A Two-Regime, Stochastic-Volatility Extension of the LIBOR Market Model - Draft

Riccardo Rebonato[+,*] & Dherminder Kainth[+]
[+]QUARC
[*]Oxford University

June 12, 2003

Abstract

We propose a two-regime stochastic volatility extension of the LIBOR market model that preserves the positive features of the recently introduced (Joshi and Rebonato 2001) stochastic-volatility LIBOR market model (ease of calibration to caplets and swaptions, efficient pricimg of complex derivatives, etc.) and overcomes most of its shortcomings. We show the improvements by analysing empirically and theoretically the real and the model-produced change sin swaption implied volatility.

1 Introduction

1.1 The Literature and Modelling Context

Rebonato and Joshi (2002) have recently presented empirical work about the changes in market implied volatility swaption matrices (see also Rebonato (2001) and Rebonato (2003)). These empirical investigations can be profitably looked at in the light of the stochastic-volatility extensions of the LIBOR market model recently introduced by Joshi and Rebonato (2001), Andersen and Andreasen (2000), Hagan et al (2003), among others. The common underlying modelling philosophy is to posit a CEV process (possibly proxied by a displaced diffusion for analytic tractability) with superimposed uncertainty in the volatility function modelled by means of one or more Brownian diffusions. The empirical findings convey a mixed picture of the adequacy of these modelling approaches.

To begin with the negative results, certain important features of the real data are not captured by the proposed models: the empirical data indicates, for instance, that the swaption matrix tends to oscillate between well-defined shape patterns, with different, and sometimes quite short, transition periods. See Rebonato and Joshi (2002). Such a behaviour is neither compatible with a stochastic volatility model with constant reversion speed, nor with a jump/diffusion process (which does not produce, in its standard formulation, stochastic smile surfaces), nor with any of the CEV extensions alluded to above. Linking the volatility in a deterministic manner to the stochastic forward rates could produce sharp moves in the level of swaption matrix, if the forward rates displayed a discontinuous behaviour (as in Glasserman and Kou (2000) and Glasserman and Merener (2001)). It is difficult, however, to see how a deterministic functional dependence on the forward rates could give rise to a sudden change in the shape of the swaption matrix. Possibly, a reversion level for the instantaneous volatility that underwent almost instantaneous transitions between a number of pre-defined values could provide a better description of the observed dynamics. The reversion speed, however, would have to change significantly (i.e. would have to display a short-lived burst) when these transitions occur if one wants to recover at the same time the diffusive behaviour of the implied volatility in 'normal times', and the quickness of the transition during 'crises', as observed, in particular, for USD.

Continuing with the 'bad news', the descriptive statistics of the empirical changes in implied volatility strongly reject the hypothesis that the instantaneous volatility should follow a diffusive (mean-reverting) behaviour. In particular, the empirical tails are far too fat when compared with the model-produced ones. Furthermore, if one orthogonalizes the covariance matrix of the changes in swaption implied volatilities (as in Rebonato and Joshi (2002)) the proportion of the variability explained by the first principal component is significantly smaller in reality (about 60%) than with the model-produced data (about 95%).

Despite these shortcomings, Rebonato and Joshi show that the mean-reverting stochastic-volatility approach they propose displays two important encouraging features: first of all, the qualitative shape of the first eigenvector turns out to bear a close resemblance with the corresponding empirical quantity. In particular, the same periodicity is observed in the real and model data. Second, the decaying behaviour of the first principal component as a function of increasing expiry, empirically observed when the real-data covariance matrix is orthogonalized, is found to be naturally recoverable and explainable by the mean-reverting behaviour for the instantaneous volatility. This feature in turn constitutes the most salient characteristic of the Joshi-Rebonato stochastic-volatility extension of the LIBOR market model. Furthermore, the values for the mean reversion that had been previously and independently obtained using static information (i.e. by fitting to the smile surface) turned out to be adequate to explain in a satisfactory way the qualitative features of such dynamic features as the shape of the eigenvectors (obtained from time series analysis). It is therefore fair to say that, despite the obvious shortcomings, a modelling approach of the type proposed by Joshi and Rebonato (2001) appears to be a useful first step in the right direction, and we propose to extend this approach in the present work.

1.2 The Relevance of the Proposed Approach

Many approaches provide a fit to current prices of similar quality. Indeed, Britten-Jones and Neuberger (2000) show that, given *any* stochastic-volatility model, it is always possible to recover an exogenous set of market caplet prices by adding a suitable local-volatility term. The main feature of the approach presented in this paper is therefore not a more accurate recovery of the empirical smile surface today, but a more convincing description of the evolution of the smile surface. It is not obvious, however, why a better stochastic-volatility description of the dynamics of the swaption matrix should be *per se* a desirable feature if, as shown below, the quality of the fit to market quantities (such as smiley caplet prices) is virtually the same in the simpler version of Joshi and Rebonato (2001). The answer lies in the practice of pricing and trading of exotic options. Typically, even if the trader uses a deterministic-volatility model, she will not simply carry out the delta-hedging predicted by the model, but will also carry out vega hedges to neutralize her exposure to the changes in the volatility surface. If the model used assumes deterministic volatility this practice is clearly logically inconsistent, but it is nonetheless universally adopted. The first requirement for the trader is therefore to recover the prices of the hedging options used at trade inception to vega-neutralize the trade, hence the importance given to the recovery of today's smile surface. During the life of the trade, however, the portfolio of the complex derivative and the plain-vanilla hedges will in general not remain vega neutral, and the trader will have to enter further vega-hedging trades. If the future state of the world at the re-hedging times were fully compatible with the model used to price the trade on day 0, this further hedging activity would have no economic effect. However, a second fundamental trading practice comes into play at this point, ie the practice of re-calibrating the model every day during the life of the trade so as to recover the then-current prices of plain-vanilla hedging options. In a way, the trader recognizes every morning the error of her ways, re-calibrates the model, and hedges *as if from now on the re-calibrated model will be 'true' for the rest of the life of the trade.*

The combined effect of the theoretically inconsistent vega (re-)hedging and re-calibration of the model therefore exposes the trader to future realizations of the smile surface (ie, to the future prices of the re-hedging options). The main criterion of success of a model, from the perspective of a complex -derivatives trader, should therefore be its ability to predict, either in a deterministic or in a stochastic manner, the current and *future* prices of the vega-hedging instruments. We therefore propose that the main criterion to choose between models which provide similar, good-quality fits to t*oday's* smile surface is their ability to predict in a realistic manner the future re-hedging costs incurred in vega hedging.

The extension of the stochastic-volatility extension of the LMM by Joshi and Rebonato presented in this work should be seen in this light: we attempt to provide a description of the smile surface dynamics that, by being more closely aligned with empirical evidence, will provide a better pricing tool for traders by implictly 'knowing' better than other approaches about future vega-re-hedging costs.

2 Description of the Stochastic-Volatility LIBOR Market Model

In order better to appreciate the changes brought about by the two-regime stochastic-volatility LMM here introduced it is useful to recall briefly the results presented in Joshi and Rebonato (2001). One can start from the usual deterministic-volatility (DV) LMM, and posit

$$\sigma(t,T) = k_T g(T-t) \tag{1}$$

$$g(T-t) = [a + b(T-t)]\exp[-c(T-t)] + d \tag{2}$$

where $\sigma(t,T)$ is the instantaneous volatility at time t of the Tmaturity forward rate, and k_T is a forward-rate specific constant needed in order to ensure correct pricing of the (at-the-money) associated caplet. Neglecting for the moment smiles, if one denotes by $\sigma_{Black}(T)$ the implied volatility of the caplet of expiry T, the caplet-pricing condition is ensured in the DV setting by imposing that

$$\frac{\sigma_{Black}(T)^2 T}{\int_0^T g(u,T)^2 du} = k_T^2 \tag{3}$$

When the instantaneous volatility is deterministic, this formulation allows to determine the most time-homogenous evolution of the term structure of volatilities and of the swaption matrix consistent with a given family of parametrized functions $g(T-t)$ simply by imposing that the idiosyncratic terms, k_T, should be as constant as possible across forward rates. See, for a detailed discussion of this point, Rebonato (2002) or Mercurio and Brigo (2001). Once this volatility function has been chosen, the arbitrage-free stochastic differential equation for the evolution of the T_i-expiry forward rate in the Q-measure associated with the chosen numeraire is given by

$$\frac{df_{T_i}(t)}{f_{T_i}(t)} = \mu^Q(\{f_{T_j}(t)\},t)dt + \sigma(t,T_i)\sum_{k=1,m} b_{ik}dz_k^Q(t) \tag{4}$$

where dz_k^Q are orthogonal increments of standard Q-Brownian motions, $\mu^Q(\{f_{T_j}(u)\},u)$ is the measure-, forward-rate- and time-dependent drifts that reflect the conditions of no arbitrage, and the coefficients $\{\mathbf{b}\}$, linked by the caplet-pricing condition $\sum_{k=1,m} b_{ik}^2 = 1$, fully describe the correlation structure given the chosen number, m, of driving factors (see Rebonato (1999)).

In order to account for smiles, this standard formulation can be extended in two ways:

i) by positing a displaced-diffusion evolution of the forward rates according to

$$\frac{d(f_{T_i}(t)+\alpha)}{f_{T_i}(t)+\alpha} = \mu_\alpha^Q(\{f_{T_j}(t)\},t)dt + \sigma_\alpha(t,T_i)\sum_{k=1,m} b_{ik}dz_k^Q(t) \tag{5}$$

and

ii) by making the instantaneous volatility non-deterministic via the following stochastic mean-reverting behaviour for the coefficients a, b, c and d, or their logarithm, as appropriate:

$$da_t = RS_a(RL_a - a_t)dt + \sigma_a(t)dz_t^a \tag{6}$$

$$db_t = RS_b(RL_b - b_t)dt + \sigma_b(t)dz_t^b \tag{7}$$

$$d\ln[c_t] = RS_c(RL_c - \ln[c_t])dt + \sigma_c(t)dz_t^c \tag{8}$$

$$d\ln[d_t] = RS_d(RL_d - \ln[d_t])dt + \sigma_d(t)dz_t^d \tag{9}$$

In Equations (6) to (9) all the Brownian increments are uncorrelated with each other and with all the Brownian increments $dz_k^Q(t)$ and the symbols

$$RS_a, RS_b, RS_c, RS_d, RL_a, RL_b, RL_c, RL_d$$

denote the reversion speeds and reversion levels, respectively, of the relative coefficients, or of their logarithms.

The introduction of the displacement coefficient α (see Rubinstein (1983) and Marris (1999) for the link with the CEV model) is intended to account for the deviation from exact proportionality with the level of the basis point move of the forward rates: this feature translates to a monotonically decaying (with strike) component of the smile surface[1]. In addition, the stochastic behaviour for the (coefficients of) the instantaneous volatility is invoked in order to account in a financially convincing way for the more recently observed 'hockey-stick' shape of the smile curves. More precisely, given the econometric interpretation that can be given to a, b, c and d, Equations 6 to 9 allow the initial slope, the long-term level and the location of the maximum of the instantaneous volatility functions to be stochastic.

The general strategy that can be followed to calibrate the stochastic-volatility (SV) LMM presented above in an efficient way rests on three simple observations:

1. Given the posited independence between the forward rates and the stochastic volatilities, *conditional on a particular volatility path having been realized* the problem looks exactly like a standard (DV) LMM problem;

2. The Black formula is, at-the-money, almost exactly linear in the root-mean square volatility;

[1] It is customary to model this feature by means of a CEV approach (see, e.g. Andersen and Andreasen (2000), or Zuehlsdorff (2001)). Marris (1999), however, shows that there exists a close correspondence between the CEV and the displaced-diffusion dynamics, and that, once the two models are suitably calibrated, the resulting caplet prices are virtually indistinguishable over a very wide range of strikes and maturities. Marris also provides a thoeretical justification as to why this should be the case. Joshi and Rebonato therefore use the displaced-diffusion setting, which, unlike the CEV case, allows simple closed-form solutions for the realization of the forward rates after a finite period of time, as a computationally simple and efficient subsititute for the theoretically more pleasing CEV framework (which does not allow negative forward rates).

3. Joshi and Rebonato (2001) show that surprisingly few volatility paths are sufficient for an accurate sampling of the volatility probability density.

By making use of this approach Joshi and Rebonato show that it is simple and computationally effective to calibrate to caplets, to obtain the prices of a swaption matrix, to calibrate in a financially desirable way to the co-terminal swaptions that underlie a given Bermudan swaption and to price complex (path-dependent) derivatives. Point 1 above (ie, the independence between the increments of the volatility processes on the one hand on the processes for the underlying on the other) is the key to obtaining such simple results[2].

It is important to stress that the forward-rate coefficients k_T required in order to ensure perfect pricing of the caplets are almost invariably found to be very close to unity. This is important, because, given the time-independence of all the coefficients above, this ensure that the evolution of the caplet and swaption surface will almost exactly retain its statistical features in the future. See the discussion in Joshi and Rebonato (2001) about this important point.

3 The Proposed Extension

How could one improve upon this approach in such a way as to retain its desirable features and to take into account the empirical evidence discussed in the introduction?

The most salient missing features are probably

i) the ability to reproduce rapid transitions of the swaption matrix from one 'mode' to another;

ii) the ability to return, after one such transition has taken place, to a similar shape;

iii) the recovery of fatter tails in the distribution of changes of implied volatilities (in agreement with empirical data); and

iv) a better apportioning of the total variance among the eigenvectors obtained from orthogonalizing the changes in swaption implied volatilities.

A simple and natural way to model these features, while retaining the simplicity and intuition behind the approach described above is the following. We posit the existence of a latent variable, y, which follows a two-state Markov-chain process between two states, x and n, with transition probabilities:

$$\begin{bmatrix} \lambda_{x\to x} & \lambda_{n\to x} \\ \lambda_{x\to n} & \lambda_{n\to n} \end{bmatrix} \tag{10}$$

and which can only take up the values 1 (if state n prevails) or 0 (if state x prevails). The pricing procedure is then as follows.

[2] This does not mean that there can be no relationship between the forward rates and the volatilities, but that this has been fully captured by the CEV (DD) treatment.

1. One can begin by choosing a simple criterion to determine whether the swaption matrix is currently in the normal or excited state. Looking at the figures in Rebonato (2002) one such criterion could be whether the n-year-into-1-year series displays a hump or not.

2. There exist two instantaneous volatility functions for each forward rate, described by these two functional forms:

$$\sigma_i^n(t, T_i) = [a_t^n + b_t^n(T - t)] \exp(-c_t^n(T - t)) + d_t^n \tag{11}$$

$$\sigma_i^x(t, T_i) = [a_t^x + b_t^x(T - t)] \exp(-c_t^x(T - t)) + d_t^x \tag{12}$$

   with different coefficients $\{a^n, b^n, c^n, d^n\}$ and $\{a^x, b^x, c^x, d^x\}$ associated with the normal (superscript n) and excited state (superscript x).

3. At any point in time the instantaneous volatility for forward rate i, $\sigma_i(t, T_i)$, is given by

$$\sigma_i(t, T_i) = y\sigma_i^n(t, T_i) + (1 - y)\sigma_i^x(t, T_i) \tag{13}$$

4. All the coefficients $\{a^n, b^n, c^n, d^n\}$ and $\{a^x, b^x, c^x, d^x\}$ are stochastic, and follow the same Ornstein-Uhlenbeck process described in the original work by Joshi and Rebonato. Their processes are all uncorrelated with the forward rates.

5. The transition of the instantaneous volatility from the normal to the excited state occurs with frequency $\lambda_{n \to x}$, and the transition from the excited state to the normal state with frequency $\lambda_{x \to n}$. Notice that both frequencies are risk-adjusted and not real-world frequencies.

6. Since the same assumption of independence between the volatility processes and the forward rate processes is enforced, once again along each volatility path the problem is exactly equivalent to the deterministic case, apart from the fact that, at random times, the coefficients would switch from one state to the other.

7. Because of 5., the evaluation of the variances or covariances along each path proceeds exactly as described in Chapter 12, with possibly different coefficients 'half-way through' some of the paths if a transition has occurred. The evaluation of caplets and European swaptions would be practically unaltered.

4 Empirical Tests

4.1 Description of the Test Methodology

In order to test the effectiveness of the procedure proposed above we began by studying the qualitative behaviour of a two-regime stochastic-volatility LMM by choosing the two instantaneous volatilities functions (normal and excited) depicted in Fig 1. For our tests we chose the swaption data described in detail in Rebonato and Joshi (2002). The reason for using swaption rather than caplet data is also explained in Rebonato and Joshi (2002). With this data we implemented an adaptation of their algorithm , as follows:

- we fitted the parameters of the normal and excited volatility curves so as to obtain a best-fit to the market caplet smile surface (see the parameters in Tab I), *given the transition probabilities* $\lambda_{n\to x}, \lambda_{x\to n}$ *required to give an acceptable recovery of the relative weights of the first three eigenvalues (see the discussion in Section 4.2.4)*;
- we evolved the $a^{n(x)}, b^{n(x)}, c^{n(x)}, d^{n(x)}$ coefficients thus obtained from today's state over a simulation period of one week;
- we evolved the forward rates over the same one week period;
- given this world state we priced the at-the-money swaptions and obtained their implied volatilities;
- we repeated we procedure over a large number of time steps, thereby creating a time series of model implied volatilities for the swaptions
- these quantities were then compared with the market data.

The following observations are in order:

1. In order to limit the number of degrees of freedom, we did not treat all the $a^{n(x)}, b^{n(x)}, c^{n(x)}, d^{n(x)}$ coefficients as fully free-fitting parameters; instead we started from 'plausible' shapes for the 'normal' and 'excited' instantaneous volatility functions and locally optimized the parameters around these initial guesses;
2. The test was run by fitting to *caplet* data and then exploring *swaption* data. No best-fit to swaption-related quantities was attempted in the choice of $a^{n(x)}, b^{n(x)}, c^{n(x)}, d^{n(x)}$. The test is therefore quite demanding, in that it requires a satisfactory description of the evolution of swap rates using parameters estimated on the basis of forward-rate information alone.
3. The levels and shape of the normal and excited volatility curves applied to the risk-neutral world and not to the real-world measure. The same consideration applies to the frequency of transition from one state to the other. Therefore no immediate conclusions can be drawn from these values. We shall nonetheless present an order-of-magnitude comparison of the transition frequency.
4. Fits to caplet prices of very similar (and very good) quality can be obtained with different parameters for the normal and excited coefficients. Therefore recovery of the caplet prices is a poor criterion to choose between different instantaneous volatility curves. We suggest below that an

analysis of the kurtosis of the changes in implied volatilities and of the eigenvalues behaviour can provide a better criterion.

5. Statistical estimates of kurtosis are in general very noisy, and, anyhow, cannot be directly compared with the model values. This is because the model kurtosis will depend on the transition probability, and the fitted quantity is risk-adjusted (ie, pertains to the pricing measure, not the real-world measure). Therefore we did not attempt a fit to the kurtoses for the various swaption series. However, we present below the real and risk-adjusted values for a qualitative comparison. A meaningful comparison across measures can however be carried out by looking at the eigenvalues and eigenvectors. See Section 4.2.4.

4.2 Results

4.2.1 The Real and Model Path of Implied Volatilities

Figs 2 and 3 display time series of the changes in instantaneous and implied volatilities for a 5 x 5 constant-maturity swaption obtained using the procedure described above and the best-fit parameters in Tab. I. It is clear that the proposed process produces regime shifts also in the implied volatility changes, with most changes being 'small', and a relatively smaller fraction very high. This result in not *a priori* obvious, since implied volatilities are linked to the root-mean-squared instantaneous volatility, and regime changes in the latter could have been 'washed out' after the integration. Figs. 4 and 5 then shows time series for one of the empirical time series of implied volatilities, showing the qualitative similarity between the model and real-world data. Also in this case, the time spent in the normal and excited states in the real and risk-adjusted world cannot be directly compared, because of the risk-adjusted nature of the quantities (the level of the volatility curves and the frequency of the jump) that enter the pricing of caplets.

4.2.2 Recovery of Market Smile Surface

The quality of the fit to the market smile surface (USD data for date here) is shown in Fig. 6. The fit was obtained by choosing beforehand values of the transition probabilities that would produce acceptable ratios for the eigenvalues obtained by orthogonalizing the implied volatility model covariance matrix. See the discussion in Section 4.2.5. (Recall that the purely diffusive Joshi and Rebonato model loaded more than 95% of the explanatory power onto the first eigenvector). With these transition probabilities the coefficients for the normal and excited volatility curves were then obtained as described above.

The fit was obtained with time-independent coefficients, and therefore the resulting swaption matrix displays a desirable time-homogenous behaviour.

This fit should be compared with the fit obtainable with a regime switch between purely deterministic volatilities. This important feature is discussed in Section 5.

4.2.3 Kurtosis

The Table Kurt shows the kurtosis for our two-regime model for the several swaptions using the parameters listed above for the two instantaneous-volatility states and the jump intensities given above. Note that these values of kurtosis are substantially larger than the kurtosis that would be obtained using a single-regime stochastic-volatility LMM, and much closer to the values observed in the market. The usual caveats about the change of measure apply. Since the true distribution of the changes in implied volatility is not known a priori it is not possible to associate statistical error bars to the experimental values. However, in order to give an idea of the possible dispersion Table KURT2 displays the real-world estimates obtained using the first and second half of the available data.

4.2.4 Skew

No attempt was made to reproduce the skew of the distribution of the change in implied volatilities. Nonetheless Table SKEW shows a good agreement between the model (theoretical) and real world quantities. The same observations about the statistical error bars and the change of measure hold, and again the estimates obtained using the first and second half of the data are presented (see Table SKEW2).

4.2.5 Eigenvalues and Eigenvectors

Rebonato and Joshi (2002) argue that a comparison between the eigenvectors and eigenvalues estimated from real-world data and simulated by the model is a powerful tool to assess the quality of a model in general, and to overcome the difficulties in comparing real-world and risk-adjusted quantities. One of the main results of the work presented in Rebonato and Joshi (2002) was that the simple stochastic-volatility model produced a good qualitative shape for the eigenvectors obtained from the orthogonalization of the covariance matrix of the changes in implied volatilities. The relative size of the eigenvalues, however, in the original model was at variance with what observed in reality.

In order to investigate the dynamics of our new model, the changes in the artificial time series generated as described above were used to generate a covariance matrix. Unfortunately, the introduction of the risk-ajusted proabilities of transition between states destroys the measure-independence of the eigenvecots and eigenvalues. This is easy to understand: a non-zero transition probability between the normal and excited state will change the relative weight of the various eigenvectors; the transition probability used for pricing, however, contains a risk adjustment, and therefore the eigenvecors/values are measure-dependent. Therefore, simply by observing that the new eigenvectors and eigenvalues are in closer agreement to the empirical data, one cannot directly conclude that the model is better. Suppose, however, that a better agreement *is* found, ie that the model eigenvectors higher than the first are found to be more important than in the purely diffusive case. This 'improvement' could be misleading only if the aversion to transition risk had the effect of making investors more 'afraid' than actuarially justifiable of parallel moves in the swaption matrix, than of other modes of deformation. This is however counterintuitive and financially difficut to justify. We therefore belive that the comparison between the eigen-structures is therefore at least qualitatively meaningful.

After diagonalization of the covariance matrix, we examined both the magnitude of the first few eigenvalues and the form of the corresponding eigenvectors. See Figures 7 and 8. Notice how the qualitative shape of the eigenvectors is recovered not only for the first but also for the second and third. (Rebonato and Joshi discuss why it is correct to compare model and real-world eigenvalues and eigenvectors).

We were also interested in obtaining a distribution of the eigenvalues such that there was a significant fraction of the spectral weight in the second/higher eigenvalues, because failure to do so was of the main shortcomings of the model in Rebonato and Joshi (2002). In the original approach almost all of the spectral weights, ϑ, defined via:

$$\vartheta_i = \frac{\lambda_i}{\sum \lambda_k} \tag{14}$$

were concentrated in the first eigenmode (ϑ_1 ~ 97%). If we were to compute the spectral weights for our two-state jumpy model under the assumption that the jump rates are equal at approximately one jump per year from the normal to the excited state and vice versa we noticed that we would obtain a very similar picture of the dynamics to that obtained within the one-state model i.e., more than 95% of the spectral weight would be concentrated in the first mode. However, upon setting the jump intensities to values more similar to the real-world jump rates, we found a very different behaviour. The relative magnitude of the real-world jump rates can be approximately estimated from the empirical observation that the majority of time is spent in the unexcited state; and that then approximately once a year the curve jumps into the excited state in which it stays for approximately 2 - 3 weeks after which it returns to the unexcited state. This suggests real-world jump intensities of approximately 1(jump/year) from the lower to the upper state and of approximately 20 from the upper state to the lower.

If we take this order-of-magnitude estimate in the real world as appropriate for the risk-adjusted frequencies, we get far more promising results: the second eigenvalue now has ~12% of the spectral weight, and it takes four eigenvectors to 'explain' more than 99% of the observed variability. This is illustrated in Figs 9 and 10.

Despite the fact that reality (see Fig. 9) is still considerably more complex, we consider this new feature a significant improvement, because it will have a direct bearing on the future possible shapes of the implied volatility surface. We have argued in Section 2 why, given the practice of vega-rehedging and model recalibration, this ability to generate complex realistic future changes in the smile surface should be of great importance.

5 How Important Is the Two-Regime Feature?

Whenever a new modelling feature is introduce it is important to ascertain to what extent the better fit to available data is simply obtained by virtue of having more parameters at one's disposal, or because new meaningful features have indeed been introduced. In order to answer this question, a very stringent test has been carried out[3]. We have *a priori* chosen two instantaneous volatility functions, one for the normal and one for the excited state, with an overall shape consistent with the financial justification discussed above: a lower-level humped curve, and a higher monotonically decaying curve. *Each volatility curve was assumed to be deterministic.* The parameters of these curves were not fit in any way to current market data. We then rescaled the precise level of each curve by a positive constant. These two scaling factors were the first two fitting parameters. The second two fitting parameters were the transition probabilities $\lambda_{n\to x}, \lambda_{x\to n}$. So, we only had at our disposal four parameters (two scaling constants and two transition probabilities to fit a full smile surface spanning maturities form 1 to 15 years.

Some of the combinations of parameters used are shown in Tab. DetVol, and Fig 11 shows the corresponding instantaneous and implied volatilities for the normal and excited states obtained from the parameters in the first column.

The results of this much-simplified fit are shown in Fig 12. The parameters in the first column in Tab DetVol were used. The fits obtained with the other sets of parameters were of very similar quality, showing that the fit does not depend on the fine features of the chosen volatility functions, but on the overall financial mechanism based on the regime switch.

It is clear that, despite the fact that the fit is far from perfect, many of the qualitative features of the real-world data are correctly recovered. Furthermore, the relative levels of the normal and excited volatility states turned out to be consistent with the financial interpretation given to these two quantities; and the (risk adjusted) transition probabilities $\lambda_{n\to x}, \lambda_{x\to n}$ were also naturally found to imply a lower probability of transition from the lower to the excited state than vice versa.

We want to stress that a better fit could clearly have been obtained by optimizing over the parameters that control the shape of the normal and excited volatility curves, but this was not the purpose of the exercise. Actually, Tab DetVol shows that, even by starting from financially plausible but otherwise rather different initial guesses for the volatility curves, fits of similar quality, and with similar scaling and transition parameters were found. This should be an indication of the robustness of the approach.

We can therefore conclude that, while further modelling flexibility is probably required to obtain an accurate fit to today's smile surface, the two-regime feature should be taken as an important part of the description of the smile surface.

[3] Numerical help by Mark Joshi is gratefully acknowledged.

6 Conclusions

We have presented a simple two-regime stochastic-volatility LIBOR market model. Its main positive features are:

- it retains the features of fast convergence, simple and efficient pricing of interest-rate derivatives, and fast calibration to the caplet and swaption markets enjoyed by the simple stochastic-volatility model introduced by Rebonato and Joshi (2001) and Joshi and Rebonato (2002);

- it produces high-quality fit to the market data in a manner that is financially justifiable. Since this fit is obtained with time-independent parameters, and the forward-rate specific coefficients necessary to ensure perfect pricing of the caplets are very close to 1, the statistical properties of the volatility surfaces are time stationary, and the future 'looks' (statistically) like the present;

- the eigenvectors are recovered as well as with the simple stochastic-volatility model (and in good agreement with the real world). The relative sizes of the eigenvalues obtained with the jumpy model displays a marked improvement: the first eigenvalue is still too big, but the second is now approximately five times larger;

- the kurtosis and the skew of the distribution of the changes in the implied volatilities show good agreement with the real-world data;

- when a simple fit to today's smile surface was attempted starting from two simple deterministic volatility function, an acceptable fit for the whole surface was found even by using as few as four parameters, suggesting that the two-regime feature should be taken as an important part of the description of the smile surface.

Overall, we believe that the proposed approach can provide a realistic and practically useful description of the dynamics of forward rates and of their volatilities.

BIBLIOGRAPHY

Andersen L, Andreasen J, (2000), 'Volatility Skews and Extensions of the LIBOR Market Model', Applied Mathematical Finance, 7, March, 1-32

Britten-Jones M, Neuberger A, (1998) 'Option Prices, Implied Price Processes and Stochastic Volatility', London Business School working paper, available at www.london.edu/ifa

Glasserman P, Kou, S.G (2000) 'The Term Structure of Simple Forward Rates with Jump Risk' – Working paper - Columbia University

Glasserman P, Merener N, (2001) 'Numerical Solutions of Jump-Diffusion LIBOR Market Models' (2001), Working paper, Columbia University

Joshi M, Rebonato R, (2001) 'A Stochastic-Volatility, Displaced-Diffusion Extension of the LIBOR Market Model', QUARC (Quantitative Research Centre) working paper available at www.Rebonato.com, and accepted fo rpublication in Quantitative Finance, 2003

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

Rebonato R (1999) "On the Simultaneous Calibration of Multi-Factor Log-Normal Interest-Rate Models to Black Volatilities and to the Correlation Matrix", Journal of Computational Finance, Vol. 2, 4, 5-27 and QUARC (Quantitative Research Centre) Working Paper available at www.Rebonato.com

Rebonato R, (2003) 'Which Process Gives Rise to the Observed Dependence of Swaption Implied Volatilities on the Underlying', accepted fro publication in International Journal of Theoretical and Applied Finance'

Rebonato R, Joshi M (2002) 'A Joint Empirical/Theoretical Investigation of the Modes of Deformation of Swaption Matrices: Implications for the Stochastic-Volatility LIBOR Market Model', International Journal of Theoretical and Applied Finance, 5, (7) 667-694 and Working Paper, QUARC (Quantitative Research Centre), available at www.rebonato.com

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

Zuehlsdorff, C, (2001), 'Extended LIBOR Market Models with Affine and Quadratic Volatility', Working Paper, Department of Statistics, Rheinische Friederich-Wilhelms-Universitaet, Bonn, Germany