A Joint Empirical and Theoretical Investigation of the Modes of Deformation of Swaption Matrices: Implications for Model Choice

Riccardo Rebonato and Mark Joshi
QUARC - Quantitative Research Centre of the Royal Bank of Scotland

July 24, 2001

Abstract

We present a joint empirical/theoretical analysis of the changes in the implied volatility swaption matrix for two currencies (USD and DEM/EUR). We recognize the existence of a small number of recognizable shape patterns, and comment about the speed of transition between them. By Principal-Component-Analyzing the associated correlation and covariance matrices we highlight a non-trivial interpretation for the leading eigenvectors. We also compare the empirically obtained eigenvectors and eigenvalues with the corresponding quantities produced by the stochastic-volatility LIBOR market model of Joshi and Rebonato (2001). This allows us to perform a measure-independent comparison that is of intrinsic interest, and that can also provide a general blueprint for analyzing the realism of and choosing among similarly-fitting stochastic models. We find that mean reversion of the instantaneous volatility is a necessary condition in order to obatin the market-observed shape of the first eigenvector associated with the covariance matrix.

1 Introduction

Caplets and European swaptions enjoy a somewhat privileged status in the interest-rate derivatives world, in that they are the liquid plain-vanilla reference instruments used by the complex derivatives traders to hedge the 'convexity' of their positions (i.e. for hedging beyond the delta level). In this respect, the exotic interest-rate option trader tends to regard cap(lets) and European swaption almost as his most natural set of 'underlying' instruments. Indeed, one of the reasons of appeal of the LIBOR market model is its ability to recover the Black prices of these benchmark derivatives. Furthermore, a significant amount of direct screen-visible information is available for these instruments in almost real time, making their prices to a large extent model independent.

When examined in detail, the situation is actually somewhat more complex, since brokers' quotes of *cap*, rather than *caplet*, implied volatilities are posted on the screen[1]. In the presence of smiles, it is not a simple and model-independent task to extract from these quotes the at-the-money volatilities of the underlying caplets. The swaption information is, on the other hand, much 'cleaner', since each screen quote refers to one single at-the-money option. The construction of the swaption implied volatility matrix, defined precisely below, is therefore a much simpler and unambiguous task, requiring as it does at most a two-dimensional interpolation/extrapolation exercise. It is for this reason that we focus the analysis of the present work on swaption, rather than caplet, implied volatilities.

The standard LIBOR market model (see, e.g., Brace Gatarek and Musiela (1995) Musiela and Rutkowski (1997), Jamshidian (1997)) and, indeed, the Black formula (Black (1976)) in terms of which the implied volatilities of swaptions are quoted, assume a deterministic volatility structure. This, in turn, can only give rise to a deterministic evolution of the term structure of volatilities and of the swaption matrix. A common useful criterion in order to constrain the parameters of the LIBOR market model has been the requirement that the evolution of these two quantities should be as time-homogeneous as possible. See, for instance, Longstaff and Santa Clara (2000a) and (2000b), Andersen and Andreasen (1999), Sibenius (2000) and (1998), Mercurio and Brigo (2001), Rebonato (2001). Empirically, however, the term structure of volatilities and the swaption matrix have been observed to display a significant degree of stochasticity. This feature has become particularly noticeable in the aftermath of the Russia crisis, (see the discussion in Section 2), and has prompted the introduction of stochastic-volatility extensions of the LIBOR market model[2]. The work by Joshi and Rebonato (2001) is one example in this direction.

The quality of the fitting of the smile surface that a given model allows is, of course, an important criterion in order to decide whether it provides a reasonable description of the market they attempt to describe. As it is well known, however, a variety of financially very different modelling approaches can produce fits of essentially the same quality. Indeed, Britten-Jones and Neuberger (1999) show how an arbitrary stochastic-volatility model can be made to fit

[1] The market adopts the following convention in quoting a cap implied volatility: given the strike and the expiry of the cap, the associated implied volatility is the single number that must be input in the Black formulae for all the associated caplets in order to obtain the desired market price for the whole cap by adding up all the resulting caplet prices. Notice that this procedure implies a much greater degree of price opaqueness than what is normally entailed by quoting an implied volatility, since, even in the absence of smiles, a *cap* implied volatility is not related in any obvious way to the root-mean-squared volatility of any of the underlying caplets.

[2] Other extensions of the standard LIBOR market model have recently been introduced, either by allowing the diffuison coefficient ot be functionally dependent on the forward rates in a CEV fashion (see, e.g. Andersen and Adreasen), or by positing a jump-diffusive behaviour for the process of the forward rates (see, e.g. Glasserman and Kou, Jamshidian). Both approaches however, have been introduced with the main goal to account for non-flat smile curves; to the extent that the latter displays deterministic coefficients, it does not allow for stochastic changes in the swaption matrix (or, for that matter, in the term structure of volatilties). CEV models, or the approach suggested by Zuehlsdorff (2001), only allow stochasticity that is perfectly correlated with the Brownian shocks that affect the forward rates. They appear therfeore unlikely to be able to explain the richness of behaviours described in Section 2.

exactly by construction any (sufficiently regular) exogenous smile surface. The need to assess the reasonableness of a model by means other than the quality of the fit to *today's* market prices is therefore acutely felt in the trading and academic community. If, because of the difficulties in making direct use of the cap data, one chose to concentrate on swaption-related information, it would be tempting to try and ascertain the realism of a proposed model by comparing the market-observed and the model-produced changes in the swaption matrix. The exercise is however made difficult by the fact that one set of quantities is related to the risk-adjusted and the other to the real world. If, as it is the case for the Joshi-Rebonato (2001) approach, the source of randomness in the volatility dynamics is introduced by positing an additional (set of) Brownian shocks, there intervene a set of drift transformations in moving between one world and the other (see, e.g. Lewis (2000)). It is therefore not *a priori* obvious to what extent the resulting evolution of the real and model swaption matrices can be compared. Despite these difficulties, it would be highly desirable to bring statistical information to bear on the financial justification of the extensions to the LIBOR market model that are being introduced.

We address this problem from a joint empirical/theoretical perspective. First of all we analyze in Section 2 market data pertaining to the US and EUR/DEM swaption matrices. At a qualitative level, we begin by recognizing different shape patterns for the swaption surface in the two currencies, and we comment on the speed of transition from one pattern to the other. We then perform a Principal Component Analysis of the changes in at-the-money implied swaption volatilities by calculating and orthogonalizing the correlation and covariance matrices among changes in the implied volatilities that make up the swaption matrix. We provide an intuitive qualitative interpretation of the results (given the complex structure of a swaption matrix, our results are more interesting than the usual level/slope/curvature decomposition). To the best of our knowledge, no study either of the qualitative patterns of change of the swaption matrix or of its Principal Component Analysis have been presented in the literature so far. The only somewhat related work we are aware of is Alexander's (2000) Principal Component study of implied equity index volatilities.

We then move to the Principal Components Analysis of the changes in the swaption matrix as predicted by the stochastic-instantaneous-volatility model proposed by Joshi and Rebonato (2001) and estimate the eigenvectors and eigenvalues of the swaption matrix correlation implied by the theoretical evolution of the instantaneous volatility. The important advantage of this PCA-based methodology is that, as we discuss in section 3.1, the statistical (real-world) and the risk-adjusted principal components are unchanged in moving between measures, since they only depend on the (deterministic) volatility and correlation of the volatility process, and can therefore be directly compared. This allows us to circumvent the problems alluded to above of comparing real-world and model-produced changes in the overall shape of the swaption matrix.

It must be stressed that there are several levels at which the 'validity' or 'reasonableness' of a modelling approach, such as the Joshi and Rebonato stochastic instantaneous volatility model, can be assessed: the most demanding test would require a substantial congruence between the relevant marginal and conditional moments of the model and statistical distributions. If this test were rigorously applied, all the modelling approaches (including the standard LIBOR market model) based on a geometric-Brownian-motion assumption for the underlying (be it an FX or interest rate or an equity price) would probably be strongly rejected. A weaker criterion of a model's adequacy and reasonableness is to explore the Principal Component Analysis results of the real-world and model-produced data. Indeed, many of the most popular implementations of the LIBOR market model are explicitly based on a re-scaling and matching of the Brownian drivers to the eigenvalues and eigenvectors obtained from orthogonalization of the real-world correlation or covariance matrices. See, e.g., Hull (2000), Rebonato (1999).We therefore explore in this work not only whether the proposed stochastic-instantaneous-volatility approach recently introduced by Joshi and Rebonato (2001) passes the more stringent test of an overall distributional match, but also whether, and under what circumstances, it can be made to satisfy the weaker criteria of a PCA test.

More generally, we think that our findings can shed some light on the necessary ingredients for a realistic description of the stochastic evolution of the implied volatility swaption matrix. The Principal Component Analysis of real-world swaption implied volatilities presented in the following yields, in fact, a first eigenvector which decays as time to expiry increases. In comparing this quantity with the corresponding quantity obtainable from a Principal Component Analysis of the 'artificial' time series generated by a stochastic-volatility model, we notice that, in this setting, this effect can only be obtained through the use of a mean-reverting stochastic volatility. The intuitive explanation for this is clear: a change in volatility will have less effect on the value of a long-dated swaption as the volatility will have plenty of time to mean-revert back to its reversion level. We take this as indirect but convincing evidence that mean reversion of volatility is an important feature when modelling changes in stochastic instantaneous volatilities.

We have found the results, reported and discussed in Sections 2 and 3, intrinsically interesting. More generally, we also propose that the methodology suggested in this work could constitute a useful blueprint for analyzing the financial desirability of different yield-curve models which can produce stochastic future swaption matrices. It must be stressed again that, at the present moment, some extensions of the LIBOR market model that attempt to account for the current smile surface are indeed available, (see, e.g. Glasserman and Kou (2000), Glasserman and Merener (2001), Jamshidian (1999), Andersen and Andreasen (1997), Zuehlsdorff (2001)), but, as mentioned in the second footnote, either they do not produce a *stochastic* term structure of volatilities or swaption matrices, or, if they do, they only allow for a very restrictive type of stochasticity (in these CEV-related models the stochastic percentage volatility must be perfectly correlated with the underlying forward rates). Conceptually straightforward extensions of some of these approaches are however possible, which would give rise to non-deterministic implied volatility structures. (This feature could be achieved, for instance, by allowing the jump intensity in Glasserman and Kou to be a stochastic variable). We believe that, if and when these extensions are undertaken, the approach employed in the present work could provide some guidance in judging their performance and desirability.

2 The empirical analysis

2.1 Description of the data

The data set used consisted of 83,136 data points (57,156 for USD and 25,980 for DEM), corresponding to all the trading dates between 1-Jan-1998 and 1-May-2001 (866 trading days, 3 years and 4 months). In the case of USD for each trading day the following at-the-money implied volatilities were available: 3m, 6m, 9m, 1y, 2y, 3y, 4y, 5y, 7y and 10y into 1y, 2y, 3y, 5y, 7y and 10y. (A remark on notation and terminology: the a x b European swaption is the swaption expiring in a years', or months', time, as appropriate, to exercise into a b-year swap. So, the 6m x 5y, read '6 month into 5 years', swaption is the option to pay or receive fixed for five years in six month's time).

In the case of DEM/EUR the following expiries were available: 3m, 6m, 1y, 2y and 3y for exercise into 2y, 3y, 4y, 5y, 7y and 10y swaps. Less than 0.12% of the data was missing, unreliable or corrupted. For these cases, in order to preserve equal time spacings between observations, rather than eliminating the trading day, a bilinear interpolation between the neighbouring cells was carried out. Care was taken to ensure that the interpolation procedure did not alter in any significant way the final results.

The data set is particularly significant because it encompasses both the Russia crisis, and the series of rate cuts carried out by the US Fed in the first months of 2001. Indeed, non-monotonic implied volatility smiles have appeared in the swaption market after the events associated with the Russia crisis. Despite the fact that, in the present study, we do not address smile issues, and restrict our attention to at-the-money volatilities, anecdotal market evidence indicates that it was the occurrence of dramatic changes in the swaption implied volatility matrix during this period that prompted the trading community to revise the shape of quoted volatility surfaces.

With this data we then calculated the absolute[3] daily changes in implied volatilities, and organized the data with the first column containing the changes for the first expiry into the first swap length, the second column containing the changes for the first expiry into the second swap length, ..., the sixty-sixth (thirtieth) column containing the changes of the eleventh (fifth) expiry into the sixth swap length for USD (DEM/EUR). In moving across columns, the data therefore naturally presents a relatively smooth behaviour until one expiry gives place to the next.

[3] Qualitatively very similar results were obtained by conducting the analysis using percentage, rather than absolute, daily changes.

Since the data presented is considerably more complex than the more-familiar equity or FX implied volatility data, it is important to familiarize oneself with the data set up in order better to appreciate the graphs and results presented in the following. More precisely, since three-dimensional graphs can be dazzling, but are rather difficult to read in detail, we present most of the graphical information in the following sections by providing on the same graph the implied volatility corresponding to the various expiries on the x-axis, with differently-marked curves referring to different swap lengths. The series corresponding to different expiries for a fixed swap length are called the 'into' series. So, the curve corresponding to the implied volatilities of the 3m x 4y, 6m x 4y, 1y x 4y, 2y x 4y and 3y x 4y will be referred to as to 'into 4y' series.

2.2 Results

2.2.1 Qualitative Patterns of the Swaption Matrix

The analysis of the data presented above can be profitably split in the USD and the EUR/DEM case. The time series of selected at-the-money implied volatility curves are shown in Figs 1 and 2 for USD and EUR/DEM, respectively. For the purpose of future discussions it is interesting to notice the more complex structure of the USD data, where time series of different 'into' series appear 'optically' less correlated. Notice, in particular, the strong *upward* change of the 1 x 3 swaption series, associated with a large *downward* move for the 10 x 10 swaption volatility during October 1998. While large sudden moves can also be observed in the case of the DEM swaption matrix, these are directionally more correlated, and, therefore, are more likely to produce a level translation of the swaption matrix, rather than a change in overall shape. These qualitative observations will be referred to later in this section.

Continuing with the analysis of the USD data the following salient features can be noticed. First of all, over the period of observation at least two main distinct patterns can be recognized: the first, that we shall call 'normal' for reasons that will become clear in the following, displays an overall humped shape for all or most of the curves. See Figs. 3 to 5); the second, that we shall dub 'excited', shows instead a monotonically-decaying behaviour. See Figs. 6 and 7.

Within the normal pattern, further systematic shape configurations can be distinguished: in the 'ordered' case, for all expiries greater than, approximately, one year, the cross-sectional curves of the swaption matrix corresponding to the different swap lengths (i.e. the into 1y, into 2y, into 3y, into 5y, into 7y and into 10y curves) are monotonically decreasing as a function of residual swap length. See Figs. 1 and 2. For expiries shorter than approximately one year the order is exactly reversed. There exists, for these 'normal' and 'ordered' shapes, a very tight range of expiries (indeed, almost a single point) where the different 'into' curves intersect. In other cases, that we shall call 'normal and scrambled', the overall shape of the surface is still humped, but the shape of some of the 'into' series is more complex (perhaps displaying several maxima). See, e.g., Fig. 5.

As for the 'excited' (monotonically decreasing) pattern, displayed in Fig. 6, the overall surface is monotonically decreasing as a function of the expiry for each 'into' series. These series, in turn, are monotonically decreasing in level from the shortest swap length (into 1 year) to the longest (into 10 years). Furthermore, all the different 'into' series occur in decreasing level from the shortest to the longest swap length (i.e., for all expiries, options to exercise into 1-year swaps have a higher implied volatility than options to exercise into 2-year swaps; these, in turn, have a higher implied volatility than options to exercise into 3-year swaps; and so on). Finally, Figure 7 displays a 'mixed' case where both humped and monotonically decreasing cross-sectional curves coexist in the same surface.

Moving to the DEM/EUR data, a much simpler set of patterns can be observed during the same period of observation, i.e. either a monotonically decreasing pattern (see Fig. 8); or a steep or shallow hump (Figs. 9 and 10). Notice that for this currency little, if any, cross-over of the different 'into' series can be observed either for the 'normal' or the 'excited' state.

Some features of the overall structure of the swaption matrix can be understood by noticing that the series corresponding to the shortest swap length should be approximately related to the term structure of caplet volatilities (a one-period swaption is clearly a caplet). It is therefore not surprising that the 'into 1y' and 'into 2y' series for USD and DEM/EUR, respectively, should bear a close resemblance to the typical humped shapes of the market-observed term structures of volatilities (See, e.g., the discussion in Rebonato (2001)). The reason for labelling the matrices 'normal' or 'excited' is that the monotonically decreasing shape is associated with periods immediately following large movements in the underlying yield curve and in the swaption matrix itself. The post-Russia environment is a prime example of this pattern. The humped, 'normal' shape, on the other hand, prevails during periods of greater stability.

It is important to point out that similar patterns (i.e. normal, excited, scrambled or mixed patterns) repeated themselves at different points in the history of the swaption matrix analyzed in this study. Generalizing considerably, one can say that, especially in the US case, the swaption matrix appears to 'oscillate' with random periodicity between a few possible fundamental shapes. The existence of these well-defined patterns is important for modelling purposes because, if the swaption matrix were to 'diffuse' freely over time, it would explore a variety of possible shapes, would quickly lose memory of where it started, and would be very unlikely to revisit time and time again very similar patterns. Similarly, if the diffusion of the swaption matrix were driven by a mean-reverting diffusion with a constant reversion level, one would observe a temporary departure from a particular 'basis' shape, but a long-term reversion *to the same* pattern. The recurrence of these well-defined patterns therefore suggests that a simple mean-reverting diffusive behaviour cannot adequately capture these fine details of the swaption matrix dynamics.

The USD swaption data turned out to be different from the DEM/EUR data not only in the sense that it showed more complex shape patterns; perhaps even more significant is, in fact, the observation that the USD implied volatilities displayed much more rapid transitions between different shapes than what observed in the case of DEM/EUR. For USD, the transition form the normal to the excited pattern sometimes occurred over a space of two or three trading days. In the DEM/EUR case, instead, a much smoother 'diffusion' from one shape to the other was observed. This observation is somewhat difficult to quantify without displaying an enormous amount of data. An idea of this feature can however be gleaned by the following analysis. If it were indeed the case that the transition between implied volatility states took place much more rapidly for USD than for DEM/EUR, one would expect to observe a different tail behaviour for the changes in implied volatilities in the two currencies. In particular, sudden transitions between one regime and the other would be associated with very large changes in the implied volatilities, which would, in turn, give rise to fat tails. If the transitions from one regime to the other were more sudden in the case of USD, one would expect to observe fatter tails for the density of changes for this currency than for DEM/EUR. We tested this conjecture by constructing a frequency histogram for the changes in implied volatilities across all expiries and swap lengths, and by comparing its overall shape, its first four moments and its right and left tails. See Figs 11 and 12.

Prima facie, the overall shapes appear quite similar for the two currencies, and both display marked leptokurtosis. See Fig. 11, which also shows the two Gaussian densities with the same expectation and standard deviation. In particular, the average change was quite similar in the two currencies (15.7 and 18.2 basis points for USD and DEM/EUR, respectively), with a (non-annulaized) standard deviation of 37.5 (USD) and 33.2 (DEM/EUR) basis points. The skewness is also very similar for the two currencies(0.044 for USD and 0.042 for DEM/EUR), but the kurtosis differs significantly, and in the direction consistent with the speed of transition between patterns highlighted above: 0.24 for DEM/EUR and 0.41 for USD. Figs. 12 illustrates this feature by displaying the left tails of the empirical frequency densities (the corresponding Gaussian tails are not shown in the graph because they do not show on the scale of the y axis).

Besides being of intrinsic interest, these results are of relevance for the modelling of the stochastic volatility of the forward or swap rates that drive the swaption matrix. Since this latter quantity is given by suitable integrals of the instantaneous volatility, these results indicate that

- the hypothesis of a purely diffusive behaviour for the instantaneous volatility is strongly rejected[4]
- very large changes in implied volatilities are more pronounced for USD than for DEM/EUR.

As mentioned in the introductory section, the rejection of the diffusive behaviour for the implied volatility components is not surprising, and consistent

[4]A series of χ^2 test to check the hypothesis whether the sample distribution of any of the empirical implied volatility series could have been drawn from a normal distribution with matching first two moment always gave a probability of less than 10^{-13}.

with similar findings for virtually all the financial-market quantities that have been investigated. A geometric-diffusion approach has nonetheless provided a useful first approximation both in the equity and FX worlds and in interest-rate modelling. In the latter field, in particular, it has become customary to calibrate a diffusive model to the principal components contained from the statistical data.

In this restricted *modus operandi* modellers have therefore accepted the fact that the real-world distributional features of the relevant underlying quantity might be poorly described by a diffusive behaviour; but have nonetheless retained a Brownian-shock modelling description as long as the shape of the eigenvectors, and the relative magnitude of the eigenvalues, could be matched in the real and risk-adjusted worlds. The results presented in the following section therefore present PCA data not only for their intrinsic interest, but also in order to explore whether a diffusive-stochastic-volatility LIBOR market model can be justified and calibrated along these lines.

2.2.2 Correlation and Principal Component Analysis results

Using the data described in Section 2.1, we constructed the correlation matrix between the (absolute) changes in the implied volatilities. The columns and rows in the matrix were organized as described in Section 2.1. Fig. 13 displays a small portion of the USD correlation matrix in order to highlight some of the main features of the whole matrix (the same considerations apply to the DEM/EUR matrix). The 36 elements in the bordered box in the top left-hand corner of the matrix contain the correlation between changes in implied volatilities of the 1y x 1y, 1y x 2y, ..., 1y x 10y swaptions. Similarly, the 36 elements in the bordered box in the top right-hand corner of the matrix contain the correlation between changes in implied volatilities for the 2y x 1y, 2y x 2y, ..., 2y x 10y swaptions; etc. There are several interesting features worthwhile noticing:

- for a given expiry (i.e. for a given 'into' series), the correlation tends to be a convex function of the swap lengths; this feature remains true for all expiries, and in both currencies;
- for a given swap length, the correlation displays less convexity as a function of swaption expiry;
- given the way the matrix has been organized, there are discontinuities with a periodicity of 6 both as one moves across columns (from one 'into' series to the next) and as one moves down rows (from one swap length to the next);
- the correlation between changes in implied volatilities of very 'distant' swaptions (i.e. swaptions with greatly different expiries and swap lengths) is very low (approximately 20% for USD and 15% for DEM/EUR).

With these preliminary considerations in mind, one can better understand the shape of the overall correlation matrices for the two currencies, reported below for the case of the DEM/EUR currency (the qualitative shape of the USD matrix is the same).

One can easily recognize the jagged structural feature of the matrix due to the transition from one expiry series to the next, or from one swap length series to the next. It is important to point out that this jagged behaviour, clearly displayed in Fig. 14, is not due to noise, but to the way the two-dimensional data must be organized along a one-dimensional axis. The individual correlation curves inside each 6 x 6 box are indeed remarkably smooth, a shown in Fig. 13 for the small sample of the USD correlation matrix displayed, indicating that *prima facie* statstical noise should not be a concern as far as the interpretation of the data is concerned.

With these correlation matrices one is in a position to obtain, by orthogonalization of the correlation matrix, the associated eigenvectors and eigenvalues. The results for USD and DEM/EUR are shown in Figs 15 and 16 below.

The first noteworthy feature is the strong qualitative similarity of the results for the two currencies. The shape of the eigenvectors also lends itself to an interesting interpretation: the first principal component, as usual, displays a virtually identical loading across the various swaption implied volatilities, and therefore describes the typical up and down rigid shift of the swaption matrix[5]. This first mode of deformation, however, only accounts for less than 60% both in USD and DEM/EUR. The second mode of deformation can be interpreted as corresponding to the first three series moving up (down) and the last three series moving down (up). Finally, the third eigenvector mainly picks up movements *within* each series, with, say, the implied volatility of swaptions with short swap length moving up and the volatility of swaptions with long swap length moving down.

Finally, it is also worthwhile pointing out the reamrkable similarity of the explanatory power of an increasing number of eigenvectors, are shown in Fig. 17: for both currencies, it takes approximately 10 eigenvectors to explain 90% of the variability across series of expiries and swap lengths. This result should be contrasted with the findings of PCA on yields or forward rates, where, typically, 90% of the variability is explained by four or fewer eigenvectors (see, e.g. Priaulet et al. for a recent survey of results). It must be stressed that the data sample included two particularly 'excited' periods, i.e. both the Russia crisis and its aftermath and the aggressive easing by the FED in early 2001. It would be interesting to see whether data in more 'normal' periods would display the same features.

Similar, but, at the same time, noticeably different, results were obtained by orthogonalizing the covariance rather than the correlation matrix. The most noteworthy difference is that the loadings onto the various implied volatilities now display a noticeable decaying behaviour as one moves across columns, indicating that the top left-hand corner of the swaption matrix is more volatile than the bottom right-hand corner. This observation, *per se* not surprising,

[5] The first eigenvector is virtually flat when the correlation, as opposed to covariance, matrix is orthogonalized. Therefore the most important mode of deformation is a parallel shift only after scaling by the volatilities. See the discussion later in the section.

will become very relevant in the discussion of the role of mean reversion in the stochastic-volatility model analyzed in Section 3. Without preempting future results, it is worthwhile pointing out that these findings will be shown to be compatible with a mean-reverting process for the instantaneous volatility. Figs. 18 and 19 show the first three eigenvectors obtainable from orthogonalizing the covariance (as opposed to the correlation) matrix. Finally, the explanatory power of an increasing number of eigenvectors was found to be quite similar to what found by orthogonalizing the correlation matrix.

3 Stochastic Extensions of the LIBOR Market Model: Implications for the Theoretical PCs

3.1 The Model Used

Joshi and Rebonato (2001) have recently presented an extension of the LIBOR market model by introducing displaced diffusion and by making the instantaneous volatility of the forward rates stochastic. They start from a standard deterministic-volatility LIBOR market model along the lines of Brace Gatarek and Musiela (1995) Musiela and Rutkowski (1997), Jamshidian (1997), which is fully specified once an arbitrary correlation function and a set of instantaneous volatilities for the state variables are given. More precisely, their chosen setting is based on the evolution of discrete-tenor forward rates, and, following Jamshidian (1997), this set-up is used for the valuation of LIBOR payoffs which can be expressed as homogeneous functions of degree one in the bond prices[6] and that satisfy the condition of being measurable with respect to the natural filtration generated by the driving Brownian processes. Such an approach neither requires the instantaneous short rate, nor the instantaneously-compounded money-market account.

Joshi and Rebonato's (2001) start from this deterministic setting, and posit

$$\begin{aligned} \sigma(t,T) &= k_T g(T-t) = && (1) \\ &= k_T([a+b(T-t)]\exp[-c(T-t)]+d) && (2) \end{aligned}$$

where $\sigma(t,T)$ is the instantaneous volatility at time t of the T-maturity 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 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 deterministic-volatility setting by imposing that

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

[6] or, equivalently, as a function homogeneous of degree zero in the forward rates, times a (linear combinaton of) bonds. In practice, this condition hardly limits the scope of typical pricing applications.

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 dsicussion, Rebonato (2001) 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 associate with the chosen numeraire is given by

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

which integrates to

$$f_{T_i}(t) = f_{T_i}(0)\exp[\int_0^t(\mu^Q(\{f_{T_j}(u)\},u) - \frac{1}{2}\sigma(u,T_i)^2)du + \int_0^t \sigma(u,T_i)\sum_{k=1,m} b_{ik}dz_k^Q(u)] \tag{5}$$

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 drift that reflects 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.

Joshi and Rebonato (2001) enrich this standard formulation in two ways:

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

$$\frac{d(f_{T_i}(t)+\alpha)}{f_{T_i}(t)+\alpha} = \mu_\alpha^Q(\{f_{T_j}(t)\},t)dt + \sigma_\alpha(t,T_i)\sum_{k=1,m} b_{ik}dz_k^Q(t) \tag{6}$$

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:

$$da_t = RS_a(RL_a - a_t)dt + \sigma_a(t)dz_t^a \tag{7}$$

$$db_t = RS_b(RL_b - b_t)dt + \sigma_b(t)dz_t^b \tag{8}$$

$$d\ln[c_t] = RS_c(RL_c - \ln[c_t])dt + \sigma_c(t)dz_t^c \tag{9}$$

$$d\ln[d_t] = RS_d(RL_d - \ln[d_t])dt + \sigma_d(t)dz_t^d \tag{10}$$

and all the Brownian increments uncorrelated with each other and with all the Brownian increments $dz_k^Q(t)$. In Equations (7) to (10) 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 coeffcients, or of their logarithms.[7]

The introduction of the displacement coefficient α (see Rubinstein (1983) for a discussion of displaced diffusions 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[8]. 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.

Since the smile dynamics are not the focus of this paper, we do not pursue this angle any further, and refer the interested reader to Joshi and Rebonato (2001), but make two observation of relevance for the future discussion: first, the stochastic evolution (under Q) of the swaption matrix is fully specified once the dynamics for the instantaneous volatility is given; second, as mentioned in the introductory section, the model-implied evolution of the swaption matrix takes place in the pricing measure, and, to the extent that investors display aversion to volatility risk, cannot be immediately related to the observable evolution of the same quantity in the real world. As it is well known, in fact, (see, e.g., Lewis (2000)), in a stochastic volatility setting the Girsanov's translation from the real-world to the risk-adjusted measure Q brings about a change in the drift of the instantaneous volatility. This change of measure, in turn, makes a naive comparison between the empirically observed and the model predicted changes in the swaption matrix impossible. It is worthwhile commenting briefly on this point: when moving between equivalent measures, Girsanov's theorem allows us to change the drifts of all Brownian motions. The requirement of no arbitrage, coupled with the fact that a forward rate times an appropriate bond becomes a tradable asset (see e.g. Jamshidian (1997) or Rebonato (2001)) forces their drifts in the martingale measure to take a unique value. However, volatility is not tradable, nor is it possible to find a simple transformation (such as the multiplication by appropriate discount bonds in the case of forward rates) that can turn it into a traded asset. It is therefore not possible to impose constraints which are strong enough to single out a unique drift of the volatility parameters in the martingale measure. As a consequence, an infinity of volatility drifts (each associated with a different equivalent pricing measure) are compatible

[7] For simplicity, and in order to enhance the time-homogeneity of the evoluation of the term structure of vlatilities and of the swaption matrix, the reversion level has always been set, in the present study, equal to the current level of the corresponding quantity. Joshi and Rebonato (2001) follow a similar procedure in their calibration to the smile surface.

[8] It is customary to model this feature by means of a CEV approach (see, e.g. Andersen and Andreasen (2000), or Zuehlsdorff (2001)). Joshi and Rebonato (2001) quote the result by Marris (1999), who 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 (1999) also provides a thoeretical justification as to why this should be the case. Joshi and Rebonato therefore use the displaced-diffusion setting, which allows simple closed-form solutions, as a computationally simple and efficient subsititute for the theoretically more pleasing CEV framework (which does not allow negative forward rates).

with absence of arbitrage. This has been one of the reasons for performing the (measure-independent) Principal-Component-Analysis study presented in this work. The PCA results, in fact, are measure-independent because they are obtained from the orthogonalization of implied volatilities which are the inverse of suitable (Black) functions[9] of the instantaneous correlation and covariance matrices. These inverse of the Black function are, in turn, linked to the instantaneous volatilities through a complex chain of Ito's lemmas. However complex this chain might be, given the assumption about the dynamics of the model instantaneous volatilies, the implied volatilies are themselves a diffusion. The volatility terms of these diffusions, which we do not calculate explictly but estimate numerically, only depend on the volatilities of, and the correlation among, the coefficients of $\sigma(t,T)$. Since we have assumed the volatilities of the coefficients, $\sigma_a(t), \sigma_b(t), \sigma_c(t)$ and $\sigma_d(t)$, to be deterministic (and their associated Brownian motions uncorrelated), they will remain unchanged under the Girsanov transformation between measures (see, e.g., Duffie (1996)). It immediately follows that the volatilities of and correlations amongst the implied volatilities will also be invariant, and hence so will their eigenvectors and eigenvalues.

3.2 The Methodology to Extract the Model-Implied Principal Components

We examine in this section the dynamics of the swaption volatility matrix when swaptions are priced using the stochastic volatility model described above. In particular, we try to answer the following questions:

1. How many principal components are required in order to account for the stochastic evolution of the swaption matrix to a given percentage of explanatory power (fraction of variance explained)?

2. How does the first principal component (eigenvector) obtained from the orthogonalization of the model covariance matrix compare with the corresponding empirically observed quantity?

3. Does this first principal component change its qualitative behaviour when the volatilities are mean reverting?

4. Has the displacement coefficient a significant effect on the overall qualitative behaviour of the first principal component?

The algorithm implemented to carry out this investigation was the following: we evolved the coefficients a, b, c and d over a time step of a week using the

[9] Note carefully that the price functions that transform the instantaneous volatilities into the prices of the European swaptions, and, utlimately, into their implied volatilities, *do* depend on the mean-reversion and, therefore, on the measure-dependent drift terms. This, however, does not invalidate in any way our results.

dynamics described in Section 3.1. Since we have chosen a Ornstein-Uhlenbeck process for the coefficients (or their logarithms) the evolution of the instantaneous volatilities can be accomplished exactly and analytically. Given this joint realization of the instantaneous volatilities we evolved the forward rates over the same time step using Equation (6)[10].

After performing this step, we 'reset' the curve so that the expiries of the various forward rates remained a constant time distance away from the new spot time. Given this state of the world we then priced all the swaptions in the different series and for the different expiries, and translated their prices into implied volatilities. By repeating this exercise many times we generated an artificial time series for the swaption implied volatilities. The changes in this time series over 1,000 weeks were finally used to generate a covariance matrix, that was then diagonalized. Despite the fact that, for numerical reasons, the time step was taken to be one week (as opposed to the daily spacing between the real-world data), the changes are fully dominated by the stochastic term: for typical values of the volatility and of the forward rates, and for the chosen time step, the volatility terms are typically 200 to 500 times larger than the drift term[11], and therefore the numerical procedure presented above accurately estimates the *instantaneous* covariance matrix.

The parameters used for this exercise are given below:

Coefficient	Initial Value	Volatility	Rev. Speed	Rev. Level
a	-0.02	10.00%	0.5	-0.02
b	0.108	10.00%	0.3	0.108
$\ln(c)$	0.8	10.00%	0.5	0.8
$ln(d)$	0.114	10.00%	0.4	.114

Tab. I: The parameters of the instantaneous volatility process used in the simulation.

In addition to the instantaneous volatility function, a functional shape for the correlation had to be chosen in order to describe the covariance matrix between the forward rates. This was taken to be given by

$$\rho_{ij} = \exp[-\beta|T_i - T_j|] \tag{11}$$

where ρ_{ij} is the correlation between the forward rate expiring at time T_i and the forward rate expiring at time T_j. This correlation function is not particularly

[10] Strictly speaking, the evolution of the forward rates over a finite step is not exact because, for a chosen numeraire, not all the forward rates can be simultaneously log-normally distributed, and the drift of all but one of the forward rates is stochastic. Their distribution is therefore not exactly log-normal, and a closed-form solution for the SDE does not exist. The size of the step is such, however, that the loss of precision introduced by this approximation is absolutely negligible. (See, e.g., Hunter, Jaeckel and Joshi (2001) for a detailed analysis of the magnitude of the effect).

[11] For this order-of-magnitude calculation the forward rate was taken to be at 6.00% and the instanatneous volatility at 15.00%. A constant correlation of 90% among all the forward rates was also assumed.

sophisticated, and can be criticized on financial grounds. However, we found in our study that the detailed shape of the correlation has a small impact on swaption prices, and we concur with De Jong et al. (1999) who state that *although swaption prices do depend on the correlation between interest rates of different maturities, this turns out to be a second order effect; swaption prices are primarily determined by the volatilities of interest rates.* Rebonato (2001) and Mercurio and Brigo (2001) provide additional evidence that, and an explanation as to why, the detailed shape of the correlation function has a small impact on the price of European swaptions.

The parameters shown in Tab. I, to which the results in this study refer, were typical of the ones found in Joshi and Rebonato (2001) by optimizing to the caplet smile surface of EUR and USD using the market data prevailing as of 8 September 2000. They were not chosen to produce a best fit to either swaption matrix, but provide a qualitatively acceptable description of the swaption matrix. The qualitative features reported below, and the trends in the eigenvectors, were found to be strongly insensitive to the details of the parametrization, and, as discussed at greater length below, mainly to depend on whether the reversion speed was assumed to be zero, or equal to a finite 'reasonable' value (where 'reasonable' means similar to the values found in the fitting to actual market data for different currencies).

Finally, the displacement coefficient α was chosen to be equal to 0.02, in close agreement with the best-fit values obtained for both currencies.

3.3 Results

We begin the discussion of the results by presenting the graphs of the first principal component obtained using the methodology described above. The x-axis displays the following swaptions, in the order

1 x 1,2,3,5,7,10
2 x 1,2,3,5,7,10
3 x 1,2,3,5,7,10
5 x 1,2,3,5,7,10
7 x 1,2,3,5,7,10
10 x 1,2,3,5.

The first observation from these results, which answers the first of the questions posed in Section 3.2 above, is that, in all cases (i.e. with and without mean-reversion and with and without a displacement coefficient), a very large fraction (approximately 90% for both currencies) of the variability across expiries and swap lengths is explained by the first principal component. This was observed despite the fact that four independent Brownian processes affect the instantaneous volatility process. The finding is not surprising, because we noticed in Joshi and Rebonato (2001) that the prices of European swaptions are mainly affected by the stochastic behaviour of the d coefficient.[12] On the other hand, the fraction of the total variability (percentage of variance) explained by the first principal component in the two currencies examined was empirically found to be slightly less than 60% (see Fig. 17). Therefore the first conclusions is that a significant degree of the real-world complexity is not accounted for by the model proposed. This is *a priori* not surprising: stochastic-volatility extensions of the LIBOR market model are after all just being introduced, and one can draw a parallel with the early simple one-factor models (such as Vasicek, CIR) that attempted to describe the evolution of the yield curve. The more relevant question is perhaps whether this initial step, albeit incomplete, is in the right direction.

In order to answer this question one can refer again to Figs. 18 and 19, which display the first principal component as estimated from the empirical data. Two features are noteworthy: first of all the clear periodicity in the values, which corresponds to the transition from one 'into' series to the next. The second important feature is that, when one uses correlation data, the average magnitude of the first eigenvector is roughly constant as one moves from one series to the next: in other words the first mode of deformation is roughly parallel across the swaption matrix. When the covariance matrix is orthogonalized, however, the first eigenvector displays a noticeable decay as one moves down the series.

The model-obtained eigenvector (obtained from the orthogonalization of the *covariance* matrix) clearly displays the same periodicity as the empirical eigenvectors. See Figs. 20, 21 and 22. When no reversion speed and zero displacement coefficient are used, the first eigenvector obtained form the model remains roughly constant as one moves down the x axis. To a limited extent, a non-zero, positive displacement coefficient, which indicates a deviation from log-normality towards normality, produces some degree of decay in the first eigenvector. It is when the reversion speed has a value similar to what obtained in the fit to the smile surface, however, that the model data displays a decay similar to what can be observed from real data. This result is particularly noteworthy, because the reversion speed used in the present study was independently determined by fitting to smile data, and without any prior knowledge of the shape of the empirical eigenvalues. Therefore, it appears that the empirical data are consistent with a mean-reverting behaviour for the instantaneous volatility (in the risk-adjusted world!) as posited by Joshi and Rebonato (2001) and as independently calibrated to static cross-sectional data (a single-day volatility surface)[13].

[12] This observation should not be taken to imply that the stochastic behaviour of the a, b and c parameters has 'no use': b and c control the location of the 'hump' in the instantaneous volatility curve, which in turn controls the degree of terminal de-correlation between forward rates. Several complex products are sensitive to different extents to this terminal decorrelation, and it would therefore be unwarranted to extrapolate conclusions from plain-vanilla to complex instruments.

[13] In their study, Joshi and Rebonato (2001) carried out optimizations of the volatility coefficients to market smile data for several different trading days. While obviously not identical, the resulting parameters displayed marked stability and robustness.

4 Discussions, Conclusions and Suggestions for Future Work

We have presented in this work some empirical work about the changes in market implied volatility swaption matrices. We have coupled this investigation with a theoretical analysis of a particular stochastic-volatility extension of the LIBOR market model. Our findings have conveyed a mixed but overall encouraging picture of the adequacy of this modelling approach.

To begin with, certain important features of the real data are not captured by the Joshi-Rebonato approach: 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. 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 which have been proposed[14]. 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. The computational and calibration problems of such a model are likely to be quite a challenge, but this modelling approach could constitute an interesting avenue for future research.

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, the proportion of the variability across the swaption matrix explained by the first principal component is significantly smaller in reality (about 60%) than with the model-produced data.

Despite these shortcomings, the modelling approach by Joshi and Rebonato has been shown to display two important encouraging features: first of all, the qualitative shape of the first eigenvector turned out to bear a close resemblance with the corresponding empirical quantity, and, in particular, the same periodicity was observed in the real and model data; second, the decaying behaviour of the first principal component as a function of increasing expiry, observed in the real data when the covariance matrix is orthogonalized, was found to be naturally recoverable and explainable by the mean-reverting behaviour for the instantaneous volatility. This feature in turn constitutes the most salient

[14] 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 *shape* of the swaption matrix.

characteristic of the stochastic-volatility extension of the LIBOR market model that we have analyzed. Furthemore, 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 therefore appears fair to say that, despite the obvious shortcomings, the modelling approach analyzed in the theoretical part of the present study appears to be a useful first step in the right direction.

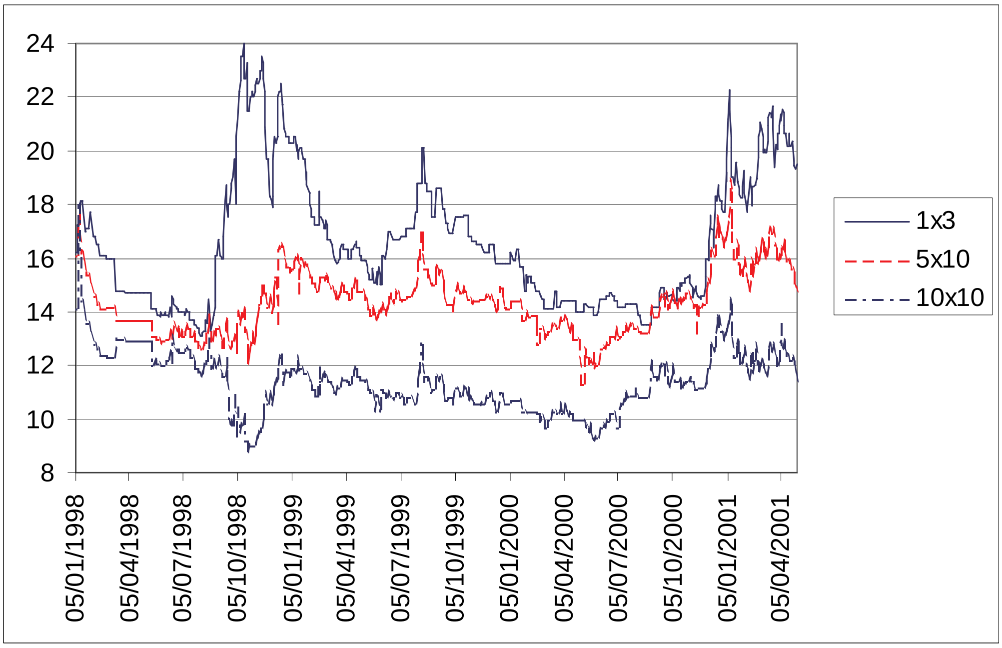

Figure 1: Selected swaption implied volatility time series for USD: 1x3 (top continuous line), 5 x 10 (middle broken line) and 10 x 10 (bottom dash-and-dot line).

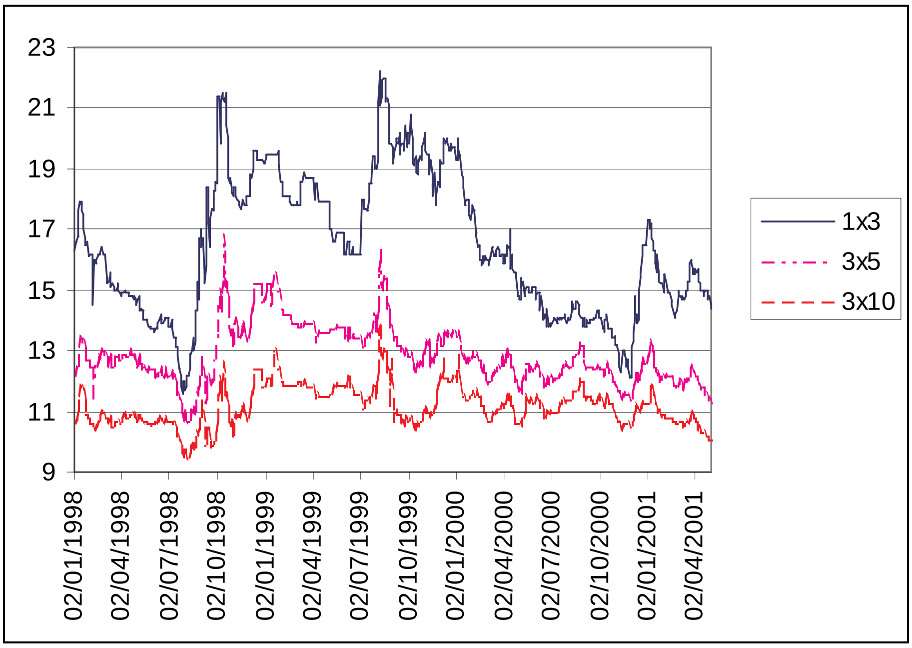

Figure 2: Selected time series of implied volatilites for DEM: 1 x 3 (top continuous line), 3 x 5 (middle dash-and-dot line) and 3 x 10 (bottom broken line)

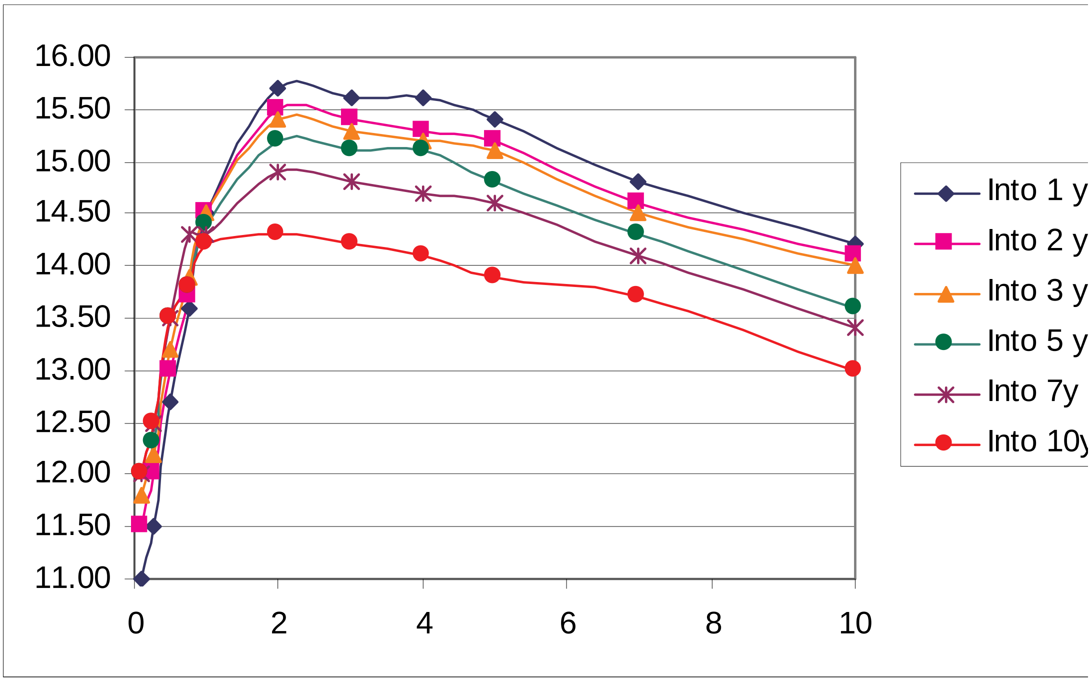

Figure 3: The 'normal' pattern for USD

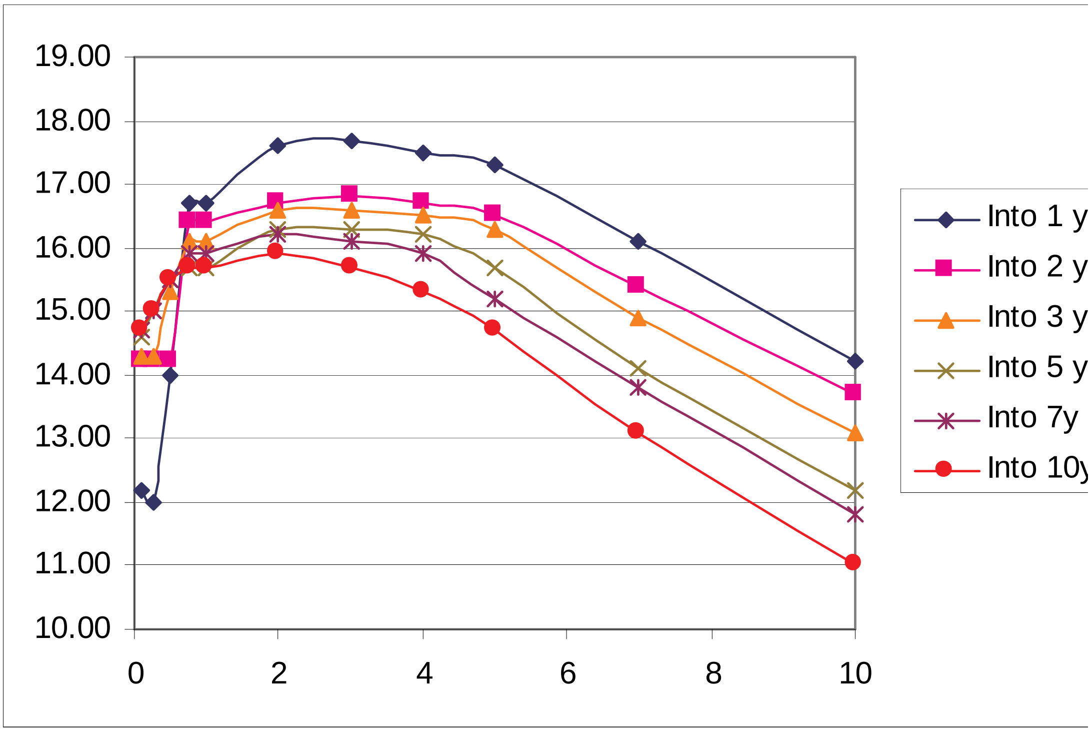

Figure 4: Another 'normal' pattern for USD with amore pronounced occurrence of th ecross-over of the different 'into' series.

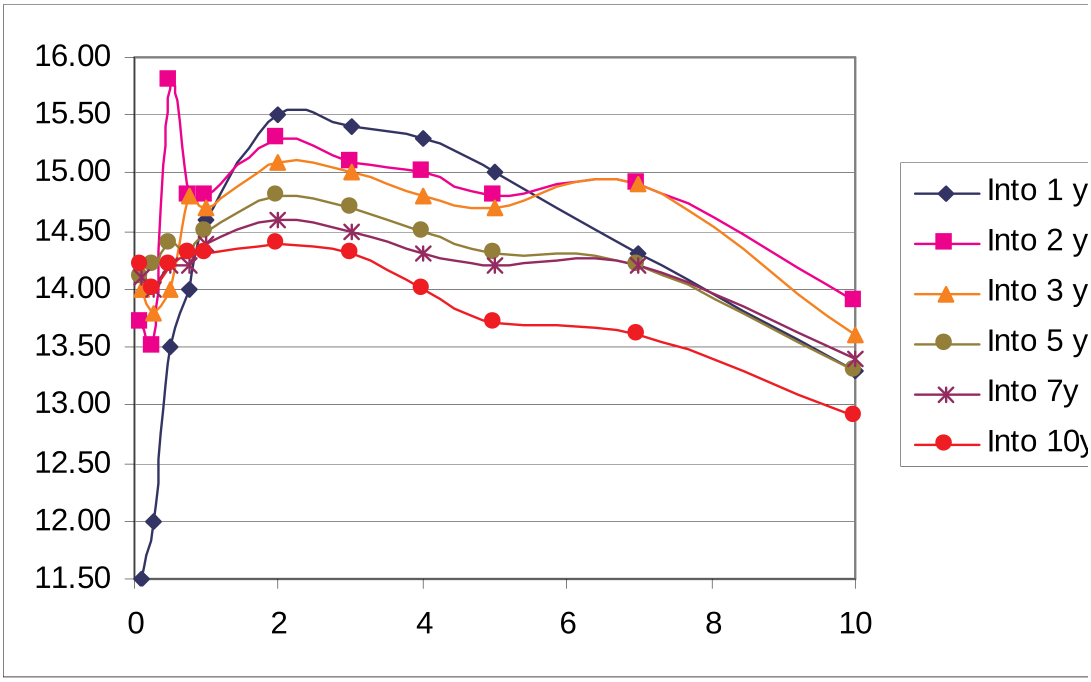

Figure 5: A more 'scrambled' pattern for the normal swaption matrix shape (USD)

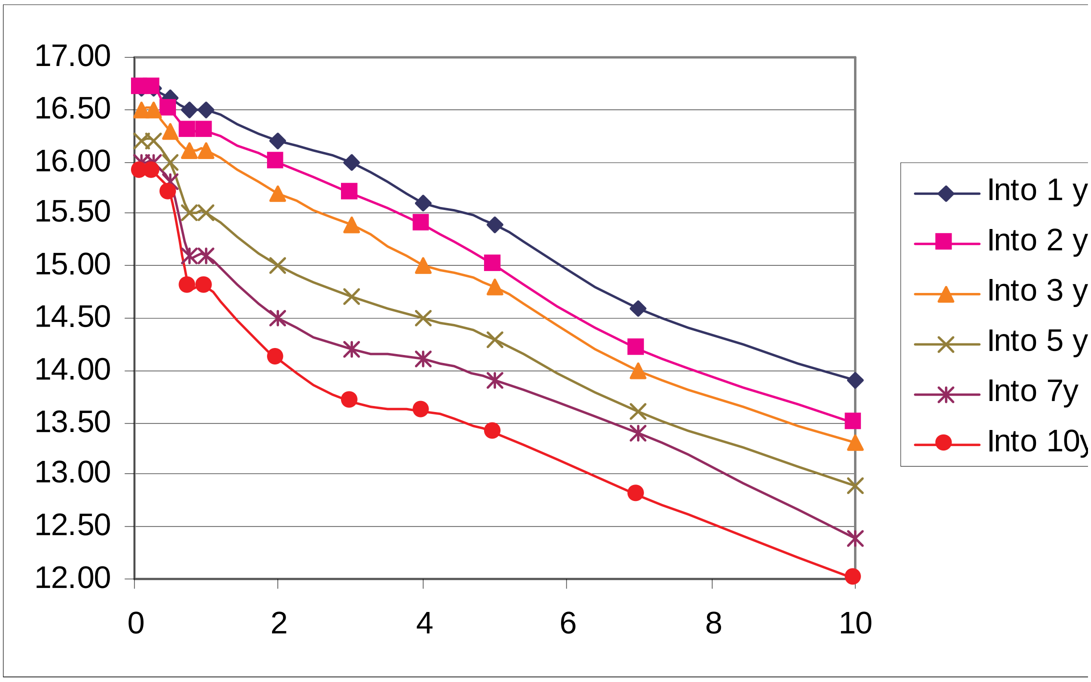

Figure 6: An 'excited' pattern for the USD swaption matrix

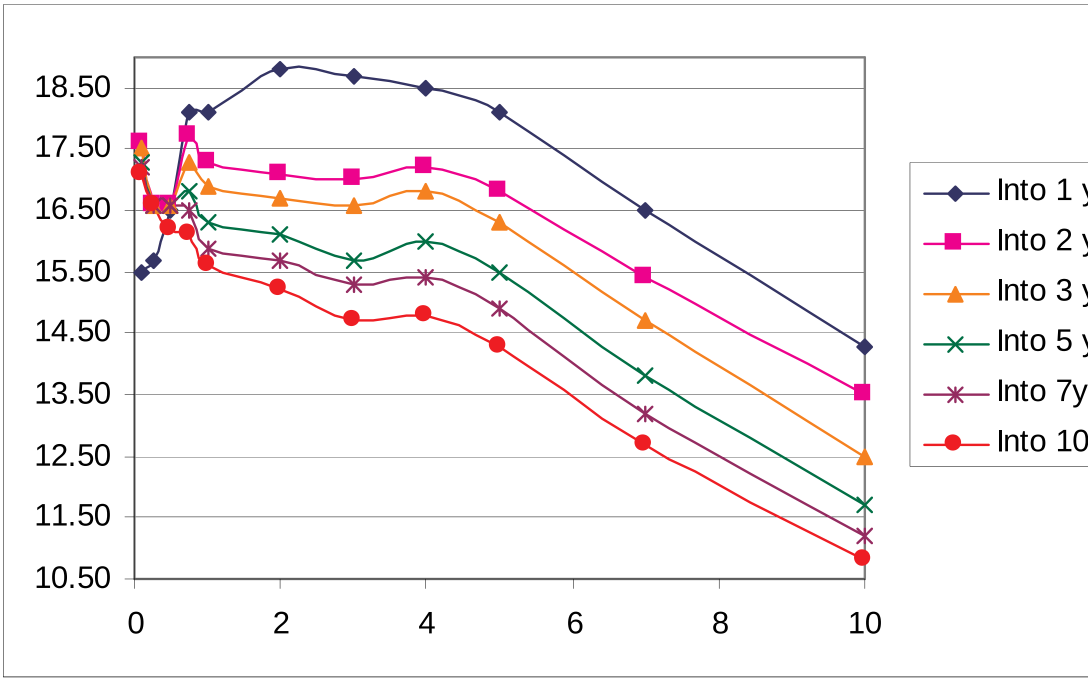

Figure 7: A 'mixed' case for the USD swaption matrix

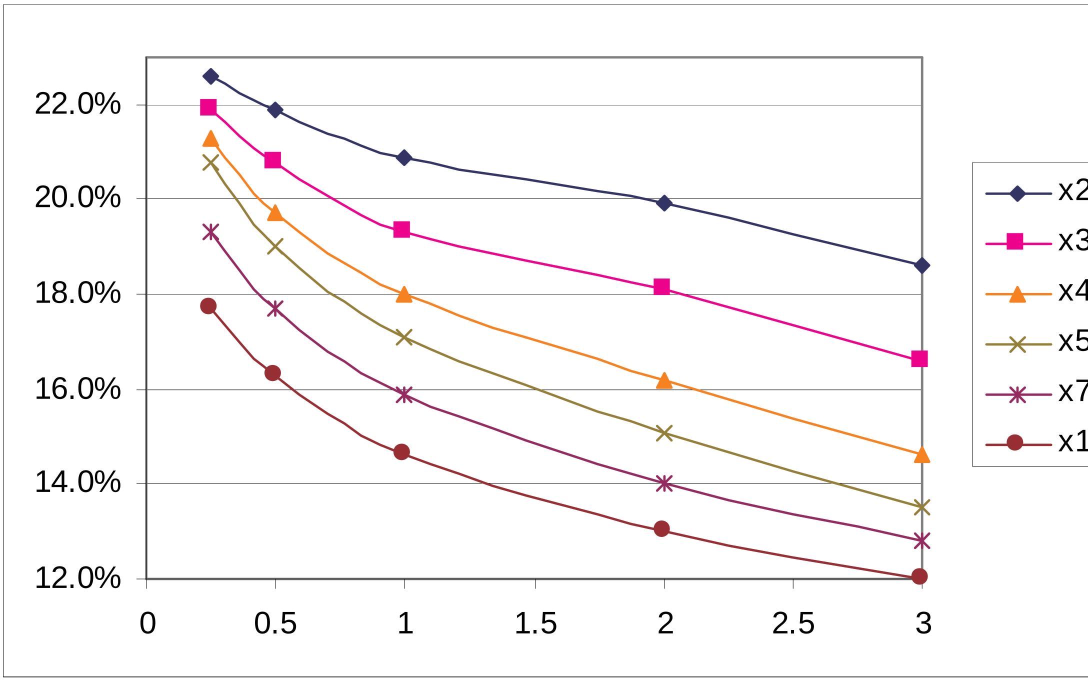

Figure 8: An 'excited' state for the DEM/EUR swaption matrix

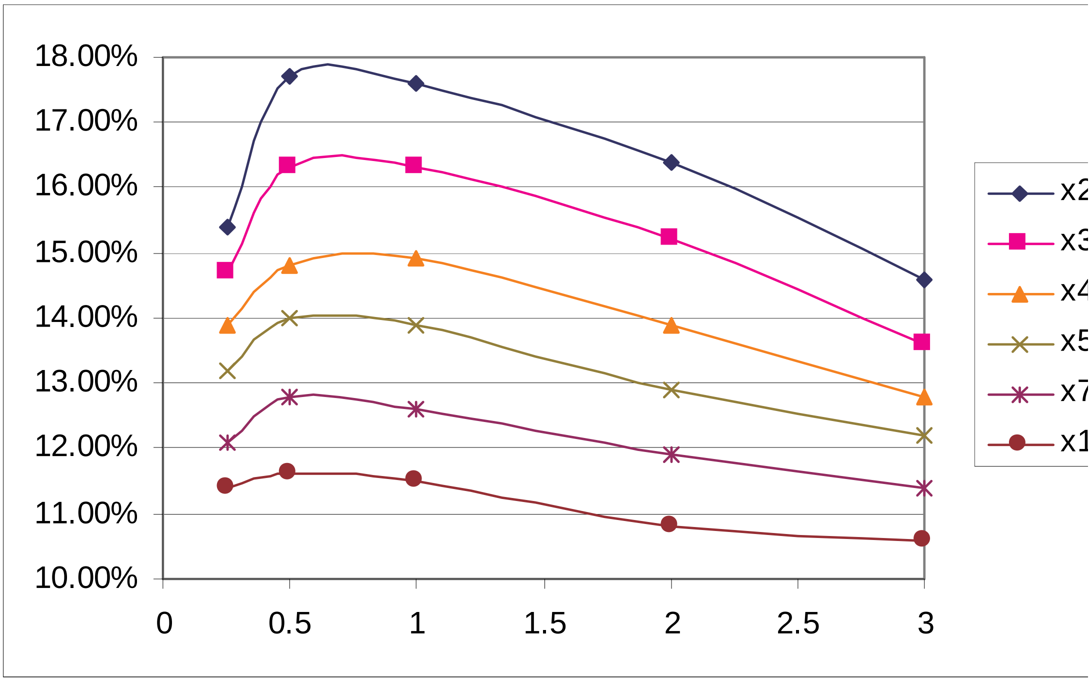

Figure 9: A 'normal' state with a sharp hump for DEM/EUR

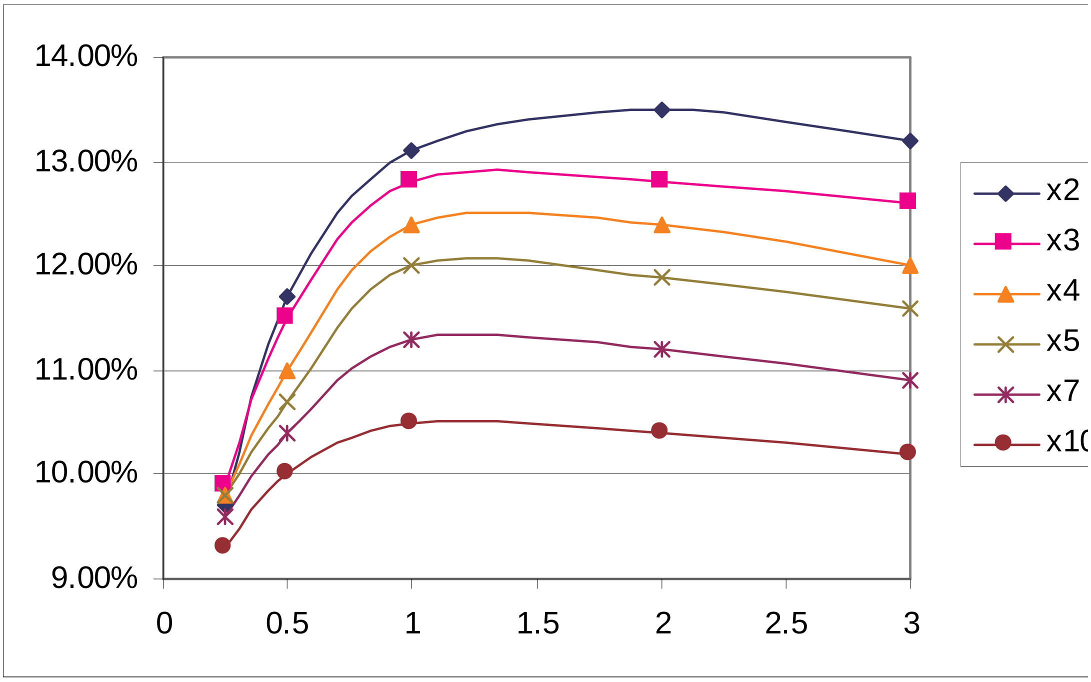

Figure 10: A 'normal' state with shallow hump for DEM/EUR

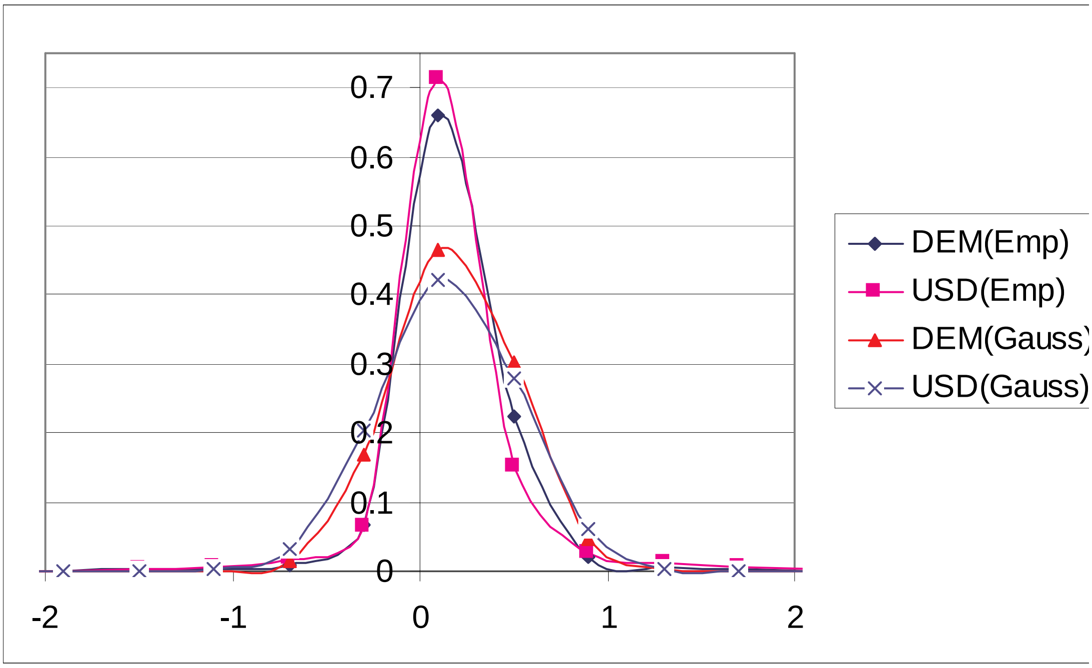

Figure 11: The empirical and fitted Gaussian densities of changes in the implied volatility swaption matrix for DEM/EUR and USD

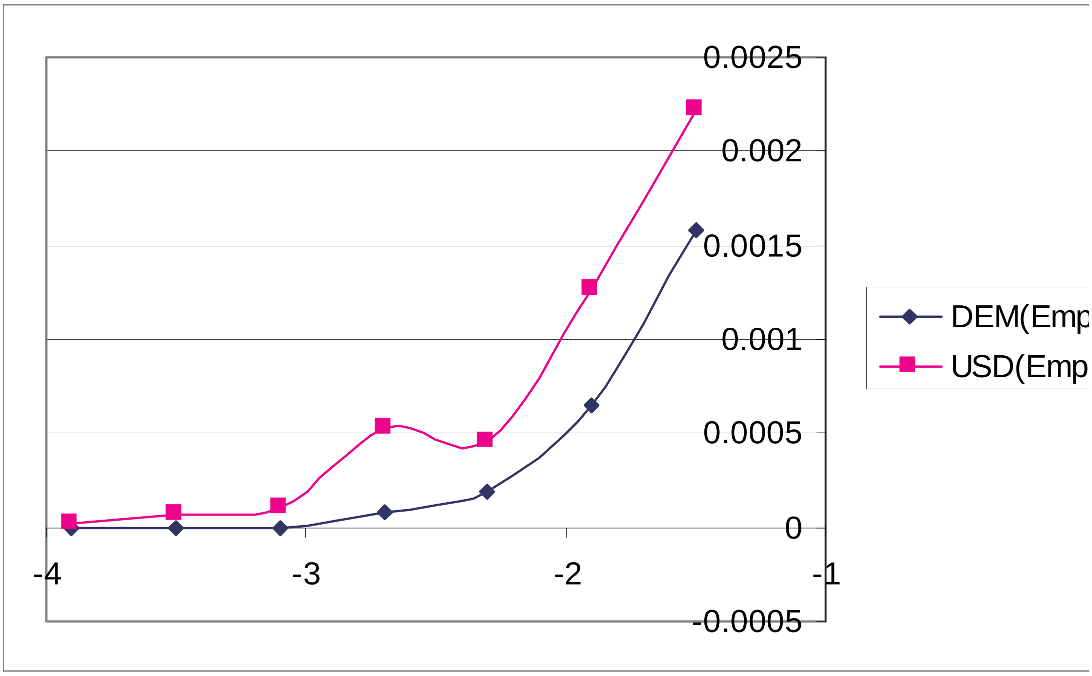

Figure 12: The left tails of the emprical frequency desnity of implied volatilit ychanges for DEM/EUR and USD. the corresponding Gaussian tails do not appear on the scale of the graph.

	1 year into						2 year into		
	1	2	3	5	7	10	1	2	3
1	1	0.7573313	0.70383312	0.70272808	0.6829902	0.66714772	0.84221759	0.72208199	0.67399784
2	0.7573313	1	0.92506863	0.87005919	0.83106392	0.79883191	0.68091232	0.88853052	0.81245114
3	0.70383312	0.92506863	1	0.90081514	0.86441631	0.84816343	0.67311099	0.88370878	0.86123263
5	0.70272808	0.87005919	0.90081514	1	0.90266159	0.87633016	0.65330642	0.87442905	0.85059915
7	0.6829902	0.83106392	0.86441631	0.90266159	1	0.89152726	0.63271429	0.84005688	0.83775257
10	0.66714772	0.79883191	0.84816343	0.87633016	0.89152726	1	0.62925843	0.83917766	0.85590297
1	0.84221759	0.68091232	0.67311099	0.65330642	0.63271429	0.62925843	1	0.72983905	0.67236483
2	0.72208199	0.88853052	0.88370878	0.87442905	0.84005688	0.83917766	0.72983905	1	0.9065737
3	0.67399784	0.81245114	0.86123263	0.85059915	0.83775257	0.85590297	0.67236483	0.9065737	1
5	0.66102105	0.83970274	0.83457423	0.87059137	0.85980908	0.84799235	0.66607135	0.89238509	0.89298015
7	0.6565227	0.7961068	0.83785346	0.83632581	0.86171228	0.87911282	0.64950144	0.86824658	0.90110557
10	0.61186077	0.71600212	0.77343124	0.81131332	0.84958077	0.88349286	0.58603538	0.80100429	0.85150274

Figure 13: A portion (top lefdt-hand corner) of the DEM/EUR correlation matrix

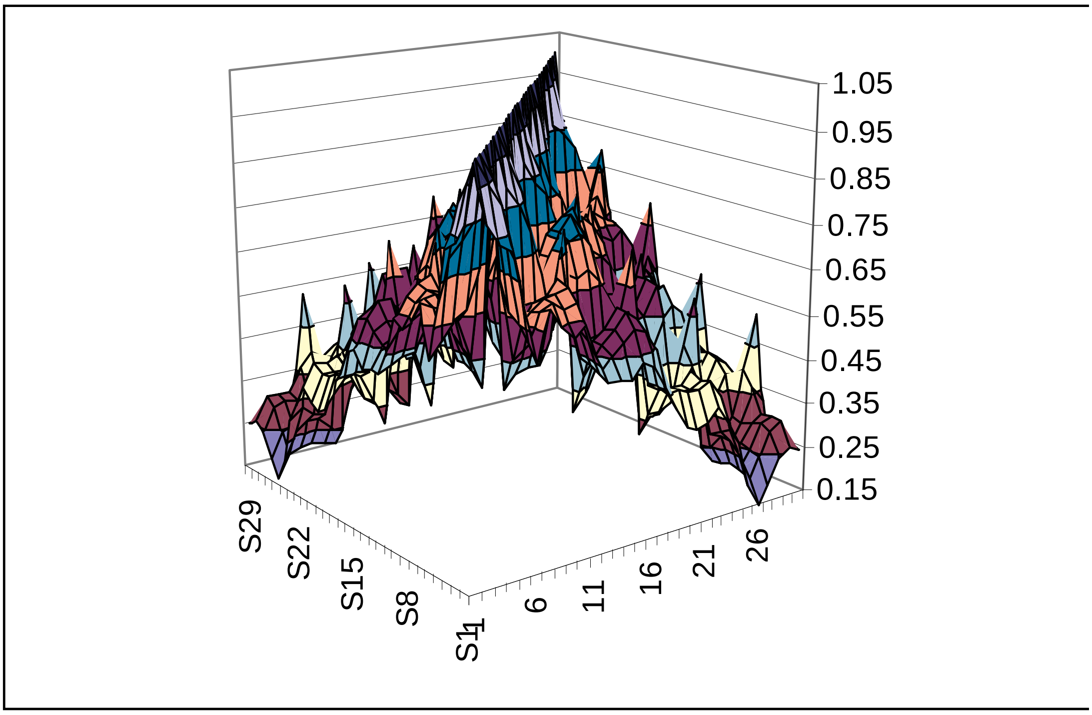

Figure 14: The correlation matrix for the changes in implied volatility in the case of DEM/EUR

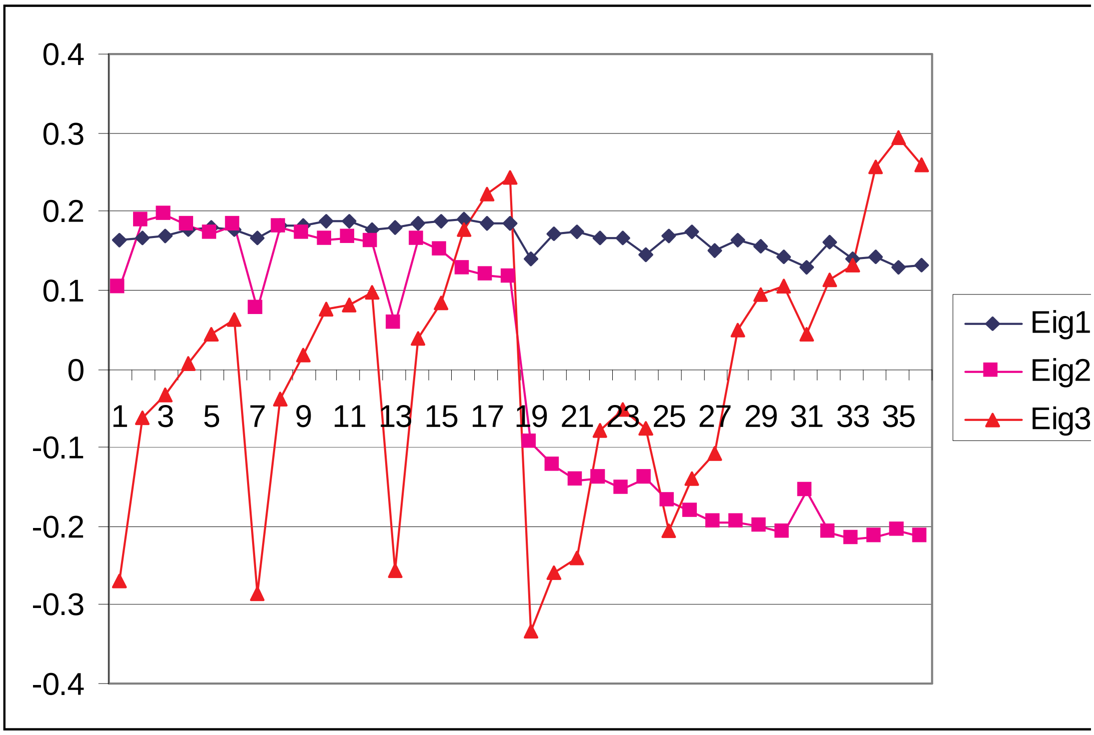

Figure 15: The first three eigenvecotors from the orthogonalization of the USD correlation matrix

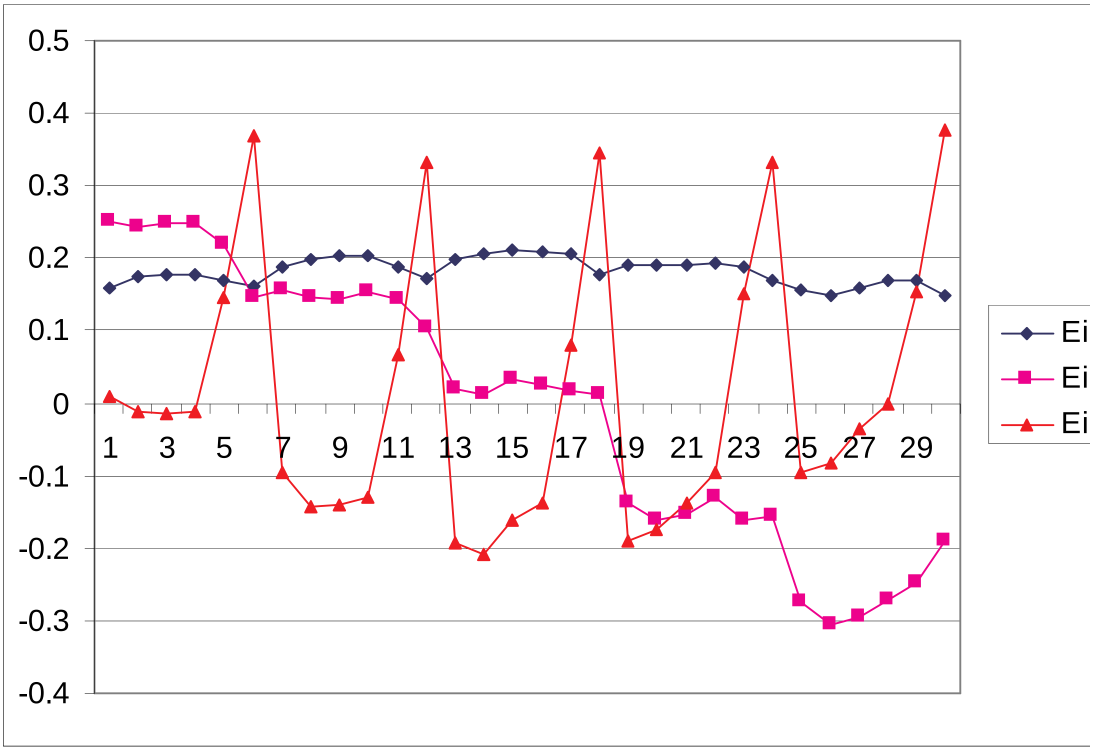

Figure 16: The first three eigenvectors from the orthogonalization of the FDEM/EUR correlation matrix

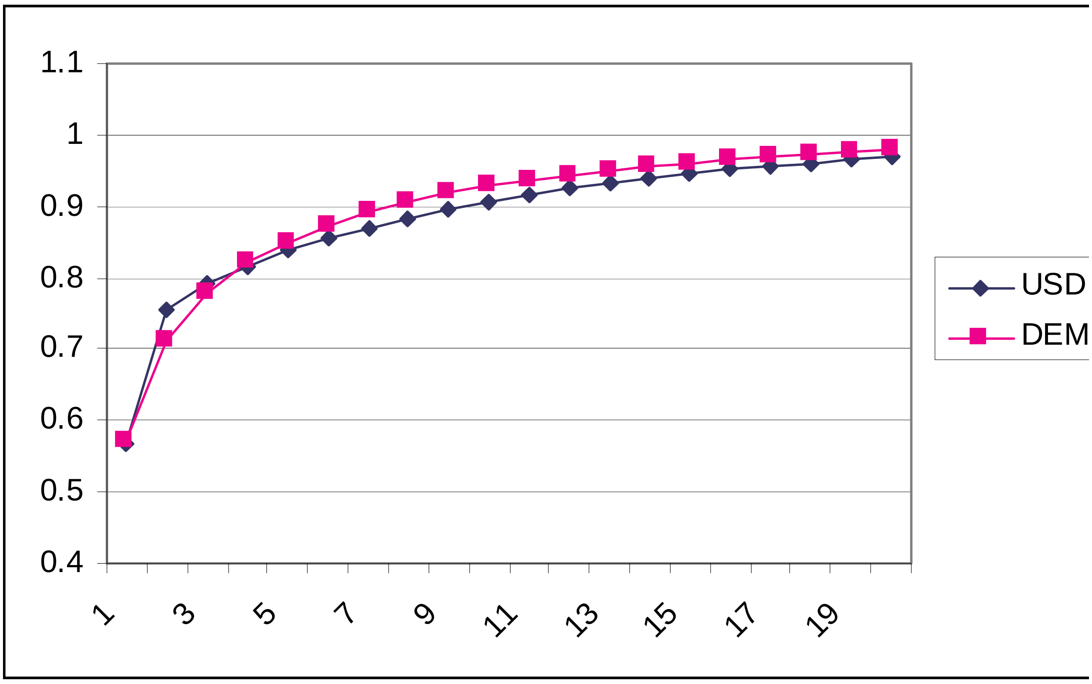

Figure 17: The proportion of the total variability explained by an increasing numberof eigenvectors (on the x axis)

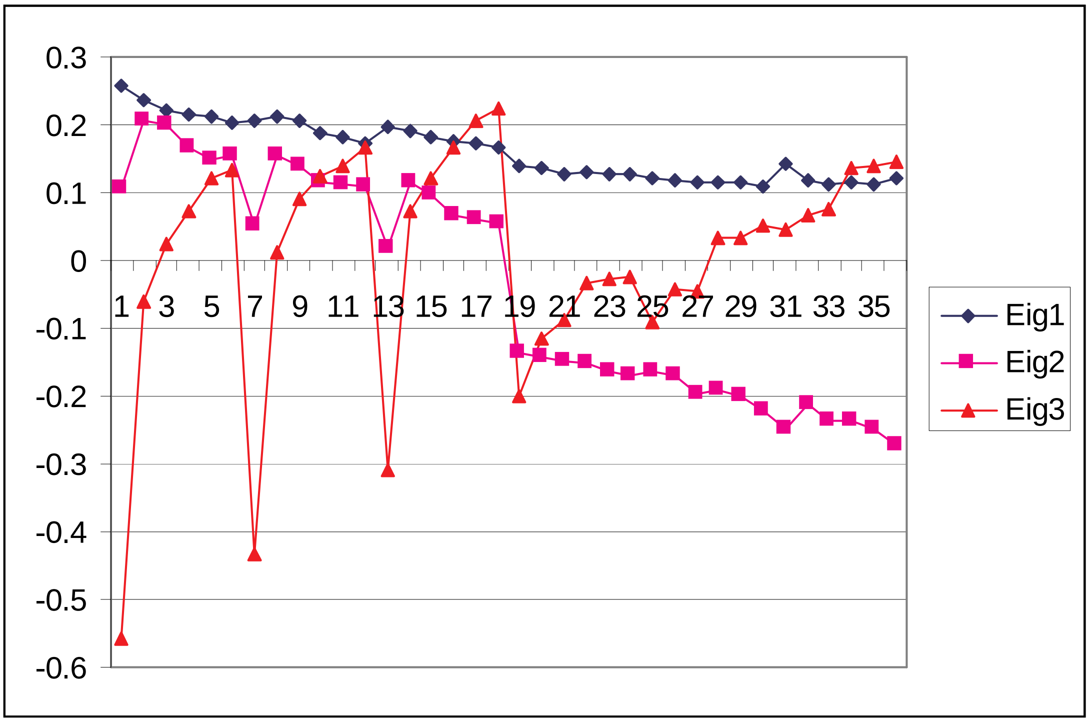

Figure 18: The foirst three eigenvectors from the orrthogonalization of the USD covariance matrix

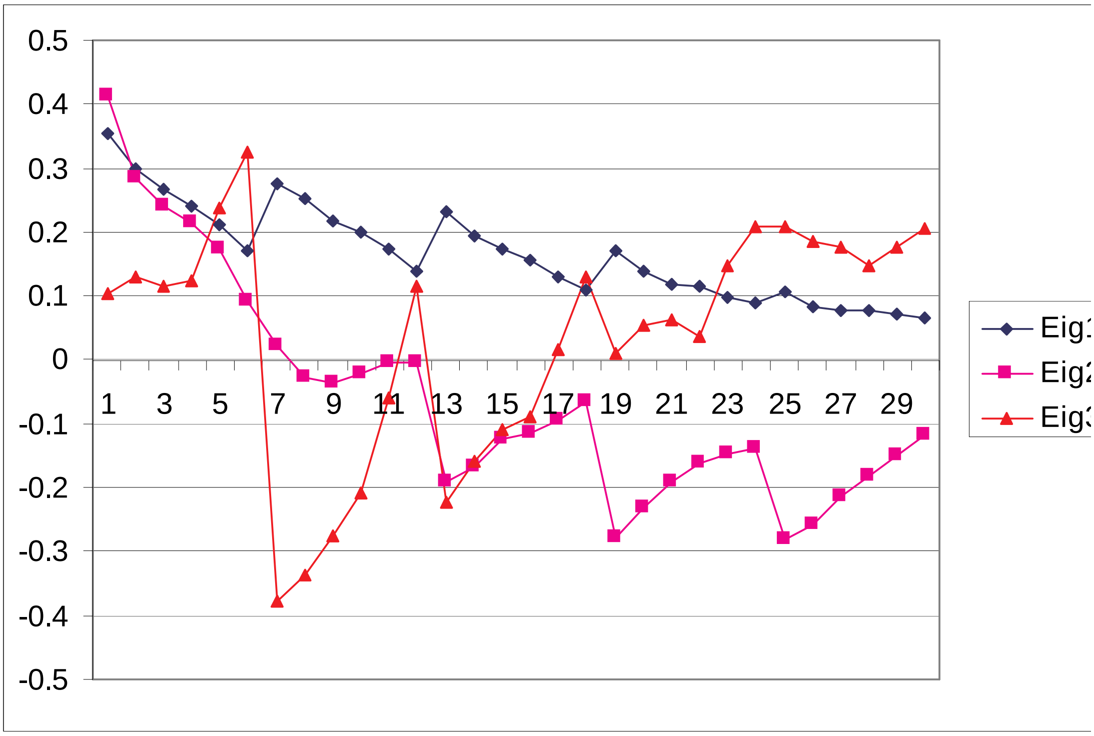

Figure 19: The first three eigenvectors from the orthogonalization of the DEM covariance matrix

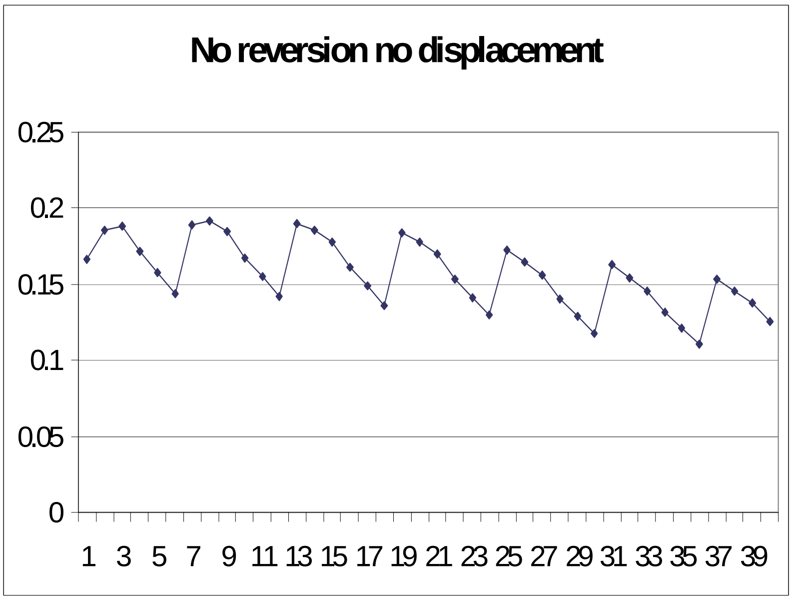

Figure 20: No reversion speed, zero displacement coefficent

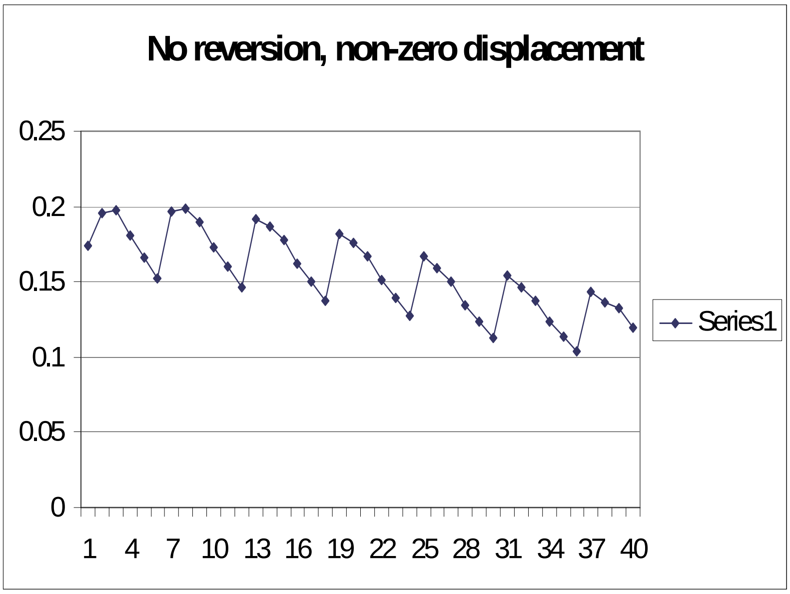

Figure 21: No reversion speed, finite displacement coefficient

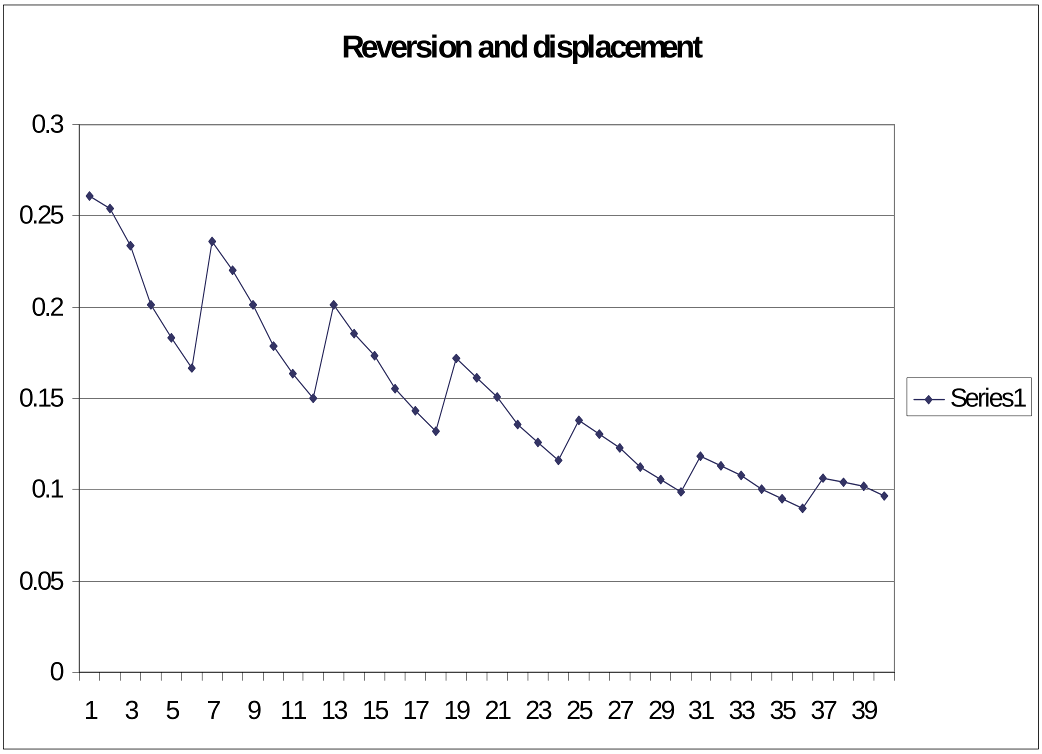

Figure 22: Non-zero reversion speed and displacement coefficient

BIBLIOGRAPHY

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

Brace A, Gatarek D, Musiela M, (1995) ‘The Market Model of Interest Rate Dynamics’, working paper, School of Mathematics, University of New South Wales, Australia

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

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

De Jong, F, Driessen L, Pelsser A (1999) ’LIBOR and Swap Market Models for the Pricing of Interest Rate Derivatives: an Empirical Comparison’, working paper

Duffie, D., (1996) ‘Dynamic Asset Pricing Theory’ , Princeton University Press, Princeton, New Jersey, Second Edition

Hunter C, Jaeckel P and Joshi M (2001) ‘Getting the Drift’ accepted for publication in Risk Magazine, December 2001, also QUARC (Quantitative Research Centre) Working Paper ‘Drift Approximations in a LIBOR Market Model’ available at www.Rebonato.com

Joshi M, Rebonato R. (2001) ‘A stochastic-volatility, displaced-diffusion extension of the LIBOR market model’, submitted to Journal of Computational Finance and QUARC (Quantitative Research Centre) Working Paper available at www.Rebonato.com

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

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

Jamshidian F (1999) ‘LIBOR Market Model with Semimartingales’, working paper, NetAnalytic Ltd., London

Longstaff F.A., Santa-Clara P, Schwartz, ES, (2000a) ‘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, (2000b) ‘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

Martellini L, Priaulet P, (2001) ‘Fixed Income Securities’, John Wiley, Chichester, New York, Weinheim, Brisbane, Singapore, Toronto.

Mercurio, D, Brigo F, (2001) ’Interest Rate Models - Theory and Practice’, Springer Verlag, Berlin, Heidelberg, New York

Miltersen, K, Sandmann K, Sondermann D, (1997) ‘Closed-Form Solutions for Term Structure Derivatives with Log-normal Interest Rates’, Journal of Finance, 52, 409-430

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

Musiela, M, Rutkowski M (1997b): ‘Martingale Methods In Financial Modelling’ Springer Verlag, Berlin

Rebonato R (1999): ’On the Simultaneous Calibration of Multifactor Lognormal Interst-Rate Models to Black Volatilities and to the Correlation Matrix’, Journal of Computational Finance, Vol2 , No 4, Summer 1999, 5-28

Rebonato R (2001) ’Modern Pricing of Interets-Rate Derivatives: the LIBOR Market Models and Beyond’, Princeton University Press, Princeton, New Jersey

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

Sidenius J (2000) ‘Evaluating New Research Exploring Practical Aspects of Multi-Factor Libor Market Models’ Presented at the 7th annual forum - Global Derivatives 2000 - Paris

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

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