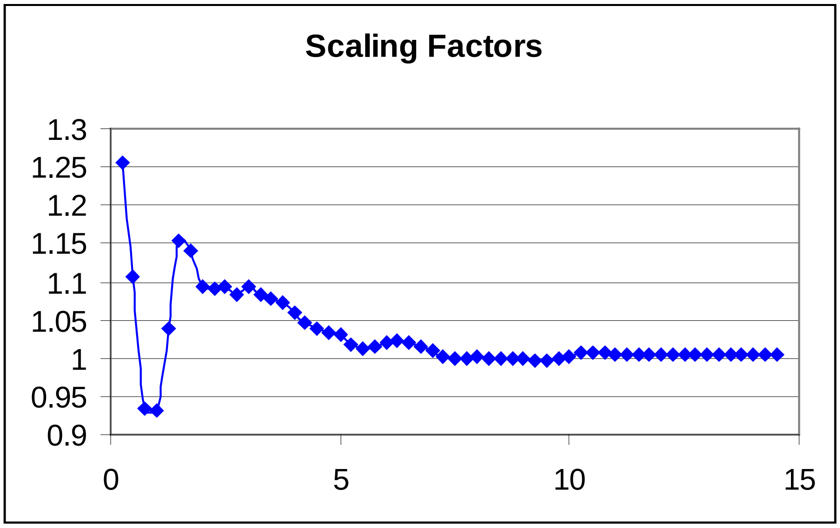

Fig. 22: The scaling factors necessary to ensure perfect pricing of the at-the-money caplets obtained in the case of only d stochastic.

a	0.02		Rsa	0.1
b	0.1		Rsb	0.1
c	1		RsLNc	0.1
d	0.14		RsLNd	0.1
Reva	0.02		Vola	0.008
Revb	0.1		Volb	0.016
RevLNc	0		VolLNc	0.12
RevLNd	-1.96611		VolLNd	0.1

Tab I: The numerical values for the a, b, c and d coefficients and their Ornstein-Uhlenbeck processes used in the calculations that produced Fig.5 to 9.

	Normal LogNormal	Initial Value	Vol	Reversion Speed	Reversion Level
a	0	-0.020	0.05	0.5	-0.020
b	0	0.108	0.1	0.3	0.108
c	1	0.800	0.1	0.5	0.800
d	1	0.114	0.2	0.4261	0.114

Tab III: The values of the stochastic parameters used for the convergence test

Approximate Vega

Length 1	5	10	Expiry
0.0002	0.0010	0.0015	0.5
0.0003	0.0013	0.0021	1
0.0004	0.0018	0.0028	2
0.0005	0.0022	0.0035	5
0.0005	0.0021	0.0034	10

Tab IV: The vega in basis points for some of the swaptions used in the tests. Following market practice, one vega has been taken as the typical bid-offer spread for an at-the-money European swaption in market size.

Paths	Length 1	5	10	Expiry
16	0.002808	0.013501	0.019882	0.5
	0.004207	0.018759	0.027485	1
	0.006743	0.025219	0.03667	2
	0.00774	0.028798	0.043695	5
	0.006901	0.026803	0.041187	10
				Expiry
32	0.002791	0.013365	0.01969	0.5
	0.00416	0.018484	0.027091	1
	0.006594	0.024692	0.035929	2
	0.007624	0.028305	0.042944	5
	0.006833	0.026712	0.041092	10
				Expiry
64	0.00279	0.013388	0.019723	0.5
	0.004155	0.018498	0.027127	1
	0.006603	0.024757	0.036062	2
	0.00767	0.028492	0.043212	5
	0.00691	0.026935	0.0414	10
				Expiry
128	0.002792	0.013386	0.019717	0.5
	0.004158	0.018492	0.027107	1
	0.006603	0.024731	0.036006	2
	0.007698	0.028563	0.0433	5
	0.006923	0.026944	0.041377	10
				Expiry
256	0.002796	0.013403	0.01974	0.5
	0.004169	0.018531	0.027158	1
	0.00662	0.024781	0.036077	2
	0.007702	0.028574	0.043329	5
	0.006927	0.026968	0.041418	10
				Expiry
512	0.002794	0.013399	0.019738	0.5
	0.004166	0.018524	0.027156	1
	0.006628	0.024817	0.036128	2
	0.00771	0.028599	0.043369	5
	0.006924	0.026965	0.041425	10
				Expiry
1024	0.002794	0.013395	0.019732	0.5
	0.004167	0.018518	0.027147	1
	0.006629	0.024809	0.036114	2
	0.007714	0.028609	0.043376	5
	0.006927	0.02697	0.041432	10
				Expiry
2048	0.002794	0.0134	0.019739	0.5
	0.004169	0.01853	0.027161	1
	0.006632	0.024809	0.036113	2
	0.007704	0.028586	0.043354	5
	0.006928	0.026973	0.041437	10

Tab V: The convergence of the swaption prices as a function of number of paths. See the text for details.

	Normal LogNormal	Initial Value	Vol	Reversion Speed	Reversion Level		
a	0	-0.020	0.05	0.500	-0.020	Beta	0.1
b	0	0.108	0.1	0.300	0.108	Strike	5.00%
c	1	0.800	0.1	0.500	0.800	Paths	16,384
d	1	0.114	0.2	0.426	0.114	Steps	32

Tab VI: The simulation parameters used in the test, together with the coefficients for the evolution of the instantaneous volatility. A displaced-diffusion coefficient α of 0.02 was used in conjunction with these parameters.

Numeraire	Caplet	Put	Forward
0	0.001155	0.000969	0.000186
1	0.001155	0.000969	0.000186
2	0.001154	0.000969	0.000185
3	0.001155	0.000969	0.000186
4	0.001155	0.000969	0.000186
5	0.001155	0.000969	0.000186
6	0.001155	0.000969	0.000186
7	0.001155	0.000969	0.000186
8	0.001155	0.000969	0.000186
9	0.001155	0.000969	0.000186
10	0.001155	0.000969	0.000186
Exact Price of a 5% FRA			0.000186

Tab VII: The values of the Caplet and Floorlet used to obtain, via call/put parity, the value of the FRA.

	Normal LogNormal	Initial Value	Vol	Reversion Speed	Reversion Level		
a	0.000	-0.020	0.050	0.500	-0.020	Beta	0.1
b	0.000	0.108	0.100	0.300	0.108	Displacement	0.0205
c	1.000	0.800	0.100	0.500	0.800		
d	1.000	0.114	0.200	0.426	0.114		

Tab VIII: The values obtained by best fitting the displacement coefficient and the parameters of the stochastic volatility processes (5) to (9) to the GBP caplet market (February 2001). These values were then used in the fitting to the caplet prices.

Expiry Times	Global K	Market ATM Vols	Global Calib ATM Vols	Rescaled Ks	ATM Calibrated Vols
0.25	1.25	**9.76%**	17.95%	0.682	**9.77%**
0.5	1.11	**9.68%**	17.51%	0.612	**9.69%**
0.75	0.93	**9.64%**	15.85%	0.569	**9.64%**
1	0.93	**11.10%**	16.59%	0.623	**11.10%**
1.25	1.04	**13.52%**	19.15%	0.733	**13.51%**
1.5	1.15	**15.95%**	21.81%	0.843	**15.95%**
1.75	1.14	**16.52%**	21.94%	0.858	**16.51%**
2	1.09	**16.43%**	21.25%	0.845	**16.42%**
2.25	1.09	**16.97%**	21.39%	0.866	**16.96%**
2.5	1.09	**17.44%**	21.48%	0.887	**17.43%**
2.75	1.08	**17.63%**	21.36%	0.895	**17.62%**
3	1.09	**18.16%**	21.55%	0.921	**18.15%**
3.25	1.08	**18.19%**	21.33%	0.924	**18.18%**
3.5	1.08	**18.11%**	21.16%	0.922	**18.10%**
3.75	1.07	**18.16%**	21.00%	0.928	**18.15%**
4	1.06	**18.07%**	20.71%	0.925	**18.06%**
4.25	1.05	**17.81%**	20.36%	0.916	**17.80%**
4.5	1.04	**17.68%**	20.11%	0.914	**17.67%**
4.75	1.03	**17.60%**	19.92%	0.914	**17.59%**
5	1.03	**17.66%**	19.85%	0.917	**17.65%**
5.25	1.02	**17.38%**	19.52%	0.906	**17.37%**
5.5	1.01	**17.30%**	19.33%	0.906	**17.29%**
5.75	1.01	**17.42%**	19.29%	0.916	**17.41%**
6	1.02	**17.74%**	19.38%	0.934	**17.73%**
6.25	1.02	**17.85%**	19.33%	0.944	**17.84%**
6.5	1.02	**17.89%**	19.24%	0.949	**17.88%**
6.75	1.01	**17.73%**	19.04%	0.945	**17.72%**
7	1.01	**17.65%**	18.90%	0.944	**17.64%**
7.25	1.00	**17.45%**	18.69%	0.937	**17.45%**
7.5	1.00	**17.39%**	18.55%	0.936	**17.38%**
7.75	1.00	**17.43%**	18.50%	0.942	**17.42%**
8	1.00	**17.53%**	18.48%	0.949	**17.52%**
8.25	1.00	**17.54%**	18.40%	0.953	**17.53%**
8.5	1.00	**17.55%**	18.34%	0.956	**17.55%**
8.75	1.00	**17.56%**	18.27%	0.960	**17.55%**
9	1.00	**17.55%**	18.21%	0.962	**17.54%**
9.25	1.00	**17.53%**	18.14%	0.964	**17.53%**
9.5	1.00	**17.58%**	18.10%	0.969	**17.57%**
9.75	1.00	**17.67%**	18.10%	0.977	**17.67%**
10	1.00	**17.82%**	18.18%	0.983	**17.81%**
10.25	1.01	**17.93%**	18.20%	0.991	**17.93%**
10.5	1.01	**17.99%**	18.21%	0.995	**17.99%**
10.75	1.01	**17.98%**	18.16%	0.996	**17.98%**
11	1.01	**17.97%**	18.13%	0.996	**17.96%**
11.25	1.01	**17.98%**	18.11%	0.998	**17.98%**
11.5	1.01	**18.00%**	18.11%	1.000	**18.00%**
11.75	1.01	**18.01%**	18.08%	1.002	**18.00%**
12	1.00	**17.98%**	18.14%	0.995	**17.97%**
12.25	1.00	**17.98%**	18.11%	0.996	**17.98%**
12.5	1.00	**17.97%**	18.11%	0.997	**17.97%**
12.75	1.00	**17.97%**	18.08%	0.999	**17.97%**
13	1.00	**17.95%**	18.08%	0.996	**17.95%**
13.25	1.00	**17.95%**	18.06%	0.998	**17.95%**
13.5	1.00	**17.92%**	18.07%	0.995	**17.92%**
13.75	1.00	**17.92%**	18.04%	0.997	**17.92%**
14	1.00	**17.89%**	18.06%	0.994	**17.89%**
14.25	1.00	**17.88%**	18.04%	0.995	**17.88%**
14.5	1.00	**17.84%**	18.05%	0.992	**17.84%**

Tab IX: The at-the-money market implied volatilities for a series if caplets (GBP curve, February 2001) (column ‘Market ATM Vols’), the model caplet prices obtained with a global coefficient fit to the market (column ‘Global Calib ATM Vols’), the re-scaling factors (column ‘Rescaled Ks’) and the fitted caplet prices (column ‘ATM Calibrated Vols’).

	5 x 10	10 x 1	10 x 5	10 x 10
Full Monte Carlo	0.0286	0.0069	0.0270	0.0440
Approximate	0.0286	0.0069	0.0269	0.0438

Tab. XI: The 'true' (row 'Full Monte Carlo') an approximate (row 'Approximate') prices for several European swaptions, as described in the text.

Expiry	Residual Length	Input Vols	Output Vols
1 X 8		13.920%	13.919%
1.5 X 7.5		14.000%	13.999%
2 X 7		14.080%	14.079%
2.5 X 6.5		14.132%	14.130%
3 X 6		14.183%	14.182%
3.5 X 5.5		14.340%	14.338%
4 X 5		14.497%	14.495%
4.5 X 4.5		14.708%	14.706%
5 x 4		14.919%	14.916%
5.5 x 3.5		15.273%	15.270%
6 x 3		15.628%	15.624%
6.5 x 2.5		16.018%	16.013%
7 x 2		16.407%	16.401%
7.5 x 1.5		16.407%	16.401%

Tab XII: The portion of the market swaption matrix corresponding to a set of co-terminal swaptions (GBP swaption matrix (February 2001). The at-the-money 'exact' (Monte Carlo) and approximate swaption implied volatilities are given in columns 'Input Vols' and 'Output Vols', respectively.

	Initial Value	Vol	Reversion Speed	Reversion Level
a	-0.089792	0.0000	0.20000	-0.08979
b	0.1083503	0.0000	0.20000	0.10835
c	0.9677668	0.0000	0.20000	0.967767
d	0.1137136	0.3976	0.42610	0.113714
a	-0.121691	0.0183	0.17018	-0.12169
b	0.1683858	0.0098	0.19798	0.168386
c	0.9512499	0.0005	0.19957	0.95125
d	0.1212483	0.1870	0.49167	0.121248

Tab XIII: The initial values a(0), b(0), c(0) and d(0) and the optimal parameters for their processes obtained by fitting globally to the whole caplet surface, and by allowing either just d (top boxes) or all the coefficients (bottom boxes) to be stochastic.