On the simultaneous calibration of multi-factor log-normal interest-rate models to Black volatilities and to the correlation matrix

Riccardo Rebonato[1]

Abstract

It is shown in this paper that it is not only possible, but indeed expedient and advisable, to perform a *simultaneous* calibration of a log-normal BGM interest-rate model to the percentage volatilities of the individual rates *and* to the correlation surface. One of the contributions of the paper it to show that the task can be accomplished in two separate and independent steps: the first part of the calibration (i.e. to cap volatilities) can always be accomplished exactly thanks to straightforward geometrical relationships; the fitting to the correlation surface, thanks to a simple theorem, can then be carried out in a numerically efficient way so that the calibration to the volatilities is not spoiled by the second part of the procedure. The ability to carry out the two tasks separately greatly simplifies the overall task.

Actual calculations are shown for a 3- and 4-factor implementation of the approach, and the quality of the overall agreement between the target and model correlation surfaces is commented upon.

Finally, the dangers of overparametrization, i.e. of forcing (near) exact fitting to certain portions of the correlation matrix, are analysed by looking at the cases of a trigger swap, a Bermudan swaption and a one-way floater (resettable cap).

1 – Introduction and motivation for the present study

Until relatively recently, the calibration to market quantities of any interest-rate option model was one of the most arduous parts of its implementation. Users of early short-rate-based models (such as the Black-Derman and Toy (1990), the Hull and White or the Black-Karasinsky (1991)) are too well aware of the difficulties one encounters when attempting to calibrate the model parameters so as to reproduce the prices of caps or swaptions. Also the more recent Heath-Jarrow-Morton (HJM) approach is, in its more general form, hardly more user-friendly when it comes to calibration of the model to market data. The common features of all these models was the fact that, explicitly or implicitly, within these traditional frameworks the stochastic behaviour was specified of *unobservable* financial quantities, such as, for instance, the

[1] The ideas expressed in this article have evolved over time thanks to discussions with several colleagues, to whom I am indebted; in particular, I would like to remember Mike Sherring and Soraya Kazziha.

instantaneous forward rates, the instantaneous short rate or its variance. The calibration of a model to a set of market quantities therefore required transforming, via the black box provided by the model itself, the dynamics of these unobservable quantities into the dynamics of observable quantities.

The recently introduced Brace-Gatrek-Musiela (BGM) approach (1995), germane to the HJM (1989) model, has radically changed this picture: now directly observable market quantities, such as discrete (Libor) forward rates or swap rates, are evolved. Given the availability from the market of the volatilities of caplets and European swaptions, calibration to *either* set of variables has become, at least for one-factor implementations, virtually immediate. Pre-empting the more precise treatment to be found later on, given a set of (forward of swap) rates, the market gives the traded prices of series of caplets and European swaptions. From these one can directly impute, via inversion of the Black formula, the value of the average variance of the appropriate log-normal rate.[2] In turn, this quantity determines the (integral of the square of) the time-dependent instantaneous volatility that must be assigned to the corresponding model rate in order to price the market instrument exactly (at least within numerical noise). See Equation (8) below. In general, the realization at time T of the k-th log-normal forward rate, of value $f_k(t_0)$ today, in terms of its time-dependent instantaneous volatility, $s_k(u)$, is given by

(1) $$f_k(T) = f_k(t_0) \exp \left[\int_{t0}^{T} \mu_k(u) - ½\, s_k(u)^2\, du \right] \exp \left[\int_{t0}^{T} s_k(u)\, dz(u)\right]$$

In the expression above, $\mu_k(u)$ is the value at time u of the time-dependent drift that ensures that the model is arbitrage free. More generally, in the case of r orthogonal[3] driving factors

(2) $$f_k(T) = f_k(t_0) \exp \left[\int_{t0}^{T} \mu_k(u) - ½\, s_k(u)^2\, du \right] \exp \left[\int_{t0}^{T} \sum_{m=1,r} s_{km}(u)\, dz_m(u)\right]$$

[2] 'Smile' effects are not taken into account in this discussion. In most interest-rate markets they tend to be of much smaller magnitude than in the FX or equity markets.

[3] The requirement of orthgonality for the driving factors entails no real loss of generality, as any non-orthogonal Brownian motions can be always transformed into an equivalent orthogonal set. Since working with orthogonal processes is computationally much simpler, the assumption will always be made in the following that $E[dz_i,dz_j] = \delta_{ij}\, dt$.

under the constraint that

(3) $$\sum_{m=1,r} s_{km}(u)^2 = s_k(u)^2$$

As long as Equations (2) and (3) are satisfied, the variance of the forward rates, and hence the caplet prices, will always be correctly recovered, irrespective of the number of driving factors. Similarly, one can write for the realization at time T of the k-th log-normal swap rate (using similar notation), $SR_k(T)$,

(4) $$SR_k(T) = SR_k(t_0) \exp [\int_{t0}^{T} \mu_{SRk}(u) - ½\, \sigma_{SRk}(u)^2 du] \exp [\int_{t0}^{T} \sigma_{SRk}(u)\, dz(u)]$$

Once again, in the case of r orthogonal driving factors, Equation (4) becomes

(5) $$SR_k(T) = SR_k(t_0) \exp [\int_{t0}^{T} \mu_{SRk}(u) - ½\, s_{SRk}(u)^2 ds]$$
$$\exp [\int_{t0}^{T} \sum_{m=1,r} s_{SRk\, m}(u)\, dz_m(u)]$$

with the constraint that

(6) $$\sum_{m=1,r} s_{SRk\, m}(u)^2 = s_{SRk}(u)^2$$

As far as the pricing of either caplets or European swaptions is concerned any number of factors can be used to obtain exactly their market (Black) prices. In particular, one single factor (used with the appropriate state variables) is perfectly adequate, since, as shown in greater detail below, the quantities to be matched are

(7) $$\sigma_{Black}^2\, (T-t_0) = \int_{t0}^{T} s\, (u)^2\, du$$

(In Equation (7) σ_{Black} is the market implied Black volatility for the forward or swap rate expiring at time T, and s(u) is the instantaneous volatility of the same rate from today (t_0) to expiry). If one wants, however, to imply the dynamics of swap rates from the assigned process for forward rates, or vice versa[4], the issue of the dimensionality of the approach is quite subtle. More strongly, one will want, in most practical applications, to use the BGM approach to value exotic interest options; in this case the model price for the exotic product can strongly depend on the dimensionality of the underlying cap- (or swaption-) fitting model. In particular, for instruments like spread options or trigger swaps the choice of an appropriate number of driving factor becomes all-important. (See Sidenius (1998) on this point). If one were to choose for the number of driving factors as many independent Brownian motions as forward rates in a given instrument, one would, of course, be giving the most general description of the problem at hand. Given the well-known results of Principal Component Analysis, however, such an extravagant increase in the dimensionality of the problem is generally deemed to be unnecessary, and computationally detrimental. With fewer factors than the number of forwards, any possible choice for the apportioning of the variance (Equations (3) or (6)), and for the time-dependence of the instantaneous volatilities will give rise to different terminal correlations, covariance elements, and, ultimately, exotic option prices. It is therefore unavoidable to optimize a very-high-dimensional problem to a target time-dependent covariance matrix. Given the complexity of the task, it is not surprising to find in the literature statements along the lines of the following:

> "With such a large number of variables a straightforward optimization is impractical at best. The problem is that finding the global best fit[…] is very difficult in high dimensions. […] The question of calibrating the correlation matrix is very interesting, but it seems impractical to undertake this calibration in parallel with the volatility calibration." (Sidenius (1998))

I propose in this paper a calibration methodology that, contrary to the statement above, does allow a simple and computationally efficient calibration designed to recover exactly and in the most general way the instantaneous volatilities of all the forward rates in the problem, and, at the same time, to fit in the 'best' possible way, given the dimensionality of the approach, the correlation matrix. Looked at in this light, the

---

[4] The issues of the consistency between the simultaneous joint log-normal assumptions for forward and swap rates and their impact on pricing are addressed in Rebonato (1998).

essential problem of calibration within the framework of the BGM approach therefore boils down to

i) choosing the most suitable number of factors given the problem at hand;

ii) choosing a suitable time-dependence for the instantaneous volatility functions;

iii) apportioning the weights necessary for the exact recovery of the desired total instantaneous volatility amongst the time-dependent volatilities of the different factors (see Equations(3) and (6)); in other words, choosing how large the contribution of the m-th factor, $\sigma_{km}(s)^2$, towards $\sigma_k(s)^2$ should be.

The present work shows how the task can be considerably simplified by means of general relationships that *must* hold true for any model of the BGM family, irrespective of the choice of instantaneous volatility; therefore the joint fitting to volatilities *and* correlations not only can, but arguably *should*, be carried out in concert. This article also illustrates important restrictions on the resulting dynamics of the state variables if a truncation of dimensionality is carried out and provides a simple and effective means towards achieving the 'best' possible simultaneous calibration to volatilities *and* to the correlation matrix obtainable given the constraints alluded to above.

2 - Definitions and results

Let $\sigma^i_{Black}(T_i)$ be implied Black volatility of forward rate[5] i of maturity T_i, $f_i(T_i)$, (often abbreviated in the following as f_i), and let $s_i(u,T_i)$ be the instantaneous volatility at time u for that forward rate. As mentioned above, the instantaneous and the implied Black volatilities are linked by the relationship

(8) $$\int_0^T s_i(u,T_i)^2\, du = \sigma^i_{Black}(T_i)^2\, T_i$$

[5] *Mutatis mutandis*, exactly the same theorems and results derived in the following apply to the case of forward swap rates. The approach described in sections 2 and 3 can therefore be used in order to price European swaptions and fit, in the best possible way, a surface of correlations amongst swap rates.

The behaviour of a log-normal forward rate is described by a SDE of the form

(9) $$df_i/f_i = \mu_i(t)\, dt + \sum_{k=1,s} a_{ik}(t)\, dz_k(t)$$

whose solution is well know to given by

(10) $$f_i(T_i) = f_i(0) \exp \left[\int_0^{T_i} \mu_i(u) - ½\, s_i(u,T_i)^2 \, du\right] \exp\left[\int_0^{T_i}\sum_{k=1,s} a_{ik}(u)\, dz_k(u) \right]$$

If

(11) $$E[dz_i,dz_j] = \delta_{ij}\, dt$$

then Equations (2) and (3) describe the most general n-factor log-normal forward-rate s-factor model. Notice that

- if condition (11) were not satisfied for a given model, a rotation of variables could always be carried out so as to ensure orthogonality; Equation (11) will therefore always be assumed to hold true without loss of generality.
- The loadings $a_{ik}(t)$ have a calendar-time dependence, and are specific to each individual forward rate via the first index. They cannot, however, be of the form $a_{ik}(f_i,t)$ and preserve the log-normal distributional feature for the forward rate f_i. Therefore they truly represent the most general specification of an s-factor model consistent with *log-normal* forward rates.
- The drift has been assumed to be forward-rate-specific but dependent at most on calendar time.

It is important to stress that the last requirement is not necessarily met for a generic choice of numeraire, since, in an arbitrary measure, the drifts may depend possibly on the full collection of forward rates. There always exists a measure, however, in which each individual forward rate displays a purely time-dependent drift. More strongly, there always exist a forward-rate-dependent measure (the terminal measure) in which that particular forward rate is driftless. We know, however, via

Girsanov's theorem, that, in moving from one measure to another, the resulting change in the process for a given forward rate is purely a drift transformation, and that, therefore, its volatility remains unchanged. Therefore, whatever result we obtain *about the volatility* in the terminal measure of each forward rate, we can rest assured that this result will hold true even when, for practical purposes, we work in a different measure.

In particular, *if we work in the terminal measure*, it is known that necessary and sufficient condition for the instantaneous volatility of a T-expiry forward rate to produce the correct Black price for a caplet (i.e. to give rise to the desired implied volatility) is that the unconditional variance of the forward rate out to its expiry T should be equal to

$$\text{var}[\ln f(T)] = \sigma^{i}_{Black}(T)^2\, T = \int_0^T s_i(u,T)^2\, du \qquad (12)$$

Therefore, in order to ensure that, in *any* measure, a particular caplet will be correctly priced it is sufficient to impose that, in the *terminal* measure (or, for that matter, in any measure with purely deterministic drifts and volatilities), the variance (12) should be recovered. Since this argument is crucial to following developments, let's look at it from a different angle: let's place ourselves in a one-factor framework, and let us assume that, in the terminal measure associated with a particular forward rate, say, forward rate k, we find a square-integrable function $s(t,T_k)$ such that condition (8) is satisfied. For that particular measure, this is tantamount to imposing the first identity in Equation (12). This would not be correct in any other measure, since the unconditional variance for any different measure will depend on the {f}-dependent drift. However, in the terminal measure the variance relationship (12) can simply be used in order to impose the constraint on the instantaneous volatility function that produces the correct caplet price. Once one has ensured that this instantaneous volatility function, by satisfying the variance constraint, gives the correct caplet price, then this function can be transferred to any other measure, and the caplet will still be exactly priced.

With these definitions clearly in mind, we can prove Theorem (1)

Theorem (1) - Given any function s(t) such that Equation (8) is satisfied, necessary and sufficient condition for the process (2) of forward rate f_i to produce a caplet price consistent with the implied Black volatility $\sigma_{Black}(T_i)$ is that

$$a_{ik}(t)= b_{ik}(t)\ \ s_i(t,T_i)$$

(13) $$b_{ik}(t) = \cos(\theta_{ik}(t))\, \Pi_{j=1,k-1} \sin(\theta_{ij}(t)), \quad k=1,s-1$$

(13') $$b_{ik}(t) = \Pi_{j=1,k-1}\sin(\theta_{ij}(t)), \quad k=s$$

Proof:

$$\mathrm{var}[\ln f_i(T_i)] = E[\ln f_i(t_i)^2] - E[\ln f_i(t_i)]^2$$

Define $f_i(0) \exp [\int_0^{Ti} \mu_i(u) - ½\, s_i(u,T)^2\ du] \equiv f_i(0,T_i)^*$

Then

$$E[\ln f_i(t_i)] = \ln f_i(0,T_i)^*$$

$$E[\ln f_i(t_i)^2] = [\ln f_i(0,T_i)^*]^2 + E[\int_0^{Ti}\Sigma_{k=1,s}\, a_{ik}(u)\, dz_k(u) \int_0^{Ti}\Sigma_{m=1,s}\, a_{im}(u)\, dz_m(u)]$$

=

$$= [\ln f_i(0,T_i)^*]^2 + E[\int_0^{Ti}\Sigma_{k=1,s}\, a_{ik}(u)^2\, dz_k(u)^2] =$$

$$[\ln f_i(0,T_i)^*]^2 + \int_0^{Ti} s_i(u,T_i)^2\, \Sigma_{k=1,s}\, b_{ik}(u)^2\, du$$

But $\Sigma_{k=1,s}\, b_{ik}(u)^2 =1\ \forall u$, because the coefficients {b} can be recognized to be the polar co-ordinates of a unit-radius (s+1)-dimensional hyper-sphere. Therefore

$$E[\ln f_i(t_i)^2] = (\ln f_i(0,T_i)^*)^2 +\int_0^{Ti} s_i(u,T_i)^2\ du$$

and

$$\int_0^{T} s_i(u,T)^2\, du = \mathrm{var}[\ln f(T)] = \sigma^i_{Black}(T)^2\, T$$

Q.E.D.

As a by-product of Theorem (1) one can immediately obtain that, if Equation (8) is satisfied, necessary and sufficient condition for the process (2) of forward rate f_i to produce a caplet price consistent with the implied Black volatility $\sigma_{Black}(T_i)$ is that

$$\sum_{k=1,s} b_{ik}(t)^2 = 1$$

Let us now consider any path-dependent option problem such that the expiries and maturities of a set of n forward rates constitute all the dates when price-sensitive events occur. Such events could be discrete trigger events, discrete sampling times for averages, resetting of stochastic strikes for ratchet caps, etc. Continuous-time barriers do not belong to this type of problem. After each price-sensitive event takes place, the number of forward rates left in the problem is reduced by one. Let us define the time interval between price-sensitive event i and price-sensitive event i+1 as the i-th time-step. Let then h(i) be the number of forward rates 'alive' at time step i. If, in the modelling of such a problem, at each time step the chosen number of factors is equal to h(i), then the problem of reproducing an arbitrary set of instantaneous volatilities and an arbitrary exogenously assigned correlation matrix can always be trivially accomplished by orthogonalizing of the time-dependent covariance matrix whose elements $Cov_{jk}(i)$ are given by (with obvious notation)

$$Cov_{jk}(i) = \int_{ti}^{ti+1} s(u,T_j)\, s(u,T_k)\, \rho_{jk}(u)\, du.$$

Since retaining as many factors as residual forward rates is practically too onerous for most applications, the interesting question is what happens to the *model* covariances when s(i) < h(i) factors are retained at each time step:

$$Covjk(i)^{mod} = \int_{ti}^{ti+1} \sum_{r=1,h(i)} a_{jr}(u)\, a_{kr}(u)\, du$$

Notice that, if one retains h(i) factors at aeach time step, then, after making use of relationships (13) and (13'), there are enough degrees of freedom in order to specify any feasible exogenously specified covariance matrix element. For an option problem with a finite number of price-sensitive events these covariance elements, in turn, fully specify the valuation problem. In a sense, therefore, if h(i) factors are reatined at each time step, the calibration problem does not exist. If, however, only s(i) < h(i) fatcors are retained, the non-diagonal model covariance elements will not coincide with the corresponding arbitrarily specified exogenous inputs. But, since, by Theorem 1 above, the instantaneous volatilities (and hence the implied Black volatilities) can always be recovered for an arbitrarily small number of factors, the calibration problem is tantamount to specification of the behaviour of the model time-dependent correlation implied by the dynamics of s(i) < h(i) factors.

As shown above, the most general h-factor implementation of a log-normal-forward-rate BGM model is fully specified by the matrix $\{b_{jk}\}$, j=1, h(i), k=1,s(i). For future reference let us denote by $\mathbf{b_r}$ the r-th column vector in the matrix **B** of elements $\{b_{jk}\}$. Since, in general, an arbitrary target (‘market’ in the following) correlation function will not be reproducible with s(i) orthogonal factors, the user is faced with the problem of determining the elements of the matrix $\{b_{jk}\}$ in such a way that

1) the orthogonality between the different vectors is retained:

$$\sum_{j=1,h(i)}\sum_{j'=1,h(i)} b_{jk}\, b_{j'k} = 0 \qquad \forall k \leq s(i)$$

2) each vector is normalized to 1:

$$\sum_{j=1,h(i)} b_{jk}^2 = 1 \qquad \forall k \leq s(i)$$

3) the sum of the coefficients $\{b_{jk}\}$ across factors also adds up to one:

$$\sum_{k=1,s(i)} b_{jk}^2 = 1 \qquad \forall j \leq h(i)$$

4) the discrepancies between the implied (model) and market correlation matrices are minimized in some precise way to be defined.

Conditions 1) and 2) ensure orthogonality of the vectors, and condition 3) ensures the correct recovery of a desired instantaneous volatility. In general, optimizing the coefficients $\{b_{jk}\}$ is a complicated exercise, given the joint constraints about sums over factors, about sums over forward rates and about the model correlation matrix. This, indeed, is the difficulty mentioned by Sidenius (1998) in the opening section. A simple and very useful theorem, however, is of great assistance in the joint fulfilment of desiderata 1) to 4).

Theorem (2) - Let B be an $[n \times s]$ real matrix such that $rank(B) = s$. Then since the matrix BB^T is real and symmetric, it can be diagonalised : i.e. there exists an orthogonal matrix $P(n \times s)$ and a diagonal matrix $\Lambda(n \times n)$ such that

$$P \Lambda P^T = BB^T$$

But $rank(B) = rank(BB^T) = s$, so that BB^T has only s eigenvalues different from zero. It is always possible to permute the columns of B such that the non-zero eigenvalues are in the top left hand corner of Λ :

$$\Lambda = \begin{bmatrix} \lambda_1^2 & & & 0 \\ & \ddots & & \\ & & \lambda_s^2 & \\ 0 & & & 0 \end{bmatrix}$$

Let

$$P = \begin{bmatrix} C_1 & C_2 & \dots & C_s & \dots & Cn \end{bmatrix}$$

where C_i is the i^{th} column of P. Define

$$A \equiv \begin{bmatrix} \lambda_1 C_1 & \lambda_2 C_2 & \dots & \lambda_s C_s \end{bmatrix}$$

then

$$\boxed{BB^T = AA^T}$$

Proof: Let us first note that if

$$M \equiv \left[\underbrace{\lambda_1 C_1 \quad \lambda_2 C_2 \quad \ldots \quad \lambda_s C_s \quad \underline{0} \quad \underline{0}}_{n} \right]$$

(where $\underline{0}$ indicates a column of zero elements) then $MM^T = AA^T$. Therefore, given that M and P are both $(n \times n)$ matrices, we can transform P to obtain M by:

1-multiplying columns $C_1, C_2, \ldots C_s$ in P by $\lambda_1, \lambda_2, \ldots \lambda_s$ respectively.
2-setting the remaining columns $C_{s+1}, \ldots C_n$ in P to zero.

These scaling operations can easily be obtained using the following class of $(n \times n)$ matrices

$$H_\alpha^i \equiv \begin{bmatrix} 1 & & & & & 0 \\ & 1 & & & & \\ & & .. & & & \\ & & & \alpha & & \\ & & & & 1 & \\ 0 & & & & & 1 \end{bmatrix} \qquad \text{i.e} \qquad \begin{matrix} \left(H_\alpha^i\right)_{ii} = \alpha \\ \left(H_\alpha^i\right)_{jj} = 1, \forall j \neq i \end{matrix}$$

Then it follows that

1- multiplying any matrix $P(n \times n)$ by H_α^i from the left amounts to multiplying row i by α
2- multiplying any matrix $P(n \times n)$ by H_α^i from the right amounts to multiplying column i by α

Therfore $M = PH_{\lambda_1}^1 H_{\lambda_2}^2 .. H_{\lambda_s}^s H_0^{s+1} .. H_0^n$ and, noting that $H_\alpha^i = \left(H_\alpha^i\right)^T$, we have

$$MM^T = PH^1_{\lambda_1}H^2_{\lambda_2}..H^s_{\lambda_s}H^{s+1}_0H^{s+2}_0..H^n_0H^n_0H^{n-1}_0..H^{s+1}_0H^s_{\lambda s}..H^2_{\lambda_2}H^1_{\lambda_1}P^T$$

$$= P\begin{pmatrix} \lambda_1^2 & & & & 0 \\ & .. & & & \\ & & \lambda_s^2 & & \\ & & & 0 & \\ 0 & & & & 0 \end{pmatrix}P^T$$

$$= P\Lambda P^T = BB^T$$

Finally

$$BB^T = MM^T = AA^T$$

Q.E.D.

The above Theorem (2) is extremely useful in the parametrization of an s-factor BGM model, and allows to reduce to a very large extent the difficulties in the *simultaneous* calibration to volatilities and to the correlation matrix alluded to in Section (1). It can be used in the following way. Let us start from a matrix $\{b_{jk}\}$ such that only condition (3) is satisfied. This can always be achieved, without any loss of generality for an arbitrary choice of the angles $\{\theta\}$, by making use of relationships (13) and (13'). Let us now vary the angles $\{\theta\}$ in such a way that the sum of the (suitably weighted) squared discrepancies between the model and market correlation surfaces are minimzed. Since by Theorem (1) we know how to calibrate exactly to the variances of the various forward rates, calibration to the correlation matrix is all that is left in order to specify completely the most general log-normal forward rate model. But this latter task can always and simply be accomplished by virtue of Theorem (2) since the model correlation matrix $\{\rho_{j,j'}\}$ is given by B B^T. More precisely, the reasoning goes as follows: after optimizing the angles $\{\theta\}$, and hence the coefficients $\{b_{jk}\}$, in the desired manner to the market correlation function, the resulting vectors are neither normalized, nor orthogonal to each other. However, after orthogonalizing the h(i) x h(i) matrix B B^T, conditions 1) and 2) will be satisfied for the resulting eigenvectors multiplied by the square root of each respective eigenvalue. Let us denote by $\{a_{jk}\}$ the elements of the new h(i) x s(i) matrix A containing the eigenvectors multiplied by the square root of the eigenvalues. The crucial point is that, by Theorem (2), one can rest assured that, after diagonalization of the non-orthogonal, non-normalized vectors, the correlation matrix obtained by $A A^T$ will be identical to the correlation matrix obtained by $B B^T$. *The tasks of optimizing to a market correlation surface and of fulfilling conditions 1) and 2) can therefore be profitably tackled separately, without having to engage in a joint constrained optimization*.

Thanks to Theorem (1) and (2) the parametrization of the most general BGM log-normal forward rate model can be accomplished in a fast and numerically efficient way so that, *at the same time,*

1) an arbitrarily specified set of instantaneous volatilities can be recovered
2) if these instantaneous volatilities are consistent with Equation (1), the the Balck volatilities are correctly reproduced within the model (and hence all the caplets are correctly priced);
3) the resulting driving factors are orthogonal to each other
4) the best fit (given a set of exogenous quality criteria) to a given market correlation function compatible with the number s of factors can easily be found.

Notice that, at time step i, the market correlation matrix is characterized by h(i) x h(i) elements. Despite the fact that there are h(i) x s(i) terms in the matrix B, in reality only at most h(i) x [s(i)-1] elements of the model correlation matrix can be forced to match exactly[6] (if one so wished) the corresponding elements of the market correlation matrix. That this is the case can be intuitively understood from the fact that there are only h(i) x [s(i)-1] angles that can be changed at will. In particular, for the case of s = 2 one can obtain Theorem (3) below.

Theorem (3) - For a 2-factor model the correlation between forward rate i and forward rate j only depends on the difference between the angles θ_{i1} and θ_{j1}.

[6] It will be argued later in this paper that forced matching of a given sub-set of the market correlation matrix can in general be less than desirable, and indeed quite dangerous.

For s=2 the correlation ρ_{ij} is given by

$$\rho_{ij} = E[\sum_{s=1,2} b_{is}\, dz_s \sum_{r=1,2} b_{jr}\, dz_r]$$

Remembering the definition of $\{b_{ir}\}$, the orthogonality of the Brownian increments $\{dw_r\}$, and dropping the redundant second subscript '1' for the angles one obtains

$$b_{i1}= \cos \theta_i$$

$$b_{i2} = \sin \theta_i$$

and therefore

$$\rho_{ij} = \cos \theta_i \cos \theta_j + \sin \theta_i \sin \theta_j = \mathbf{\cos (\theta_i-\theta_j)}$$

Q.E.D.

3 - Numerical results

All the results reported in the following pertain to the case of a collection of 12 twelve-month forward rates. The task was undertaken of fitting simultaneously and exactly to all the market volatilities, and to obtain the best possible fit to a given ('target' or 'market') correlation matrix by using the procedure described in Section 2. As we know, any combination of the angles $\{\theta\}$ in Equations (13) and (13') is by construction compatible with the exact recovery of the volatilities of all the individual forward rates. We also know, however, that each of these combinations gives rise to a different correlation surface. Given an arbitrarily chosen quality function[7], the various angles can therefore be varied in such a way as to minimize this exogenoulsy specified function. One such quality function, for instance, could be the sum of the squares of the 12 x 12 differences between the whole model and 'market' correlation matrices. The exact procedure and the results are described in detail below.

---

[7] A quality function is defined as the sum of the squared differences of specified subsets of the model and the target correlation matrices

3-a: Fitting the correlation surface with a 3-factor model

Random numbers were first of all chosen for the 2 x 12 angles θ_{ij} (i=1,2, ...,12, j=1,2). The coefficients $\{b_{ij}\}$ (i=1,2,...,12, j=1,2,3) were created using Equations (13) and (13’). Let **B** be the 12 x 3 matrix made up by the vectors b_{ij}. The model correlation matrix was then constructed by means of Theorem (2) as $\{\boldsymbol{\rho}^{\mathbf{mod}}\} = \mathbf{B}\,\mathbf{B}^T$. These random numbers were varied until the sum of the squared discrepancies $[\rho^{mod}_{ij} - \rho^{market}_{ij}]^2$ *over the whole matrix* was reduced to a minimum. The ‘market’ correlation function was assumed to be given by the following expression:

$$\rho^{market}_{ij} = \text{LongCorr} + (1 - \text{LongCorr}) \exp[\ \beta\, |t_i - t_j|\ ]$$

$$\beta = d_1 - d_2 \,\text{Max}\,(t_i, t_j)$$

and LongCorr = 0.3, d_1 =- 0.12, d_2 = 0.005

B(1)	B(2)	B(3)
0.802	-0.490	0.342
0.848	-0.391	0.357
0.899	-0.248	0.361
0.948	-0.014	0.318
0.971	0.140	0.196
0.978	0.199	0.057
0.977	0.195	-0.085
0.965	0.120	-0.233
0.934	-0.048	-0.354
0.899	-0.177	-0.401
0.873	-0.250	-0.420
0.851	-0.305	-0.426

A(1)	A(2)	A(3)
0.839	**0.436**	**0.33**
0.874	**0.432**	**0.22**
0.909	**0.410**	**0.07**
0.934	**0.323**	**-0.15**
0.944	**0.175**	**-0.28**
0.949	**0.027**	**-0.31**
0.953	**-0.111**	**-0.28**
0.954	**-0.242**	**-0.18**
0.944	**-0.329**	**0.01**
0.924	**-0.351**	**0.15**
0.906	**-0.356**	**0.23**
0.891	**-0.353**	**0.28**

Tab I : The vectors $\{b_{ij}\}$ and $\{a_{ij}\}$

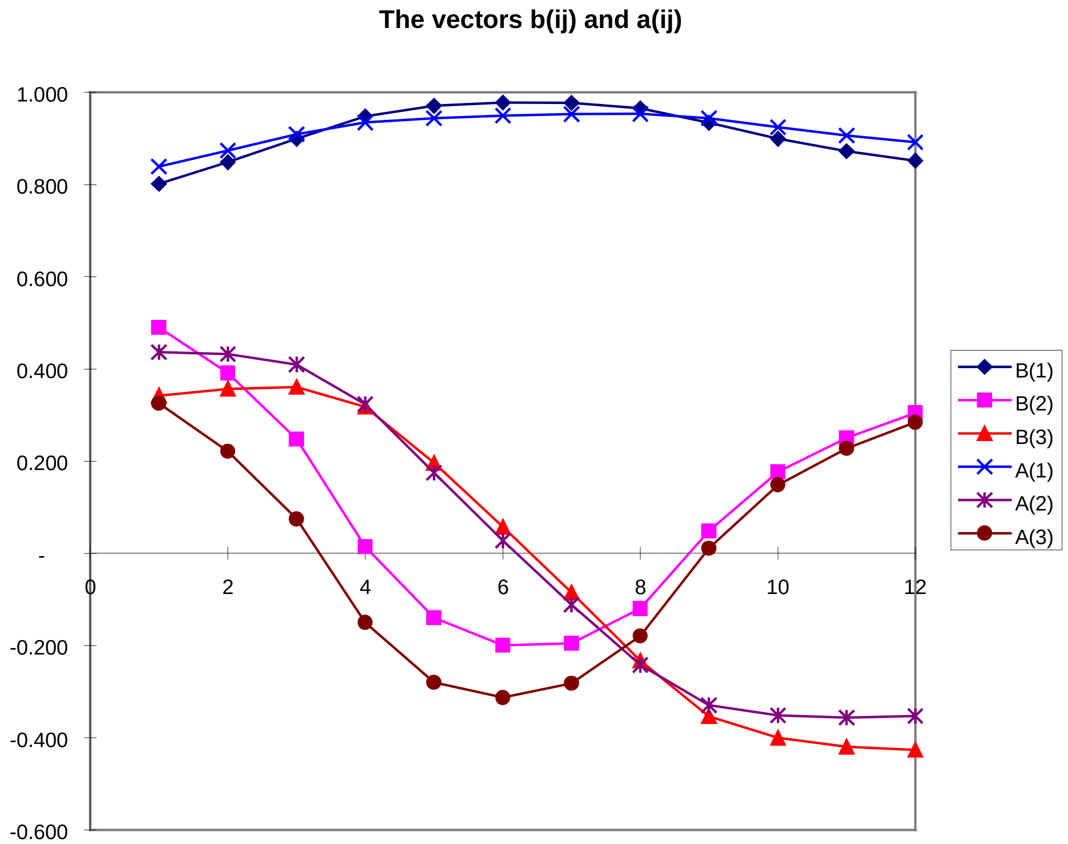

Fig 1: The vectors {b_{ij}} and {a_{ij}} for the 3-factor case (all-matrix fit)

Theta	Phi
0.640174	-0.96091
0.557733	-0.83127
0.453171	-0.6008
0.323572	-0.04526
0.243088	0.618934
0.208983	1.290709
0.214491	1.980806
-0.26474	-0.47569
-0.36511	0.135639
-0.45315	0.415468
-0.51038	0.537576
-0.55197	0.620986

Tab Ib: The angles obtained after the optimization described in the text for the 3-factor fit to the whole 'market' correlation matrix in Tab II.

When the optimal vectors {b_{ij}}were found, the resulting 12 x 12 matrix was orthogonalized, giving rise to new vectors {a_{ij}}. Given the rank of the $\mathbf{B}\,\mathbf{B}^T$ matrix, only 3 of the resulting eigenvalues were different from zero. The vectors $\{b_{ij}\}$, $\{a_{ij}\}$, the model ($\mathbf{B}\,\mathbf{B}^T$) and target correlation matrices, the eigenvectors and the eigenvalues resulting from the orthgonalization are shown below, either in tabular form or as figures. It is interesting to notice that, despite the fact that no orthogonality constraints were imposed in the optimazation, the solution qualitatively turned out to be very similar to what found using Principal Component Analysis. Indeed, as shown in Fig. 1, the rotation induced by the orthogonalization was minor.

ModelCorrelMatrix

1.0000	0.993951	0.9658	0.8761	0.7771	0.7065	0.6590	0.6358	0.6517	0.6705	0.6787	0.6864
0.9940	1.0000	0.9884	0.9233	0.8388	0.7725	0.7224	0.6891	0.6852	0.6890	0.6885	0.6896
0.9658	0.9884	1.0000	0.9707	0.9089	0.8509	0.7995	0.7541	0.7240	0.7074	0.6949	0.6870
0.8761	0.9233	0.9707	1.0000	0.9805	0.9428	0.8966	0.8395	0.7739	0.7277	0.6976	0.6762
0.7771	0.8388	0.9089	0.9805	1.0000	0.9886	0.9590	0.9079	0.8305	0.7694	0.7297	0.7002
0.7065	0.7725	0.8509	0.9428	0.9886	1.0000	0.9899	0.9547	0.8838	0.8213	0.7796	0.7477
0.6590	0.7224	0.7995	0.8966	0.9590	0.9899	1.0000	0.9862	0.9333	0.8780	0.8393	0.8086
0.6358	0.6891	0.7541	0.8395	0.9079	0.9547	0.9862	1.0000	0.9780	0.9397	0.9098	0.8845
0.6517	0.6852	0.7240	0.7739	0.8305	0.8838	0.9333	0.9780	1.0000	0.9900	0.9756	0.9610
0.6705	0.6890	0.7074	0.7277	0.7694	0.8213	0.8780	0.9397	0.9900	1.0000	0.9968	0.9903
0.6787	0.6885	0.6949	0.6976	0.7297	0.7796	0.8393	0.9098	0.9756	0.9968	1.0000	0.9982
0.6864	0.6896	0.6870	0.6762	0.7002	0.7477	0.8086	0.8845	0.9610	0.9903	0.9982	1.0000

MarketCorrelMatrix

0	1	2	3	4	5	6	7	8	9	10	11
1.0000	0.9240	0.8618	0.8109	0.7692	0.7353	0.7079	0.6861	0.6691	0.6564	0.6476	0.6424
0.9240	1.0000	0.9271	0.8674	0.8186	0.7787	0.7463	0.7203	0.6998	0.6842	0.6728	0.6654
0.8618	0.9271	1.0000	0.9302	0.8731	0.8264	0.7884	0.7576	0.7331	0.7141	0.6998	0.6900
0.8109	0.8674	0.9302	1.0000	0.9334	0.8789	0.8344	0.7982	0.7692	0.7463	0.7288	0.7162
0.7692	0.8186	0.8731	0.9334	1.0000	0.9366	0.8847	0.8424	0.8083	0.7811	0.7599	0.7441
0.7353	0.7787	0.8264	0.8789	0.9366	1.0000	0.9398	0.8906	0.8506	0.8186	0.7933	0.7739
0.7079	0.7463	0.7884	0.8344	0.8847	0.9398	1.0000	0.9430	0.8965	0.8590	0.8290	0.8058
0.6861	0.7203	0.7576	0.7982	0.8424	0.8906	0.9430	1.0000	0.9462	0.9025	0.8674	0.8397
0.6691	0.6998	0.7331	0.7692	0.8083	0.8506	0.8965	0.9462	1.0000	0.9494	0.9086	0.8760
0.6564	0.6842	0.7141	0.7463	0.7811	0.8186	0.8590	0.9025	0.9494	1.0000	0.9527	0.9147
0.6476	0.6728	0.6998	0.7288	0.7599	0.7933	0.8290	0.8674	0.9086	0.9527	1.0000	0.9559
0.6424	0.6654	0.6900	0.7162	0.7441	0.7739	0.8058	0.8397	0.8760	0.9147	0.9559	1.0000

Tab II: The model and market correlation matrices.

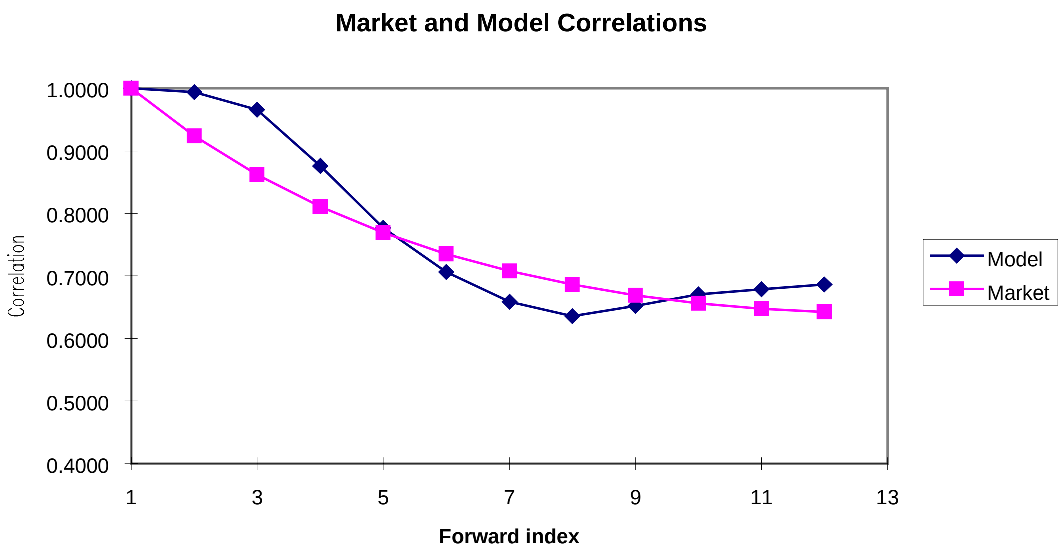
Market and Model Correlations
1.0000
0.9000
0.8000
0.7000
0.6000
0.5000
0.4000
Correlation
1
3
5
7
9
11
13
Forward index
Model
Market

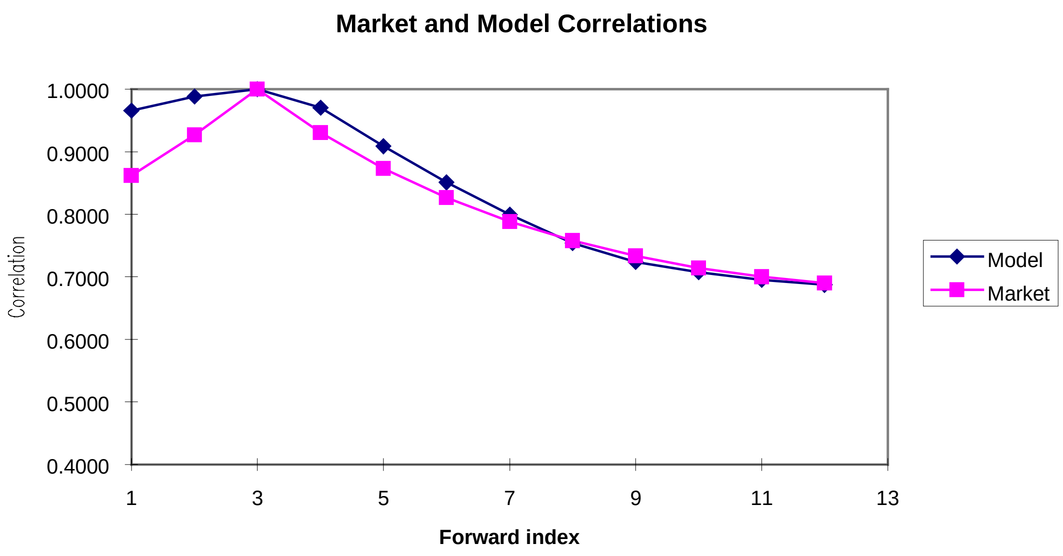
Market and Model Correlations
1.0000
0.9000
0.8000
0.7000
0.6000
0.5000
0.4000
Correlation
1
3
5
7
9
11
13
Forward index
Model
Market

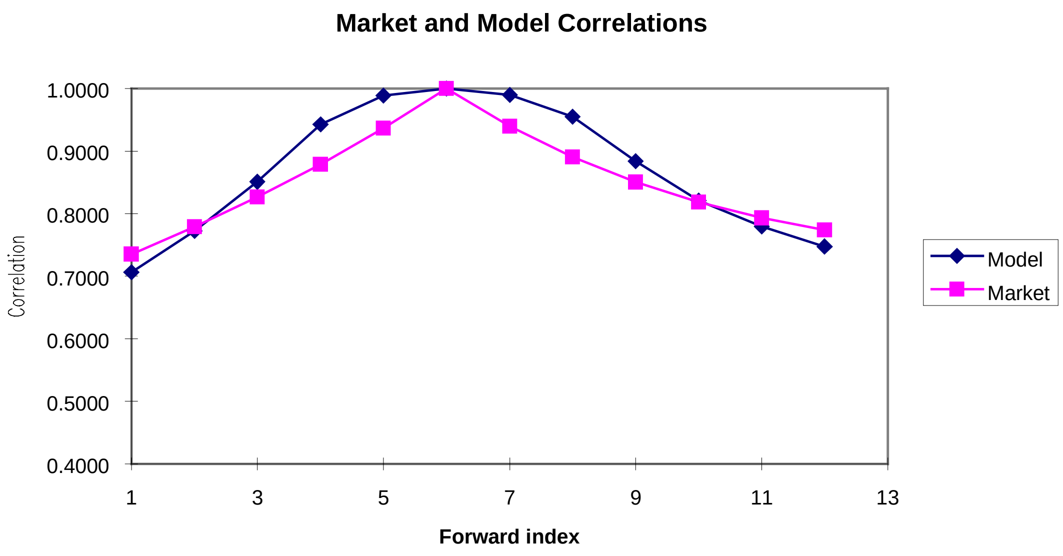

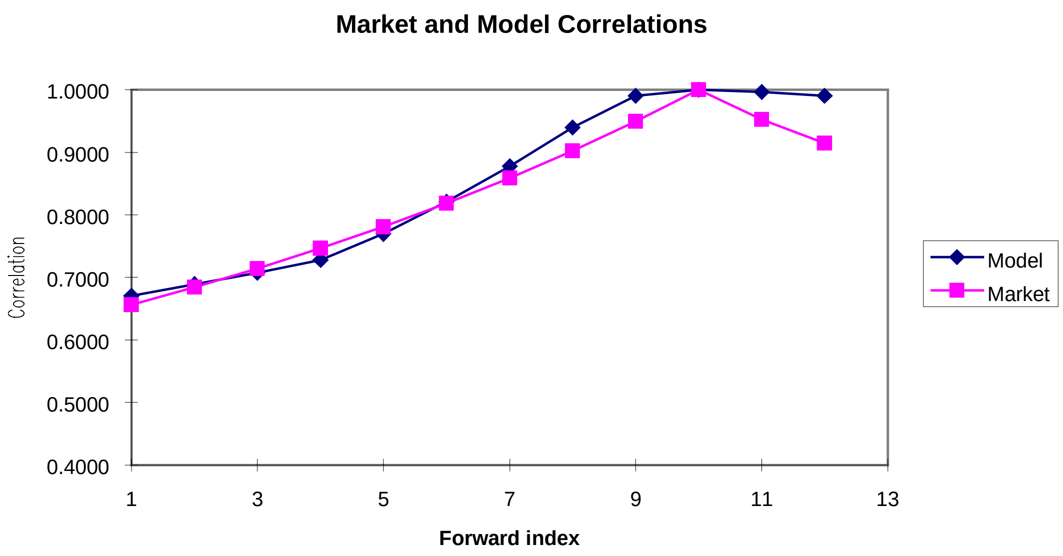

Fig 2 to 5: The market and model correlations between the first, the third, the fifth and the eighth forward rates and all the other forwards obtained using 3 factors, and imposing an overall best fit to the whole correlation matrix.

The well known shortcomings of low-dimensionality models in reproducing correlation functions with positive convexity at the origin are well known (see, e.g. Rebonato and Cooper (1995) or Rebonato (1998), where the implications for pricing are discussed at length), and were indeed observed again in this study. Without repeating material presented elsewhere, it will suffice to say that, by creating too *strong* a correlation between adjacent forward rates, and too *weak* a correlation between distant forward rates, the model correlation surface obtained with a small number of factors systematically misprices swaptions ('short' swaptions too 'expensive', and 'long' swaptions too 'cheap'). With this proviso in mind, the overall agreement found by using the procedure described above is however good, and shown more clearly by the model correlation surface shown in Fig 6 below:

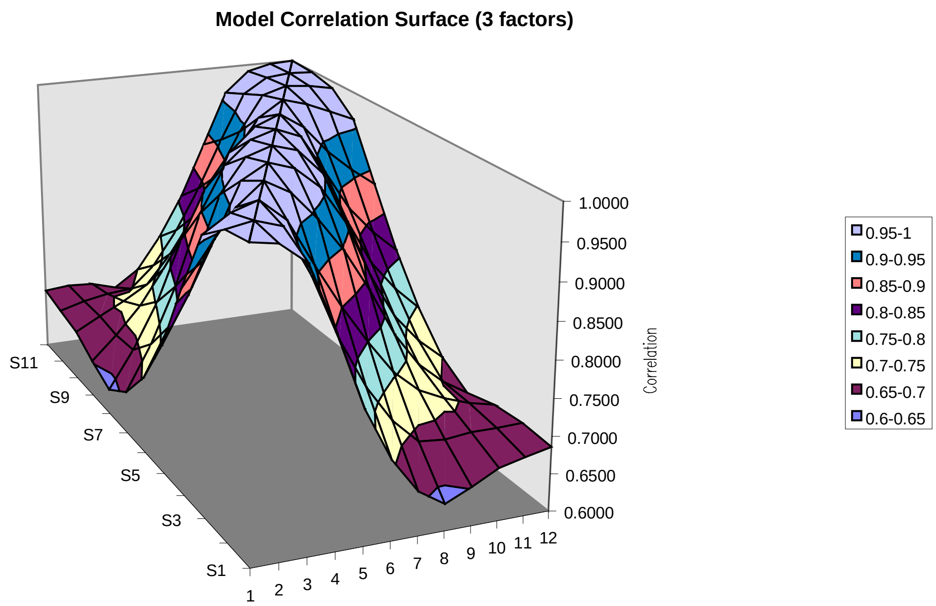

Fig 6 : The correlation surface obtained by overall best fit to the whole correlation surface with 3 factors.

3-b: Fitting the correlation surface with a 4-factor model

In order to see to what extent the agreement between the model and market correlation surfaces is sensitive to the number of factors, a similar investigation was carried out for a four-factor approach. The same procedure described in Section 3-a was followed by adding a further column of 12 angles to the set used for the 3-factor case. After randomizing all the angles and optimizing again to the whole-matrix quality function, the vectors $\{b_{ij}\}$ and $\{a_{ij}\}$ shown in the figure below were obtained. Once again, the vectors $\{b_{ij}\}$ thus obtained bear a close resemblance with the first four eigenvectors usually found with PCA. In particular, qualitatively it seems that the vectors $\{b_{ij}\}$ and their orthgonalized counterparts are linked by a phase rather than frequency transformation. This would indicate that the findings of Theorem (3) above can be extended to higher dimensions.

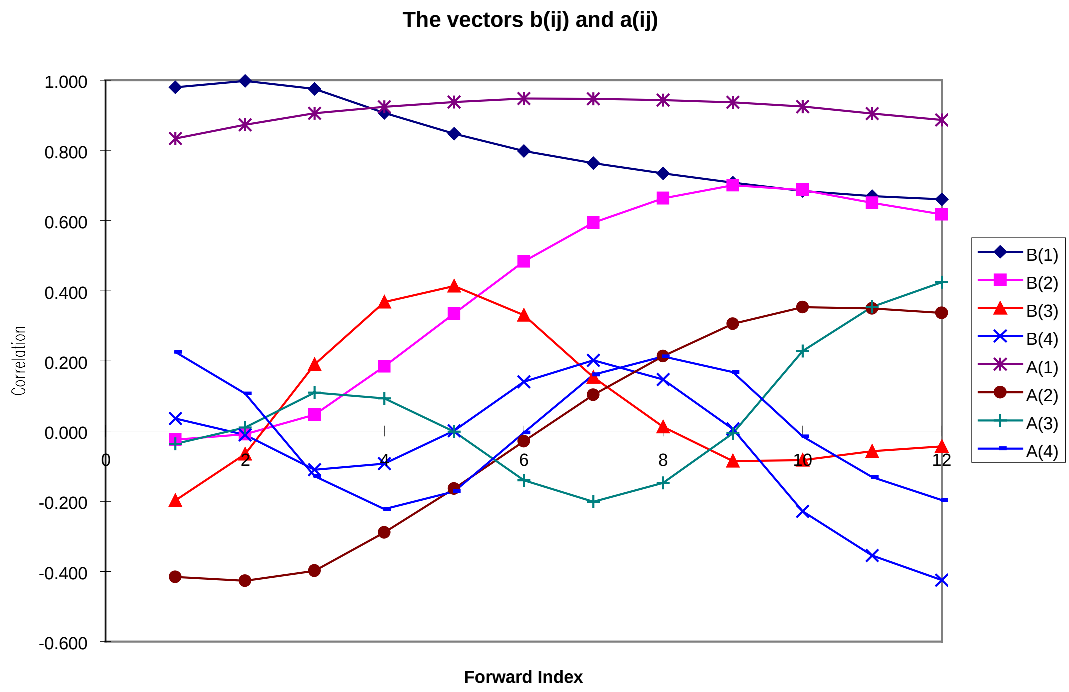

Fig 7: The vectors $\{b_{ij}\}$ and $\{a_{ij}\}$ for the 4-factor case (all-matrix fit)

It is interesting to compare the ‘improvement’ in the fitting of the same columns of the correlation matrix in going from 3 to 4 factors. As is apparent from Figs. 8 to 10, the greatest changes take place for the first series, i.e. for the correlation between the first and second forward rates and all the others. Overall, however, the improvement is rather limited, confirming previous findings (Rebonato and Cooper (1995)) that he convergence to an exponentially decaying target correaltion surface is very slow. In particular, the same qualitative features concerning negative convexity at the origin are still observed.

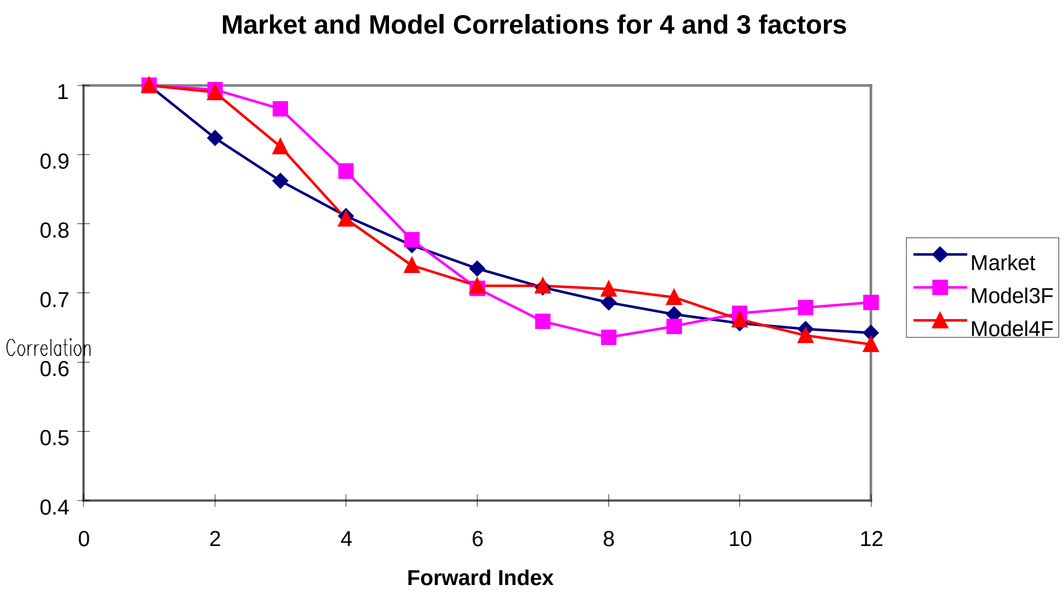
Market and Model Correlations for 4 and 3 factors
1
0.9
0.8
0.7
Correlation
0.6
0.5
0.4
0
2
4
6
8
10
12
Forward Index
Market
Model3F
Model4F

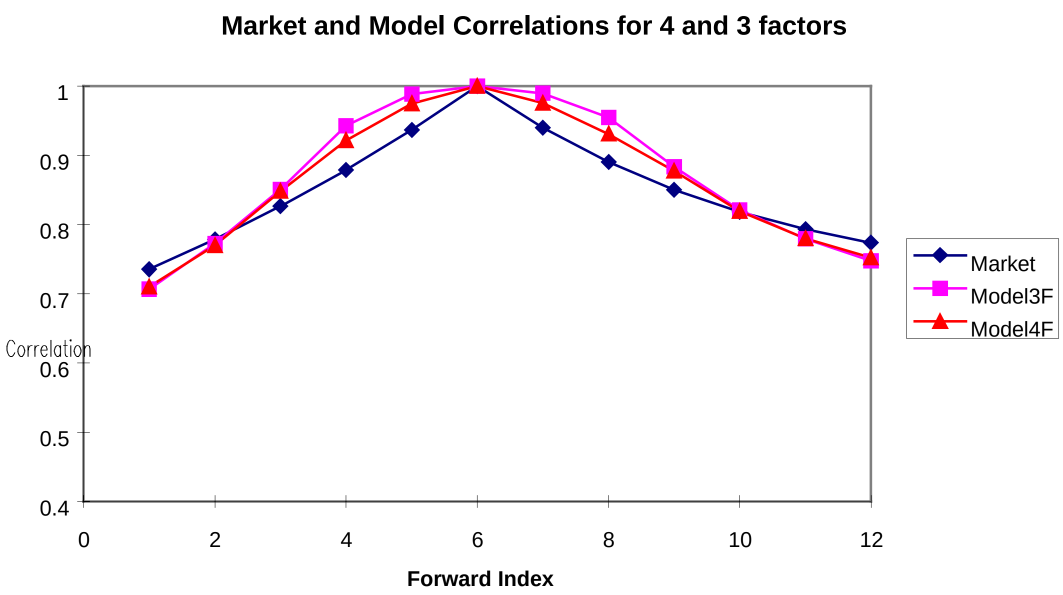
Market and Model Correlations for 4 and 3 factors
1
0.9
0.8
0.7
Correlation
0.6
0.5
0.4
0
2
4
6
8
10
12
Forward Index
Market
Model3F
Model4F

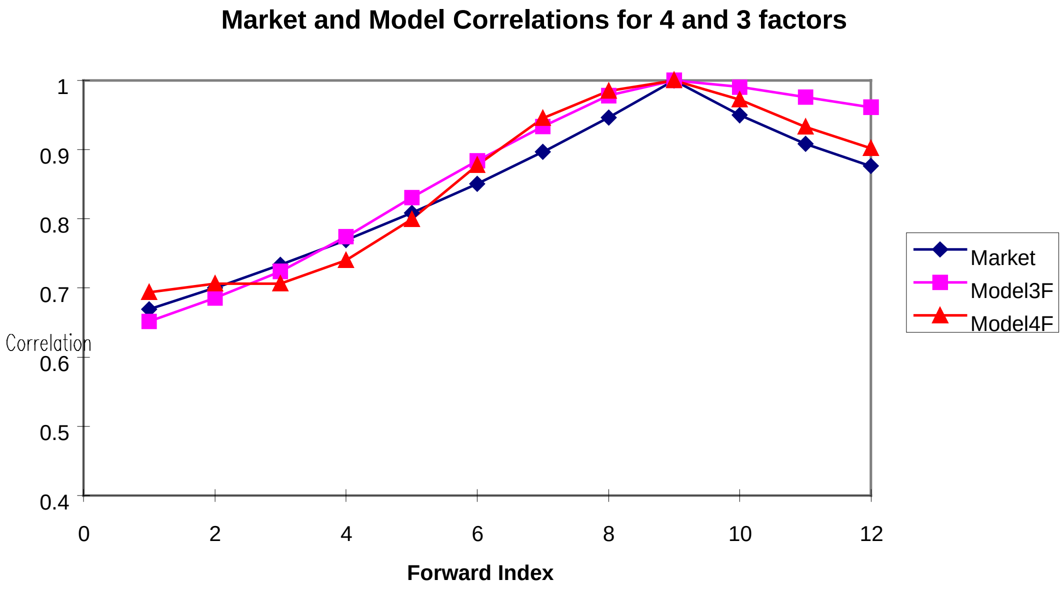

Fig 8 to 10: Comparison between the market and model correlations between the first, the fifth and the eighth forward rates and all the other forwards obtained using 3 or 4 factors (Model3F and Mode4F, respectively) and imposing an overall fit to the correlation matrix.

The analysis so far has shown the results obtained by using as quality function the sum of the squared discrepancies over the all correlation matrix. For specific option problems, however, it might be argued that specific subsections of the correlation matrix might be more important for pricing, and that the attempt should therefore be made to strive for a closer fit to these particular areas. It is therefore instructive to examine in detail the correlation surfaces obtained by using different quality functions. With 4 factors there are enough degrees of freedom to match exactly up to the first 3 columns of the correlation matrix. This choice for the quality function could be motivated by the desire to capture exactly (or as well as possible) the correlation between the first short-maturity LIBOR index rates and the residual swap rates to maturity in the case of a trigger swap. Whilst achieving this goal is indeed possible, as displayed below, if one so did the price would be paid of an increasingly unsatisfactory fit (see Figs. 11 to 14) to the remaining columns of the correlation matrix (giving therefore rise to highly unsatisfactory correlation between later indices and the residual swaps).

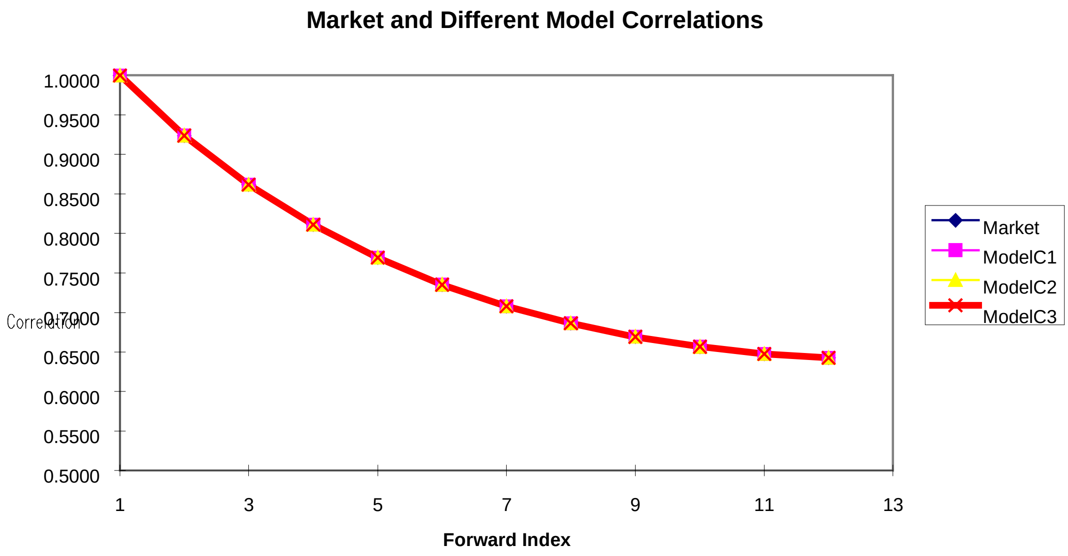
Market and Different Model Correlations
1.0000
0.9500
0.9000
0.8500
0.8000
0.7500
0.7000
0.6500
0.6000
0.5500
0.5000
Correlation
1
3
5
7
9
11
13
Forward Index
Market
ModelC1
ModelC2
ModelC3

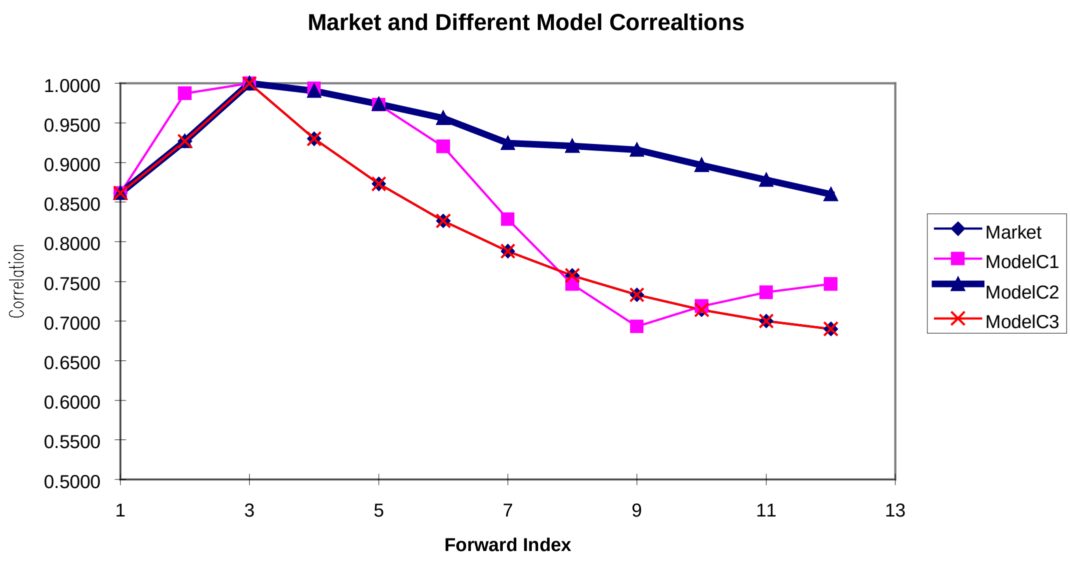
Market and Different Model Correaltions
1.0000
0.9500
0.9000
0.8500
0.8000
0.7500
0.7000
0.6500
0.6000
0.5500
0.5000
Correlation
1
3
5
7
9
11
13
Forward Index
Market
ModelC1
ModelC2
ModelC3

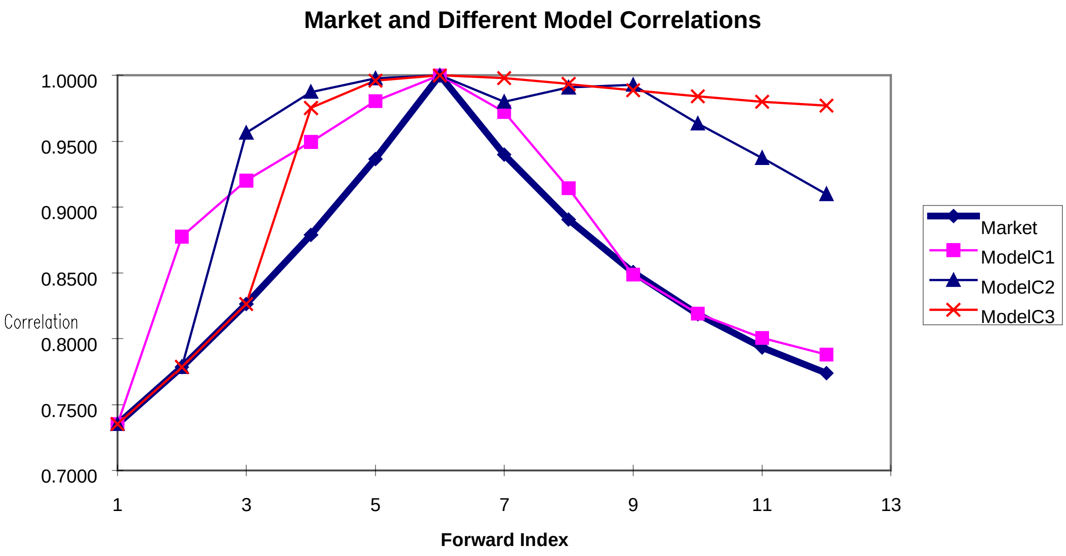

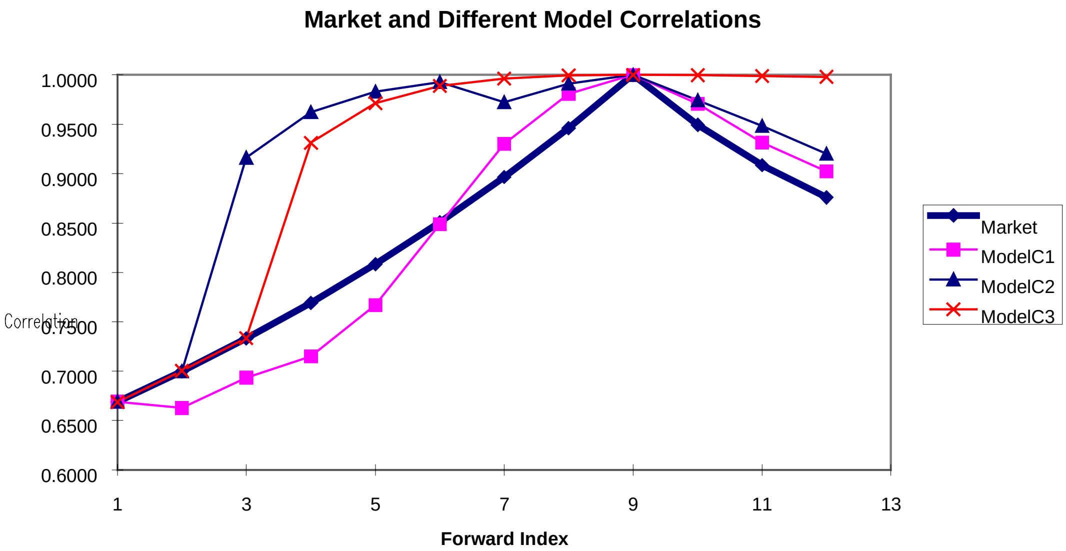

Fig 11 to 14: Comparison between the market and model correlations between the first, the third, the fifth and the eighth forward rates and all the other forwards obtained using 4 factors and imposing an exact fit to the first (ModelC1), the first and second (ModelC2), the first, second and third(ModelC3) column of the correlation matrix.

Even more interesting is to explore how the whole correlation matrix behaves when use is made of all the available degrees in order to fit to as closely as possible to the elements of the matrix itself which influence the value of a particular Bermudan swaption (a 9-non-call-2[8] was chosen for the example). Despite the fact that the fit to

[8] A N-non-call-J Bermudan swaption is a swaption with a maturity of N years, that cam be called or put on every reset date after year J

the desired portion of the correlation surface is successfully achieved (see Figs. 15 to 18), it is clear that

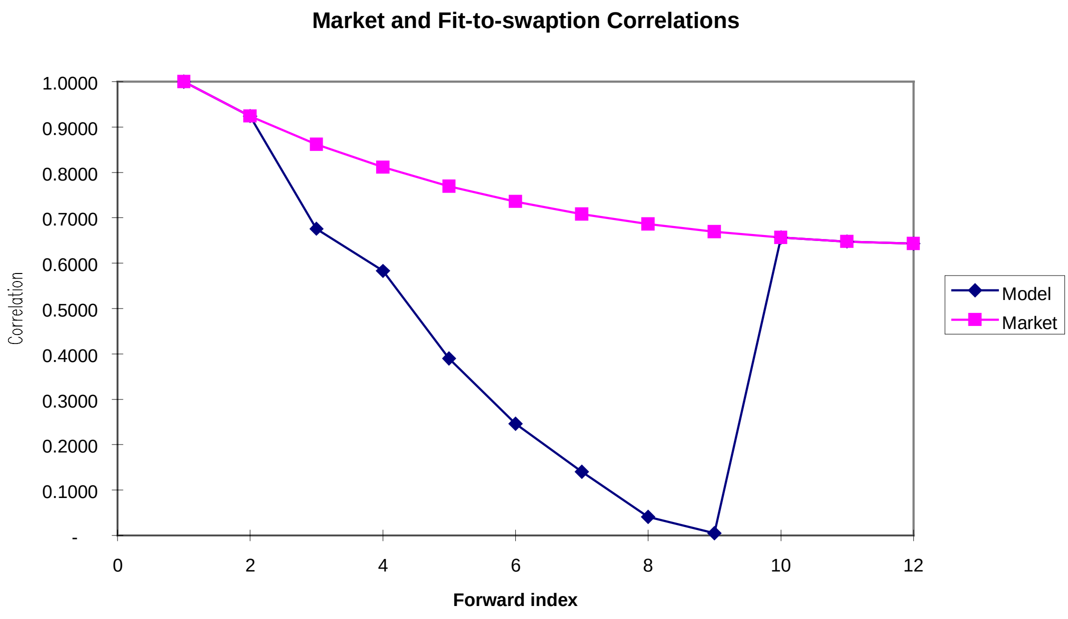

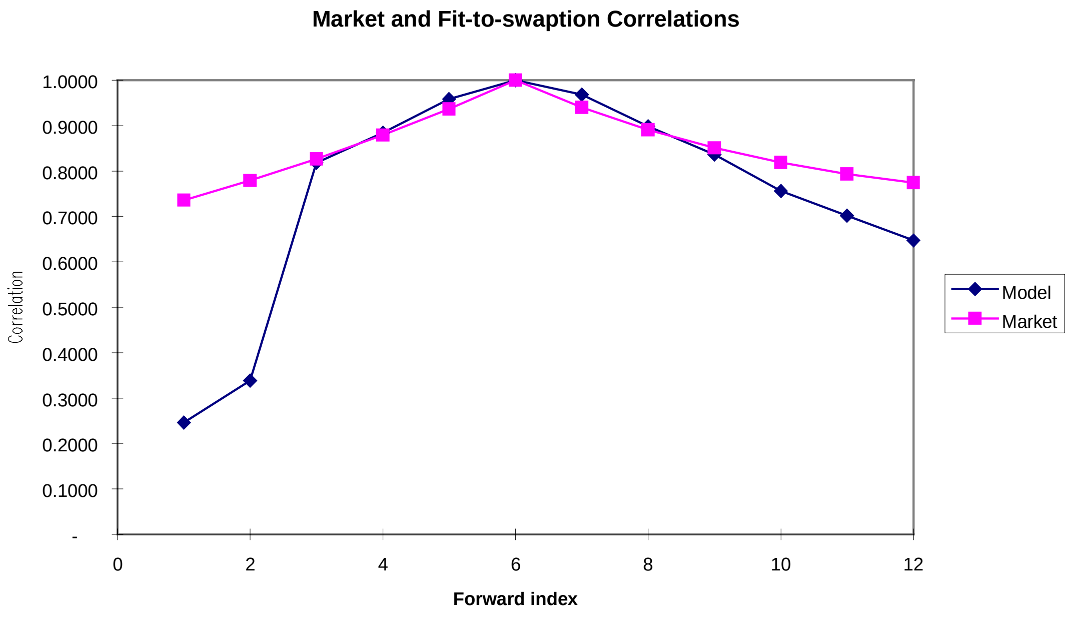

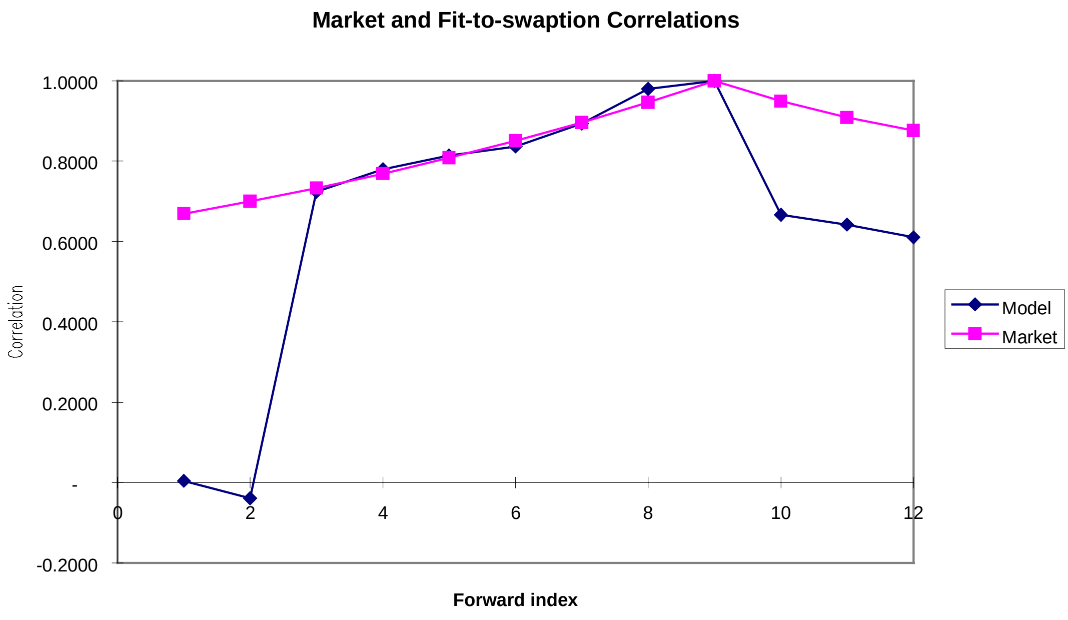

Fig 15 to 18: Comparison between the market and model correlations between the first, the third, the fifth and the eighth forward rates and all the other forwards obtained using 4 factors and imposing an best fit to the elements of the correlation matrix which affect the value of a 9NC5 Bermudan swaption.

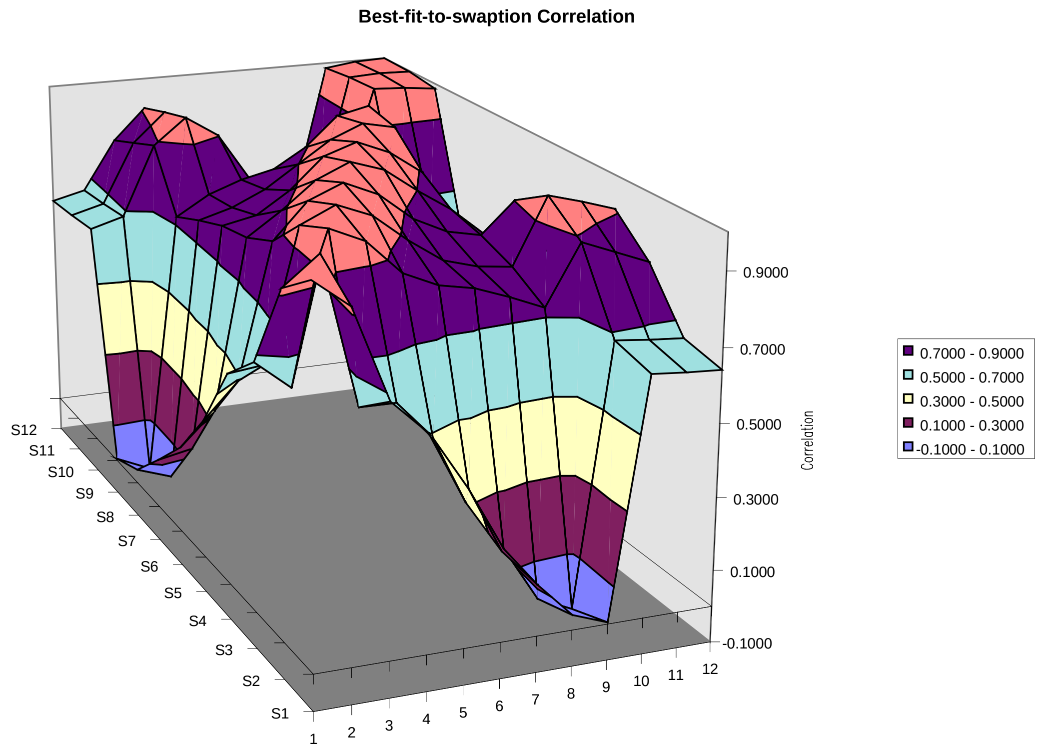

Fig 19: The model correlation surface for the 9NC5 Bermudan case.

the resulting overall surface is completely unrealistic. (See Fig. 19.) In addition, the resulting eigenvectors bear hardly any resemblance with the familiar results of PCA; this fact leaves the user without any intuitional insight about the possible interpretation of the factors as level, slope, curvature, etc. . The set of data pertaining to the Bermudan-swaption fit, in addition, caution strongly against using a low-dimensionality forward-rate-based model to price Bermudan swaptions by forcing *exact* pricing of a subset of the underlying European swaptions (perhaps the most valuable of the underlying European swaptions). Notice that, if the user could afford exactly to match *all* the underlying Euroepan swaptions, rather than a subset, then he would be in a position not to need any reduction in the dimensionality of the problem, which is the very motivation of this study, and of most practical applications. Swap-rate-based models, instead, can reproduce exactly the prices of all the European swaptions for *any* number of factors and *any* implied correlation (see Equations (4) and (5), and Rebonato (1998) for a detailed discussion of risk-neutral drifts in the swap-rate-based framework). It is therefore very dangerous, in so far as the overall correlation between different forward rates is concerned, to impose too strictly a good (or even perfect) fit to specific portions of the correlation matrix.

Along similar lines, one could explore the choice for the quality function of the sum squared errors along the tri-diagonals of the model and market correlation matrices. This choice could be seen as an attempt to reflect the importance of the correlation amongst contiguous forward rates in the pricing of instruments of the resettable-cap family (instruments, that is, where a stochastic strike on a given forward rate is determined by the reset of the immediately preceding forward rate). Without going into a detailed analysis, the resulting model correlation (see Fig. 21) matrix displays, once again, the shortcomings of imposing over-fitting to particular subsections of the correlation surface.

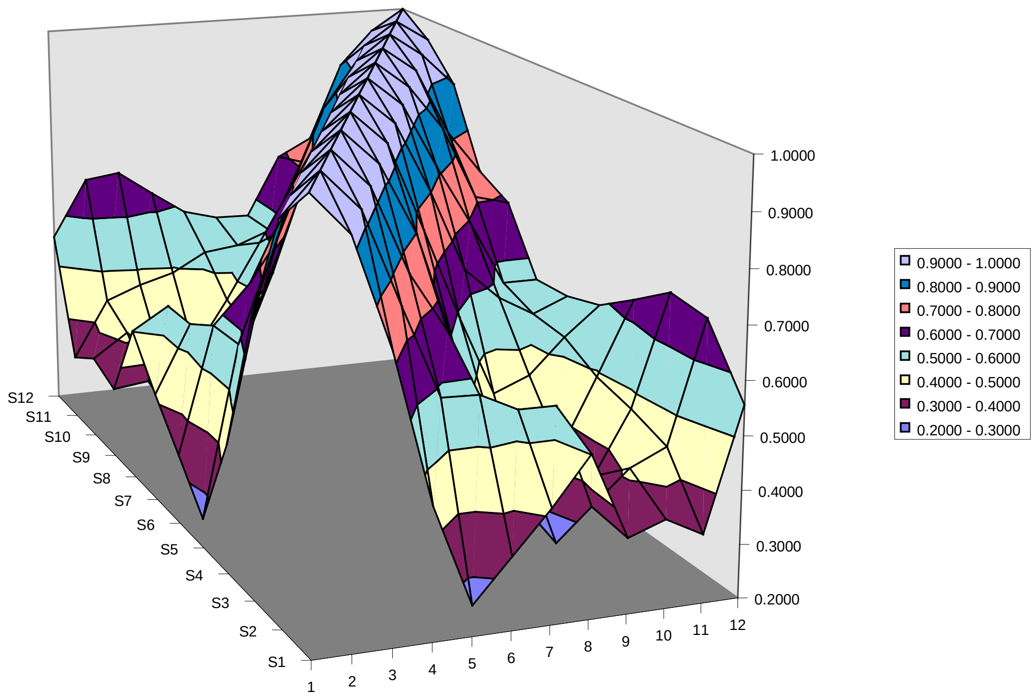

Fig 20: The model correlation surface for the resettable-cap case.

As is clearly visible from the figure above, the effect of imposing an (almost) exact correlation along the three main diagonals has the effect of imposing an unrealistic decorrelation amongst forward rates farther apart. This, in turn, will produce undesirable tilts and twists in the yield curve, that, after a sufficiently high number of resets, will ultimately have a strong effect on the price of the resettable product.

4 - Conclusions

This paper has shown that it is not only possible, but indeed expedient and advisable, to perform a *simultaneous* calibration of a log-normal BGM interest-rate model to the percentage volatilities of the individual rates *and* to the correlation surface. It has been shown that the task can be accomplished in two separate steps. More precisely, the first part of the calibration (i.e. to volatilities) can always be accomplished exactly thanks to straightforward geometrical relationships. After this step has been carried out, the fitting to the correlation surface, thanks to a simple theorem, can be efficiently carried out in such a way that the calibration to the volatilities is not spoiled by the second part of the procedure. The ability to carry out the two tasks separately greatly simplifies the overall task.

Actual calculations were then shown for a 3- and 4-factor implementation of the approach. Notwithstanding fundamental limitations of low-dimensioanlity models to recover the positive convexity at the origin of a market correlation function, the overall agreement between target and model correlation surfaces was shown to be very good.

Finally, the dangers of overparametrization, i.e. of forcing (near) exact fitting to certain portions of the correlation matrix, were highlighted by analysing the cases of a trigger swap, a Bermudan swaption and a one-way floater (resettable cap).

References

Black F, Derman E, Toy W, "A one-factor model of interest rates and its application to Treasury bond options", Fin. An. Jour., 1990, 33-339

Black F, Karasinski P, “Bond and option pricing when short rates are lognormal’, Fin. An. Journ, July-August 1991

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

Heath D, Jarrow RA, Morton A, "Bond pricing and the term structure of interest rates: a new methodology", working paper, 1989, (revised edition) , Cornell University

Rebonato R‘Interest Rate Option Models’ John Wiley Publ. (Chichester) (1998) Second Edition

Rebonato R, “On the pricing implications of the joint log-normality assumption for the cap and swaption markets”, submitted to Net Exposure (1998)

Rebonato R, Cooper I, ‘Limitations of simple two-factor interest rate models’ Journ Fin Engnin,5, 1-16, (1995)

Sibenius, J, “LIBOR Market Models in Practice” submitted to Journ. of Comp. Fin.