

Common Correlation and Calibrating the Lognormal Forward Rate Model

Carol Alexander, ISMA Centre, UK

Abstract

This paper examines some challenging problems for calibrating the lognormal forward rate (LFR) model. For interest rate options, the calibration of forward rate correlations to the swaption implied volatility surface requires careful selection of a parametric form, otherwise over-fitting can produce prices that are unstable over time. This paper introduces some useful new full-rank and reduced-rank correlation parameterizations that are very parsimonious. First two full-rank parameterizations that connect semi-annual with annual forward rate correlations are derived, and then the main focus of the paper is on a reduced rank parameterization which is based on the *common eigenvector* model of Flury (1988).

In the common eigenvector framework the same eigenvectors are used for all correlation matrices of the same dimension. Although the eigenvalues are fully time-varying and are specific to each tenor and maturity, the eigenvectors are constant and will be common to all swaptions of the same tenor. The 'trend', 'tilt', 'curvature' interpretation is shown to be a very accurate interpretation to place on these common

eigenvectors. Therefore only four parameters are required, for three orthogonal eigenvectors which can be calibrated to the market using, simultaneously, implied volatilities of different maturity swaptions.

Several academics and practitioners advocate using historical forward rate data when calibrating reduced rank correlations, and use the longer term swaption data for calibrating volatilities. We discuss the model risk that arises from using historical forward rate data, and the reasons why one should seek to use market data alone. However, lack of market data, particularly at the long end, means that a very parsimonious parameterization of correlation is required. Since the common eigenvector parameterization is much more parsimonious than individual eigenvector parameterizations, and is based simultaneously on swaptions of several maturities, it is likely to give more stable results. Common eigenvector parameterizations also have more general applications to multi-factor pricing models where forward correlations play an important role, e.g., for pricing and hedging OTC options in the power market.

1. Introduction

In 1997 three papers that introduced very similar lognormal diffusion processes for interest rates appeared virtually simultaneously. These models, now commonly called the 'Libor models' are based on either lognormal diffusions of forward rates as in Brace, Gatarek & Musiela (1997) and Miltersen, Sandmann & Sondermann (1997) or lognormal diffusions of swap rates, as in Jamshidian (1997). The consequent research interest in the calibration of these models has engendered a growing empirical literature, and in particular there are very many papers by Brigo and Mercurio, and Riccardo Rebonato. Further details are given in Brigo and Mercurio (2001, 2002a), Rebonato (1999b, 2002) and many of their papers are downloadable from www.fabiomercurio.it and www.damianobrigo.it and www.rebonato.com.

The Lognormal forward rate version of the Libor model is very useful for pricing and hedging exotic OTC products such as auto-caps and in-arrear swaps. However, the model is not easy to calibrate because its parameters include forward rate correlations that are difficult to forecast, particularly at the short end. To gain some idea of how these correlations change, Figure 1 shows the average daily correlations between semi-annual forward rates in the UK during four three month periods between June 2001 and June 2002. It is clear that changes in correlation are not as significant at the long end as they are at the short end. Figure 2 shows the correlation between the semi-annual forward rate starting in six months, and several longer maturity semi-annual forward rates starting in one year, 18 months, 2 years and so forth. Each line in Figure 2 represents the average daily correlation during a quarterly period. The correlations between the six month forward rate and

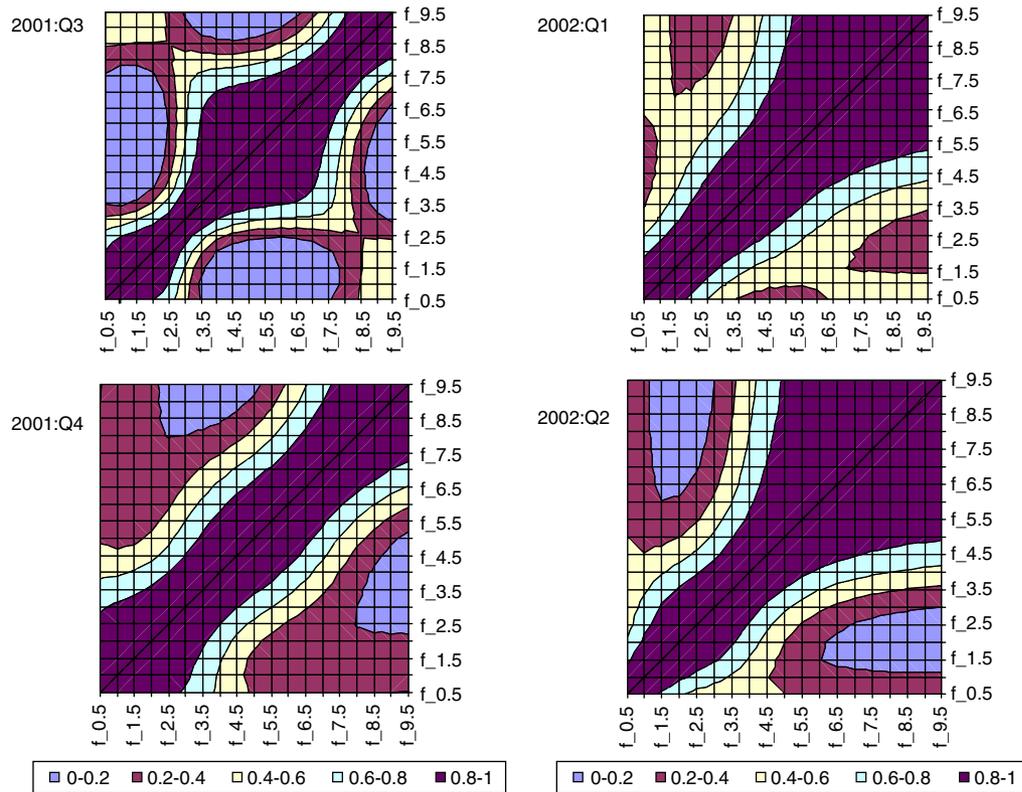


Figure 1: Semi-Annual Forward Rate Correlations. Semi-annual forward rates obtained by fitting the Svensson (1994) model to the UK yield curve. Unconditional correlation estimates for the quarter were based on daily first difference in log rates. Source: Alexander and Lvov (2003).

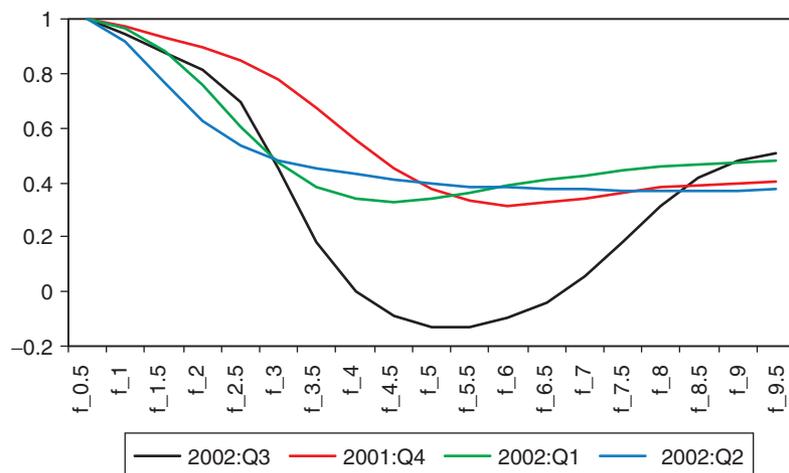


Figure 2: Correlations of Semi-Annual Forward Rates with Six Month Forward Rate.

the medium maturity forward rates are particularly unstable, and in the third quarter of 2001, when the terrorist attacks in the US occurred, they even became negative.

Based on the variability of the historical correlations observed in Figures 1 and 2, correlation forecasts must be difficult to obtain, particularly for short maturity forward rates. The forecasts have a great deal of uncertainty. This is one of many reasons why we should seek to use market data rather than historical data for calibrating correlations. The current swaption implied volatility surface captures market expectations of future correlation, and forecast uncertainty can be ignored if one uses instruments that are linear in the correlation. An at-the-money swaption is approximately linear in the forward rate volatility which is itself approximately linear in the factors of the forward rate correlation matrix. This correlation matrix should be calibrated to the swaption implied volatility surface, but one must be very careful how the factors of the correlation matrix are specified.

Many authors, including Rebonato (1999a), Rebonato and Joshi (2001) Hull and White (1999, 2000) and Longstaff, Santa-Clara and Schwartz (1999) advocate the use of historical data on forward rates for calibrating correlations. Long term swaptions are used for calibrating forward rate volatilities, since the cap market is typically only short to medium term, and so there is a lack of reliable data at the long end for calibrating also the correlations.

But historical correlations are not stable over time, and even if they were, there is another rather important problem with the use of historical data on forward rates. Long dated forward rates are unobservable. The market convention for stripping out long dated forward rates is straight-forward, and simply assumes a flat semi-annual rate between observed swap rates at annual maturities. Although this convention is necessary, for example to obtain cap prices from the Black cap volatilities that are quoted in the market, it is not informative of the instantaneous forward rate correlations in the Libor model.

If historical data on forward rates were used to calibrate these correlations, the forward rates would have to be obtained via a yield curve fitting model. Those in Figures 1 and 2 were obtained using the Svensson (1994) model, but there are many alternative yield curve models: for example see McCulloch (1975), Nelson and Siegel (1987) and Stealy (1991). Alexander and Lvov (2003) show that the forward rate correlations that are estimated from historical data can be quite different depending on the yield curve model chosen.¹ Therefore, the model risk arising from (a) the choice of yield curve model from which the forward rates are derived,

and (b) the choice of historical period over which observations are taken, may be considerable.

On the other hand, when trying to calibrate forward rate correlations to the swaption implied volatility surface, one encounters other problems. These are similar in spirit to the problems of calibrating local volatility surfaces to option implied volatilities. It is well-known that, if no parameterization for local volatility is used, as advocated by Dupire (1994), Derman, Kani and Chriss (1996) and Andersen and Brotherton-Ratcliffe (1997) amongst others, the local volatilities will be very sensitive to extrapolation and interpolation of implied volatilities and the volatility surface will lack robustness. Anyone who has tried to calibrate local volatilities directly from the implied surface will know that exotic path dependent option pricing models which depend directly on local volatilities do not always give stable results, even when regularization is used to introduce smoothness to the local volatility function as in Avellaneda et. al. (1997) and Bouchouev and Isakov (1997, 1999). So a functional form for local volatility can be chosen in order to calibrate local volatility to only the observed market data, which may be quite sparse, and to use the local volatility surface for hedging. Many different parsimonious functional forms for local volatility have been suggested in the literature, by Coleman, Li and Verma (1998), Dumas, Fleming and Whaley (1998), Brown and Randall (1999), McIntyre (2001), Brigo and Mercurio (2002b) and others. But the subjective choice of functional form is open to question: how should we determine which is the best choice?

The same type of problems arise when attempting to calibrate forward rate correlations to market data. Direct calibration can lead to unstable results, where correlations are also counter-intuitive, or even lie outside $[-1, 1]$. Therefore one should seek to impose a parametric form that extends to all the correlation matrices underlying all swaptions, which is parsimonious enough to give stable results but at the same time sufficiently general to capture the 'typical' behaviour of forward rate correlations—for example, where the correlation between adjacent forward rates increases with their maturity. Many different functional forms have been proposed and a useful review of these is given in Brigo (2001).

The aim of this paper is to derive some new parsimonious parameterizations of forward rate correlation matrices that enable the lognormal forward rate model to be calibrated to market implied swaption volatilities, without reverting to historical data for calibrating the correlations. These include full-rank parameterizations of the correlation matrices for semi-annual and annual forward rates, and reduced rank parameterizations that are based on the key concept of common principal components for the covariance matrices that relate to all swaptions.

The outline of this paper is as follows: Section provides a brief outline of the lognormal forward rate version of the Libor model and the calibration of forward rate volatilities. Section three proposes some full rank parameterizations of a semi-annual correlation matrix that are very parsimonious, involving only one or two parameters, and that

yield approximate correlations between annual forward rates which are determined by the same parameters. These are useful because they link the semi-annual cap market forward rate correlations with swap market forward rate correlations which, when the fixed leg has annual payments, are normally derived from annual rates.² Section four explains the orthogonal transformation of the LFR model that is commonly used for pricing path dependent products. Rebonato (1999a), Rebonato and Joshi (2001) Hull and White (1999, 2000) and Longstaff, Santa-Clara and Schwartz (1999) reduce the rank of forward rate covariance matrices by setting all but the three largest eigenvalues to zero. The implication of zeroing eigenvalues is a transformation of the lognormal forward rate model, where each forward rate is driven by three orthogonal factors that are derived from a principal component analysis (PCA). Section five introduces a new method for calibrating reduced rank correlation matrices to market data. Instead of using historical data on forward rates to estimate the three eigenvectors in the PCA, calibration to the market can be based on the common principal components model of Flury (1988). That is, the same eigenvectors, with only four parameters, are calibrated to all swaptions of the same tenor. The number of correlation parameters is very substantially reduced and the simultaneous use of market data on swaptions with a range of maturities will ensure that calibrated correlation matrices are stable. Section six summarizes the main results and concludes.

2. The lognormal forward rate model

To ease notation we assume that day counts are constant. That is, year-fractions between payment (or reset) dates are constant for all forward rates, and the basic forward rate is a semi-annual rate. Denote by f_i the semi-annual forward rate that is fixed at time t_i but stochastic up to that point in time. Each forward rate has its own 'natural' measure, which is the measure with numeraire P_{i+1} where P_i is the value of a zero coupon bond maturing at date t_i . Under its natural measure each forward rate is a martingale and therefore has zero drift in its dynamics. The log normal forward rate (LFR) model is thus:

$$df_i(t)/f_i(t) = \sigma_i(t)dW_i \quad [i = 1, \dots, m; 0 < t < t_i] \quad (1)$$

where dW_1, \dots, dW_m are Brownian motions with correlations $\rho_{ij}(t)$. That is,

$$E[dW_i dW_j] = \rho_{ij}(t)dt \quad (2)$$

Calibration of the model requires estimation of the parameters of the instantaneous volatilities $\sigma_i(t)$ and instantaneous correlations $\rho_{ij}(t)[i, j = 1, \dots, m]$.

Now consider a T maturity cap with strike K as a set of caplets from t_i to $t_{i+1}[i = 1, \dots, m - 1$ and $t_m = T]$. Each caplet pay-off = $\max(L_i - K, 0)$ where L_i is the LIBOR rate revealed at time t_i and at t_i we have $L_i = f_i$. Assume for the moment that $\sigma_i(t) = \sigma_i$. Then the Black-Scholes

type formula for each caplet value at time 0 is:

$$\text{B-S(Caplet, } t_i, K, \sigma_i) = P_{i+1} [f_i(0)\Phi(x) - K\Phi(y)] = C_i(\sigma_i), \quad \text{say}$$

$$\text{and} \quad x = [\ln(f_i(0)/K)/\sigma_i\sqrt{t_i}] + \sigma_i\sqrt{t_i}/2; \quad y = x - \sigma_i\sqrt{t_i}$$

The B-S cap price is the sum $\sum C_i(\sigma_i)$ of all prices of the caplets.

If we further assume that $\sigma_i = \sigma$ for every i then we can 'back-out' a single implied 'flat volatility' $\theta(K, T)$ from the observed market price of a cap with strike K and maturity T . Similarly, we can 'back-out' the implied caplet volatilities from the observed market prices of several caps of different maturities, using some interpolation between prices because each fixed strike caplet in a cap with strike K has a different moneyness. We normally define the 'at-the-money' (ATM) cap strike as the current value of the swap rate for the period of the cap, so the different caplets in an ATM cap are only approximately ATM.

Figure 3 shows the typical 'humped shaped' term structure for flat volatilities of caps. In the simple version of the LFR model, with $\sigma_i(t) = \sigma_i\forall i$, the implied forward rate volatilities are the caplet volatilities that are backed-out from interpolated caplet prices. These volatilities will display a more pronounced humped shape than the cap volatilities—an example is given in the Appendix, which re-examines the relationship between cap and caplet volatilities.

Caplet implied volatilities may be used for calibrating the LFR model (1) with time-varying but deterministic volatilities $\sigma_i(t)$ for the forward rates. It is standard to assume that $\sigma_i(t) = \eta_i h(t)$ where the time-varying parametric form $h(t)$ that is common to all volatilities is a simple 'hump' introduced by Rebonato (1999a) and since used by many others, given by:

$$h(t) = [(a + b(T - t)) \exp(-c(T - t)) + d] \quad (4)$$

Note that $h(t)$ —and therefore also $\sigma_i(t)$ —depends on maturity T and the parameters a, b, c, d which define a common volatility structure. The individual parameters η_i are there so that instantaneous volatilities can be adjusted upwards or downwards according to the prices in the cap market. The parameters may be calibrated to implied caplet volatilities, θ_i

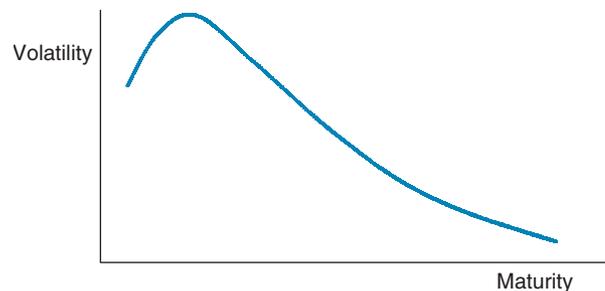


Figure 3: Typically 'humped shaped' term structure of flat cap implied volatility: $\theta(T)$ vs T .

and a simple calibration objective is:

$$\min \sum_i \omega_i \left(\theta_i \sqrt{t_i} - \eta_i \sqrt{\int_0^{t_i} h(t)^2 dt} \right)^2$$

The weights ω_i are there to account for the uncertainty that surrounds different option prices and vegas. To place more weight on the more certain volatilities, the weights should be directly proportional to vega and inversely proportional to the bid-offer spread.³ The weights should also be a decreasing function of maturity (and tenor, in the case of a swaption), and since gamma is directly proportional to vega and decreases with maturity, we suggest taking the weights as gamma divided by the bid-ask spread.

With so many parameters, of course the model could fit the market prices of caps perfectly, and there is a danger in over-fitting as shown by De Jong, Dreissen and Pelsser (1999). Moreover, since the usefulness of the LFR lies in pricing exotic OTC products, we shall also need to calibrate forward rate correlations, and this cannot be done using cap market data alone—unless one is happy to use historical forward rate data, which is prone to the difficulties outlined in section one.

3. Calibration to the swaption market

Unlike caps, a swaption pay-off [$\max(\text{value of swap}, 0)$] cannot be written as a simple sum of options, so forward rate correlations as well as their volatilities will affect the value of a swaption. We need to work in a single measure, but forward rates are only drift-less under their natural measure. Hull and White (1999, 2000) take as numeraire the discretely reinvested money market account. This gives the 'spot labor measure', where the appropriate rate for discounting an expected cash flow at time t_{i+1} to time t_i is the forward rate f_i . This is intuitive and it leads to a (relatively) tractable specification of the drift terms, so we shall adopt this here.⁴ Under the spot labor measure the forward rate f_i has dynamics given by

$$\left. \begin{aligned} df_i(t)/f_i(t) &= \mu_i(t)dt + \sigma_i(t)dW_i \\ \mu_i(t) &= \sigma_i(t) \sum_{j=m(t)}^i \frac{\rho_{ij}(t)\sigma_j(t)f_j(t)}{1+f_j(t)} \end{aligned} \right\} \quad (5)$$

where $m(t)$ is the 'number' of the accrual period.

The drift becomes important when using the LFR model to price path dependent options, where the resolution method will, most likely, be Monte Carlo simulation, and the drift will need updating every time step. We shall return to this in the next section, but in this section we are considering the calibration of forward rate volatilities and correlations to swaption volatilities, where the drift is not important.

We now derive some new approximations for the instantaneous volatility of an annual swap rate in terms of the instantaneous volatilities and correlations of the semi-annual forward rates underlying the swap.

These approximations allow the calibration of instantaneous volatilities and correlations for semi-annual rates to the average swap rate volatilities that are implied from the market, so that historical data on forward rates are not required. The results are useful in the US market when the fixed leg is based on annual forward rates.

Following Rebonato (1998) we write the swap rate as an approximate linear function of annual forward rates. For an m -year swap starting at t_n (that is, an “ n into m year” swap) we have the swap rate:

$$SR_{n,m} = w_1 F_n + \dots + w_m F_{n+m-1}$$

where F_i is the annual forward rate starting at year i and the weights $w_i = P_{n+i}/[P_{n+1} + \dots + P_{n+m}]$ are assumed constant at their current value.⁵ Consequently the variance of the swap rate may be written as a quadratic form in the 3×3 covariance matrix of annual forward rates and thus the swap rate volatilities are linked to annual forward rate volatilities and correlations.

To see this, write $\mathbf{w} = (w_1, \dots, w_m)'$ and let $\mathbf{V}_{n,m}(t)$ be the instantaneous covariance matrix of the forward rates F_n, \dots, F_{n+m-1} for $0 < t < t_n$. Then $SR_{n,m}$ has instantaneous variance given by

$$\left. \begin{aligned} \sigma_{n,m}(t)^2 &= \mathbf{w}' \mathbf{V}_{n,m}(t) \mathbf{w} \\ \text{and the average swap volatility during the interval } [0, t_n] \text{ is} \\ \sigma_{n,m}^{\text{model}} &= \sqrt{\frac{1}{t_n} \int_0^{t_n} \sigma_{n,m}(t)^2 dt} \end{aligned} \right\} \quad (6)$$

We know the weights \mathbf{w} from current values of discount bonds and we want to infer the parameters of $\mathbf{V}_{n,m}(t)$ by equating the model volatilities that are defined by equations (6) to the smoothed⁶ implied swaption volatilities $\theta_{n,m}$. The covariance matrix of annual forward rates $\mathbf{V}_{n,m}(t)$ contains instantaneous volatilities and correlations of annual rates. Indeed it may be factored as

$$\mathbf{V}_{n,m}(t) = \mathbf{D}_{n,m}(t) \boldsymbol{\Sigma}_{n,m}(t) \mathbf{D}_{n,m}(t) \quad (7)$$

where $\mathbf{D}_{n,m}(t)$ is the diagonal matrix of the instantaneous volatilities and $\boldsymbol{\Sigma}_{n,m}(t)$ is the correlation matrix of the m annual forward rates F_n, \dots, F_{n+m-1} underlying the n into m year swap.

It is simple to express the i th annual forward rate variance $\sigma_{i_ann}(t)^2$ in terms of the two underlying semi-annual rate variances $\sigma_{2i-1}(t)^2$ and $\sigma_{2i}(t)^2$. Since

$$(1 + F_i(t)) = (1 + f_{2i-1}(t)/2)(1 + f_{2i}(t)/2)$$

we have $F_i(t) = (f_{2i-1}(t) + f_{2i}(t))/2 + f_{2i-1}(t)f_{2i}(t)/4$. With this identity, and assuming lognormal dynamics for the semi-annual rates, it is shown in Brigo and Mercurio (2001) that

$$\begin{aligned} \sigma_{i_ann}(t)^2 &= u_{2i-1}(t)^2 \sigma_{2i-1}(t)^2 + u_{2i}(t)^2 \sigma_{2i}(t)^2 \\ &\quad + 2u_{2i-1}(t)u_{2i}(t)\rho_{2i-1,2i}\sigma_{2i-1}(t)\sigma_{2i}(t) \end{aligned} \quad (8)$$

where $\rho_{2i-1,2i}$ is the correlation between the two semi-annual rates and

$$u_{2i}(t) = 1 - [f_{2i-1}(t)/2F_i(t)] \quad \text{and} \quad u_{2i-1}(t) = 1 - [f_{2i}(t)/2F_i(t)];$$

Brigo and Mercurio (2001) assume $u_{2i}(t) = u_{2i}(0)$ and $u_{2i-1}(t) = u_{2i-1}(0)$ for all time t , and use the consequent approximation based on (8) to connect annual (swaption) volatilities to semi-annual (caplet) volatilities.

The point to note is that the correlation of adjacent semi-annual forward rates also enters the approximation based on (8). Therefore, if such an approximation is used with (7) for calibration to the swaption market, the semi-annual correlation parameters will also enter $\mathbf{D}_{n,m}(t)$. However we can use some further assumptions to express $\boldsymbol{\Sigma}_{n,m}(t)$ in terms of semi-annual rate correlations only. In that case all the parameters in $\mathbf{V}_{n,m}(t)$ will relate only to semi-annual rates and not to annual rates.

Some researchers, including Brigo and Mercurio (2001) and Rebonato and Joshi (2001), assume that the annual rate correlation matrix is $\boldsymbol{\Sigma}_{n,m} = \{\rho_{i,j}\}$ where

$$\rho_{i,j} = \exp(-\varphi|i-j|) \quad (9)$$

Variants of this, such as $\rho_{i,j} = \exp(-\varphi|i^{\gamma} - j^{\gamma}|)$ are discussed in Brigo (2001) but these parameterizations do not lend themselves to a straightforward relationship between the semi-annual rate correlations and the annual correlations. We write $\rho = \exp(-\varphi)$, so that the assumption (9) may be written $\rho_{i,j} = \rho^{|i-j|}$ and, if we use the parameterization (9) for semi-annual forward rates, we now derive a simple approximate relationship between semi-annual correlations and annual correlations.

Let ρ be the correlation between any two adjacent semi-annual forward rates, and assume that the correlation between two semi-annual rates f_i and f_j is $\rho^{|i-j|}$ so that the $2m \times 2m$ correlation matrix of semi-annual forward rates underlying an n into m year swaption, denoted $\boldsymbol{\Sigma}_{n,m_semi}$, is given by the *circulant* correlation matrix:

$$\begin{pmatrix} 1 & \rho & \rho^2 & \dots & \dots & \rho^{2m-1} \\ \rho & 1 & \rho & \rho^2 & \dots & \rho^{2m-2} \\ \rho^2 & \rho & 1 & \rho & \dots & \rho^{2m-3} \\ \vdots & & & \vdots & & \\ \rho^{2m-1} & & & & \dots & 1 \end{pmatrix}$$

Then, if the two semi-annual forward rates f_{2i-1} and f_{2i} in the annual forward rate F_i have equal volatilities (though not necessarily the same volatilities for each i) the $m \times m$ correlation matrix of the annual forward rates underlying the same swaption, denoted $\boldsymbol{\Sigma}_{n,m}$, is approximated by the single parameter full rank correlation matrix:

$$\begin{pmatrix} 1 & \phi & \rho^2 \phi & \dots & \rho^{2(m-2)} \phi \\ \phi & 1 & \phi & \rho^2 \phi & \dots & \rho^{2(m-3)} \phi \\ \rho^2 \phi & \phi & 1 & \phi & \dots & \rho^{2(m-4)} \phi \\ \vdots & & & \vdots & & \\ \rho^{2(m-2)} \phi & & & & \dots & 1 \end{pmatrix}$$

where $\phi = \rho(1 + \rho)/2$. That is, the correlation between two annual forward rates F_i and F_j is approximately $\rho^{2|i-j|-1}(1 + \rho)/2$.

To prove this we must assume that the two semi-annual forward rates f_{2i-1} and f_{2i} in the annual forward rate F_i have equal volatilities, though not necessarily the same volatilities for each i . Denote the (equal) volatility of f_{2i-1} and f_{2i} by $\sigma_{2i}(t)$. We must also introduce a further approximation, that $u_j(t) = u$, a constant, for all time t and for all maturities j . Then the annual forward rate dynamics are approximately

$$dF_i(t)/F_i(t) \approx (\dots)dt + u\sigma_{2i}(t)(dW_{2i-1}(t) + dW_{2i}(t))$$

where the two Brownian motions driving the semi-annual rates are correlated:

$$E[dW_{2i-1}dW_{2i}] = \rho_{2i-1,2i}dt$$

This gives an approximation for the annual forward rate variance:

$$\sigma_{i_ann}(t)^2 \approx 2u^2\sigma_{2i}(t)^2.(1 + \rho_{2i-1,2i})$$

and an approximation for the covariance between the i th and the j th annual rates:

$$\sigma_{ij_ann}(t) \approx u^2\sigma_{2i}(t)\sigma_{2j}(t)(\rho_{2i-1,2j-1} + \rho_{2i-1,2j} + \rho_{2i,2j-1} + \rho_{2i,2j})$$

Now, when the semi-annual rate correlation matrix Σ_{n,m_semi} takes the circulant form given above we have:

$$\sigma_{i_ann}(t)^2 \approx 2u^2\sigma_{2i}(t)^2.(1 + \rho)$$

and

$$\sigma_{ij_ann}(t) \approx u^2\sigma_{2i}(t)\sigma_{2j}(t)\rho^{2|i-j|-1}(1 + \rho)^2$$

so the result follows.

Figure 4 illustrates the correlations obtained using the circulant structure with $\rho = 0.9$. Although this leads to a simple form for annual rate correlations, it is a little inflexible. In particular it does not allow for the correlation of adjacent forward rates to increase with maturity, as has been observed recently in some markets (see Alexander and Lvov, 2003). Schoenmakers and Coffey (2000) point out the importance of allowing for this, and consequently propose a full-rank semi-parametric form based on an increasing sequence of M real numbers. In fact the increasing correlation of adjacent forward rates with maturity may be more parsimoniously captured by a two parameter correlation matrix for semi-annual rates, as shown in Figure 5.

Suppose the $2m \times 2m$ correlation matrix Σ_{n,m_semi} of semi-annual forward rates underlying an n into m year swaption is given by the two parameter correlation matrix (ρ_{ij}) where:

$$\rho_{ij} = \psi^{j-1}\rho^{|i-j|}i < j, 0 < \rho < \text{ and } \rho < \psi^{2m-1}\rho < 1.$$

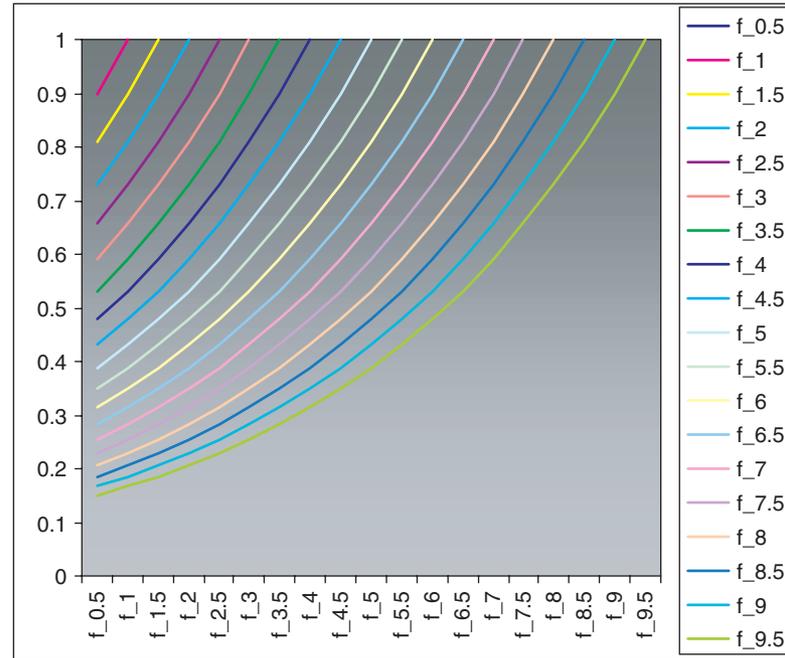


Figure 4: Semi-Annual Correlations with the Circulant Form [$\rho = 0.9$].

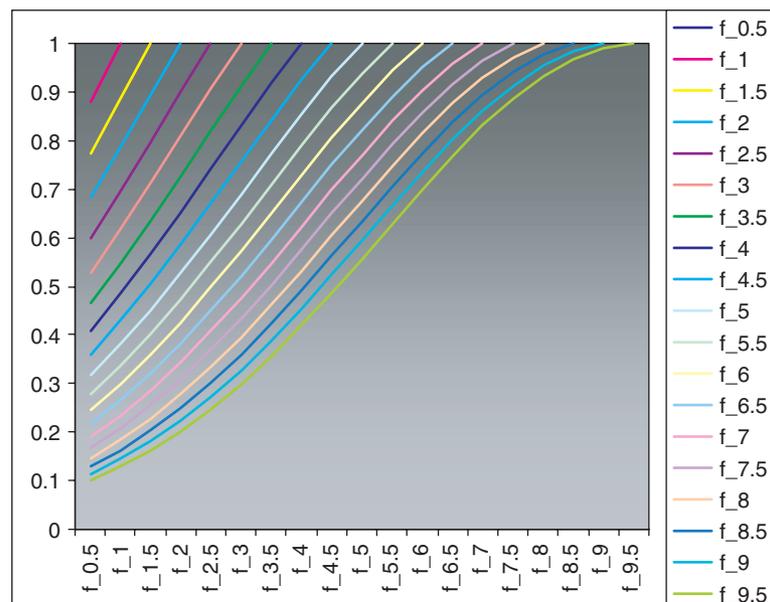


Figure 5: Semi-annual Correlations with a Two Parameter Form [$\rho = 0.88, \psi = 1.007$].

In Figure 5, the case $m = 10$, $\rho = 0.9$, $\psi = 1.007$ is illustrated. Then the $m \times m$ correlation matrix of the annual forward rates underlying the same swaption, denoted $\Sigma_{n,m}$ is approximated by the correlation matrix (ρ_{ij}) where:

$$\rho_{ij} = \frac{\rho^{2|i-j|-1} \psi^{2j-1} (1 + \rho)(1/\psi + \rho)}{2\sqrt{(1 + \psi^{2(i-1)}\rho)(1 + \psi^{2(j-1)}\rho)}}$$

This result follows after some algebra using similar assumptions and methods as for the circulant parameterization.

Thus each correlation matrix $\Sigma_{n,m}(t)$ is now defined in terms of at most two parameters, ρ and ψ . Also, by (8), the diagonal matrix of the annual forward rate volatilities $\mathbf{D}_{n,m}(t)$ is expressed in terms of the semi-annual forward rate volatilities and these same parameters. Thus from (7), the only parameters in the annual rate covariance matrix $\mathbf{V}_{n,m}(t)$ are: a, b, c, d, η, ρ and ψ . Note that we may suppose that η is itself parameterized in terms of a, b, c , and d , serving only to equate model and market prices for caps, as advocated by Brigo and Mercurio (2001). In that case only five or six parameters need to be determined by the swaption data: these are a, b, c and d , for the volatilities; ρ and possibly also ψ for the correlations. If there are sufficient market data to calibrate more parameters, one could assume time-varying parametric forms for ρ and ψ , or that ρ and ψ depend on maturity and tenor so they are different for each $\mathbf{V}_{n,m}(t)$. However this should be attempted with caution—too many correlation parameters may lead to results that are not sufficiently stable as the implied volatility surfaces changes from day to day.

4. Calibration for pricing path dependent interest rate options

The LFR model may be used to price path dependent options, but no analytic pricing formulae exist and some sophisticated resolution methods, often Monte Carlo simulation, are used. Whichever method is chosen the drifts that are induced by the change of numeraire will be important, and they will need updating with the current value of the instantaneous volatilities and correlations at every time step. The instantaneous drift term in the i th forward rate f_i depends on the forward rates of lower maturity f_{i-1}, f_{i-2} , etc., but only if these are still random variables at the time that the drift is estimated. The correlation matrix that one needs to consider for the calibration of the drift term therefore decreases in dimension (by 1) as every payment/reset date passes. This is not the only complication for the simulation. If m is large the dimension of the simulation will be large and computationally burdensome. However it is possible to reduce dimensions by reducing the rank of the correlation matrix. A review of rank reduction methods is given in Brigo (2001).

A standard method for rank reduction is advocated by Rebonato (1999c) and Hull and White (1999, 2000) and many others. They use an orthogonal transformation of the correlated Brownian motions in (5).

The forward rate dynamics are expressed in terms of three uncorrelated stochastic processes that are common to all forward rates:

$$df_i(t)/f_i(t) = \mu_i(t)dt + \lambda_{i,1}(t)dZ_1 + \lambda_{i,2}(t)dZ_2 + \lambda_{i,3}(t)dZ_3 \quad (10)$$

where dZ_1, dZ_2, dZ_3 are uncorrelated Brownian motions and:

$$\sigma_i(t)dW_i = \lambda_{i,1}(t)dZ_1 + \lambda_{i,2}(t)dZ_2 + \lambda_{i,3}(t)dZ_3 \quad (11)$$

From this it follows that:

$$\sigma_i(t) = \sqrt{[\lambda_{i,1}(t)^2 + \lambda_{i,2}(t)^2 + \lambda_{i,3}(t)^2]} \quad (12)$$

and

$$\rho_{ij}(t) = [\lambda_{i,1}(t)\lambda_{j,1}(t) + \lambda_{i,2}(t)\lambda_{j,2}(t) + \lambda_{i,3}(t)\lambda_{j,3}(t)]/\sigma_i(t)\sigma_j(t) \quad (13)$$

Thus the forward rate volatilities and correlations are completely determined by three volatility ‘components’ for each forward rate, which are $\lambda_{i,1}(t)$, $\lambda_{i,2}(t)$ and $\lambda_{i,3}(t)$ for the i th forward rate. Denote by $\sigma_i(t)$ the instantaneous i th annual forward rate volatility. Given these, the volatility components may then be determined from the spectral decomposition of the covariance matrix of the first difference in the logarithm of the m annual forward rates F_n, \dots, F_{n+m-1} underlying the n into m year swap.⁷ The spectral decomposition is

$$\mathbf{V}_{n,m}(t) = \mathbf{A}_{n,m} \mathbf{\Lambda}_{n,m}(t) \mathbf{A}'_{n,m}$$

where $\mathbf{\Lambda}_{n,m}(t)$ is the diagonal matrix of eigenvalues and $\mathbf{A}_{n,m}$ is the $m \times m$ matrix of eigenvectors of $\mathbf{V}_{n,m}(t)$. Note that the eigenvalues are time-varying whereas the eigenvectors are assumed constant. To derive the volatility components, denote by $\Lambda_1(t), \Lambda_2(t), \Lambda_3(t)$, the three largest eigenvalues of $\mathbf{V}_{n,m}(t)$ and denote their eigenvectors by $\alpha_1, \alpha_2, \alpha_3$. Set $= \sigma_i(t)/\sqrt{V_i(t)}$ where

$$V_i(t) = \alpha_{i,1}^2 \Lambda_1(t) + \alpha_{i,2}^2 \Lambda_2(t) + \alpha_{i,3}^2 \Lambda_3(t)$$

Then, to satisfy (12), simply set the volatility components to be:

$$\lambda_{i,1}(t) = M(t)\alpha_{i,1}\sqrt{\Lambda_1(t)}$$

$$\lambda_{i,2}(t) = M(t)\alpha_{i,2}\sqrt{\Lambda_2(t)}$$

$$\lambda_{i,3}(t) = M(t)\alpha_{i,3}\sqrt{\Lambda_3(t)}$$

5. A common principal components model

A number of researchers, notably Rebonato (1999a), Rebonato and Joshi (2001) Hull and White (1999, 2000) and Logstaff, Santa-Clara and Schwartz (1999) have advocated the use of volatility components derived from eigenvectors that are estimated using historical forward rate data, and eigenvalues that are calibrated to the market. That is, calibration of the volatility

components in (11) is based on a principal component analysis (PCA) of historical data on forward rates. However, following the remarks made in section one, a LFR model calibration that is based on PCA of historical forward rates will give prices that depend on (a) which forward rates are used, and (b) the length of the historical period chosen for the PCA.

A preferred approach would be to parameterize the eigenvectors of the correlation matrices in the natural way: the first eigenvector should be constant and have only one parameter which represents a parallel shift in the forward curve, the second eigenvector will have two parameters and represent a tilt in the forward curve, and the third eigenvector will have three parameters and represents the curvature of the forward curve. The orthogonality of these vectors induces certain restrictions and reduces the number of parameters required. For example, the orthogonality between the constant and tilt (parameterized as $c_1 + c_2k$, say, for $k = 1, 2, \dots, m$) implies that $c_2 = -2c_1/(m + 1)$. In fact three orthogonal trend-tilt-curvature eigenvectors can be parameterized by only four parameters.

Nevertheless, this construction still gives rather too many parameters to be calibrated: four parameters for the eigenvectors of every covariance matrix $\mathbf{V}_{n,m}(t)$ ⁸ and there are nm of these matrices! With so many parameters, the current swaption implied volatility surface will of course be fit very closely—but over fitting will lead to prices that are too sensitive and unstable over time.

How, then, should the parameters be reduced? One possibility is to deduce a parsimonious representation of the eigenvectors and eigenvalues from a full-rank parameterization such as those given in section 2. Another possibility, described here, is to use the *same* eigenvectors for all correlation matrices of the same dimension. This is a special case of the common principal components model introduced by Flury (1988).

Alexander and Lvov (2003) show that the implementation of the common eigenvector tests introduced by Flury (1988) on historical forward rate data yields results that are very robust indeed: to different sets of forward rates (i.e. those that are obtained from several different yield curve models), and to the choice of historic data period. Their results are that there is very strong evidence for common eigenvectors in correlation matrices of the same dimension, for both annual and semi-annual forward rates, and moreover these common eigenvectors are very close indeed to the archetypal ‘trend, tilt, curvature’ form—much closer than the eigenvectors obtained from a single covariance matrix.

The common eigenvector parameterization of the covariance matrix is

$$\mathbf{V}_{n,m}(t) = \mathbf{A}_m(t)\mathbf{\Lambda}_{n,m}(t)\mathbf{A}_m'(t)$$

where $\mathbf{\Lambda}_{n,m}(t)$ is the 3×3 matrix of the three largest eigenvalues of $\mathbf{V}_{n,m}(t)$ and $\mathbf{A}_m(t)$ is a matrix of eigenvectors, scaled up from the common eigenvectors of the corresponding correlation matrix using the forward rate volatilities (see footnote 8), where each common eigenvector of the correlation matrices has only four parameters. Moreover these parameters are the same for all n . The strength of the common eigenvector approach is that it substantially reduces the number of correlation parameters, but

at the same time allows the eigenvectors to be time-varying, so that one can still model the instantaneous volatilities in the flexible manner.

Without common eigenvectors, one needs to calibrate four parameters (for three orthogonal eigenvectors) to each swaption smile of maturity n and tenor m . For longer maturity, not enough smile data will be available, and this may be one reason why several authors favour using historical data to estimate these eigenvectors, instead of calibrating them to the market. However, with common eigenvectors \mathbf{A}_m has only four correlation parameters and applies to all swaptions of tenor m . The total number of eigenvector parameters is therefore reduced by a factor of n , from $4mn$ to $4m$, where m is the number of different tenors of the swaptions used to calibrate the model. Since swaptions with tenors of more than 5 or 6 years are often illiquid, and since those with tenors of less than three years have too small dimension for three eigenvectors, we seek to calibrate only \mathbf{A}_3 , \mathbf{A}_4 , and \mathbf{A}_5 where each matrix has only four correlation parameters.

The great advantage of using common PCA to parameterize forward rate correlations is that the common eigenvectors can be calibrated to the swaption market as constants, whilst the forward rate volatilities and the eigenvalues of the covariance matrix can be fully time varying and specific to each tenor and maturity, and these can be calibrated to the

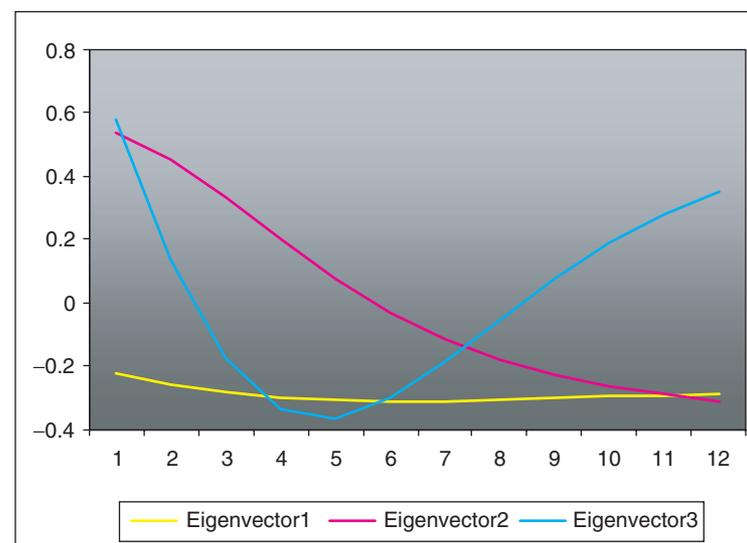


Figure 6: Parameterizing Common Eigenvectors.

The figure shows the three largest common eigenvectors for all semi-annual forward rate correlation matrices of dimension 12; the data are based on Svensson yield curve fitting for UK rates between Jan 1st and Dec 31st 2002 and a total of eight matrices were diagonalized simultaneously in the common principal component algorithm. The eigenvalues are different for each matrix, but they indicate that in every case these three common principal component explain over 99% of the variation of forward rates during 2002. See Alexander and Lvov (2003) for further details.

cap and/or swaptions market. Moreover we have reason to believe that parameterized common eigenvectors are likely to fit market data better than parameterized individual eigenvectors. This is because when common principal components are estimated using historical forward rate data we find that it is definitely appropriate to parameterize common eigenvectors of the correlation matrices so that their associated principal components represent almost exactly the orthogonal trend-tilt-curvature movements in the forward rate curve. The common eigenvectors shown in Figure 6, for example, certainly justify the use of the parsimonious common eigenvector parameterization described in this section.

6. Summary and conclusion

In the lognormal forward rate option pricing model, forward rate volatilities and correlations may be calibrated to both cap volatility term structures and a smoothed swaption volatility surface. Often this requires linking annual forward rate correlations with semi-annual forward rate correlations, for example in the US swap market when the fixed leg is paid annually. The first part of this paper has introduced some parsimonious parametric forms for the semi-annual forward rate correlations which imply approximations for the annual forward rate correlations that depend on the same parameter(s).

Pricing of exotic path dependent interest rate options normally involves numerical resolution (e.g. simulation) for the forward rate dynamics, where the drift terms that are induced by the change to a single measure will become important. Since these change at every time step, the computations are burdensome and it is desirable to reduce dimensions in the numerical method. To achieve this, several authors have used principal component analysis in conjunction with historical time series on forward rate data for calibrating the eigenvectors of the forward rate covariance matrix.

However historical data on the unobservable long maturity forward rates are subject to substantial model risk arising from the choice of yield curve model, the choice of historical observation period, and the statistical model used to estimate the correlations. Therefore the final part of this paper introduces an orthogonal transformation based on common eigenvectors of the forward rate correlation matrices which allows calibration to swaption market data alone. The parametric form for common eigenvectors is based on the common principal component model of Flury (1988), and the resultant parsimony allows calibration of these common eigenvectors simultaneously to all swaptions of the same tenor.

Appendix

6 ‘Rule-of-Thumb’ relationship between cap and caplet volatilities

It is interesting to consider the relationship between the flat volatilities quoted in the market for caps of different maturities and strikes and the implied volatilities of the caplets that make up these caps. In option markets where the underlying is a single asset whose dynamics are governed

by a standard lognormal geometric Brownian motion, the variance of log returns over a period $T = T_1 + T_2$ is the sum of the variance of log returns over period T_1 and the variance of log returns over period T_2 . The additivity of variances is a consequence of the independent increments in the stochastic process for a single underlying asset, and because variances are additive implied volatilities $\{\theta_1, \theta_2, \dots, \theta_T\}$ can be derived from a term structure $\theta(T)$ using the iterative method: Set $\theta_1 = \theta(1)$; then solve for θ_2 from $[\theta_1^2 + \theta_2^2]/2 = \theta(2)^2$ and so forth.

However, for a cap there is no single stochastic process for the underlying. The underlying forward rate changes for every caplet in the cap and therefore an iteration on variances is not appropriate. In fact the iteration should be performed on the cap volatilities. Denote by $\theta(K, T)$ the market quote of a ‘flat’ implied volatility for a cap of strike K and maturity T . Denote by $\theta_i(K, T)$ the Black-Scholes implied volatility and by $v_i(K, T)$ the vega of the i th caplet in the cap of strike K and maturity T . Thus $v_i(K, T) = \partial C_i(\sigma_i)/\partial \sigma_i$ evaluated at $\sigma_i = \theta_i(K, T)$. Then

$$\theta(K, T) \approx \Sigma[v_i(K, T)\theta_i(K, T)]/\Sigma v_i(K, T) \quad (A)$$

That is, the flat cap implied volatility $\theta(K, T)$ is approximately equal to the vega weighted sum of the caplet implied volatilities.

To explain the derivation of (A), for ease of notation we shall drop the explicit mention of the dependence of implied volatilities on strike and maturity. For a fixed maturity T and strike K , the flat volatility $\theta(K, T)$ is defined as the volatility such that the B-S cap price, (denoted $C(\sigma)$ and written as a function just of volatility) is equal to the sum of the caplets priced at the caplet constant volatilities, σ_i . That is, at $\sigma = \theta(K, T)$, we have $C(\sigma) = \Sigma C_i(\sigma_i)$. Now expanding each $C_i(\sigma_i)$ using a first order Taylor approximation about σ gives:

$$C(\sigma) = \Sigma C_i(\sigma_i) \approx \Sigma [C_i(\sigma) + (\sigma_i - \sigma)\partial C_i(\sigma_i)/\partial \sigma_i]$$

Where the derivative $\partial C_i(\sigma_i)/\partial \sigma_i = v_i$ is evaluated at $\sigma_i = \sigma$. This implies

$$\Sigma(\sigma_i - \sigma)\partial C_i(\sigma_i)/\partial \sigma_i \approx 0$$

and so it follows that the flat volatility is approximately a vega weighted sum of each caplet implied volatility θ_i . That is, $\theta \approx \Sigma w_i \theta_i$ where $w_i = v_i/\Sigma v_i$ which proves (A).

The cap vega is the sum of the caplet vegas, so caplet vegas can be backed out from estimates of the cap vegas and they will normally decrease with maturity. Given these, and the relationship (A) between caplet and cap volatilities, caplet implied volatilities can be approximated as follows: Set $\theta_1 = \theta(1)$; then solve for θ_2 from $[v_1\theta_1 + v_2\theta_2]/[v_1 + v_2] = \theta(2)$, and so forth. Caplet volatilities that are approximated by a variance iteration may be negative, or even complex. However there will be no problem with imaginary volatilities when a volatility iteration is used and, although they could be negative, negative forward rate volatilities are unlikely, as illustrated by the following simple example.

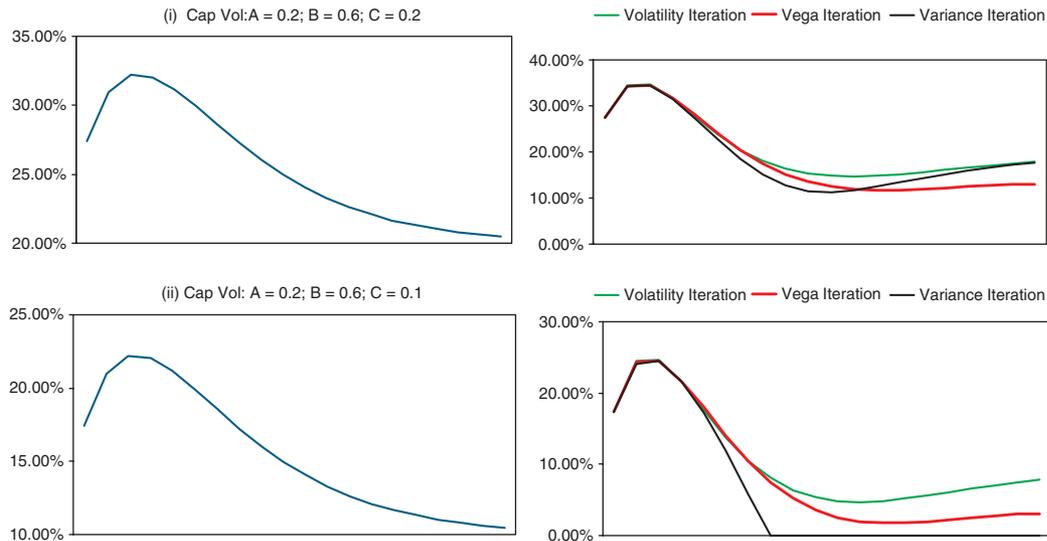


Figure A: Relationship between Caplet and Cap Implied Volatilities.

Suppose that a flat cap volatility term structure is given by $AT[\exp(-BT)] + C$, which is a simple hump shaped curve, and $T = 0.5, 1, 1.5, \dots, 10$ years. We consider the two cases (i) $A = 0.2$; $B = 0.6$; $C = 0.2$ and (ii) $A = 0.2$; $B = 0.6$; $C = 0.1$. Figures A(i) and 2(ii) illustrate the cap volatility curve and the caplet volatilities that are obtained by three different methods: the variance iteration, the volatility iteration (which is based on (A) with all vegas equal) and the vega iteration based on (A). The cap volatilities are uniformly greater in case (i), and all three methods lead to real, non-negative caplet volatilities; however in case (ii)—with lower cap volatilities which are perhaps more realistic of current market conditions in the US and Europe—the variance weighted method gives imaginary volatilities from 4 years onwards. Note that the effect of vega weighting is to reduce the longer maturity caplet volatilities in both cases.

FOOTNOTES & REFERENCES

- When estimating volatilities and correlations from historical data, one has to consider whether the first difference of the forward rates, or the first difference of the log of forward rates, is the variable. The choice should be consistent with the assumed diffusion process. If a normal geometric Brownian process were assumed, then the appropriate stationarity transform would be the difference in the forward rates. But a lognormal geometric Brownian process is assumed in the LFR model described in section 2, so the appropriate stationarity transform is the difference in the log forward rates.
- Most Sterling and Euro swap have semi-annual fixed leg payments, but in the US the fixed leg sometimes has annual payments.
- The bid-offer spread is a good indication of the uncertainty in pricing; it normally increases with maturity. Uncertainty in pricing also translates to uncertainty in volatility through the caplet vega: high vega means that large price errors will induce small volatility errors and low vega means that small price errors will induce large volatility errors.
- See Brigo and Mercurio (2001) for the specification of the drifts under other choices of numeraire.
- Other relationships between swap and forward rates have been developed by several authors. Hull and White (1999) write the log swap rate as a difference in logs of products of forward rates to derive an exact expression for the volatility of the swap rate. Longstaff, Santa-Clara and Schwartz (1999) use a least square regression technique to write the swap rate as an approximate linear function of forward rates. Jäckel and Rebonato (2000) express the swap rate variance as a weighted sum of forward rate covariances and thus derive an approximation for the volatility of the swap rate. Andersen and Andreasen (2000) refine the relationship between forward rates and swap rates in the presence of a volatility skew.
- If reliable (B-S) swaption prices are available in the market, from these we can immediately imply the average swaption volatility $\theta_{n,m}$. Otherwise we shall need to smooth the

volatility surface using a 2-dimensional smoothing algorithm. It is well known that calibration results will be sensitive to the choice of smoothing algorithm (see for example Brigo and Mercurio, 2001).

- Alternatively, the $2m \times 2m$ covariance matrix of semi-annual forward rates can be used. Note that the correlation or covariance matrix should be calculated on the first differences in the logarithms of the forward rates, as is standard for lognormal diffusions.
- The eigenvectors of a covariance matrix V can be obtained from those of the associated correlation matrix C . Writing a_i and m_i as the i th eigenvectors of V and C , and σ_j as the volatility of the j th variable in the system, the required scaling is $a_i = \sigma \otimes m_i$ where \otimes denotes the element by element (Hadamard) product. We therefore seek to parameterize the eigenvectors of the correlation matrix using four parameters, and then scale up these eigenvectors using the forward rate volatilities.

- Alexander, C. and D. L'vov (2003) Statistical Properties of Forward Libor Rates *ISMA Centre Discussion Papers in Finance* available from www.isamacentre.rdg.ac.uk
- Andersen, L. and R. Brotherton-Radcliffe (1997) The Equity Option Volatility Smile: An Implicit Finite Difference Approach *Journal of Computational Finance* Vol. 1, pp. 5–37.
- Andersen, L. and J. Andreasen (2000) Volatility Skews and Extensions of the LIBOR Market Model *Applied Mathematical Finance* Vol. 7, pp. 1–32
- Avellaneda, M., C. Friedman, R. Holmes and D. Samperi (1997) Calibrating volatility surfaces via relative entropy minimization *Appl. Math. Finance* 4 37–64
- Bouchouev, I. and V. Isakov (1997) The inverse problem of option pricing *Inverse Problems* 13 L11–7
- Bouchouev, I. and V. Isakov (1999) Uniqueness, stability and numerical methods for the inverse problem that arises in financial markets *Inverse Problems* 15 R95–116
- Brace, A., D. Gatarek and M. Musiela (1997) The Market Model of Interest Rate Dynamics *Mathematical Finance* Vol. 7, pp. 127–155
- Brigo, D. (2001) A Note on Correlation and Rank Reduction *Working Paper for Banca IMI, Milan* available at <http://www.damianobrigo.it/correl.pdf>

- Brigo, D. and F. Mercurio (2001) *Interest Models, Theory and Practice* Springer-Verlag
- Brigo, D. and F. Mercurio (2002a) Calibrating LIBOR *RISK* Vol. 15:1 pp. 117–121
- Brigo, D. and F. Mercurio (2002b), “Lognormal-Mixture Dynamics and Calibration to Market Volatility Smiles”, *International Journal of Theoretical & Applied Finance* 5(4), pp. 427–446.
- Brown, G. and C. Randall (1999) If the Skew Fits *RISK* Vol. 12:4 pp. 62–65.
- Coleman, T.F., Yuying Li and A. Verma (1998) Reconstructing the Unknown Local Volatility Function. In *Proceedings of Computational and Quantitative Finance '98*, New York.
- De Jong, F., J. Driessen and A. Pelsser LIBOR and Swap Market Models for the Pricing of Interest Rate Derivatives: An Empirical Analysis *Working Paper* http://center.kub.nl/phd_stud/jdries/index.htm.
- Derman, E., I. Kani and N. Chriss (1996) Implied Trinomial Trees of the Volatility Smile *Journal of Derivatives* Vol. 3 pp. ??
- Dumas, B., F. Fleming and R. Whaley (1998) Implied Volatility Functions: Empirical Tests *Journal of Finance* Vol. 53:6 pp. 2059–2106.
- Dupire, B. (1994) Pricing with a Smile *RISK* Vol. 7:1 pp. 18–20.
- Flury, B. (1988) *Common Principal Components and Related Multivariate Models* Wileys, NY.
- Hull, J. and A. White (1999) Forward Rate Volatilities, Swap Rate Volatilities and the Implementation of the LIBOR Market Model *Joseph L. Rotman School of Management, University of Toronto, preprint*
- Hull, J. and A. White (2000) The Essentials of the LMM *RISK* Vol. 13:12 pp. 126–129
- Jäckel, P. and R. Rebonato (2000) Linking Caplet and Swaption Volatilities in a BGM/J Framework: Approximate Solutions *QUARC Working Paper* available at www.rebonato.com forthcoming in the *Journal of Computational Finance*.
- Jamshidian, F. (1997) LIBOR and Swap Market Models and Measures *Finance and Stochastics* 1, 293–330.
- Longstaff F., P. Santa-Clara and E. Schwartz (1999) The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence *UCLA Working Paper*
- McCulloch, J.H. (1975) The Tax Adjusted Yield Curve *Journal of Finance* Vol. 30:3 pp. 811–30.
- McIntyre, M. (2001) Performance of Dupire’s Implied Diffusion Approach Under Sparse and Incomplete Data *Journal of Computational Finance* Vol. 4 (Summer) pp. 33–84.
- Miltersen K., K. Sandmann and D. Sondermann (1997) Closed Form Solutions for Term Structure Derivatives with Lognormal Interest Rates *Journal of Finance* Vol. 52:1 pp. 409–430
- Nelson, C.R. and Siegel, A.F. (1987) Parsimonious Modelling of Yield Curves *Journal of Business* Vol. 60:4 pp. 473–89.
- Rebonato, R. (1998) *Interest Rate Option Models* Wileys, Chichester.
- Rebonato, R. (1999a) *Volatility and Correlation* Wileys, Chichester.
- Rebonato, R. (1999b) Calibrating the BGM Model *RISK* Vol. 12:3.
- Rebonato, R. (1999c) On the Simultaneous Calibration of Multifactor Lognormal Interest rate Models to Black Volatilities and to the Correlation Matrix *Journal of Computational Finance*, Vol. 2:4 pp. 5–27.
- Rebonato, R. and M. Joshi (2001) A Joint Empirical and Theoretical Investigation of the Modes of Deformation of Swaption Matrices: Implications for Model Choice *QUARC Working Paper* available at www.rebonato.com
- Rebonato, R. (2002) *Modern Pricing of Interest-rate Derivatives: the LIBOR Market Model and beyond*, Princeton University Press, Princeton (New Jersey) and Woodstock (UK).
- Schoenmakers, J. and C. Coffey (2000) Stable Implied Calibration of a Multi-Factor LIBOR Model via a Semi-Parametric Correlation Structure *Weierstrass Institute Preprint* no. 611.
- Steely, J.M. (1991) Estimating the Gilt-Edged Term Structure: Basis Splines and Confidence Intervals *Journal of Business Finance and Accounting* Vol. 18:4 pp. 512–29.
- Svensson, L. (1994) Estimating and Interpreting Forward Interest rates: Sweden 1992–94 *International Monetary Fund Working Paper* No. 114.

ACKNOWLEDGEMENT

I would like to thank Damiano Brigo and Fabio Mercurio of Banca IMI, Milan, Italy; Riccardo Rebonato of the Royal Bank of Scotland, London, UK; Hyung-Sok Ahn of Constellation Power, Baltimore, USA; and Jacques Pézier of the ISMA Centre, UK all of whom have provided very helpful comments and suggestions on preliminary drafts of this paper. Particular thanks go to my PhD student Dmitri Lvov for providing the historical forward rate data and the common principal components used to construct the figures.