

**FORWARD RATE VOLATILITIES, SWAP RATE VOLATILITIES,
AND THE IMPLEMENTATION OF THE LIBOR MARKET MODEL**

**John Hull and Alan White
Joseph L. Rotman School of Management
University of Toronto**

August, 1999

Abstract

This paper is concerned with the implementation of the LIBOR market model and its extensions. It develops and tests a simple analytic approximation for calculating the volatilities used by the market to price European swap options from the volatilities used by the market to price interest rate caps. The approximation is found to be very accurate for the range of market parameters normally encountered. It enables swap option volatility skews to be implied from cap volatility skews. It also allows the LIBOR market model to be easily calibrated to broker quotes on caps and European swap options so that a wide range of non-standard interest rate derivatives can be valued.

FORWARD RATE VOLATILITIES, SWAP RATE VOLATILITIES, AND THE IMPLEMENTATION OF THE LIBOR MARKET MODEL

The most popular over-the-counter interest rate options are interest rate caps/floors and European swap options. The standard market models for valuing these instruments are versions of Black's (1976) model. This model was originally developed for valuing options on commodity futures, but has found many other applications in financial engineering. When Black's model is used to value a caplet (one element of an interest rate cap), the underlying interest rate is assumed to be lognormal. When it is used to value European swap options, the underlying swap rate is assumed to be lognormal. Researchers such as Jamshidian (1997) have shown that the cap/floor market model and the European swap option market model are each internally consistent in the sense that they do not permit arbitrage opportunities. However, they are not exactly consistent with each other.

In recent years Brace, Gatarek, and Musiela (1997), Jamshidian (1997), and Miltersen, Sandmann, and Sondermann (1997) have developed what has become known as the LIBOR market model. This is an extension of the Heath, Jarrow, and Morton (HJM) (1992) model. Whereas the HJM model describes the behavior of instantaneous forward rates expressed with continuous compounding, the LIBOR market model describes the behavior of the forward LIBOR rates underlying caps and floors, with the usual market conventions being used for compounding (that is, the compounding period equals the tenor of the rate). Brace, Gatarek, and Musiela (1997) show that the LIBOR market model overcomes some technical existence problems associated with the lognormal version of the HJM model.

Like the HJM model, the LIBOR market model must be implemented using Monte Carlo simulation. Fortunately, the development of the LIBOR market model has been accompanied by the development of a number of new numerical procedures for valuing American and Bermudan options using Monte Carlo simulation. Andersen (1998) proposes an approach where the decision on early exercise is assumed to depend on the intrinsic value of the option and the values of the still-active European options. Broadie and

Glasserman (1997) propose an approach involving a stochastic recombining tree. Longstaff and Schwartz (1998) assume that the value of an option, if not exercised, is a simple function of variables observed at the exercise date in question. These approaches make it feasible to use the LIBOR market model to value instruments such as Bermudan swap options, which are becoming increasingly popular.

This paper develops and tests an analytic approximation for pricing European swap options under the LIBOR market model and its extensions. This approximation makes it possible to translate the volatility skews observed for caps into volatility skews for European swap options. It also simplifies the process of calibrating the LIBOR market model to the market prices of caps and swap options.

The rest of the paper is organized as follows. Section 1 describes the standard market models for pricing interest rate caps and European swap options. Section 2 explains the LIBOR market model and how it can be implemented. Section 3 presents a very fast and accurate procedure for pricing European swap options when the LIBOR market model is used. Section 4 explains the extended LIBOR market model proposed by Andersen and Andreasen (1997) and develops a procedure for translating the volatility skews observed for caps into volatility skews for swap options. Section 5 discusses the implementation of the model and its applications. Conclusions are in Section 6.

1. THE STANDARD MARKET MODELS FOR CAPS AND SWAPTIONS

In this section we explain the standard market models for pricing caps and European swap options. We will use the result of Harrison and Kreps (1979) that, in any market where there is no arbitrage, for any given numeraire security whose price is $g(t)$, there exists a measure for which $f(t)/g(t)$ is a martingale for all security prices $f(t)$. We will denote this measure by $\mathcal{M}\{g(t)\}$.

1.1 Interest Rate Caps

Consider a cap with strike rate R_c and reset dates at times t_1, t_2, \dots, t_N with a final payment date t_{N+1} . Define $\delta_i = t_{i+1} - t_i$ and R_i as the δ_i -maturity rate observed at time t_i and expressed with a compounding period of δ_i ($1 \leq i \leq N$). The cap is a portfolio of N caplets. The i th caplet provides a payoff at time t_{i+1} equal to

$$L_c \delta_i \max[R_i - R_c, 0] \tag{1}$$

where L_c is the principal.

Define $P(t, T)$ as the price of a discount bond paying off \$1 at time T . To value the i th caplet we use $P(t, t_{i+1})$ as the numeraire. Under the measure, $\mathcal{M}\{P(t, t_{i+1})\}$

$$\frac{f(t)}{P(t, t_{i+1})}$$

is a martingale for all security prices $f(t)$. Hence

$$\frac{f(0)}{P(0, t_{i+1})} = E_{i+1} \left[\frac{f(t)}{P(t, t_{i+1})} \right] \tag{2}$$

where E_{i+1} denotes expectations under $\mathcal{M}\{P(t, t_{i+1})\}$.

By setting $f(t) = P(t, t_i) - P(t, t_{i+1})$ in equation (2), we see that

$$F_i(0) = E_{i+1}(R_i) \tag{3}$$

where $F_i(t)$ is the forward interest rate for the period (t_i, t_{i+1}) observed at time t . Equation (3) shows that, under the assumed measure, the value at time zero of this forward rate equals the expected future spot rate for (t_i, t_{i+1}) .

By setting $f(t)$ equal to the price of the i th caplet and noting that $P(t_{i+1}, t_{i+1}) = 1$, we see from equation (2) that

$$f(0) = P(0, t_{i+1})E_{i+1}[f(t_{i+1})]$$

or

$$f(0) = P(0, t_{i+1})L_c\delta_i E_{i+1}[\max(R_i - R_c, 0)]$$

Assuming R_i is lognormal, with the standard deviation of $\ln(R_i)$ equal to $\sigma_i\sqrt{t_i}$, this becomes

$$f(0) = P(0, t_{i+1})L_c\delta_i[E_{i+1}(R_i)N(d_1) - R_cN(d_2)]$$

where

$$d_1 = \frac{\ln[E_{i+1}(R_i)/R_c] + \sigma_i^2 t_i/2}{\sigma_i\sqrt{t_i}}$$

$$d_2 = d_1 - \sigma_i\sqrt{t_i}$$

and $N(x)$ is the cumulative normal distribution function. Substituting from equation (3) gives the standard market model for valuing the caplet:

$$f(0) = P(0, t_{i+1})L_c\delta_i[F_i(0)N(d_1) - R_cN(d_2)] \quad (4)$$

where

$$d_1 = \frac{\ln[F_i(0)/R_c] + \sigma_i^2 t_i/2}{\sigma_i\sqrt{t_i}}$$

$$d_2 = d_1 - \sigma_i\sqrt{t_i}$$

Similarly the standard market model for valuing the i th element of a floor is

$$f(0) = P(0, t_{i+1})L_c\delta_i[R_cN(-d_2) - F_i(0)N(-d_1)] \quad (5)$$

The variable σ_i is often referred to the volatility of the (t_i, t_{i+1}) forward rate. In fact, there is no requirement that the volatility of this forward rate be constant during the life of the caplet. All that is required is that the variable σ_i^2 be the average variance rate of

the forward rate during the life of the caplet. This point is important for understanding of the model in the next section. To avoid confusion with variables introduced in the next section, we will refer to σ_i as the spot volatility of the (t_i, t_{i+1}) rate, or the spot volatility of the i th caplet.

In practice brokers usually quote what is known as a flat volatility for a cap. This is a volatility which, if used as the spot volatility for all the underlying caplets, reproduces the cap's market price. When flat volatilities for all cap maturities are available or can be estimated, spot volatilities can be calculated. The procedure is to use the flat volatilities and equation (4) to calculate cap prices, deduce caplet prices by subtracting one cap price from the next, and then imply the spot volatilities from these caplet prices using equation (4).

1.2 European Swap Options

We now move on to consider a European swap option. Suppose that the underlying swap lasts from time t_n to time t_{N+1} and the strike rate is R_s . The swap reset dates are t_n, t_{n+1}, \dots, t_N and the corresponding payment dates are $t_{n+1}, t_{n+2}, \dots, t_{N+1}$, respectively. We define

$$A_{n,N}(t) = \sum_{i=n}^N \delta_i P(t, t_{i+1})$$

where as before $P(t, T)$ is the price of a discount bond paying off \$1 at time T and $\delta_i = t_{i+1} - t_i$. We will denote by $S_{n,N}(t)$ ($0 \leq t \leq t_n$) the forward swap rate. This is the fixed rate which, when it is exchanged for floating in a forward start swap, causes the value of the swap to be zero. An expression for $S_{n,N}(t)$ is

$$S_{n,N}(t) = \frac{P(t, t_n) - P(t, t_{N+1})}{A_{n,N}(t)} \quad (6)$$

Define $R_{n,N} = S_{n,N}(t_n)$ as the swap rate observed at time t_n . Both $R_{n,N}$ and R_s are expressed with compounding frequencies reflecting the number of swap payments per year. A European swap option where the holder has the right to pay fixed and receive floating

can be viewed as providing payments of

$$L_s \delta_i \max(R_{n,N} - R_s, 0)$$

at times t_{i+1} for $n \leq i \leq N$ where L_s is the swap principal. The value of the swap option at time t_n is

$$L_s A_{n,N}(t_n) \max(R_{n,N} - R_s, 0) \tag{7}$$

We will value the European swap option at time zero using $A_{n,N}(t)$ as the numeraire. Under the measure $\mathcal{M}\{A_{n,N}(t)\}$

$$\frac{f(t)}{A_{n,N}(t)}$$

is a martingale for all security prices $f(t)$. Hence

$$\frac{f(0)}{A_{n,N}(0)} = E_A \left[\frac{f(t_n)}{A_{n,N}(t_n)} \right] \tag{8}$$

where E_A denotes expectations under the measure. By setting $f(t) = P(t, t_n) - P(t, t_{N+1})$ in equation (8) we see that

$$S_{n,N}(0) = E_A[S_{n,N}(t_n)] = E_A(R_{n,N}) \tag{9}$$

The forward swap rate, therefore, equals the expected future swap rate under the measure.

We now set $f(t)$ equal to the value of the swap option. Equation (8) gives

$$f(0) = A_{n,N}(0) E_A \left[\frac{f(t_n)}{A_{n,N}(t_n)} \right]$$

Substituting for $f(t_n)$ from equation (7) we obtain

$$f(0) = L_s A_{n,N}(0) E_A[\max(R_{n,N} - R_s, 0)]$$

Assuming $R_{n,N}$ is lognormal with the standard deviation of $\ln(R_{n,N})$ equal to $\sigma_{n,N} \sqrt{t_n}$ this becomes

$$f(0) = L_s A_{n,N}(0) [E_A(R_{n,N}) N(d_1) - R_s N(d_2)]$$

where

$$d_1 = \frac{\ln[E_A(R_{n,N})/R_s] + \sigma_{n,N}^2 t_n / 2}{\sigma_{n,N} \sqrt{t_n}}$$

$$d_2 = d_1 - \sigma_{n,N} \sqrt{t_n}$$

Substituting from equation (9) gives the standard market model for valuing the swap option

$$f(0) = L_s A_{n,N}(0) [S_{n,N}(0) N(d_1) - R_s N(d_2)] \quad (10)$$

where

$$d_1 = \frac{\ln[S_{n,N}(0)/R_s] + \sigma_{n,N}^2 t_n / 2}{\sigma_{n,N} \sqrt{t_n}}$$

$$d_2 = d_1 - \sigma_{n,N} \sqrt{t_n}$$

Similarly the standard market model for valuing a swap option that gives the holder the right to receive fixed and pay floating is

$$f(0) = L_s A_{n,N}(0) [R_s N(-d_2) - S_{n,N}(0) N(-d_1)] \quad (11)$$

The variable $\sigma_{n,N}$ is usually referred to as the forward swap rate volatility. As in the case of interest rate caps, we emphasize that the model does not require the forward swap rate's volatility to be constant between times 0 and t_n . It is sufficient for its average variance rate to be $\sigma_{n,N}^2$ during this period. We will refer to $\sigma_{n,N}$ as the spot volatility of the swap rate or the spot volatility of the European swap option.

2. THE LIBOR MARKET MODEL

In this section we discuss the development and implementation of the LIBOR market model. Our notation is consistent with Section 1.1. We consider a cap with reset dates at times t_1, t_2, \dots, t_n and a final payment date t_{n+1} . We define $t_0 = 0$, $\delta_i = t_{i+1} - t_i$ ($0 \leq i \leq n$), and

$F_i(t)$: Forward rate observed at time t for the period (t_i, t_{i+1}) , expressed with a compounding period of δ_i

$P(t, T)$: Price at time t of a zero-coupon bond that provides a payoff of \$1 at time T

$m(t)$: Index for the next reset date at time t . This means that $m(t)$ is the smallest integer such that $t \leq t_{m(t)}$.

p : Number of factors

$\zeta_{i,q}$: q th component of the volatility of $F_i(t)$ ($1 \leq q \leq p$)

$v_{i,q}$: q th component of the volatility of $P(t, t_i)$ ($1 \leq q \leq p$)

We assume that the components to the volatility are independent. (This is not restrictive since they can be orthogonalized.) The processes followed by $F_i(t)$ and $P(t, t_i)$ are:

$$dF_i(t) = \dots dt + \sum_{q=1}^p \zeta_{i,q}(t) F_i(t) dz_q \quad (12)$$

$$dP(t, t_i) = \dots dt + \sum_{q=1}^p v_{i,q}(t) P(t, t_i) dz_q$$

where the dz_q are independent Wiener processes and the drifts depend on the measure. In this section and the next, we will assume a model where $\zeta_{i,q}(t)$ is a function only of time. The bond price volatility, $v_{i,q}(t)$, is in general stochastic in this model.

We will use as numeraire a money market account that is invested at time t_0 for a period δ_0 , reinvested at time t_1 for a period δ_1 , reinvested at time t_2 for a period δ_2 , and so on. This is equivalent to using a numeraire at time t equal to $P(t, t_{m(t)})$. Under the

chosen measure

$$\frac{f(t)}{P(t, t_{m(t)})}$$

is a martingale for all security prices $f(t)$ when $t_{m(t)-1} \leq t \leq t_{m(t)}$ so that

$$\frac{f(t_{m(t)-1})}{P(t_{m(t)-1}, t_{m(t)})} = E_{m(t)} \left\{ \frac{f(t_{m(t)})}{P(t_{m(t)}, t_{m(t)})} \right\}$$

or

$$f(t_{m(t)-1}) = P(t_{m(t)-1}, t_{m(t)}) E_{m(t)} \{ f(t_{m(t)}) \}$$

where $E_{m(t)}$ denotes expectations under $\mathcal{M}\{P(t, t_{m(t)})\}$. This equation shows that the chosen measure allows us to discount expected values “one accrual period at a time” when security prices are valued. Since cash flows and early exercise opportunities usually occur only on reset dates, this is an attractive feature of the measure.

It can be shown that the q th component of the market price of risk under the measure $\mathcal{M}\{P(t, t_j)\}$ is $v_{j,q}$ for all j . [See for example, Jamshidian (1997).] As shown in Section 1.1, under the measure $\mathcal{M}\{P(t, t_{i+1})\}$ $F_i(t)$ is a martingale so that its drift is zero. It follows that the expected growth rate of $F_i(t)$ under the $\mathcal{M}\{P[t, t_{m(t)}]\}$ measure is

$$\sum_{q=1}^p \zeta_{i,q}(t) [v_{m(t),q}(t) - v_{i+1,q}(t)]$$

so that

$$\frac{dF_i(t)}{F_i(t)} = \sum_{q=1}^p \zeta_{i,q}(t) [v_{m(t),q}(t) - v_{i+1,q}(t)] dt + \sum_{q=1}^p \zeta_{i,q}(t) dz_q \quad (13)$$

The relationship between forward rates and bond prices is

$$\frac{P(t, t_j)}{P(t, t_{j+1})} = 1 + \delta_j F_j(t)$$

Using this in conjunction with Ito's lemma leads to

$$v_{j,q}(t) - v_{j+1,q}(t) = \frac{\delta_j F_j(t) \zeta_{j,q}(t)}{1 + \delta_j F_j(t)} \quad (14)$$

Repeated application of this result gives

$$v_{m(t),q}(t) - v_{i+1,q}(t) = \sum_{j=m(t)}^i \frac{\delta_j F_j(t) \zeta_{j,q}(t)}{1 + \delta_j F_j(t)}$$

and substituting into equation (13) we see that the process followed by $F_i(t)$ under the measure $\mathcal{M}\{P(t, t_{m(t)})\}$ is

$$\frac{dF_i(t)}{F_i(t)} = \sum_{j=m(t)}^i \frac{\delta_j F_j(t) \sum_{q=1}^p \zeta_{j,q}(t) \zeta_{i,q}(t)}{1 + \delta_j F_j(t)} dt + \sum_{q=1}^p \zeta_{i,q}(t) dz_q \quad (15)$$

When we take limits allowing the δ_i 's to approach zero, this becomes

$$dF(t, T) = \sum_{q=1}^p \zeta_q(t, T) F(t, T) \int_{\tau=t}^T \zeta_q(t, \tau) F(t, \tau) d\tau + \sum_{q=1}^p \zeta_q(t, T) F(t, T) dz_q$$

where the notation is temporarily changed so that $F(t, T)$ is the instantaneous forward rate at time t for maturity T and $\zeta_q(t, T)$ is the q th component of the volatility of $F(t, T)$. This is the HJM result. The HJM model is, therefore, a limiting case of the LIBOR market model.

2.1. Determining Forward Rate Volatilities

We now simplify the model by assuming that $\zeta_{i,q}(t)$ is a function only of the number of whole accrual periods between the next reset date and time t_i . Define $\lambda_{j,q}$ as the value of $\zeta_{i,q}(t)$ when there are j such accrual periods. This means that

$$\zeta_{i,q}(t) = \lambda_{i-m(t),q} \quad (16)$$

Define Λ_j as total volatility of a forward rate when there are j whole accrual periods until maturity so that

$$\Lambda_j = \sqrt{\sum_{q=1}^p \lambda_{j,q}^2} \quad (17)$$

The Λ 's can be calculated from the spot volatilities, σ_i , used to value caplets in equation (4). Equating variances we must have:

$$\sigma_i^2 t_i = \sum_{j=1}^i \Lambda_{i-j}^2 \delta_{j-1} \quad (18)$$

This allows the Λ 's to be obtained inductively. For example, if σ_1 , σ_2 , and σ_3 are 20%, 22% and 21%, respectively and the δ_i are all equal, $\Lambda_0 = 20\%$, $\Lambda_1 = 23.83\%$ and $\Lambda_2 = 18.84\%$

Similarly to Rebonato (1999), we favor a two step approach to obtaining the $\lambda_{j,q}$. The first stage is to determine the Λ_j from market data; the second stage is to use historical data to determine the $\lambda_{j,q}$ from the Λ_j . Rebonato (1999) proposes a way of determining the $\lambda_{j,q}$ from the Λ_j to provide as close a fit as possible to the correlation matrix for the $F_j(t)$ ($1 \leq j \leq N$). As Rebonato points out, the approach gives similar results to a principal components analysis. If output from a principal components analysis on the $F_j(t)$ ($1 \leq j \leq N$) is available, we can use it in a direct way to determine values of $\lambda_{j,q}$. Suppose that $\alpha_{i,q}$ is the factor loading for the i th forward rate and the q th factor, and s_q is the standard deviation of the q th factor score.¹ If the number of factors being used, p , were equal to N , the number of forward rates, it would be correct to set

$$\lambda_{j,q} = \alpha_{j,q} s_q$$

for $1 \leq j, q \leq N$. When $p < N$, the $\lambda_{j,q}$ can be scaled so that the relationship in equation (17) holds. This involves setting

$$\lambda_{j,q} = \frac{\Lambda_j s_q \alpha_{j,q}}{\sqrt{\sum_{q=1}^p s_q^2 \alpha_{j,q}^2}} \quad (19)$$

¹ The principal components analysis model is

$$\Delta F_j = \sum_{q=1}^N \alpha_{j,q} x_q$$

where

$$\sum_{j=1}^N \alpha_{j,q_1} \alpha_{j,q_2}$$

equals 1 when $q_1 = q_2$ and zero when $q_1 \neq q_2$. The variable s_q is the standard deviation of the random variable x_q .

Table 1 shows the $\lambda_{j,q}/\Lambda_j$ values used for the analyses reported in this paper. These are based on equation (19) and data that is typical of the results obtained from a principal components analysis. The first factor corresponds to a roughly parallel shift in the yield curve; the second factor corresponds to a “twist” in the zero curve where short maturity rates move in one direction and long maturity rates move in the opposite direction; the third factor corresponds to a “bowing” of the zero curve where short and long maturity rates move in one direction and intermediate maturity rates move in the opposite direction.

2.2. Implementation of the Model

From equation (15), the process for $F_i(t)$ under the measure $\mathcal{M}\{P(t, t_{m(t)})\}$ is

$$\frac{dF_i(t)}{F_i(t)} = \sum_{j=m(t)}^i \frac{\delta_j F_j(t) \sum_{q=1}^p \lambda_{j-m(t),q} \lambda_{i-m(t),q}}{1 + \delta_j F_j(t)} dt + \sum_{q=1}^p \lambda_{i-m(t),q} dz_q \quad (20)$$

or

$$d \ln F_i(t) = \sum_{j=m(t)}^i \left[\frac{\delta_j F_j(t) \sum_{q=1}^p \lambda_{j-m(t),q} \lambda_{i-m(t),q}}{1 + \delta_j F_j(t)} - \sum_{q=1}^p \frac{\lambda_{i-m(t),q}^2}{2} \right] dt + \sum_{q=1}^p \lambda_{i-m(t),q} dz_q \quad (21)$$

An approximation that simplifies the Monte Carlo implementation of the model is that the drift of $\ln F_i$ is constant between times t_k and t_{k+1} so that

$$F_i(t_{k+1}) = F_i(t_k) \exp \left[\left(\sum_{j=k+1}^i \frac{\delta_j F_j(t_k) \sum_{q=1}^p \lambda_{j-k-1,q} \lambda_{i-k-1,q}}{1 + \delta_j F_j(t_k)} - \sum_{q=1}^p \frac{\lambda_{i-k-1,q}^2}{2} \right) \delta_k + \sum_{q=1}^p \lambda_{i-k-1,q} \epsilon_q \sqrt{\delta_k} \right] \quad (22)$$

where ϵ_q are independent random samples from standard normal distributions.

The LIBOR market model is designed to price caplets exactly consistently with the standard market model. We can test the approximation in equation (22) by using Monte Carlo simulation in conjunction with equation (22) to value caplets and comparing the prices to those given by the standard market model in equation (4). Each Monte Carlo trial consists of using the equation to generate a path for each forward rate under the

$\mathcal{M}\{P(t, t_{m(t)})\}$ measure. The value of $F_i(t_i)$ is the realized rate for the time period between t_i and t_{i+1} and enables the caplet payoff at time t_{i+1} to be calculated. This payoff is discounted to time zero using $F_j(t_j)$ as the discount rate for the interval (t_j, t_{j+1}) . The estimated caplet value is the mean of the discounted payoffs.

We have carried out the test just described for a variety of term structures, volatility structures, strike rates, and number of factors. Typical results are shown in Table 2. This table is based on an upward sloping term structure, a “humped” volatility structure similar to that observed in the market, and at-the-money caplets (that is, caplets where the strike rate equals the forward rate). Three factors and 200,000 Monte Carlo trial were used. The δ_i were set equal to one year.² This table shows the mean and standard error of the difference between the true spot volatility of a caplet and the volatility implied from the caplet price calculated using the LIBOR market model, with both volatilities being measured in percent. The results shown are for the fifth caplet. Similar results are obtained for other caplets, other term structures, other volatility structures, and other choices for the strike rate and the number of factors. The average interest rates reported in Table 2 are the arithmetic average the zero rates for maturities of 1, 2, 3, . . . , 10 years. The average volatilities are similarly the arithmetic average of the spot volatilities between one and ten years.

Table 2 shows that, for the volatility and interest rate environments that are typically encountered in North America and Europe, the implementation of the model in equation (22) gives very accurate results. Consider, for example, the case where the average interest rate is 5% and the average volatility is 20%. If the true spot volatility for a particular caplet is $V\%$, Table 2 shows that we can be 95% certain that the equation (22) implementation of the LIBOR market model is implicitly assuming a spot volatility between $V - 0.12\%$ and

² The error in equation (22) is likely to increase with the length of the accrual period. Setting the accrual period equal to one year is therefore providing a relatively tough test of equation (22).

$V + 0.08\%$.³ As the volatility increases the approximation in equation (22) works less well. In some countries (for example, Japan), very low short-term interest rates are sometimes coupled with spot volatilities as high as 100% for short-maturity caplets. Table 2 shows that the approximation in equation (22) may not be appropriate in these cases.⁴

³ To put this result in perspective, the bid–offer spread for the flat volatilities quoted by brokers in the U.S. is typically 0.5%.

⁴ The lognormal LIBOR market model described here may itself be inappropriate in very-low interest rate environments. This is because it assumes a stationary volatility structure. When short rates are in the region of 1%, their volatilities are very high. But, if as time passes short rates increase, more normal volatilities are likely to be observed.

3. APPROXIMATE PRICING OF EUROPEAN SWAP OPTIONS

In this section we present an approximate, but very accurate, procedure for pricing European swap options in the LIBOR market model. Other approximations have been suggested by Brace, Gatarek, and Musiela (1997) and Andersen and Andreasen (1997). Our approach is similar in spirit to that of Andersen and Andreasen, but is much easier to implement than either of the other two approaches. It is motivated by the observation that when forward LIBOR rates are lognormal, swap rates are approximately lognormal and approximately linearly dependent on forward LIBOR rates.

As in Section 1.2 we consider an option on a swap lasting from t_n to t_{N+1} with reset dates at times t_n, t_{n+1}, \dots, t_N . Initially we assume that the reset dates for the swap coincide with the reset dates for caplets underlying the LIBOR market model. Later we relax this assumption.

The relationship between bond prices and forward rates is

$$\frac{P(t, t_k)}{P(t, t_n)} = \prod_{j=n}^{k-1} \frac{1}{1 + \delta_j F_j(t)}$$

for $k \geq n + 1$. It follows that the formula for the forward swap rate in equation (6) can be written

$$S_{n,N}(t) = \frac{1 - \prod_{j=n}^N \frac{1}{1 + \delta_j F_j(t)}}{\sum_{i=n}^N \delta_i \prod_{j=n}^i \frac{1}{1 + \delta_j F_j(t)}}$$

(We employ the convention that empty sums equal zero and empty products equal one.)

Equivalently

$$S_{n,N}(t) = \frac{\prod_{j=n}^N [1 + \delta_j F_j(t)] - 1}{\sum_{i=n}^N \delta_i \prod_{j=i+1}^N [1 + \delta_j F_j(t)]} \quad (23)$$

or

$$\ln S_{n,N}(t) = \ln \left\{ \prod_{j=n}^N [1 + \delta_j F_j(t)] - 1 \right\} - \ln \left\{ \sum_{i=n}^N \delta_i \prod_{j=i+1}^N [1 + \delta_j F_j(t)] \right\}$$

so that

$$\frac{1}{S_{n,N}(t)} \frac{\partial S_{n,N}(t)}{\partial F_k(t)} = \frac{\delta_k \gamma_k(t)}{1 + \delta_k F_k(t)}$$

where

$$\gamma_k(t) = \frac{\prod_{j=n}^N [1 + \delta_j F_j(t)]}{\prod_{j=n}^N [1 + \delta_j F_j(t)] - 1} - \frac{\sum_{i=n}^{k-1} \delta_i \prod_{j=i+1}^N [1 + \delta_j F_j(t)]}{\sum_{i=n}^N \delta_i \prod_{j=i+1}^N [1 + \delta_j F_j(t)]}$$

From Ito's lemma the q th component of the volatility of $S_{n,N}(t)$ is

$$\sum_{k=n}^N \frac{1}{S_{n,N}(t)} \frac{\partial S_{n,N}(t)}{\partial F_k(t)} \zeta_{k,q}(t) F_k(t)$$

or

$$\sum_{k=n}^N \frac{\delta_k \zeta_{k,q}(t) F_k(t) \gamma_k(t)}{1 + \delta_k F_k(t)}$$

The variance rate of $S_{n,N}(t)$ is therefore

$$\sum_{q=1}^p \left[\sum_{k=n}^N \frac{\delta_k \zeta_{k,q}(t) F_k(t) \gamma_k(t)}{1 + \delta_k F_k(t)} \right]^2$$

Assuming, as in Section 2.1 that $\zeta_{k,q}(t) = \lambda_{k-m(t),q}$, the variance rate of $S_{n,N}(t)$ is

$$\sum_{q=1}^p \left[\sum_{k=n}^N \frac{\delta_k \lambda_{k-m(t),q} F_k(t) \gamma_k(t)}{1 + \delta_k F_k(t)} \right]^2$$

This expression is in general stochastic showing that when the forward rates underlying caplets are lognormal, swap rates are not lognormal. For the purposes of evaluating the expression we assume that $F_k(t) = F_k(0)$. This means that the volatility of $S_{n,N}(t)$ is constant within each accrual period and the swap rate is lognormal. The average variance rate of $S_{n,N}(t)$ between time zero and time t_n is

$$\frac{1}{t_n} \sum_{j=0}^{n-1} \left\{ \delta_j \sum_{q=1}^p \left[\sum_{k=n}^N \frac{\delta_k \lambda_{k-j-1,q} F_k(0) \gamma_k(0)}{1 + \delta_k F_k(0)} \right]^2 \right\}$$

The spot swap option volatility is therefore

$$\sqrt{\frac{1}{t_n} \sum_{j=0}^{n-1} \left\{ \delta_j \sum_{q=1}^p \left[\sum_{k=n}^N \frac{\delta_k \lambda_{k-j-1,q} F_k(0) \gamma_k(0)}{1 + \delta_k F_k(0)} \right]^2 \right\}} \quad (24)$$

The swap option price can be calculated by substituting this volatility into equation (10) or (11).

3.1 Non-Matching Accrual Periods

The accrual periods for the swaps underlying broker quotes for European swap options do not always match the accrual periods for the caps and floors underlying broker quotes. For example, in the United States the benchmark caps and floors have quarterly resets while the swaps underlying the benchmark European swap options have semiannual resets. To accommodate this, we now extend the analysis to the situation where each swap accrual period includes M cap accrual periods where $M \geq 1$ is an integer.

Assume that $\delta_n, \delta_{n+1}, \dots, \delta_N$ are the swap accrual periods and that within the i th swap accrual period the cap accrual periods are $\epsilon_{i,1}, \epsilon_{i,2}, \dots, \epsilon_{i,M}$ with

$$\delta_i = \sum_{m=1}^M \epsilon_{i,m}$$

Assume that $G_{i,m}(t)$ is the forward rate observed at time t for the $\epsilon_{i,m}$ accrual period. Since

$$1 + \delta_i F_i(t) = \prod_{m=1}^M [1 + \epsilon_{i,m} G_{i,m}(t)]$$

$$\frac{\partial F_i(t)}{\partial G_{i,m}(t)} = \frac{\epsilon_{i,m} [1 + \delta_i F_i(t)]}{\delta_i [1 + \epsilon_{i,m} G_{i,m}(t)]}$$

The q th component of the volatility of $S_{n,N}(t)$ is

$$\sum_{k=n}^N \sum_{m=1}^M \frac{1}{S_{n,N}(t)} \frac{\partial S_{n,N}(t)}{\partial G_{k,m}(t)} \zeta_{k,m,q}(t) G_{k,m}(t)$$

where $\zeta_{k,m,q}$ is the q th component of the volatility of $G_{k,m}(t)$. Using

$$\frac{\partial S_{n,N}(t)}{\partial G_{i,m}(t)} = \frac{\partial S_{n,N}(t)}{\partial F_i(t)} \frac{\partial F_i(t)}{\partial G_{i,m}(t)}$$

we find that the estimate in equation (24) of the spot swap option volatility is

$$\sqrt{\frac{1}{t_n} \sum_{j=0}^{n-1} \left\{ \delta_p \sum_{q=1}^p \left[\sum_{k=n}^N \sum_{m=1}^M \frac{\delta_k \lambda_{k-j-1,m,q} G_{k,m}(0) \gamma_{k,m}(0)}{1 + \delta_k F_k(0)} \right]^2 \right\}} \quad (25)$$

where

$$\gamma_{k,m}(t) = \gamma_k(t) \frac{\epsilon_{k,m}[1 + \delta_k F_k(t)]}{\delta_k[1 + \epsilon_{k,m} G_{k,m}(t)]}$$

and $\lambda_{j,m,q}$ is the q th component of the volatility of $G_{i,m}(t)$ when $i - m(t) = j$.

3.2. Accuracy of Approximation

To assess the accuracy of the assumption that forward rates are constant in equations (24) and (25) we performed tests analogous to those described in section 2.2 for testing the accuracy of caplet pricing. Typical results are shown in Table 3. These are based on 200,000 Monte Carlo trials of equation (22) using $\delta_i = 1$, an upward sloping term structure, a “humped” volatility structure similar to that observed in the market, and three factors. The table shows the mean and standard error of the difference between the estimates of the spot volatility of swap options calculated from the LIBOR market model and estimates of the spot volatilities of swap options calculated from the approximation in Section 3.1. The results shown are for a 5×5 swap option (that is, a five year option to enter into a five year swap) with the underlying swap being reset annually. The strike rates are at-the money (that is, the strike rate equals the forward swap rate). Similar results are obtained for other swap options, other strike rates, other interest rate term structures, other volatility structures, and other choices for the number of factors. Average interest rates and average volatilities are defined as in Table 2.

The results in Table 3 are a little more difficult to interpret than those in Table 2. This is because we do not know what the true swap option volatility is. To obtain the “true” swap option volatility for the LIBOR market model we used Monte Carlo simulation and equation (22) to calculate the swap option price, and equation (10) to convert the price into a volatility. There are therefore three sources of error in the difference reported in Table 3: the Monte Carlo error, the error arising from the approximation in equation (22), and the error arising from the approximations in equations (24) and (25). Tests using equation (21) and large number of time steps indicate that for a given term structure and volatility structure, the errors in estimating swap option volatilities are similar in magnitude to those

reported for caps in Table 2. For situations where the errors in both Tables 2 and 3 are low, we can therefore assume that the approximation in Section 3.1 works well.

The results show that the approximation in Section 3.1 works well for the interest rates and volatilities normally encountered in North America and Europe. Consider, for example, the situation where interest rates average 5% and volatilities average 20%. Even after the “equation (22) error” is taken into account, the error in the approximation of the 5×5 swap option volatility is likely to be less than 0.1%. This translates into a pricing error of less than \$2 per \$1 million of principal.⁵

⁵ Brace, Gatarek, and Musiela (1997) report errors for a 5×5 swap option of about \$5 per \$1 million of principal using their approximation. Although their errors are higher than ours, they are acceptably low. A key advantage of our approach over theirs is that it is much easier to implement.

4. MODELING VOLATILITY SKEWS

The volatility skew for equities has been well documented by authors such as Jackwerth and Rubinstein (1996). As pointed out by Andersen and Andreasen (1997), caps and floors exhibit a similar volatility skew to equities. The lower the strike price, the higher the volatility implied from the standard market model in equations (4) and (5). Some brokers do provide quotes for caps that are not at the money, but there is very little data on the pricing of non-at-the-money European swap options. It is, therefore, of interest to investigate how cap volatility skews can be converted to swap option volatility skews.

As shown by Andersen and Andreasen (1997), the LIBOR market model can be extended to incorporate volatility skews. We consider the version of their model where the process for forward rates in equation (12) is replaced by a CEV model

$$dF_i(t) = \dots + \sum_{q=1}^p \zeta_{i,q}(t) F_i(t)^\alpha dz_q \quad (26)$$

where α is a positive constant. We generalize the notation in Section 2.1 to define $\lambda_{j,q}$ as the value of $\zeta_{i,q}(t)$ in this model when there are j whole accrual periods between time t and time t_i . This means that

$$\zeta_{i,q}(t) = \lambda_{i-m(t),q}$$

The process for forward rates in equation (20) becomes

$$\frac{dF_i(t)}{F_i(t)^\alpha} = \sum_{j=m(t)}^i \frac{\delta_j F_j(t)^\alpha \sum_{q=1}^p \lambda_{j-m(t),q} \lambda_{i-m(t),q}}{1 + \delta_j F_j(t)} dt + \sum_{q=1}^p \lambda_{i-m(t),q} dz_q$$

or

$$dQ_i(t) = \sum_{j=m(t)}^i \left[\frac{\delta_j F_j(t)^\alpha \sum_{q=1}^p \lambda_{j-m(t),q} \lambda_{i-m(t),q}}{1 + \delta_j F_j(t)} - \sum_{q=1}^p \frac{\alpha F_i(t)^{\alpha-1} \lambda_{i-m(t),q}^2}{2} \right] dt + \sum_{q=1}^p \lambda_{i-m(t),q} dz_q$$

where

$$Q_i(t) = \frac{1}{1-\alpha} F_i(t)^{1-\alpha}$$

Analogously to the approach taken in section 2.2 in arriving at equation (22), we assume that the drift of $Q_i(t)$ is constant between times t_j and t_{j+1} to get

$$Q_i(t_{k+1}) = Q_i(t_k) + \delta_k \sum_{j=k+1}^i \left[\frac{\delta_j F_j(t_k)^\alpha \sum_{q=1}^p \lambda_{j-k-1,q} \lambda_{i-k-1,q}}{1 + \delta_j F_j(t_k)} - \sum_{q=1}^p \frac{\alpha F_i(t_k)^{\alpha-1} \lambda_{i-k-1,q}^2}{2} \right] + \sum_{q=1}^p \lambda_{i-k-1,q} \epsilon_q \sqrt{\delta_k} \quad (27)$$

We define Λ_j by

$$\Lambda_j = \sqrt{\sum_{q=1}^p \lambda_{j,q}^2} \quad (28)$$

and σ_i by:

$$\sigma_i^2 = \frac{1}{t_i} \sum_{j=1}^i \Lambda_{i-j}^2 \delta_{j-1} \quad (29)$$

These definitions are consistent with equations (17) and (18) in Section 2. However, our notation is now more general than in Section 2. In equations (17) and (18) the Λ 's and σ 's were volatilities whereas here they are volatilities only when $\alpha = 1$.

Andersen and Andreasen (1997) show that the process in equation (26) implies that the price of the i th caplet is

$$P(0, t_{i+1}) L_c \delta_i [F_i(0) - F_i(0) \chi^2(a, b + 2, c) - R_c \chi^2(c, b, a)] \quad (30)$$

when $0 < \alpha < 1$ and

$$P(0, t_{i+1}) L_c \delta_i [F_i(0) - F_i(0) \chi^2(c, -b, a) - R_c \chi^2(a, 2 - b, c)] \quad (31)$$

when $\alpha > 1$ where

$$a = \frac{R_c^{2(1-\alpha)}}{(1-\alpha)^2 \sigma_i^2 t_i}; \quad b = \frac{1}{1-\alpha}; \quad c = \frac{F_i(0)^{2(1-\alpha)}}{(1-\alpha)^2 \sigma_i^2 t_i}$$

and $\chi^2(z, v, k)$ is the cumulative probability that a non-central chi-squared distribution with non-centrality parameter v and k degrees of freedom is less than z . The price of the i th floorlet is

$$P(0, t_{i+1}) L_c \delta_i [R_c - F_i(0) \chi^2(a, b + 2, c) - R_c \chi^2(c, b, a)]$$

when $0 < \alpha < 1$ and

$$P(0, t_{i+1})L_c\delta_i[R_c - F_i(0)\chi^2(c, -b, a) - R_c\chi^2(a, 2 - b, c)]$$

when $\alpha > 1$. When $\alpha = 1$, equations (4) and (5) give the caplet and floorlet prices, respectively.

Andersen and Andreasen (1997) also show that skews for swap rates can be analyzed in an analogous way to skews for forward rates. Suppose that the swap rate considered in Section 1.2 follows a process of the form

$$dS_{n,N}(t) = \dots + \sum_{q=1}^p \eta_{q,n,N}(t)S_{n,N}(t)^\beta dz_q \quad (32)$$

where β is a positive constant. We define $\sigma_{n,N}$ by

$$\sigma_{n,N}^2 = \frac{1}{t_n} \int_{\tau=0}^{t_n} \sum_{q=0}^p \eta_{q,n,N}(\tau)^2 d\tau$$

In the special case where $\beta = 1$ this is the same as the $\sigma_{n,N}$ in Section 1.2. The value of the European swap option that gives the holder the right to pay fixed is

$$L_s A_{n,N}(0)[S_{n,N}(0) - S_{n,N}(0)\chi^2(e, f + 2, g) - R_s\chi^2(g, f, e)] \quad (33)$$

when $0 < \alpha < 1$ and

$$L_s A_{n,N}(0)[S_{n,N}(0) - S_{n,N}(0)\chi^2(g, -f, e) - R_s\chi^2(e, 2 - f, g)] \quad (34)$$

when $\alpha > 1$ where

$$e = \frac{R_s^{2(1-\beta)}}{(1-\beta)^2\sigma_{n,N}^2 t_n}; \quad f = \frac{1}{1-\beta}; \quad g = \frac{S_{n,N}(0)^{2(1-\beta)}}{(1-\beta)^2\sigma_{n,N}^2 t_n}$$

The value of the European swap option that gives the holder the right to pay floating is

$$L_s A_{n,N}(0)[R_s - S_{n,N}(0)\chi^2(e, f + 2, g) - R_s\chi^2(g, f, e)] \quad (35)$$

when $0 < \alpha < 1$ and

$$L_s A_{n,N}(0) [R_s - S_{n,N}(0) \chi^2(g, -f, e) - R_s \chi^2(e, 2 - f, g)] \quad (36)$$

when $\alpha > 1$. When $\beta = 1$, equations (10) and (11) give the swap option price.

In Section 3 we showed that, for the volatility and interest rate environments normally encountered, when $\alpha = 1$ in equation (26) it is reasonable to assume that swap rates follow the process in equation (32) with $\beta = 1$. We now hypothesize that a more general version of this result holds. We suppose that, when rates follow the process in equation (26), to a reasonable approximation swap rates follow the process in equation (32) with $\beta = \alpha$. This is an attractive hypothesis. A similar analysis to that in Section 3.1 shows that it leads to $\sigma_{n,N}$ being approximately equal to⁶

$$S_{n,N}(0)^{1-\beta} \sqrt{\frac{1}{t_n} \sum_{j=0}^{n-1} \left\{ \delta_p \sum_{q=1}^p \left[\sum_{k=n}^N \sum_{m=1}^M \frac{\delta_k \lambda_{k-j-1,m,q} G_{k,m}(0)^\alpha \gamma_{k,m}(0)}{1 + \delta_k F_k(0)} \right]^2 \right\}} \quad (37)$$

where $S_{n,N}(0)$ is given by equation (23).

To test the hypothesis we compared European swap option prices calculated from:

1. A Monte Carlo simulation of the extended LIBOR market model using the approximation in equation (27); and
2. Equation (35) or (36) combined with the estimate of $\sigma_{n,N}$ in equation (37) and $\beta = \alpha$.

Typical results, based on 100,000 antithetic simulation trials, are shown in Table 4. These are for a 3×3 European swap option in which the holder has the right to pay floating. The average ten year interest rate is 5% and the average ten year volatility is

⁶ A more general version of the hypothesis is that swap rates follow the process in equation (32) for some value of β not necessarily equal to α . We tested whether we should use the more general hypothesis by defining a “best fit” value of β as the value for which the expression in equation (37) is least sensitive to perturbations in forward rates. In a wide range of situations this value of β was found to be very close to α . Encouragingly, the best fit value of β also led to the expression in equation (37) being almost exactly constant.

20%. The term structure is upward sloping and the volatility structure is humped. The tenor of the caplets underlying the model is three months and the swap underlying the European swap option is reset semiannually.

The results for $\alpha \neq 1$ are not quite as good as those for $\alpha = 1$, but they are still quite acceptable. The largest errors (about 0.4%) are for deep-in-the-money swap options. The prices of these options are relatively insensitive to volatilities and so these errors are not a cause for serious concern.

Figures 1 and 2 use the results in this section to compare cap volatility skews with swap option volatility skews in a variety of situations. The figures show flat volatilities for a five year cap and spot volatilities for 3×3 European swap option. In Figure 1 the term structure is flat at 5%; in Figure 2 it is upward sloping. Three different values of α are considered. The cap is reset quarterly; the swap underlying the swap option is reset semiannually. In Figures 1(a) and 2(a) the volatility structure is flat at 20%. In Figures 1(b) and 2(b) it has a hump similar to that observed in the market.

The figures show that care must be exercised in interpreting cap volatility skews. The latter are affected by the way in which the flat volatilities quoted by brokers are calculated. A flat volatility is the volatility that, when applied to all caplets, produces the market price. It is, therefore, a complex non-linear function of true caplet spot volatilities. As shown by the cap volatilities for the $\alpha = 1$ case in Figure 1(b), the function depends on the strike price. (If the function did not depend on the strike price, these cap volatilities would be the same for all strike prices.)

When the term structure is upward sloping (Figure 2), the cap results are affected by the fact that the moneyness of the early caplets is less than that of the later caplets. Consider, for example, the case where the volatility structure is flat and the cap is at the money (Figure 2a). The short maturity caplets are out of the money and the long maturity caplets are in the money. When $\alpha = 1$ they are all priced with the same volatility. When $\alpha = 0.5$, the short maturity caplets are priced with relatively low volatilities and long

maturity caplets are priced with relatively high volatilities. This increases the value of the cap. When $\alpha = 1.5$ the short maturity caplets are priced with relatively high volatilities and long maturity caplets are priced with relatively low volatilities. This decreases the value of the cap.

The figures show that, once an appropriate model has been chosen 3×3 swap option volatilities are relatively insensitive to the shape of the volatility term structure. Implied volatilities are slightly higher for the humped volatility structure. This is simply a reflection of the fact the volatilities during the first three years are slightly higher for the humped volatility structure we used than for a flat 20% volatility structure.

5. APPLICATIONS OF THE MODEL

There are two main applications of the model in Section 4. The first is to valuation and marking to market of non-standard interest rate derivatives. The second is to the valuation and marking to market of non-at-the-money European swap options. We will discuss both of these applications in this section. We will use for illustration U.S. data on caps and European swap options for August 12, 1999, kindly supplied to us by a major investment bank. The data was compiled by averaging mid-market quotes from a number of different brokers.⁷ It was used by the bank to mark to market its portfolio of interest rate derivatives on August 12, 1999. The cap data consists of 133 flat volatility quotes where the strike prices range from 3% to 10% and the cap lives range from one to ten years. The swap data consists of at-the-money volatility quotes for a range of different option maturities and swap lives. The caps were reset quarterly and the swaps were reset semiannually. The zero curve on August 12, 1999 was upward sloping with short rates approximately 5.3% and long rates approximately 7.2%.

The first step in implementing the model in Section 4 is to estimate α . This can be done by simultaneously estimating α and the Λ_j from cap prices.⁸ We search for the values of α and the Λ_j that minimize the root mean square cap pricing error

$$\sqrt{\frac{1}{K} \sum_{i=1}^K (u_i - v_i)^2} \quad (38)$$

where K is the number of caps for which market data is available, u_i is the market price of the i th cap calculated from the broker (flat) volatility quotes using equation (4), and v_i is the model price of the i th cap calculated using equations (30) and (31).

⁷ Because different brokers quote data in different ways, some judgment must be used in averaging quotes.

⁸ Note that cap prices depend only on the Λ_j in equation (28) — not on the way in which Λ_j is divided into its component parts, $\lambda_{j,q}$.

The best fit value of α on August 12, 1999 was found to be 0.716. This value of α reflects a situation where quoted cap volatilities declined, in absolute terms, by 3% to 4% as the strike price increased from 3% to 10%. The best fit values of the Λ_j obtained from this analysis of cap prices are shown in the Caps column of Table 5. As indicated by the table, we assumed a step function for Λ_j .⁹ This reflects the fact that the only cap maturities for which we had data were 1, 2, 3, 4, 5, 7, and 10 years.

A natural application of the LIBOR market model, once it has been calibrated to plain vanilla caps in the way just described, is to the valuation of non-standard cap products. Hull (2000) examines three such products: ratchet caps, sticky caps, and flexi caps. In a ratchet cap the strike rate for a caplet equals the LIBOR rate at the previous reset date plus a spread. In a sticky cap it equals the previous capped LIBOR rate plus a spread. In a flexi cap there is an upper bound to the total number of caplets that can be exercised. Hull finds that the prices of all three types of nonstandard caps are dependent on the number of factors. This is because their payoffs, unlike those of a plain vanilla cap, depend on the joint behavior of two or more forward rates. Hull's analysis used Monte Carlo simulation in conjunction with equation (22) to price the instruments. To investigate the impact of volatility skews we could carry out a similar analysis using Monte Carlo simulation in conjunction with equation (27).

To value non-standard swap option products, such as Bermudan swap options, the most appropriate procedure would seem to be to set α equal to its best fit value calculated from caps and choose the Λ_j so that they fit broker quotes on at-the-money European swap options. We did this using 30 swap option quotes where the option maturity ranged from 0.5 years to 5 years and the swap life ranged from 1 to 5 years. We used a three factor model, determining the $\lambda_{j,q}$ from the Λ_j as indicated in Table 1. The results are

⁹ When finding best fit Λ_j values we imposed some "smoothness" constraints to avoid large changes in the volatility of a forward rate from one period to the next.

shown in the Swap Option column of Table 5.

If the three-factor extended LIBOR market model were perfect, we would of course obtain the same parameter values regardless of the calibrating instruments being used. The reality is that all derivative models — even those involving several factors — have to be fine tuned to reflect the use to which they are being put. Table 5 shows that the Λ_j values obtained by fitting the model to swap options are different from those obtained by fitting it to caps. Specifically, the Λ 's implied from swap options are less humped than those implied from caps. This appears to be because the pricing of options on one and two year swaps does not fully reflect the hump observed in cap volatilities.

The root mean square pricing error for caps in Table 5 is greater than that for swap options. This reflects the fact that we are fitting the model to 133 caps and only 30 swap options. The relatively high average absolute percentage error for caps reflects both this and the existence of some errors in deep-out-of-the-money caps that, although small in absolute terms, were large when measured as percentages.

5.1 Incorporating Volatility Skews when European Swap Options are Priced

As already mentioned broker quotes are available only for at-the-money European swap options. This is probably because, when they are initiated, most European swap options are close to the money. However, financial institutions are faced with the problem of marking to market their portfolios each day. Although they are usually close to the money when they first trade, European swap options are liable to be significantly in or out of the money when marked to market at a later date.

The analysis in Section 4 makes it possible to calculate a complete volatility skew for European swap options from broker quotes on at-the-money instruments. The procedure is as follows.

1. Calculate the price of the at-the-money swap option from its volatility using equation (10) or (11)
2. Imply a value of $\sigma_{n,N}$ from the price of the at-the-money European swap option using

equation (33) or (35) with α set equal to its best fit value

3. Use this value of $\sigma_{n,N}$ to calculate the price of non-at-the-money European swap option using equation (33) or (35)
4. Imply a Black volatility for the non-at-the-money swap options using equation (10) or (11)¹⁰

As an example we consider a 5×5 European swap option on August 12, 1999. (For the purposes of calculating the volatility skew it does not matter whether it is an option to receive fixed or receive floating.) The broker quote for the volatility of this instrument is 17.58% and the at-the-money strike swap rate is 7.47%. The results of the above procedure are displayed in Figure 3.

Note that the calculation of a volatility skew for European swap options depends only on the best fit value of α (0.716 for the case considered) and the at-the-money volatilities. It does not require an estimate of the Λ_j or the $\lambda_{j,q}$. It should be relatively easy for financial institutions to store a value of α and update it periodically so that very fast volatility skew calculations can be made whenever required.

¹⁰ In theory, this last step is unnecessary since the third step provides the price. In practice it is likely to be desirable as most traders and analysts naturally think in terms of Black volatilities.

6. CONCLUSIONS

Caps and European swap options are quite different instruments. A cap is a portfolio of options; a European swap option is an option on an annuity. This paper has provided a simple robust procedure for relating the volatilities used by the market to price European swap options to the volatilities used by the market to price caps. Given the popularity of caps and European swap options, the procedure presents traders with opportunities to fine tune their pricing and search for arbitrage opportunities. A key contribution of the results in the paper is to make it very easy for traders to quickly calculate volatility skews for European swap options (which are not provided by brokers) from volatility skews for caps (which are provided by brokers).

The model underlying our results is the multifactor extended LIBOR market model proposed by Andersen and Andreasen (1997). The results we have produced make it possible to calibrate the model very quickly to either caps or European swap options. New numerical procedures developed by authors such as Andersen (1998), Broadie and Glasserman (1997), and Longstaff and Schwartz allow American-style as well as European-style non-standard interest rate derivatives to be priced in a multifactor setting once the calibration has been accomplished.

REFERENCES

- Andersen, L. "A Simple Approach to the Pricing of Bermudan Swaptions in the Multifactor LIBOR Market Model," Working paper, Gen Re Financial Products, 1998, forthcoming, *Journal of Computational Finance*.
- Andersen L. and J. Andreasen, "Volatility Skews and Extensions of the LIBOR Market Model," Working paper, Gen Re Financial Products, 1997, forthcoming, *Applied Mathematical Finance*.
- Black, F. "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3 (March 1976), 167–79.
- Brace, A, D. Gararek, and M. Musiela. "The Market Model of Interest Rate Dynamics," *Mathematical Finance*, 7, 2 (1997), 127–55.
- Broadie, M. and P. Glasserman. "A Stochastic Mesh Method for Pricing High Dimensional American Options," Working Paper, Columbia University, New York, 1997.
- Harrison, J. M. and D. M. Kreps. "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20 (1979), 381–408.
- Heath, D., R. Jarrow, and A. Morton. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology," *Econometrica*, 60, 1 (1992), 77–105.
- Hull, J. C. *Options, Futures, and Other Derivatives*, Fourth Edition, Prentice Hall, Englewood Cliffs, NJ, 2000.
- Jackwerth, J.C. and M. Rubinstein, "Recovering Probability Distributions from Option Prices," *Journal of Finance*, 51 (December 1996), 1611–31
- Jamshidian, F. "The LIBOR and Swap market Models," *Finance and Stochastics*, 1 (1997), 293–330.
- Longstaff, F. and E. Schwartz. "Valuing American Options by Simulation: A Least Squares Approach," Working Paper 25-98, 1998, Andersen School at UCLA.
- Miltersen, K., K. Sandmann, and D. Sondermann. "Closed Form Solutions for Term Structure Derivatives with Lognormal Interest Rates," *Journal of Finance*, 52, 1 (March 1997), 409–30.
- Rebonato, R. "On the Simultaneous Calibration of Multifactor Lognormal Interest rate Models to Black Volatilities and to the Correlation Matrix," *Journal of Computational Finance*, 2, 4 (Summer 1999), 5–27.

Table 1**Values of $\lambda_{j,q}/\Lambda_j$ Used**

The table shows values for $\lambda_{j,q}/\Lambda_j$, the ratio of the q th component of the volatility of a forward rate to the total volatility of the forward rate. These were used to produce the results in Tables 3, 4, and 5.

q	Forward Rate Maturity (yrs)								
	1	2	3	4	5	6	7	8	9
1	0.852	0.917	0.964	0.982	0.964	0.982	0.964	0.917	0.852
2	0.467	0.377	0.264	0.135	0.000	-0.135	-0.264	-0.377	-0.467
3	0.236	0.127	0.000	-0.136	-0.267	-0.136	0.000	0.127	0.236

Table 2**Accuracy of Caplet Volatilities**

The table shows the amount by which the true spot caplet volatility exceeds the volatility given by the Monte Carlo implementation of the LIBOR market model in equation (22). For example, when the average rate is 1%, the average volatility is 50%, and the true caplet volatility is $V\%$, the Monte Carlo simulation produces a price consistent with a volatility of $V + 0.08\%$. The standard error of the volatility difference is in parentheses. Results are for 200,000 antithetic simulation trials, an upward sloping term structure, and a humped volatility structure. The caplet has a maturity of five years and tenor of one year. All volatilities are measured as percentages.

Average Rate	Average Volatility			
	10%	20%	50%	100%
1%	-0.01 (0.02)	-0.01 (0.05)	-0.08 (0.25)	0.25 (1.19)
5%	-0.02 (0.02)	-0.02 (0.05)	-0.56 (0.17)	2.57 (0.60)
10%	-0.02 (0.02)	-0.01 (0.04)	-0.02 (0.12)	9.42 (0.44)
20%	-0.01 (0.02)	-0.02 (0.03)	0.39 (0.08)	14.64 (0.43)
50%	0.00 (0.01)	0.01 (0.02)	1.65 (0.11)	14.55 (0.52)

Table 3**Accuracy of Swap Option Volatilities**

The table shows the amount by which 5×5 swap option spot volatilities, calculated using the approximation in Section 3, exceed those calculated using the Monte Carlo implementation of the LIBOR market model in equation (22). For example, when the average rate is 1%, the average volatility is 50%, and the swap option volatility calculated from the approximation in section 3 is $V\%$, the Monte Carlo simulation produces a price consistent with a volatility of $V - 0.05\%$. The standard error of the Monte Carlo estimate is in parentheses. Results are for 200,000 antithetic simulation trials, an upward sloping term structure, and a humped volatility structure. Both the caps underlying the LIBOR market model and the swaps underlying the swap option are reset annually. All volatilities are measured as percentages.

Average Rate	Average Volatility			
	10%	20%	50%	100%
1%	-0.01 (0.01)	-0.02 (0.02)	0.05 (0.02)	1.00 (0.06)
5%	-0.01 (0.01)	0.00 (0.02)	0.14 (0.04)	3.29 (0.05)
10%	0.00 (0.02)	0.02 (0.03)	0.34 (0.08)	5.11 (0.16)
20%	-0.01 (0.02)	0.04 (0.04)	1.01 (0.14)	4.13 (0.35)
50%	0.01 (0.03)	0.24 (0.07)	-0.65 (0.35)	-71.24 (4.13)

Table 4

Accuracy of Swap Option Volatilities When There is a Volatility Skew

The table shows the amount by which 3×3 swap option volatilities calculated from the approximation in Section 4 exceed those calculated using the Monte Carlo implementation of the extended LIBOR market model in equation (27). For example, when α is 0.50, the strike rate is 80% of the forward par yield 1%, and the lognormal swap volatility calculated using Section 4 is $V\%$, the Monte Carlo simulation produces a price consistent with a lognormal volatility of $V - 0.19\%$. The standard error of the Monte Carlo price is in parentheses. Results are for 100,000 antithetic simulation trials, an upward sloping term structure and a humped volatility structure. Caps used to define the LIBOR market model are reset quarterly and the swap underlying the swap option is reset semiannually. All volatilities are measured as percentages.

α	Strike Rate (% of Forward Par Yield)				
	60	80	100	120	140
0.10	0.15 (0.05)	0.03 (0.05)	-0.03 (0.04)	-0.02 (0.04)	-0.01 (0.05)
0.25	0.34 (0.05)	0.14 (0.04)	-0.11 (0.04)	-0.11 (0.04)	-0.01 (0.04)
0.50	0.38 (0.04)	0.19 (0.04)	0.09 (0.04)	0.09 (0.03)	0.12 (0.03)
0.75	0.17 (0.04)	0.05 (0.04)	0.01 (0.04)	0.01 (0.04)	0.05 (0.02)
1.00	0.09 (0.04)	0.02 (0.04)	-0.01 (0.04)	0.01 (0.03)	0.06 (0.02)
1.50	0.21 (0.04)	0.04 (0.04)	0.01 (0.03)	0.05 (0.02)	0.09 (0.02)

Table 5**Best Fit Values of Volatility Parameters on August 12, 1999**

The table show best fit values for the volatility parameters, Λ_j when the model defined by equations (26) to (28) is fitted to 133 caps (column 2) and 30 European swap options (column 3). The CEV parameter α was set equal to its best fit value of 0.716. Principal=\$1,000 and $\delta_j = 0.25$ for all j .

	Values of Λ_j for	
	Caps	Swap Options
$1 \leq j \leq 3$	0.0795	0.0866
$4 \leq j \leq 7$	0.1300	0.1076
$8 \leq j \leq 11$	0.1274	0.0890
$12 \leq j \leq 15$	0.0795	0.0890
$16 \leq j \leq 19$	0.0756	0.0890
$20 \leq j \leq 27$	0.0756	0.0808
$28 \leq j \leq 39$	0.0756	0.0730
RMS Pricing Error (\$)	0.57	0.21
Mean Abs % Pricing Error	5.72	1.27

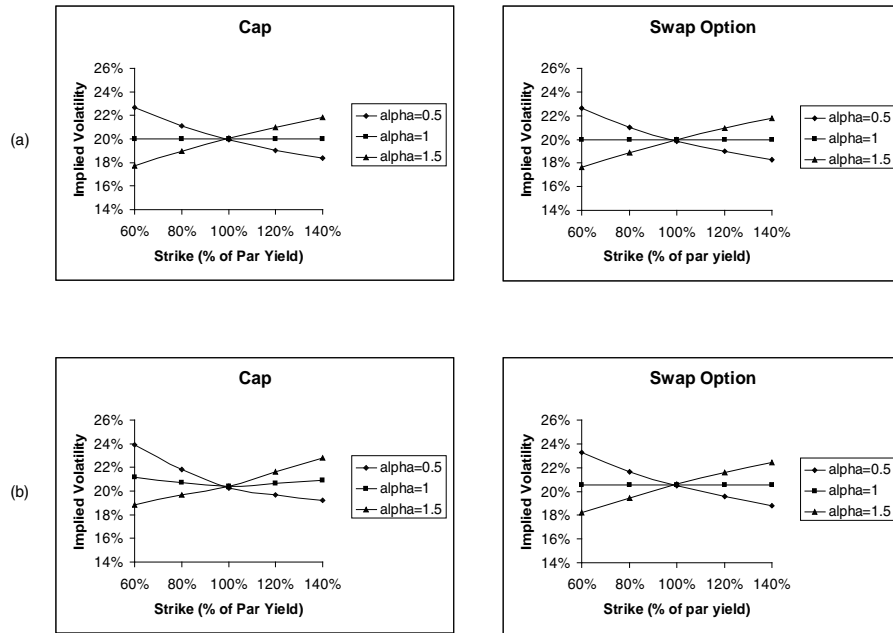


Figure 1

Volatility skews for a flat term structure. The cap is reset quarterly and has a life of 3 years. The European swaption is a 3 year option into a 3 year swap that is reset semiannually. Results are based on the model in Section 4. In (a) the volatility structure is flat. In (b) it is humped.

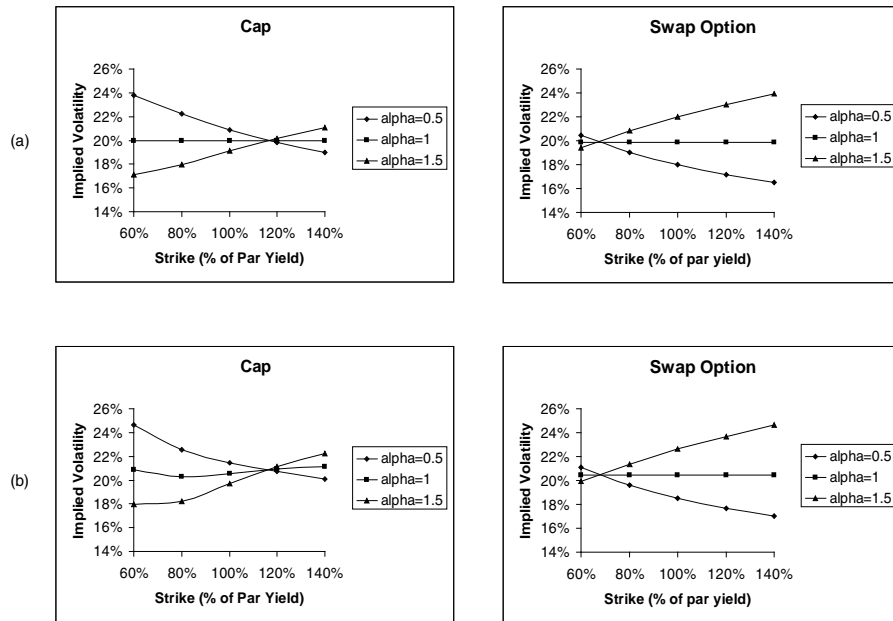


Figure 2

Volatility skews for an upward sloping term structure. The cap is reset quarterly and has a life of 3 years. The European swaption is a 3 year option into a 3 year swap that is reset semiannually. Results are based on the model in Section 4. In (a) the volatility structure is flat. In (b) it is humped.

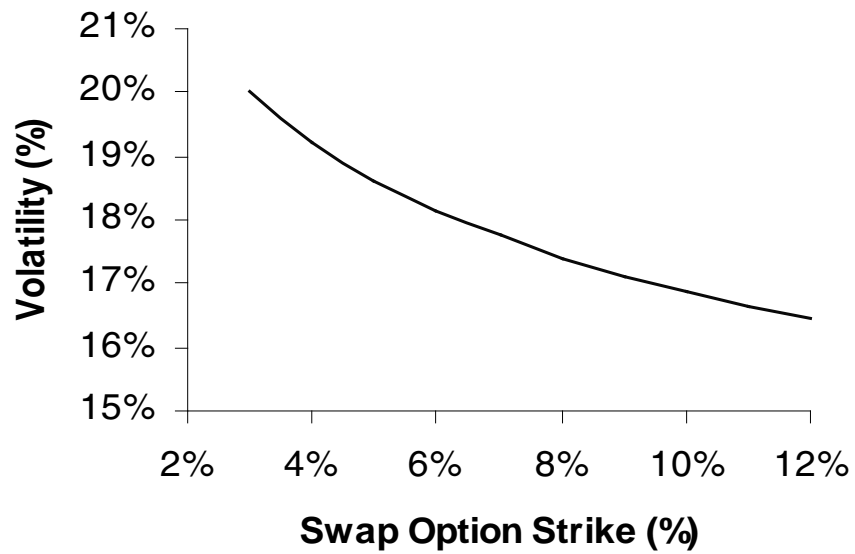


Figure 3

Volatility skew for 5 × 5 swap options on August 12, 1999. This is based on a value of 0.716 for the CEV parameter, α (estimated from cap data) and a broker quote of 17.58 % for the volatility of an at-the-money 5 × 5 swap option