On the Information in the Interest Rate Term Structure and Option Prices

Frank de Jong

Joost Driessen

Antoon Pelsser

This version: October, 2003

Frank de Jong, Finance Group, University of Amsterdam, Roetersstraat 11, 1018 WB, Amsterdam, The Netherlands. Phone: +31-20-5255815. E-mail: fdejong@fee.uva.nl.
Corresponding author: Joost Driessen, Finance Group, University of Amsterdam, Roetersstraat 11, 1018 WB, Amsterdam, The Netherlands. Phone: +31-20-5255263. E-mail: j.j.a.g.driessen@uva.nl.
Antoon Pelsser, Econometric Institute, Erasmus University Rotterdam, and Actuarial Dept, Nationale-Nederlanden. Phone: (+31)-10-4081259, E-mail: pelsser@few.eur.nl.

We thank Wolfgang Buhler, David Lando, Bertrand Melenberg, Theo Nijman, Marti Subrahmanyam (the editor), and two anonymous referees for helpful comments and suggestions, as well as seminar participants at University of Leuven, Aarhus University, Odense University, and Copenhagen Business School.

On the Information in the Interest Rate Term Structure and Option Prices

Abstract

Cap and swaption prices contain information on interest rate volatilities and correlations. In this paper, we examine whether this information in cap and swaption prices is consistent with realized movements of the interest rate term structure. To extract an option-implied interest rate covariance matrix from cap and swaption prices, we use Libor market models and discrete-tenor string models as a modelling framework. We propose a flexible parameterization of the interest rate covariance matrix, which cannot be generated by standard low-factor term structure models. The empirical analysis is based on weekly US data from 1995 to 1999. Our empirical results show that the option prices imply a covariance matrix of interest rates that is significantly different from the covariance matrix implied by realized interest rate changes. In particular, if one uses the latter covariance matrix to price caps and swaptions, one significantly underprices these options. We discuss and analyze several explanations for our findings.

JEL Codes: G12, G13, E43.

Keywords: Term Structure Models; Interest Rate Derivatives; Volatility Hump.

One of the central questions in option pricing is whether the information reflected in option prices is consistent with the time series behavior of the underlying security. One set of articles in this field focuses on analyzing whether implied volatility is an efficient and unbiased predictor of realized volatility (Amin and Ng (1997), Canina and Figlewski (1993), Christensen and Prabhala (1998)). Another series of articles compares the risk-neutral density implied by option prices with the density estimated using time series data on the underlying security (Ait-Sahalia, Wang, and Yared (2001), Jackwerth and Rubinstein (1996)). In general, the conclusion is that there are significant differences between the time series behavior of the underlying security and the cross section information in option prices.

All these articles focus on the case where there is one underlying security, typically a stock, futures price, or currency rate. We contribute to this literature by analyzing the interest rate option market, specifically, the market for caps and swaptions. This market is one of the largest OTC option markets. In contrast to the equity and currency option markets, there are many underlying securities in the interest rate option market (swaps and bonds of different maturities) which are strongly interrelated. In addition to the volatility of interest rates, their mutual correlations are crucial for derivative prices (as noted by, for example, Rebonato (1996)). Our focus is thus the covariance matrix of interest rates of different maturities. The goal of this paper is to examine what cap and swaption prices imply regarding the covariance matrix of interest rate changes of different maturities, and, second, whether these option price implications are consistent with the covariance matrix estimated from the time series of interest rates.

To 'invert' cap and swaption prices to option-implied interest rate variances and correlations, we use the Libor market models (Brace, Gatarek, and Musiela (1997), Miltersen, Sandmann, and Sondermann (1997), and Jamshidian (1997)), which are equivalent to the discrete-tenor case of the string term structure models (Longstaff, Santa-Clara, and Schwartz (2001) and Santa-Clara and Sornette (2001)).

This framework implies that forward Libor rates of different forward maturities have a joint lognormal distribution.

An important aspect of this analysis is the parameterization of the interest rate covariance matrix associated with the joint lognormal distribution. In line with Rebonato (1996), we show that standard low-factor term structure models do not generate satisfactory shapes for the correlation structure of interest rates of different maturities. Therefore, we avoid imposing a factor structure on the model, and use a full-factor model. We directly parameterize the covariance matrix of this full-factor model using a flexible specification.

Next, based on this model specification, we derive moment conditions that interest rate data and option price data should satisfy. These moment restrictions involve variances and covariances of forward Libor rate changes of different maturities, as well as average cap and swaption prices. We use the Generalized Method of Moments (Hansen (1982)) for estimation and testing, and allow for the presence of measurement error in the cap and swaption prices. The moment restrictions are estimated using weekly US data on Libor and swap rates and prices of caps and swaptions from 1995 until 1999.

We compare the information in option prices and interest rates in two ways. First, we estimate the parameters in the covariance matrix specification using interest rate data only ('historical estimation'), and analyze the implications for cap and swaption prices. This leads to cap and swaption prices that are, on average, significantly lower than the observed prices. A test on the joint pricing restrictions for caps and swaptions strongly rejects that caps and swaptions are priced consistently with the historical interest rate covariance matrix.

Second, we estimate the interest rate covariance matrix using cap and swaption data only, and compare this option-implied covariance matrix to the historically estimated covariance matrix. The results show that the option-implied interest rate volatility term structure is hump shaped, which is in line with the shape of the historically estimated volatility term structure. However, the option-implied hump is higher and steeper than the historically estimated hump. Furthermore, the option-implied interest rate correlations are much lower than the realized (historically estimated) interest rate correlations for short maturities, and higher for longer maturities. We also statistically test whether the option-implied covariance matrix is equal to the historically estimated covariance matrix . The test results show a rejection of these restrictions.

We analyze and reject two explanations for our results: (i) the high option prices during the Russia/LTCM crisis, and (ii) the presence of measurement error in the forward Libor rates. In addition, our results are robust to an alternative specification of the forward Libor rate probability distribution. We discuss other potential explanations in the final section of the paper.

Our paper is related to Longstaff, Santa-Clara, and Schwartz (LSS, 2001), who use a two-step estimation procedure to estimate a four-factor model for caps and swaptions. They extract eigenvectors from a historically estimated correlation matrix, and subsequently the eigenvalues are calibrated to swaption prices at each day in the dataset. LSS focus on whether cap prices are consistent with swaption prices. They do not study the consistency of the interest rate volatility term structure and cap prices. A by-product of their estimation results is that the correlations, implied by their two-step procedure, are all lower than the historically observed correlations. A potential problem with the approach of LSS is that the daily re-calibration is inconsistent with the model (which has constant parameters over time) and implies that the option-implied correlations change from day to day. Also, the LSS two-step procedure can lead to strange shapes for the correlation matrix, since it results from linear combinations of the eigenvectors. In contrast, we do not rely on a two-step procedure and directly parameterize the interest rate covariance matrix. Also, we do not re-calibrate our model every day, but base our analysis on moment conditions that involve the time series averages of cap and swaption prices.

Jagannathan, Kaplin, and Sun (JKS, 2001) and Han (2001) also find discrepancies between interest rate data and prices of caps and swaptions. JKS estimate three-factor affine models using interest rate data and cap prices, and report large pricing errors for caps and swaptions. These results may be driven by the restrictive covariance matrix structure implied by the three-factor CIR model. Han (2001) estimates a two-factor model with stochastic bond price volatilities using interest rate data and swaption prices. He shows that including stochastic bond price volatilities and correlations is important to explain swaption prices and interest rate covariances. Still, cap prices implied by this model are not completely consistent with observed cap prices.

The remainder of this paper is organized as follows. Section 1 discusses and motivates the modeling framework. Section 2 describes the interest rate data and option price data, as well as the estimation methodology. Section 3 contains the results on the comparison of option and interest rate data, and discusses a trading strategy based on the results. Section 4 concludes and discusses possible explanations for our results.

1 Extracting Information from Cap and Swaption Prices

1.1 Setup

We start with a short review of caps and swaptions. We use a finite set of dates $T_1 < T_2 < ... < T_N$, the so-called *tenor structure*. We also define $\delta_i = T_{i+1} - T_i$, $i = 1,..,N-1$ as the so-called daycount fraction, which is equal to the maturity of the Libor rate that is used to determine caplet payoffs and is most often equal to 3 or 6 months. Associated with each tenor date T_n is a bond that matures at this date, and its time t price is denoted with $P(t,T_n)$. These N bond prices, with maturities $T_1,...,T_N$, determine $(N-1)$ forward Libor rates. The forward Libor rate $L(t,T_n)$ is defined by $L(t,T_n) = \left(P(t,T_n)/P(t,T_{n+1}) - 1\right)/\delta_n$.

A *caplet* with strike rate k and maturity date T_n pays off $\delta_n(L(T_n,T_n) - k)^+$ at time T_{n+1}. In general, the price of this caplet at time t can be calculated using the expectation of the discounted payoff under the so-called forward martingale measure Q^{n+1}

$$Caplet(t, T_n, k) = P(t, T_{n+1}) \, E_t^{n+1} [\, \delta_n (L(T_n, T_n) - k)^+ \,] \tag{1}$$

A *cap* is a sum of caplets of different maturities. The expression in (1), which is of course similar to the price equation for equity options, implies that the volatility of the forward Libor rate $L(t,T_n)$ is the most important determinant of the caplet price.

A *swaption* is an option on a swap. Consider a forward swap, with principal 1, where two parties agree to exchange at dates $\{T_{n+1},...,T_{n+m}\}$ the floating Libor rates $\{L(t,T_n),...,L(t,T_{n+m-1})\}$ for a fixed rate. The forward swap rate is the fixed rate that gives this contract zero initial value and is given by

$$S(t,T_n,T_m) = \frac{\sum_{j=1}^{m} \delta_{n+j-1} P(t,T_{n+j}) \, L(t,T_{n+j-1})}{\sum_{j=1}^{m} \delta_{n+j-1} P(t,T_{n+j})} = \frac{P(t,T_n) - P(t,T_{n+m})}{\sum_{j=1}^{m} \delta_{n+j-1} P(t,T_{n+j})} \tag{2}$$

A payers swaption with strike rate k, maturity date T_n and m payment dates gives right to enter into a swap at date T_n, where floating Libor payments are received and fixed payments k are paid. Equivalently, a payers swaption gives the right to receive a series of cash flows $\delta_{n+j-1}(S(T_n,T_n,T_m) - k)^+$ at dates T_{n+j}, $j=1,..,m$ (see Musiela and Rutkowski (1997)).

Equation (2) shows that a forward swap rate depends on several forward Libor rates, so that the variance of a swap rate is a function of both the variances and covariances (or correlations) of forward Libor rates. Swaption prices thus contain information on both the variances and covariances of forward Libor rates of different maturities, whereas caplet prices only contain information on the variance of a single forward Libor rate.

1.2 Libor Market Models

As outlined in the introduction, our aim is to compare the historically estimated interest rate variances and covariances with the variances and covariances implied by caps and swaptions. In order to 'invert' the cap and swaption prices to interest rate variances and covariances, we choose the Libor market model (LMM) as modelling framework, which is described in this subsection.

The LMM assumes lognormal processes for forward Libor rates. As shown by Kerkhof and Pelsser (2002), the LMM framework is equivalent to the discrete-tenor string model of Longstaff, Santa-Clara, and Schwartz (2001). Our option price data consist of implied Black (1976) volatility quotes for caps and swaptions, and the LMM implies simple Black-type pricing formulas for caps (and approximate pricing formulas for swaptions). This facilitates the estimation of the model. In De Jong, Driessen, and Pelsser (2001) other advantages of the market models are mentioned.

We analyze a LMM where each forward Libor rate is driven by its own factor. These factors are allowed to be correlated across forward maturities. Such a LMM implies that the forward Libor rate $L(t,T_n)$ satisfies the following Itô process under the true probability measure

$$\frac{dL(t,T_n)}{L(t,T_n)} = \mu(t,T_n)\,dt + \sigma(t,T_n)\,dW_n(t), \quad n=1,...,N-1 \tag{3}$$

The function $\mu(t,T_n)$ is the drift function of the forward Libor rate, and $\sigma(t,T_n)$ is a one-dimensional function (the volatility parameter) for the forward Libor rate with maturity date T_n. $W_n(t)$ is a standard Brownian motion. The Brownian motions that drive the different forward Libor rates are allowed to be correlated: the correlation between $W_i(t)$ and $W_j(t)$ is denoted by $\rho(t,T_i,T_j)$.

By choosing one of the N bonds as the numeraire asset, we can obtain the process of the forward Libor rates under the equivalent martingale measure associated with this numeraire choice. Under such an equivalent martingale measure, the drift of the forward Libor rates is completely determined by the volatility and correlation parameters, see Jamshidian (1997). For example, if we take the longest maturity bond $P(t,T_N)$ as the numeraire, we obtain the so-called terminal measure Q^N, under which forward Libor rates follow the process

$$\frac{dL(t,T_n)}{L(t,T_n)} = -\sum_{i=n+1}^{N-1} \frac{\delta_i L(t,T_i)\,\sigma(t,T_i)\,\sigma(t,T_n)\,\rho(t,T_i,T_n)}{1+\delta_i L(t,T_i)}\,dt + \sigma(t,T_n)\,dW_n^*(t), \quad n=1,\ldots,N-1 \qquad (4)$$

where $W_n^*(t)$ is a one-dimensional Brownian motion under the terminal measure. Note that, when changing the probability measure, the correlation structure of the Brownian motions remains unchanged.

We refer to Brace, Gatarek, and Musiela (1997) and Jamshidian (1997) for the pricing formulas for caps and swaptions. Most importantly, these formulas show that cap prices depend on conditional variances of forward Libor rates, whereas swaption prices both depend on conditional variances and covariances of forward Libor rates. The pricing formula for swaptions given by Brace, Gatarek, and Musiela (1997) is an approximate pricing formula. For our empirical analysis, we use this approximate pricing formula for estimation. To calculate swaption prices at the final parameter estimates, we simulate the LMM using an Euler discretization of (4).[1] Our results indicate that the approximate pricing formula of Brace, Gatarek, and Musiela (1997) is quite accurate: the maximum difference between the simulation price and the analytical approximation in our analysis is 0.10 Black volatility points (which turns out to be roughly 0.8% of the price). On average, the analytical approximation yields prices that are slightly lower than the prices obtained by simulation, but the average difference is only 0.04 volatility points. These errors seem small compared to the bid-ask spread of around 6% that is typically found in the swaptions market (see Longstaff, Santa-Clara, and Schwartz (2001)).

Equations (3) and (4) imply a simple structure on the variances and covariances of instantaneous changes in log-forward Libor rate changes

$$Cov(\, d\ln L_i(t),\, d\ln L_j(t)) \;=\; \rho(t,T_i,T_j)\,\sigma(t,T_i)\,\sigma(t,T_j)\,dt, \quad i,j = 1,...,N-1 \tag{5}$$

This relation will later be used to derive moment conditions for the historically estimated variances and covariances of forward interest rates.

1.3 Specification of Volatility and Correlation Parameters

Equation (4) implies that, in order to price and hedge interest rate derivatives, only the volatility parameters $\sigma(t,T_n)$, $n=1,...,N-1$ and the correlation parameters $\rho(t,T_i,T_j)$, $i,j=1,...,N-1$ have to be determined. In this subsection we propose a parameterization for the volatility and correlation parameters. Following Santa-Clara and Sornette (2001), we directly parametrize the volatilities and correlations, as opposed to the usual approach that specifies the dynamics of the underlying factors.

If the volatility and correlation parameters explicitly depend on time t, these parameters become time-inhomogenous, which would make the comparison between interest rate data and option prices problematic. Furthermore, all standard models of the term structure, such as the affine models of Duffie and Kan (1996), imply time-homogenous volatility and correlation parameters (i.e., volatility and correlation parameters than only depend on the time to maturity). Therefore, we impose time-homogeneity on the specification for these parameters

$$\sigma(t,T_i) \;=\; \sigma(T_i-t), \quad \rho(t,T_i,T_j) \;=\; \rho(T_i-t,T_j-t),\;\; i,j=1,...,N-1 \tag{6}$$

where $\sigma(\cdot)$ and $\rho(\cdot,\cdot)$ are functions that remain to be specified. The specification of these functions should be such that it allows for large variety of volatility and correlation shapes across maturities. In particular, the functions should be able to generate some of the particular features of interest rate volatilities and correlations, such as a humped shape term structure of volatilities, as well as interest rate correlations that decrease with the difference between the two associated maturities.

For the volatility parameters, we use the following parameterization[2]

$$\sigma(T-t) = \sigma_0 + \sigma_1 \exp(-\kappa_1 (T-t)) + \sigma_2 \exp(-\kappa_2 (T-t)) \tag{7}$$

This parameterization can generate both decreasing, increasing, and hump shaped volatility structures. If $\sigma_0 = \sigma_2 = 0$, we obtain the volatility function implied by an one-factor Vasicek (1977) model, that is also used by Santa-Clara and Sornette (2001). Our specification can be seen as an extension of this specification that allows for a humped volatility structure. More precisely, a humped shape can be obtained if σ_1 and σ_2 are of opposite sign, given that κ_1 and κ_2 are both positive.

We use the following flexible form for the correlation structure[3]

$$\rho(T_i-t,T_j-t) = \exp(-\gamma_1|T_i-T_j| - \frac{\gamma_2|T_i-T_j|}{\max(T_i-t,T_j-t)^{\gamma_3}} - \gamma_4|\sqrt{T_i-t}-\sqrt{T_j-t}|), \quad \gamma_1,\gamma_2,\gamma_4>0 \tag{8}$$

This specification extends the correlation functies studies by Santa-Clara and Sornette (2001), who analyze the case $\gamma_2 = \gamma_4 = 0$ and the case $\gamma_1 = \gamma_2 = 0$. Our specification captures the effect that correlations decrease with the maturity difference, and that this maturity decay differs for short and long forward maturities. In Figure 1, we graph the influence of the several parameters in (8) on the correlation structure, by plotting the partial derivatives of the correlations with respect to the parameters (at the parameter estimates obtained by so-called joint estimation, see Section 3.2). Increasing the parameter γ_1 causes correlations to decrease, where the size of the decrease is positively related to the maturity difference. Through the parameters γ_2 and γ_3, the specification in (8) allows for a correlation decay that differs across maturities. In particular, if $\gamma_3>0$, the correlation decay for longer maturities is smaller than for short maturities, and if $\gamma_3<0$, the converse is true. Figure 1 shows that increasing this parameter leads to lower correlations for short maturities, and higher correlations for longer maturities. The parameter

γ_4 implies a decay of the correlation function that is, relative to the influence of γ_1, stronger for small maturity differences and smaller for larger maturity differences. This turns out to be important to fit interest rate and option price data.

The above specification is based on a structure where the number of Brownian motions is equal to the number of forward Libor rates. In a large part of the term structure literature, models with two or three factors are estimated. In the empirical analysis we provide a comparison with a three-factor model. We choose the following time-homogeneous specification for this three-factor model

$$Cov(d\ln L_i(t), d\ln L_j(t)) =$$

$$\begin{bmatrix} \alpha_1 e^{-\beta_1(T_i-t)} & \alpha_2 e^{-\beta_2(T_i-t)} & \alpha_3 e^{-\beta_3(T_i-t)} \end{bmatrix} \begin{bmatrix} 1 & \theta_{12} & \theta_{13} \\ \theta_{12} & 1 & \theta_{23} \\ \theta_{13} & \theta_{23} & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 e^{-\beta_1(T_j-t)} \\ \alpha_2 e^{-\beta_2(T_j-t)} \\ \alpha_3 e^{-\beta_3(T_j-t)} \end{bmatrix} dt, \quad i,j = 1,\ldots,N-1 \qquad (9)$$

We allow for an unrestricted correlation matrix, with elements θ_{ij}, $i,j = 1,\ldots,3$, for the Brownian motions[4], since Dai and Singleton (2000) illustrate that allowing for nonzero correlations between factors in (affine) term structure models is important for accurately describing US interest rate behavior. The three-factor model in (9) has exponentially downward sloping volatility functions. These volatility functions are very similar to the volatility functions implied by affine term structure models of Duffie and Kan (1996), in particular, the *K*-factor version of the Vasicek (1977) model. Note that the number of parameters of this model is exactly equal to the number of parameters in our full-factor specification in equations (7) and (8).

2 Data and Testing Methodology

2.1 Data

We use two data sets: one data set containing US money-market rates and swap rates and another data set containing implied Black (1976) volatilities of US caps and swaptions.

The derivatives data that we use are weekly quotes for the implied Black (1976) volatilities of at-the-money-forward (ATMF) US caps and swaptions. For these data we have 232 weekly observations from January 1995 until June 1999. The caps have maturities ranging from 1 to 10 years, and their payoffs are defined on 3-month Libor rates. The 1-year cap consists of 3 caplets with maturities of 3, 6, and 9 months, and the 10-year cap consists of 39 caplets, with maturities ranging from 3 months to 9 years and 9 months. The other caps are constructed in a similar way. The strike of each ATMF cap is equal to the corresponding swap rate with quarterly compounding. Note that this implies that the caplets of the cap are not exactly at-the-money-forward. The Black implied volatilities for caps are so-called flat volatilities. This implies that the price of the cap is obtained by applying the Black formula to each caplet using the same volatility parameter (the quoted volatility for the cap) for all caplets. In Figure 2 we plot the time series average of the implied volatilities of the caps. There is evidence for a hump shaped volatility structure.

For swaptions, we use three option maturities, 3 months, 1 year, and 5 years, and four swap maturities, 1 year, 3 years, 5 years, and 7 years. Since the 5x7 swaption is not in our dataset, we end up with 11 different swaptions.[5] The strike of an at-the-money-forward swaption is equal to the corresponding forward swap rate. In Figure 3, we plot the time series averages of the swaption implied volatilities. Again, there is evidence for a volatility hump at the short end of the maturity axis.

The second dataset involves US interest rates. We use US money-market rates with maturities of 3, 6, 9, and 12 months, and data on US swap rates with maturities ranging from 2 to 10 years to estimate the forward Libor rate curve using a piecewise constant specification for this forward Libor rate curve (for a 3-month Libor maturity of each forward Libor rate), where the forward Libor rates are constant between the maturities of the observed money-market and swap rates. This way, we obtain a perfect fit of the observed money-market and swap rates. In Section 3.4, we examine whether it is likely that there is measurement error in these interest rates.

For the interest rate data, we use weekly data from January 1995 until June 2000. Because options are forward-looking, in the sense that they contain information on the 'expected' interest rate variances and covariances, we extend the interest rate data period one year beyond the last observation on derivative prices in June 1999. This is similar to Christensen and Prabhala (1998) and others, who use realized volatility of the underlying equity price to assess the predictive value of option-implied equity volatility.

In Figure 4 the annualized standard deviations of changes in the log-forward Libor rates are plotted. In line with results presented in Amin and Morton (1994), and Moraleda and Vorst (1997), there is evidence for a humped volatility structure for forward Libor rate changes. In Figure 5, we graph the correlation matrix of weekly changes in the logarithm of these forward Libor rates, estimated directly from the data without restrictions as in equation (8). Correlations are typically quickly decreasing in the maturity difference, and Libor rates with longer forward maturities are somewhat more interrelated than short-maturity forward Libor rates.

2.2 Estimation and Testing Methodology

In this subsection, we derive moment conditions for both the caps and swaptions data and the interest rate data. The moment conditions are used to estimate the parameters in the specification of the volatilities and correlations in equations (7) and (8), and allow us to analyze the consistency of the information in cap and swaption prices and interest rates. We use two sets of moment restrictions:

1. Variances of log forward Libor rate changes and covariances between log forward Libor rate changes of different forward maturities.

2. Expected (squared) cap implied volatilities and swaption implied volatilities.

We refer to estimation on the basis of the first set of moments as *interest-rate-based estimation* or *historical estimation*, and to estimation on the basis of the second set of moments as *option-based estimation* or *implied estimation*. The use of both sets of moment restrictions is referred to as *joint estimation*. All moment restrictions are formulated under the true probability measure.

The first set of moment restrictions is based on a time-discretization of equation (5), which gives

$$Cov(\Delta \ln L_i(t), \Delta \ln L_j(t)) = \rho(T_i - t, T_j - t)\,\sigma(T_i - t)\,\sigma(T_j - t)\,\Delta t, \quad i,j = 1,\dots,N-1 \qquad (10)$$

By approximation, this relation holds for small time intervals Δt. This approximate relation is only exact if the drift of the log forward Libor rates is deterministic. Since we use weekly time intervals, the variation in the drift rate is most likely low, so that this approximation is accurate.

Using data on the log forward Libor rates, we can estimate the left hand side of equation (10) and confront these estimates with the model-implied (co)variances on the right hand side. For estimation, we annualize the (co)variances by multiplying (10) with $1/\Delta t$ such as to obtain the same scaling for these restrictions as the implied volatilities (see below).[6]

To derive the moment restrictions for derivative prices, we assume that the (square of the) observed implied Black volatility quote for a cap or swaption is equal to the (square of the) Black volatility that corresponds to the model price, plus an independent zero-expectation error term, that represents measurement error in the observed implied volatility quote. For caps, we thus get

$$[IV^{C}(t,T_i)]^2 = [IV^{C,Model}(t,T_i))]^2 + \eta_i(t), \quad E(\eta_i(t)) = 0 \tag{11}$$

where $IV^{C}(t,T_i)$ is the observed implied flat volatility for the cap with maturity T_i. $IV^{C,Model}(t,T_i)$ is the Black flat volatility for this cap implied by the model. At a given time t, the latter volatility is obtained as follows. First, given certain parameter values for the model in (7) and (8), we calculate model-implied prices for all caplets. Summing over all relevant caplet prices gives the price of the cap. This model-implied cap price is then inverted to a single Black flat volatility parameter $IV^{C,Model}(t,T_i)$, which can be directly compared to the quoted flat volatility that is observed in the data.

For swaptions, the moment restriction are completely similar to (11). Again, the model-implied swaption price is inverted to a Black volatility parameter, which can be compared with the quoted Black volatility for the swaption. We take the square of the implied volatilities so that these moment restrictions are measured with the same scale as the restrictions in (10). By taking the unconditional expectation on both sides of equation (11) we obtain moment restrictions for caps and swaptions[7].

As noted above the prices of caps depend on the conditional variances of Libor rates, whereas the prices of swaptions depend both on conditional variances and covariances of forward Libor rates of different forward maturities. Thus, both sets of moment restrictions involve (conditional) variances and covariances of forward Libor rates, and from both sets of moment restrictions it is possible to identify all volatility and correlation parameters.

We use the Generalized Method of Moments (Hansen (1982)) to estimate the parameters in the volatility and correlation specification in equations (7) and (8). For the forward Libor rate variance restrictions, we choose the following forward maturities (in years): 0.25, 0.5, 0.75, 1.5, 2.5, 3.5, 4.5, 6.5, and 8.5, in total 9 moment restrictions. The first three of these maturities are equal to the money-market rate maturities. Since we use a piecewise constant forward Libor rate curve, we choose for the remaining maturities the midpoints between the available swap maturities. For the covariance restrictions, we take the covariances between forward Libor rate changes of all these forward Libor rate maturities, in total 36 moment restrictions. Below, we discuss how we weight these moment restrictions.

For the cap moment restrictions, we use all 7 option maturities that are available in the cap data, ranging from the 1-year cap to the 10-year cap. For the swaption moment restrictions, we include the 11 swaptions discussed in Section 2.1.

Applying the first step of GMM, we choose a diagonal weighting matrix. Recall that we formulated all moment restrictions such that they all refer to annualized variances and covariances. We choose the diagonal weights such that the four sets of moment conditions (interest rate variances, interest rate covariances, cap volatilities, and swaption volatilities) contribute equally to the GMM goal function. The GMM weighting matrix W then becomes

$$W = \begin{bmatrix} \frac{1}{N_1^2} I_{N_1} & 0_{N_1 x N_2} & 0_{N_1 x N_3} & 0_{N_1 x N_4} \\ 0_{N_2 x N_1} & \frac{1}{N_2^2} I_{N_2} & 0_{N_2 x N_3} & 0_{N_2 x N_4} \\ 0_{N_3 x N_1} & 0_{N_3 x N_2} & \frac{1}{N_3^2} I_{N_3} & 0_{N_3 x N_4} \\ 0_{N_4 x N_1} & 0_{N_4 x N_2} & 0_{N_4 x N_3} & \frac{1}{N_4^2} I_{N_4} \end{bmatrix} \tag{12}$$

where I_N is an N-dimensional identity matrix, and 0_{NxM} is an N-by-M matrix with zeros. This implies that each variance restriction is weighted with 1/9 (N_1=9), each covariance restriction with 1/36 (N_2=36), each cap restriction with 1/7 (N_3=7), and each swaption restriction with 1/11 (N_4=11). This way, none of these four sets of restrictions dominates the estimation results. Within each of the four sets of moment restrictions, we use constant diagonal weights. As an alternative to constant diagonal weights, we also perform an option-based estimation where we weight each option moment restriction using the inverse of the variance of the option implied volatility over the sample period.[8] This way, options for which the implied volatility is very stable over time obtain a higher weight in the estimation. In our sample, swaption volatilities turn out to be more stable than cap volatilities.

It turns out that the covariance matrix of these estimated moment restrictions is close to singularity[9]. The efficient, second step of GMM requires that the inverse of the covariance matrix of the estimated moment restrictions is used as the weighting matrix. Hansen (1982) shows that this is the optimal choice for a correctly specified model as it yields the lowest asymptotic variance for the GMM parameter estimates. However, as noted by Cochrane (2001), using a near-singular covariance matrix as weighting matrix implies that one fits the parameters to linear combinations of the original moment restrictions that have very large positive and negative weights on the original moment restrictions. Using these linear combinations of moment restrictions to estimate the model may be statistically optimal for a correctly specified model (that is, asymptotically), but one can question whether these extreme linear combinations are the most interesting moment restrictions from an economic point of view (see Cochrane (2001)).

We find that, when using the optimal weighting matrix, the model is essentially fitted to the differences between the moment restrictions rather than to the level of the moment restrictions. Due to the high correlations between the estimated moment restrictions, the standard errors of these differences are much lower than the standard errors of the levels. When performing two-stage GMM estimation, we find that the shape of the Libor variance term structure, the Libor covariance structure, and the cap and swaption implied volatility term structures are fitted quite accurately, whereas the level of these term structures is not fitted well. Therefore, we use in the empirical analysis only the first-stage GMM estimator, that is obtained using a diagonal weighting matrix. Of course, if the model is correctly specified, this estimator is still consistent and asymptotically normal, and standard errors and tests are constructed in a straightforward way.[10]

3 Empirical Results

We estimate the covariance matrix parameterization in (7) and (8) three times: on the basis of interest-rate-based estimation, option-based estimation, and joint estimation.

3.1 Implications of interest-rate based estimation

In this subsection, we study the results obtained by interest-rate based estimation. In this case, the parameters in the volatility and correlation functions in (7) and (8) are estimated using variances and covariances on changes in forward Libor rates of different forward maturities.

First, we check whether our parameterization in (7) and (8) is flexible enough to provide a satisfactory fit to the interest rate moments. In Figure 4, we graph the fit to the interest rate variance moments. This graph shows a good fit of the term structure of interest rate variances. In Figure 6, we graph the correlation matrix implied by interest-rate based estimation. Again, the fit is quite good: the average absolute difference between the correlations implied by (8), and the historically observed correlations (Figure 5) is equal to 0.044. The good fit of the covariance matrix parameterization is also shown in Table 1, where we tabulate the p-values of joint tests of (subsets of) moment restrictions. The p-value for the test that the covariance matrix parameterization correctly describes the interest rate variance and covariance moment restrictions is 0.142. In Table 2 we report average absolute t-values for the individual moment restrictions. This table shows that none of the interest rate variance and covariance restrictions in (10) are significantly misfitted.

Next, we analyze the implications for prices of caps and swaptions (measured in Black volatilities). The results are shown in Figures 2 and 3. First of all, Figure 2 shows that the cap prices implied by the interest-rate based covariance matrix are much lower than the average observed prices for all cap maturities. For the 2-year and 3-year caps the difference amounts to almost 3 volatility points, averaged over the 1995-1999 period. Not surprisingly, Table 2 shows that each cap price restriction in (11) is individually rejected, and Table 1 reveals that the joint test of the cap price restrictions also leads to a rejection. Interestingly, the Black volatility term structure implied by interest-rate based estimation does exhibit a hump shaped form, but the hump is too low and too 'flat' to fit the cap price data.

In Figure 3 we plot the implications of the interest-rate based covariance matrix for swaption Black volatilities. For 9 out of the 11 swaptions, the interest-rate based covariance matrix gives too low prices for swaptions. Also, for short swap maturities, the hump in the fitted Black volatilities is too pronounced compared to the observed hump, whereas the fitted hump shape is too flat for swaptions with longer swap maturities. Tables 1 and 2 show that the mispricing of swaptions, when using the interest-rate based covariance matrix, is jointly significant and individually significant for six swaptions.

A possible explanation for these results might be the fact that our 1995-1999 data period for the option data contains the Russia/LTCM crisis, during which option prices were higher than under 'normal' market conditions (see also Longstaff, Santa-Clara, and Schwartz (2001)). Therefore, we have recalculated the time series averages of the observed Black volatilities in Figures 2 and 3, now excluding the 4-month period from August 1998 until November 1998. The results, depicted in Figures 7 and 8, show that this leads to slightly lower averages for the observed Black volatilities, but these averages are still very different from the Black volatilities implied by the interest-rate based covariance matrix. Thus, even when excluding this period of extreme market conditions, the mispricing remains.

Summarizing, the results in this subsection point at a significant difference between the covariance information in cap and swaption prices and the (co)variability of the term structure of interest rates. Buhler et al. (1999) and Driessen, Klaassen, and Melenberg (2003) test particular low-factor term structure models on the basis of option pricing performance, using interest rate data to estimate the parameters. Although they do not compare the information in interest rate data directly with the information in option price data, these articles report substantial option pricing errors, which is consistent with our results.

3.2 Implications of option-based estimation

In this subsection, we study the results obtained by option based estimation. In this case, the parameters in (7) and (8) are estimated using the moment restrictions for cap and swaption prices in equation (11). This estimation procedure is related to previous literature, where option price data are used to estimate and test particular term structure models.[11]

We first analyze whether option-based estimation leads to a good fit of average prices of caps and swaptions. Figure 2 shows that the fit to cap Black volatilities is almost perfect. The fit to swaption Black volatilities, shown in Figure 3, is reasonably good, except perhaps for swaptions with short swap maturities. In total, Table 2 shows that 16 of the 18 options are not significantly mispriced, and a joint test of all cap and swaption moment restrictions does not lead to a rejection. Therefore, we conclude that our covariance matrix specification provides a reasonable fit to the option price data. Table 2 also shows that these results are robust to using the variance-weighted GMM procedure. Given that this procedure gives more weight to swaptions and less to caps, it is not surprising to see in Table 2 that variance-weighted GMM leads to lower t-ratios for swaptions and higher t-ratios for caps. The difference with constant-weights GMM is however small. This can also be seen in Table 3 which presents the parameter estimates for both estimation approaches, which are quite similar to each other.

Next, we investigate whether the option-implied estimates lead to a good fit of the interest-rate moment restrictions. First, we look at the interest rate variances, or, equivalently, the term structure of forward Libor rate volatility. Figure 4 graphs these term structures. Compared to the realized (historically estimated) Libor rate volatilities, the option-implied forward Libor rate volatility term structure is higher at the short end and lower for long forward maturities. It is important to understand the relation between these differences and the underpricing of caps in case of interest-rate based estimation. Since a cap is a portfolio of caplets of different maturities, the Black volatility for a cap is roughly the average of the caplet's Black volatilities. In turn, the caplet Black volatility is the option-implied volatility of the corresponding forward Libor rate. Therefore, although for long forward maturities the realized Libor volatility is higher than the option-implied Libor volatility, the cap Black volatilities implied by interest-rate based estimation (which uses the realized volatility) are all lower than the observed cap Black volatilities, because for short and intermediate maturities the option-implied Libor volatility is higher than realized Libor volatility.

In Figure 9 we graph the correlation matrix implied by the option prices. The option-implied correlations of short-maturity forward Libor rates with other short-maturity forward Libor rates are much lower than the realized interest rate correlations, while all other correlations, that involve longer forward maturities, are higher in case of option-based estimation. For example, the correlation between the 3-month forward Libor rate and 6-month forward Libor rate is equal to 0.697 based on the historical interest rate data, while the option-implied estimate is 0.362. This is confirmed by the parameter estimates for the correlation structure in equation (8), which are given in Table 3. Compared to interest-rate based estimation, option-based estimation leads to higher estimates for γ_2 and γ_3, and a lower estimate for γ_4, which in total decreases short-maturity correlations and increases long-maturity correlations. These results on the correlation matrix are slightly different from Longstaff, Santa-Clara, and Schwartz (2001), who find that option-implied correlations are always lower than correlations estimated from interest rate data.

The formal tests of the interest rate variance and covariance moment restrictions in case of option-based estimation (Tables 1 and 2) indicate that the difference between the option-implied and realized interest rate covariance matrix is statistically significant in most cases. Again, these results are robust to using a variance-weighted GMM procedure for option-based estimation instead of using constant diagonal weights.

We also use a different test to analyze whether the information in interest rate data is consistent with the option price data. We test whether the parameters in the covariance matrix specification in (7) and (8), estimated using either interest rate data or option price data, are equal to each other. Since the joint set of interest rate and option price moment restrictions has asymptotically a normal distribution, it is easy to show that the interest rate based parameter estimator and the option based parameter estimator have a joint normal asymptotic distribution, so that a simple chi-square test can be performed to test this hypothesis. The p-value for this test turns out to be 0.0023, so that the hypothesis that the option-implied parameters are equal to the interest-rate implied parameters is rejected.

In Section 1, we also presented a three-factor model (equation (9)) as a comparison to our specification of the covariance matrix in the full-factor model in (7) and (8). To avoid an overload of tables and figures, we only present results for this three-factor model in case of option-based estimation. In Table 4, we give the pricing errors in terms of volatility points for this three-factor model. It follows that, although the three-factor model and the full-factor model contain an equal number of parameters, the fit of the three-factor model is less good.[12] In particular, while the full-factor model yields pricing errors that are on average very close to zero for caps and swaptions, the three-factor model on average underprices caps and overprices swaptions. The main reason for this result is that the three-factor model cannot generate the correlation structure that is implicit in swaption prices. This is shown in Figure 10, where we graph the covariance matrix implied by the three-factor model and option-based estimation. Compared to Figure 9, the three-factor model cannot generate low correlations between short-maturity interest rates, and it implies very high correlations between near-maturity interest rates. Rebonato (1996) also discusses this property of low-factor term structure models using principal components analysis.

Finally, we note that we have also performed a joint estimation of the full-factor model that uses both the interest rate moment restrictions and the cap and swaption moment restrictions. As expected, the jointly estimated covariance matrix is roughly the average of the historically estimated covariance matrix and the option-implied covariance matrix, and there is a clear trade off in the fit of the interest rate moments and the fit of the cap/swaption moments. In total, a joint test of all moment restrictions (the GMM 'J-test') again leads to a rejection.

3.3 Trading Strategies

So far, we have provided evidence for a significant difference between the historical and option-implied interest rate covariance matrix. In order to assess the economic significance of this difference, we set up trading strategies that try to exploit this difference. Given that, on the basis of interest-rate-based estimation, almost all observed option prices seem to be too high on average, each trading strategy takes a short position in a cap or swaption (i.e., writing each option). We calculate for each cap and swaption the LMM-implied deltas with respect to the hedge instruments (zero-coupon bonds of different maturities), and construct for each derivative instrument a delta-hedged portfolio. After a week, we compute the return on this hedge portfolio, using the observed prices for the hedge instruments and the derivatives.[13] This procedure is repeated each week. Given that we observe option quotes for fixed option maturities, we do not follow an option until maturity. The convexity and time value of the hedged short option position imply that this strategy generates a profit if interest rates do not move much, and a loss in case of large interest rate movements. Given the positive difference between option-implied and historical volatilities, we expect on average positive returns for these strategies.

Our interest-rate hedging method follows the so-called bucket hedging approach of Driessen, Klaassen, and Melenberg (2003). This procedure uses zero-coupon bonds for hedging options, where the bond maturities match all tenor dates of the cap or swaption. Given the full-factor model that is used in this paper, this bucket hedging procedure is appriopriate since it does not impose any restriction on the movements of bond prices across different maturities. The hedge instruments for each cap or swaption are zero-coupon bonds with maturities that correspond to all payment dates relevant to the particular derivative. For example, for a 2-year cap, that consists of 7 quarterly caplets, we use as hedge instruments zero-coupon bonds with maturities of 3 months, 6 months, and so on, up to 2 years. For a 1-year option on a 5 year swap with annual payments, we use zero-coupon bonds with maturities of 1 year, 2 years, and so on, up to 6 years, as hedge instruments. The time-*t* price *V(t)* of a cap or swaption is a function of precisely these zero-coupon bonds. In formulas we have

$$dV(t) = ...dt + \frac{\partial V(t)}{\partial P(t,T)} dP(t,T) + \sum_{j=1}^{n} \frac{\partial V(t)}{\partial P(t,T_j)} dP(t,T_j) \tag{13}$$

where V(t) is the price of the derivative at time *t*, *T* is the maturity date of the caplet or swaption, and $T_1, ..., T_n$ are the payment dates. Thus, bonds of all relevant tenor dates are included on the right-hand side. We leave the drift unspecified in (13), since it is not relevant for the hedge strategy. Following Brace et al. (1999) and Driessen, Klaassen, and Melenberg (2003), we use the approximating swaption valuation formula of Brace, Gatarek, and Musiela (1997) to calculate the hedge ratios.

We analyze a separate hedge strategy for each option, and calculate the Sharpe ratio using the time series of returns of each strategy. Table 5 presents these Sharpe ratios, which are annualized for convenience (assuming i.i.d. returns). To put the results into perspective, common values for the equity premium of 5% and annual equity volatility of 20% give an annual Sharpe ratio of 0.25. Our results show that especially the strategies for caps and swaptions with short option maturities have economically very significant Sharpe ratios of about 0.2. This is in line with Figures 2 and 3, which show that short-maturity caps and swaptions have the largest mispricing. For swaptions with longer option maturities, the Sharpe ratios are close to zero, in line with the small mispricing in Figure 3. In sum, these results show that the difference between the historical and option-implied covariance matrices is also economically significant.

3.4 Possible Explanations

The moment restrictions for the cap and swaption implied volatilities in (11) explicitly allow for the presence of measurement error. So far, we have neglected the possible presence of measurement error in the interest rate data. We now include measurement error on the interest rates in our model.

There are several reasons to include the error term in the log forward Libor rate. First of all, the underlying money-market and swap data might contain measurement error due to illiquidity and time-of-the-day effects. Also, the first-order autocorrelation of weekly changes in the log-forward Libor rate is, averaged over all forward maturities, equal to -0.185, whereas the higher-order autocorrelations are close to zero or even positive. This is an indication of the presence of measurement error, since it is easy to show that, abstracting from the drift of forward Libor rates that is implied by the model, measurement error in the level of interest rates leads to negative first-order autocorrelation for discrete-time changes in the forward Libor rate, and zero higher-order autocorrelations.

In line with previous research on term structure models (for example, De Jong (2000) and Duan and Simonato (1999)), we assume that the observed log forward rate, $\ln L^*(t,T_i)$, is equal to the true rate $\ln L(t,T_i)$, plus a zero-expectation error term $\varepsilon_i(t)$, which is independently distributed over time and across forward maturities, and independent of the true log forward Libor rate $\ln L(t,T_i)$

$$\ln L^*(t,T_i) = \ln L(t,T_i) + \varepsilon_i(t), \quad E(\varepsilon_i(t)) = 0, \quad i=1,..,N-1 \tag{14}$$

We impose a very simple structure on the measurement error variance

$$\begin{aligned} V(\varepsilon_i(t)) &= \sigma_\varepsilon^2, \; i=1,...,N-1 \\ Corr(\varepsilon_i(t), \varepsilon_j(t)) &= 0, \quad i,j=1,...,N-1, \; i \neq j \end{aligned} \tag{15}$$

This way, the moment restrictions for variances and covariances are now given by

$$\begin{aligned} V(\Delta \ln L_i^*(t)) &= \sigma(T_i-t)^2 \Delta t + 2\sigma_\varepsilon^2, \quad i=1,...,N-1 \\ Cov(\Delta \ln L_i^*(t), \Delta \ln L_j^*(t)) &= \rho(T_i-t,T_j-t)\,\sigma(T_i-t)\,\sigma(T_j-t)\,\Delta t, \quad i,j=1,...,N-1 \end{aligned} \tag{16}$$

By approximation, the measurement error variance of the forward Libor rates does not enter the moment restrictions for caps and swaptions.[14] Thus, by combining the forward Libor rate moment restrictions and the cap and swaption restrictions, the measurement error variance can be identified.

Using these modified moment restrictions in (16) and the moment restrictions for caps and swaptions, we re-estimate the covariance matrix specification (i.e., using joint estimation). It turns out that the estimate for the measurement error variance is given by a corner solution where $\sigma_\varepsilon=0$. Thus, including some measurement error on the forward Libor rates only decreases the fit of the moment restrictions. The reason for this is the following. Consider the interest rate moments first. Starting from the parameter estimates without measurement error, adding measurement error increases the total variance of forward Libor rates. Figure 4 shows that, without measurement error, the jointly estimated covariance matrix already yields interest rate variances that are higher than realized variances, and adding measurement error only increases this differences. Of course, one could try to lower the underlying variances of the model (i.e., $\sigma(T_i-t)$), to compensate for this effect, but this would lead to lower cap Black volatilities, since cap Black volatilites are not influenced by interest rate measurement error. Since caps are already underpriced in case of joint estimation and no measurement error, this also decreases the fit. Therefore, the best fit is obtained when the measurement error variance is set equal to zero. Since we argued above that it is not unlikely that there is some measurement error in the forward Libor rate data, these estimation results only increase the puzzling difference between cap and swaption information and the information in term structure data.[15]

A second possible explanation for the discrepancy between interest rate and option data is the fact that some of the options used for estimation have maturity dates that exceed the final date of our interest rate sample (June 2000).[16] For example, a 10-year caplet contains information on the volatility of the 3-month forward rate over 1995 to 2005. This may impact our results if market participants expect volatility to increase over the 2001-2005 period. Therefore, we re-run our option-implied estimation using only options for which all maturity dates are within our interest rate sample period. This leaves us with 9 options: the 1-year cap, and all 3-month and 1-year swaptions. Since our model in (7) and (8) contains 9 parameters, we obtain essentially a perfect fit of the moment restrictions for these instruments. This implies that, to compare the option-implied and interest-rate-based option prices, we can use Figures 2 and 3 to compare the average Black volatilities in the data (fitted perfectly by option-based estimation), with the Black volatilities implied by interest-rate-based estimation. This shows again a clear difference between option-based and interest-rate-based prices, and in 7 out of 9 cases interest-rate-based estimation implies too low prices for the options. In Figure 11 we present the option-implied correlation matrix in case of this subset estimation. This matrix is quite similar to the option-implied matrix in Figure 9 that is based on all option data, and quite different from the interest-rate-based correlation matrix.

Finally, another possible explanation for our results is that the lognormal distribution for forward Libor rates is not appropriate. Although our analysis is based on at-the-money options, which may not be too sensitive to misspecification of the tail of the distribution, we have performed interest-rate based estimation for a market model with normally distributed forward Libor rates. The results, available on request, are qualitatively and quantitatively very similar to the results for the lognormal model.

4 Summary and Conclusions

In this paper, we examine whether the information in cap and swaption prices on interest rate variances and correlations is consistent with realized movements of the interest rate term structure. We use lognormal market models for forward Libor rates to invert cap and swaption prices to an interest rate covariance matrix, using a full-factor model with a flexible parameterization for this covariance matrix. We show that this model performs better than a standard three-factor model.

We document clear inconsistencies between the option and interest rate data. Both the option-implied and realized interest rate variances and the corresponding correlations differ. If one uses interest rate data for estimation, the resulting option prices are much lower than the observed option prices. Especially for caps, these differences are economically large. For some caps the fitting error is almost 3 Black volatility points, averaged over the 1995-1999 period. These results are particularly striking given that caps and swaptions are subject to some counterparty risk. The presence of counterparty default risk will lower option prices, to compensate for the risk that option payoffs are not received, so that the high option prices in our sample are even more puzzling.

A first possible explanation for our results is a peso-problem interpretation: option prices incorporated the (small) possibility of large interest rate movements, which did not occur and are, therefore, not observed in the interest rate data. A second possible explanation is the presence of transaction costs on the underlying assets. In the OTC market for caps and swaptions, banks typically sell options to other institutions, and then hedge the obtained risk exposure. If banks cannot perfectly hedge the option price risk due to the presence of transaction costs, they may require a premium for this residual risk. Related to this is the presence of supply and demand imbalances in the cap and swaption markets (Rebonato (2003)). A third possible explanation is the fact that our data consist of option price quotes, and not transaction prices. This may explain some of the mispricing of options when using interest rate data, but it seems less likely that the large and systematic mispricing is caused by this effect. Finally, a possible explanation for our results is that bond and swap markets are incomplete. Our results are based on the assumption that the market defined by the underlying securities is complete. In an incomplete market, generated for example by stochastic volatility processes for forward Libor rates, the volatility process under the true probability measure can differ from the volatility process under an equivalent martingale measure. For example, in the Heston (1993) model, the conditional variance (over a discrete-time interval) is not the same under the true and risk-neutral measure if volatility risk is priced. In particular, if the volatility risk premium is negative, the conditional variance is on average higher under the risk-neutral measure than under the true probability measure. This may be an explanation for the apparent inconsistencies found in this paper. Bakshi and Kapadia (2003) provide evidence for a negative volatility risk premium in equity options. As an alternative to stochastic volatility, (priced) jump risk might be another candidate to explain our findings. For the bond market, the current evidence on market completeness is inconclusive. Collin-Dufresne and Goldstein (2002) and Heiddari and Wu (2001) conclude that bond markets are incomplete, whereas Fan, Gupta, and Ritchken (2003) conclude that swaptions can be hedged well using bonds only.

References

Ait-Sahalia, Y., Y. Wang, and F. Yared. (2001). "Do Options Markets Correctly Price the Probabilities of Movement of the Underlying Asset?," *Journal of Econometrics* 102, 67-110.

Amin, K.I., and A. Morton. (1994). "Implied Volatility Functions in Arbitrage-Free Term Structure Models," *Journal of Financial Economics* 35, 141-180.

Amin, K.I., and V.K. Ng. (1997). "Inferring Future Volatility from the Information in Implied Volatility in Eurodollar Options: A New Approach," *Review of Financial Studies* 10(2), 333-367.

Bakshi, G.S., and N. Kapadia. (2003). "Delta Hedged Gains and the Pricing of Volatility Risk," *Review of Financial Studies,* 527-566.

Black, F. (1976). "The Pricing of Commodity Contracts," *Journal of Financial Economics* 3, 167-179.

Brace, A., D. Gatarek, and M. Musiela. (1997). "The Market Model of Interest Rate Dynamics," *Mathematical Finance* 7, 27-155.

Buhler, W., M. Uhrig, U. Walter, and T. Weber. (1999). "An Empirical Comparison of Forward- and Spot-Rate Models for Valuing Interest-Rate Options," *Journal of Finance* 54, 269-305.

Canina, L. and S. Figlewski. (1993). "The Informational Content of Implied Volatility," *Review of Financial Studies* 6, 659-681.

Cochrane, J.H. (2001). *Asset Pricing*, Princeton University Press.

Collin-Dufresne, P., and R.S. Goldstein. (2002). "Do Bonds Span Fixed Income Markets? Theory and evidence for `Unspanned' Stochastic volatility," *Journal of Finance*, 57 (4), 1685-1730.

Christensen, B.J., and N.R. Prabhala. (1998). "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics*, 125-150.

De Jong, F. (2000). "Time-series and Cross-section Information in Affine Term Structure Models," *Journal of Economics and Business Statistics* 18 (3), 300-318.

De Jong, F., J. Driessen, and A.A.J. Pelsser. (2001). "Libor Market Models versus Swap Market Models for Pricing Interest Rate Derivatives: An Empirical Analysis," *European Finance Review* 5 (3), 201-237.

Driessen, J., P. Klaassen, and B. Melenberg. (2003). "The Performance of Multi-Factor Term Structure Models for Pricing and Hedging Caps and Swaptions," *Journal of Financial and Quantitative Analysis* 38(3), 635-672.

Duan, J., and J.G. Simonato. (1999). "Estimating Exponential-affine Term Structure Models by Kalman Filter," *Review of Quantitative Finance and Accounting* 13(2), 11-135.

Duffie, D., and R. Kan. (1996). "A Yield-Factor Model of Interest Rates," *Mathematical Finance* 64, 379-406.

Fan, R., A. Gupta, and P. Ritchken. (2001). "On Pricing and Hedging in the Swaption Market: How Many Factors, Really?," working paper, Case Western Reserve University.

Fan, R., A. Gupta, and P. Ritchken. (2003). "Hedging in the Possible Presence of Unspanned Stochastic Volatility: Evidence from Swaption Markets?," *Journal of Finance*, forthcoming.

Flesaker, B. (1993). "Testing the Heath-Jarrow-Morton/Ho-Lee Model of Interest Rate Contingent Claims," *Journal of Financial and Quantitative Analysis* 28, 483-495.

Gourieroux, C., and A. Monfort. (1995). *Statistics and Econometric Models*, Cambridge University Press.

Gupta, A. and M. G. Subrahmanyam. (2001). "An Examination of the Static and Dynamic Performance of Interest Rate Option Pricing Models in the Dollar Cap-Floor Markets," working paper, Case Western Reserve University.

Han, B. (2001). "Stochastic Volatilities and Correlations of Bond Yields," Working Paper, UCLA.

Hansen, L.P. (1982). "Large Sample Properties of Generalized Methods of Moments Estimators," *Econometrica* 50, 1029-1054.

Heidari, M. and L. Wu. (2001). "Are Interest Rate Derivatives Spanned by the Term Structure of Interest Rates," *Journal of Financial and Quantitative Analysis*, forthcoming.

Heston, S. (1993). "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies* 6, 327-343.

Jackwerth, J., and M. Rubinstein. (1996). "Recovering Probability Distributions from Option Prices," *Journal of Finance* 51, 1611-1631.

Jagannathan, R., A. Kaplin, and S.G. Sun. (2001). "An Evaluation of Multi-Factor CIR Models Using LIBOR, Swap Rates, and Cap and Swaption Prices," *Journal of Econometrics* 116, 113-146.

Jamshidian, F. (1997). "Libor and Swap Market Models and Measures," *Finance and Stochastics* 1, 293-330.

Kerkhof, J., and A.A.J. Pelsser. (2002). "Observational Equivalence Of Discrete String Models And Market Models," *Journal of Derivatives* 10(1), 55-61.

Longstaff, F.A., P. Santa-Clara, and E.S. Schwartz. (2001). "The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence." *Journal of Finance* 56, 2067-2109..

Miltersen, K., K. Sandmann, and D. Sondermann. (1997). "Closed Form Solutions for Term Structure Derivatives with Lognormal Interest Rates," *Journal of Finance* 52, 407-430.

Moraleda, J.M., and A.C.F. Vorst. (1997). "Pricing American Interest Rate Claims with Humped Volatility Models," *Journal of Banking and Finance* 21, 1131-1157.

Musiela, M., and M. Rutkowski. (1997). *Martingale Methods in Financial Modeling*, Springer Verlag.

Newey, W.K., and K.D. West. (1987). "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55, 703-708.

Rebonato, R. (1996). *Interest-Rate Option Models*, John Wiley & Sons.

Rebonato, R. (1999). "On the Simultaneous Calibration of Multi-Factor Lognormal Interest-rate Models to Black Volatilities and to the Correlation Matrix." *Journal of Computational Finance*, 2 (3).

Rebonato, R. (2003). "Term Structure Models: a Review", Working Paper Oxford University.

Santa-Clara, P., and D. Sornette. (2001). "The Dynamics of the Forward Interest Rate Curve with Stochastic String Shocks," *Review of Financial Studies* 14(1), 149-185 .

Schoenmakers, J.G.M., and B. Coffey. (2003). "Systematic Generation of Correlation Structures for the Libor Market Model," *International Journal of Theoretical and Applied Finance*, 6(4), 1-13.

Vasicek, O. (1977). "An Equilibrium Characterization of the Term Structure." *Journal of Financial Economics* 5, 177-188.

1.We use the cap prices, for which an exact analytical formula exists, as control variates.

2.We have experimented with other parsimonious specifications for the volatility structure, such as the specification of Moraleda and Vorst (1996), but the specification in (7) provides the best fit of both option and interest rate data.

3.Rebonato (1999) and Schoenmakers and Coffey (2003) propose other specifications for correlation structures.

4.The instantaneous covariance matrix of the Brownian motions has to be positive definite. This restriction is imposed when estimating the model parameters.

5.Our initial dataset contains 56 swaptions. We use a subset of 11 swaptions in order to avoid multi-collinearity.

6.To perform GMM on variance and covariance restrictions, we add auxiliary moment restrictions of the form $E(\Delta \ln L_i(t)) = \alpha_i$, $i=1,...,N-1$, where the α_i`s are free coefficients that are estimated along with the other parameters. Even if the true means (i.e., the $\alpha_i$`s) are equal to zero, which would be the case if forward Libor rates are stationary, Cochrane (2001) notes that, in small samples, better estimates are obtained if one uses variances and covariances instead of uncentered second moments. In our case, the sample means are very small relative to the variance of the forward Libor rates, so that imposing that the α_i`s are equal to zero would hardly affect the GMM parameter estimates.

7.This expectation is taken under the true probability measure, since the option prices are observed under this measure. Of course, to calculate the option prices implied by the model, one uses an equivalent martingale measure.

8.We thank an anonymous referee for this suggestion.

9.We use the method of Newey and West (1987) to correct this covariance matrix for heteroskedasticity and autocorrelation.

10.The near-singular covariance matrix of the moment restrictions also causes the GMM J-statistic, that can be used to jointly test the overidentifying restrictions, to be very large for all models that we estimate. Therefore, to calculate the inverse of this covariance matrix of the J-statistic we use an eigenvector decomposition of this covariance matrix. The (approximate) inverse of the covariance matrix is calculated using only those eigenvectors whose eigenvalues sum up to more than 99% of the total sum of eigenvalues. In case of the variance and covariance moment restrictions, we exploit the fact that the model implies that the changes in log-forward interest rates are normally distributed. More precisely, let V be the estimated covariance matrix of the normally distributed log-forward Libor rate changes of different maturities. Then the covariance matrix of the sample counterparts of the variance and covariance moment restrictions is equal to $2V \otimes V/T$, where T is the number of observations. The eigenvector decomposition is applied to V. Besides a joint test of moment restrictions, we also report t-values for individual moment restrictions, that are unaffected by the correlations across moment restrictions and the eigenvector decomposition.

11.Flesaker (1993), Amin and Morton (1994), Amin and Ng (1997), Fan, Gupta, and Ritchken (2001), Gupta and Sybrahmanyam (2001), De Jong, Driessen, and Pelsser (2001), and Driessen, Klaassen, and Melenberg (2003),.

12. In case of interest-rate based estimation and joint estimation the full-factor model also outperforms the three-factor model. These results are available on request.

13. Prices for the zero-coupon bonds are directly obtained from the forward Libor rates data. We only observe the prices of at-the-money caps and swaptions at each trading day. Clearly, an option that is at-the-money at a particular trading day will not be exactly at-the-money one week later. To be able to calculate the price of an off-market cap or swaption after one week, we follow Driessen, Klaassen, and Melenberg (2003) and assume that there is no implied Black volatility smile, i.e., we assume that the observed implied Black volatility for a cap or a swaption is the same for all strike rates.

14. This depends on the dependence of the LMM Black volatility $IV^{C,LMM}(t, T_i)$ on the underlying forward Libor rates. If this dependence would be linear, the presence of measurement error in the forward Libor rates would not change the unconditional expectation of $IV^{C,LMM}(t, T_i)$. In reality, this dependence is not linear, so that the expectation of $IV^{C,LMM}(t, T_i)$ will depend on the variance of the measurement error in the forward Libor rates (and higher-order moments of the measurement error distribution). A Taylor expansion shows that, for at-the-money-forward caps and swaptions, this is a second order effect, and we will therefore neglect this effect when estimating the model.

15. We have also analyzed a more sophisticated measurement error structure, where we allowed the measurement error variance to depend on the forward maturity and also allowed for correlation between the measurement errors across forward maturities. The results remain the same: the best fit is obtained if all measurement error variances are equal to zero.

16. We thank an anonymous referee for this suggestion.

Tables

Table 1. Tests of Moment Restrictions.[a]

	Interest-Rate-Based Estimation	Option-Based Estimation	Option-Based Estimation (Variance-weighted)
Libor Variances	-	0.000	0.000
Libor Covariances	-	0.134	0.119
Variances + Covariances	0.142	0.000	0.000
Caps	0.004	-	-
Swaptions	0.010	-	-
Caps + Swaptions	0.004	0.579	0.419
All	0.000	0.000	0.000

[a]The table reports p-values of joint tests of the moment restrictions in equations (10) and (11), for different estimation setups: interest-rate-based and option-based estimation described in Section 3.2, both using a identity weighting matrix, and option-based estimation that uses a variance-weighted diagonal weighting matrix. We use a standard Wald test to jointly test the moment restrictions. This test-statistic uses the asymptotic covariance matrix of moment restrictions in case of GMM estimation (see Gourieroux and Monfort (1995)), and has asymptotically a chi-square distribution. We correct the covariance matrix for heteroskedasticity and 8th-lag autocorrelation using Newey-West (1987).

Table 2. Average Absolute T-ratios Moment Restrictions.[a]

	Interest-Rate-Based Estimation	Option-Based Estimation	Option-Based Estimation (Variance-weighted)
Libor Variances (9)	0.677 (0)	1.474 (2)	1.395 (2)
Libor Covariances (36)	0.423 (0)	1.440 (12)	1.659 (14)
Caps (7)	3.432 (7)	0.227 (0)	0.270 (2)
Swaptions (11)	2.004 (6)	1.164 (2)	0.931 (1)

[a]For all moment restrictions, the t-ratios of the individual moment restrictions are calculated using the asymptotic covariance matrix of moment restrictions in case of GMM estimation (see Gourieroux and Monfort (1995)), correcting for heteroskedasticity and 8th-lag autocorrelation using Newey-West (1987). The table reports for each set of moments the average of the absolute value of these t-ratios, and the number of moment restrictions that is individually rejected at the 5% significance level. As in Table 1, results are reported for interest-rate-based estimation, option-based estimation, and variance-weighted option-based estimation.

Table 3. Parameter Estimates.[a]

	Interest-Rate-Based Estimation	Option-Based Estimation	Option-Based Estimation (Variance-weighted)
σ_0	0.147 (0.008)	0.137 (0.004)	0.148 (0.004)
σ_1	0.427 (0.496)	0.808 (0.997)	0.468 (0.803)
σ_2	-0.565 (0.457)	-0.972 (0.953)	-0.572 (0.492)
κ_1	1.062 (0.978)	1.089 (0.519)	0.901 (0.488)
κ_2	1.726 (1.308)	1.688 (1.097)	1.472 (0.981)
γ_1	0.000 (-)	0.009 (0.012)	0.008 (0.009)
γ_2	0.480 (0.099)	1.127 (0.693)	1.031 (0.539)
γ_3	1.511 (0.289)	1.849 (0.386)	1.724 (0.318)
γ_4	0.186 (0.127)	0.024 (0.098)	0.025 (0.065)

[a]Interest-rate based estimation, option-based estimation, and variance-weighted option-based estimation of the parameters in the covariance matrix parameterization in equations (7) and (8) is performed as described in Section 3.2. Each estimation setup is the first-step of GMM. The table reports the parameter estimates and associated standard errors, calculated using Newey-West (1987).

Table 4. Option Pricing Errors.[a]

	Caps		*Swaptions*	
	Avg. Volatility Point Error	Avg. Abs. Volatility Point Error	Avg. Volatility Point Error	Avg. Abs. Volatility Point Error
Interest Rate Based Estimation	-2.01	2.01	-0.70	0.79
Option Based Estimation	-0.01	0.14	0.01	0.40
Option Based Estimation (Variance-weighted)	-0.08	0.21	0.03	0.34
3-Factor Model Option Based Estimation	-0.22	0.31	0.15	0.57

[a]The table reports option pricing errors for caps and swaptions, measured in Black volatility percentages (also referred to as Black volatility points). Results are given for the full-factor model in equations (7) and (8), for three sets of estimation results: interest-rate-based estimation, option-based estimation, and variance-weighted option-based estimation. The table also includes results for option-based estimation of the three-factor model in equation (9).

Table 5. Sharpe Ratios of Trading Strategies.[a]

	Annual Sharpe Ratio		Annual Sharpe Ratio
1Y Cap	0.230	5Yx1Y Swaption	0.025
2Y Cap	0.264	3Mx3Y Swaption	0.163
3Y Cap	0.249	1Yx3Y Swaption	0.143
4Y Cap	0.232	5Yx3Y Swaption	-0.023
5Y Cap	0.206	3Mx5Y Swaption	0.154
7Y Cap	0.151	1Yx5Y Swaption	0.128
10Y Cap	0.141	5Yx5Y Swaption	-0.020
3Mx1Y Swaption	0.110	3Mx7Y Swaption	0.207
1Yx1Y Swaption	0.091	1Yx7Y Swaption	0.083

[a]The table reports Sharpe ratios of the trading strategies discussed in Section 3.3, which consist of a short position in a cap or swaption, that is delta-hedged for interest-rate movements using discount bonds of different maturities (see Section 3.3). The hedge portfolio is rebalanced every week. Using the time series of returns on these hedge portfolios over the 1995-1999 data period, Sharpe ratios are calculated. The Sharpe ratios are annualized assuming i.i.d. returns.

Figure 1. Impact of Correlation Parameters. This graph shows the change in the correlation of a 3-month forward Libor rate with forward Libor rates of other forward maturities, due to a change in one of the four parameters in equation (8) of 1% at the estimated parameter values. For γ_1 we use a change of 0.5% in order to have similar scales.

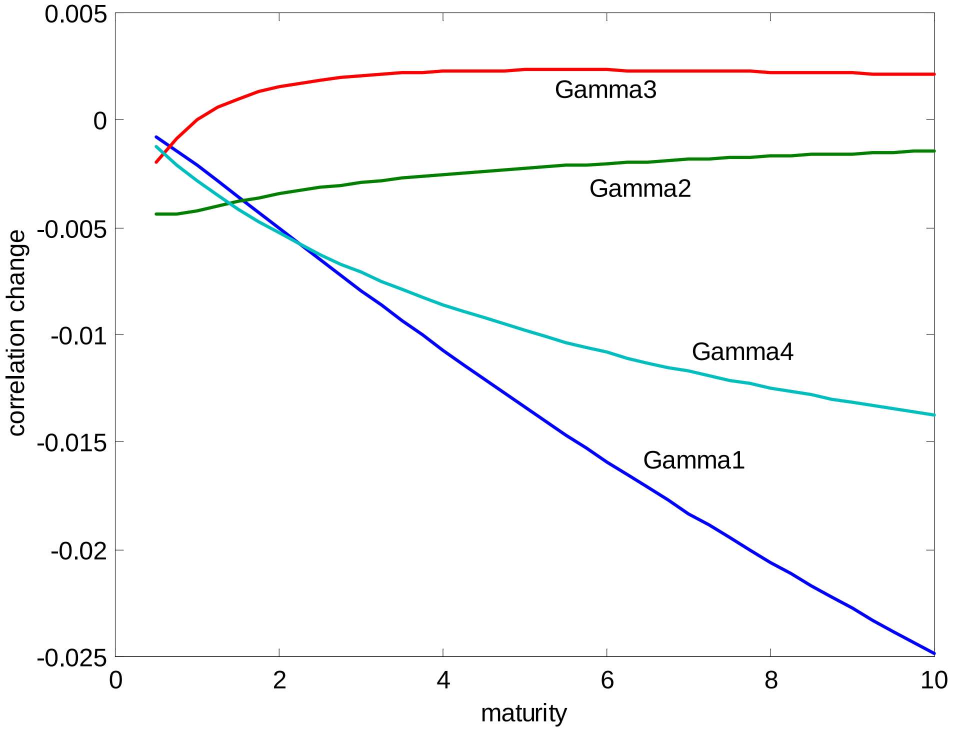

Figure 2. Term Structure of Cap Black Volatilities. The figure reports term structures of Black volatilities for caps. The solid line represents the time series average of the cap volatility data over the period January 1995 - June 1999. The other lines represent the time-series averages of cap Black volatilities, which are implied by model-based cap prices from interest-rate-based estimation and option-based estimation of the parameterization in (7) and (8).

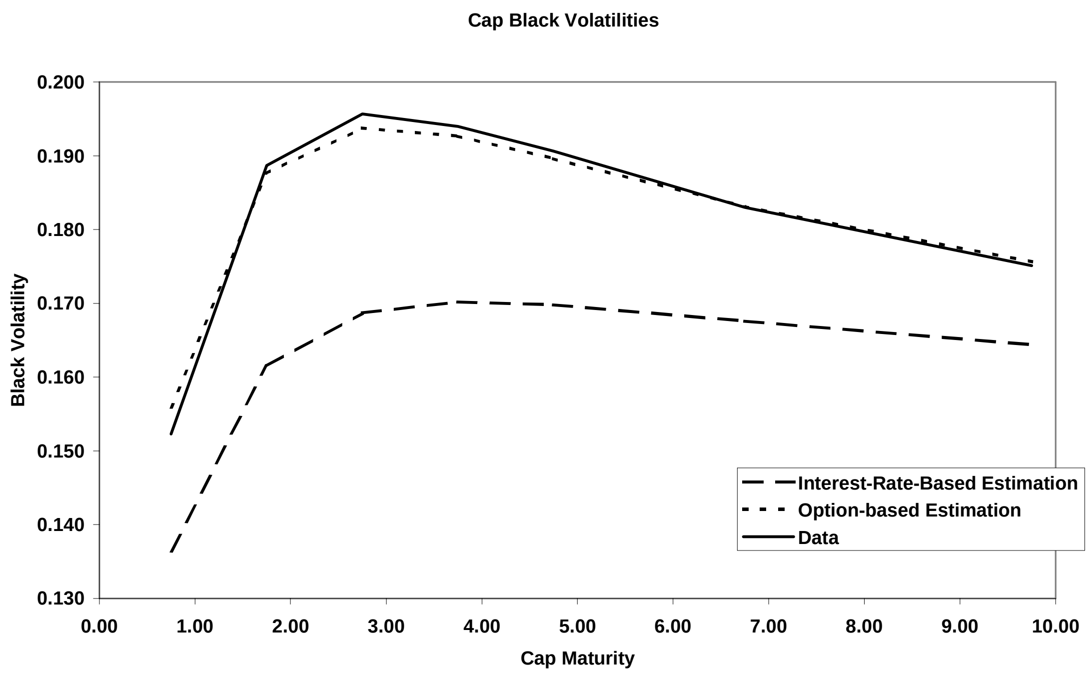

Figure 3. Term Structures of Swaption Black Volatilities. The figure reports term structures of Black volatilities for swaptions. The solid lines represent the time series average of the swaption volatility data over the period January 1995 - June 1999. The other lines represent the time-series averages of swaption Black volatilities, which are implied by model-based swaption prices from interest-rate-based estimation and option-based estimation of the parameterization in (7) and (8).

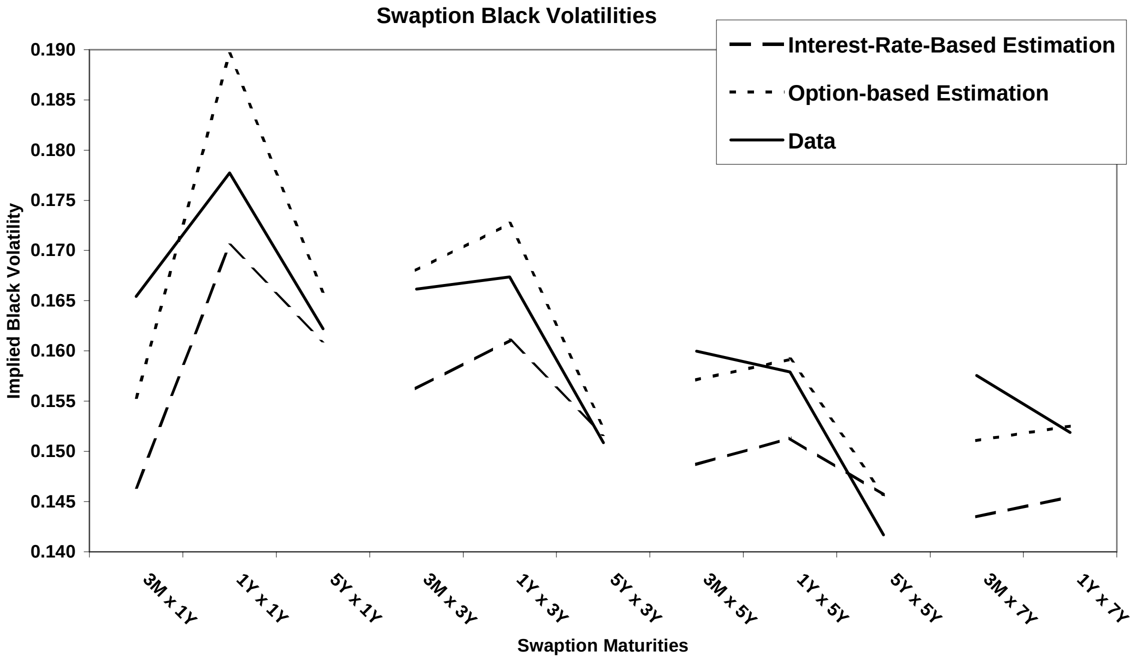

Figure 4. Term Structure of Forward Libor Volatilities. The figure reports term structures of forward Libor rate volatilities. The solid line represents annualized standard deviations of weekly forward Libor rate changes as observed in the 1995-2000 sample period. The other lines represent the model-implied annualized standard deviations of forward Libor rate changes, for respectively interest-rate-based estimation and option-based estimation of the parameterization in (7) and (8).

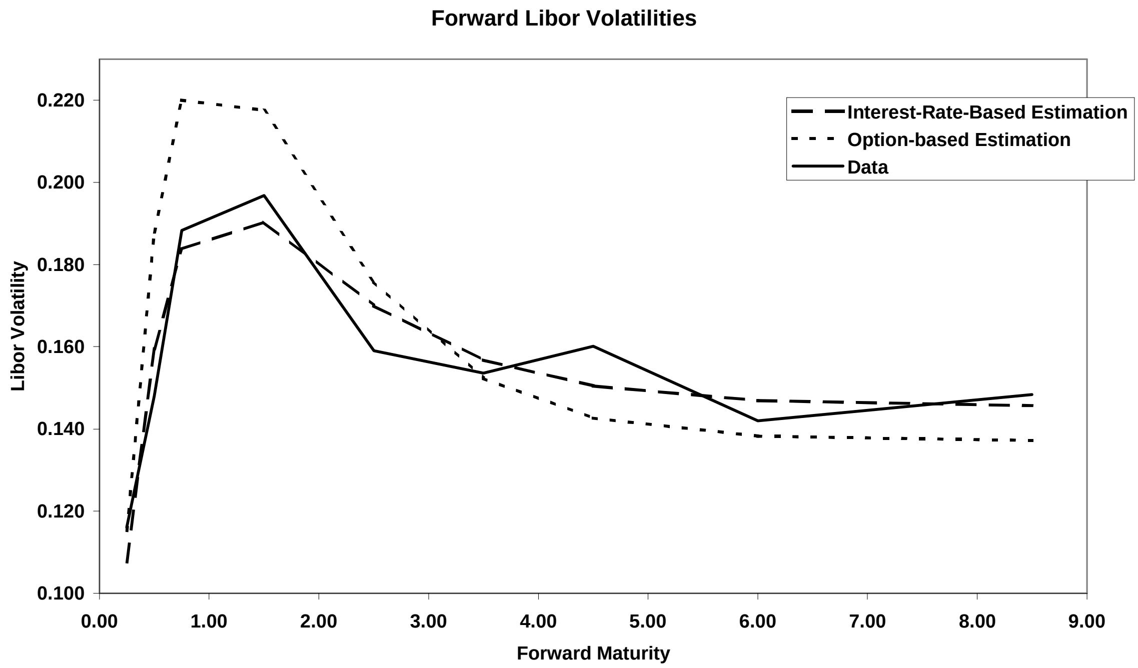

Figure 5. Data Correlation Matrix. The figure graphs the correlations between forward Libor rate changes of different forward maturities, estimated directly from weekly data on these rates from 1995-2000.

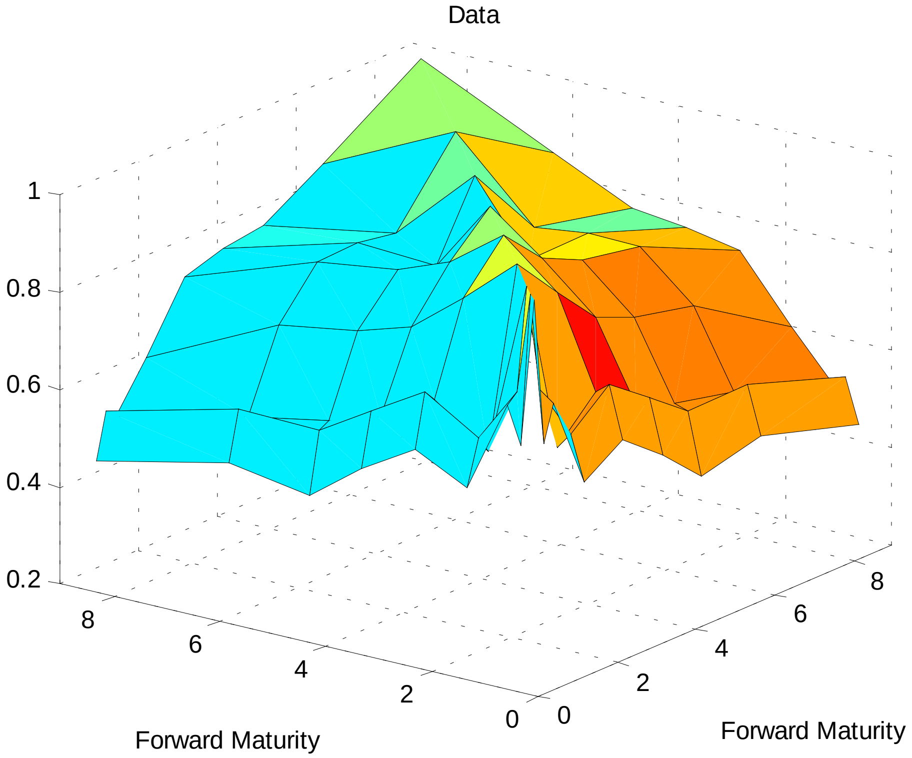

Figure 6. Interest-Rate Based Correlation Matrix. Interest-rate-based estimation is performed as described in Section 3.2. The figure graphs the correlations between forward Libor rate changes of different maturities, as implied by the correlation parameterization in equation (8) and interest-rate based parameter estimates (see Table 3).

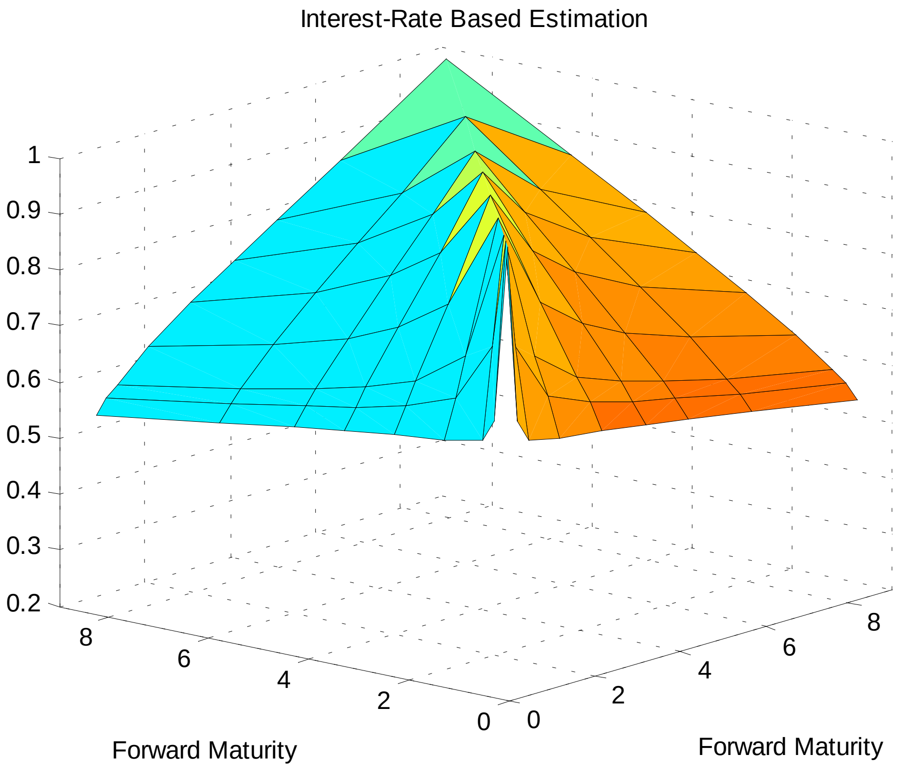

Figure 7. Term Structure of Cap Black Volatilities. The figure reports the time series average of the cap volatility data, using the full data set (solid line), and using a data set that excludes the Russia/LTCM crisis period, Aug 1998-Nov 1998 (dotted line). For comparison, the dashed line represents the time-series averages of cap Black volatilities implied by interest-rate based estimation.

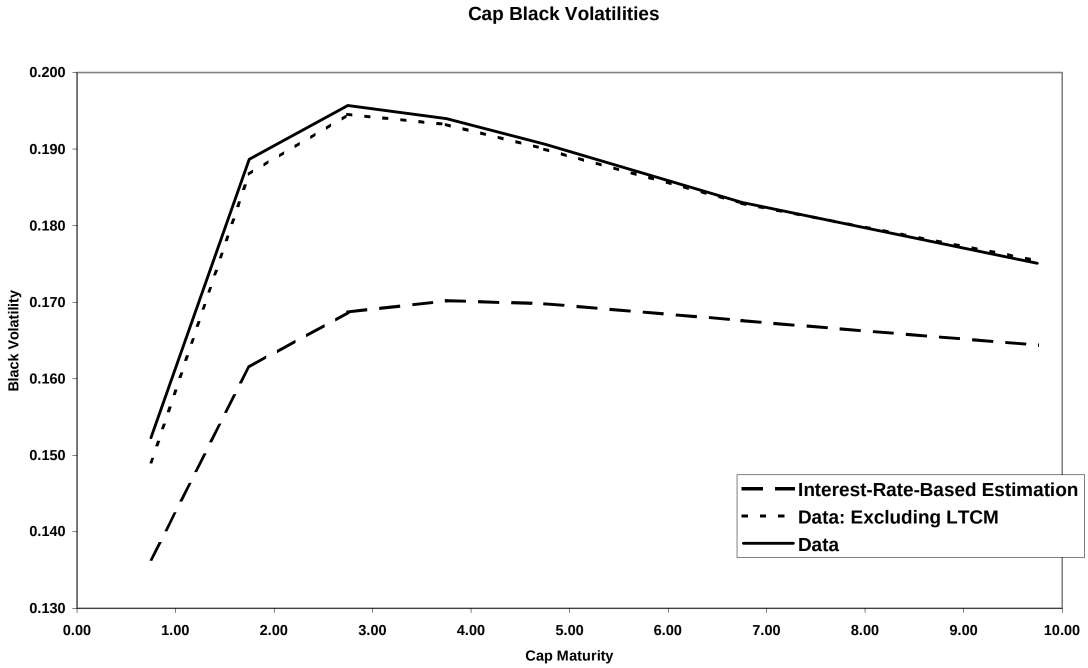

Figure 8. Term Structures of Swaption Black Volatilities. The figure reports the time series average of the swaption volatility data, using the full data set (solid line), and using a data set that excludes the Russia/LTCM crisis period, Aug 1998-Nov 1998 (dotted line). For comparison, the dashed line represents the swaption Black volatilities implied by interest-rate based estimation.

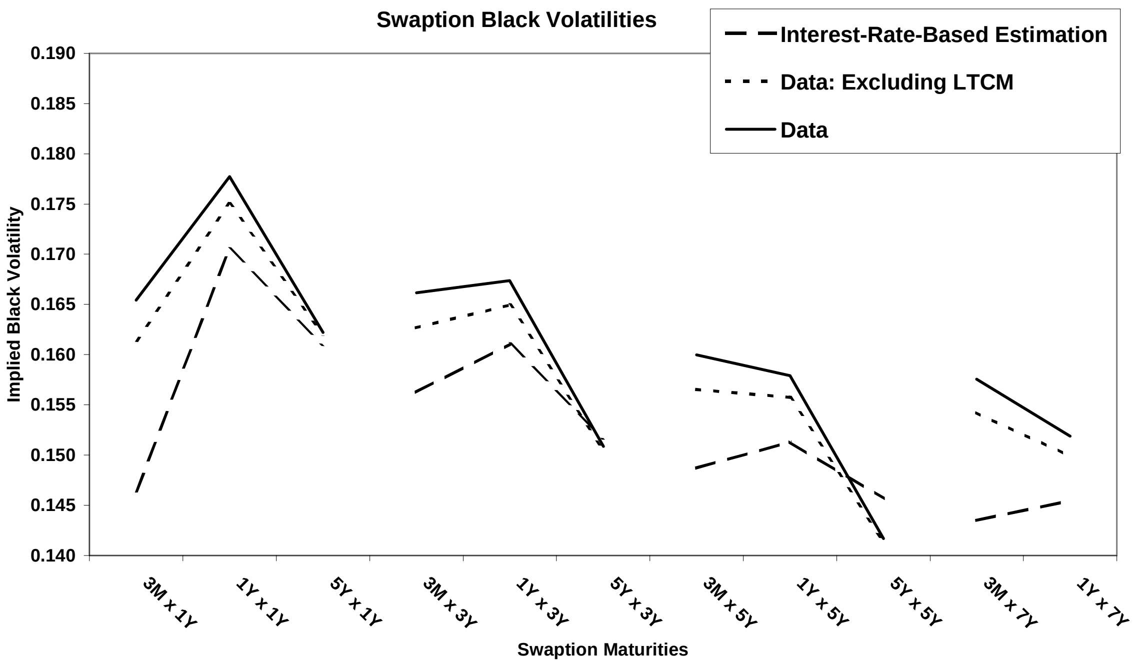

Figure 9. Option Based Correlation Matrix. Option-based estimation is performed as described in Section 3.2. The figure graphs the correlations between forward Libor rate changes of different maturities, as implied by the correlation parameterization in equation (8) and option based parameter estimates (see Table 3).

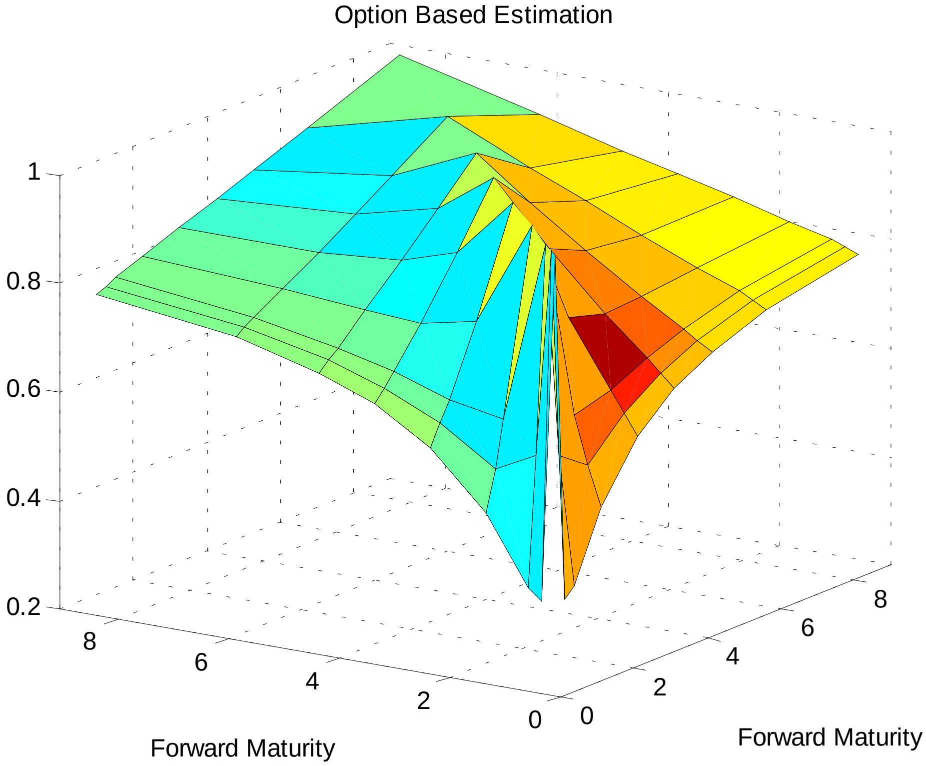

Figure 10. Option Based Correlation Matrix: Three-Factor Model. The figure graphs the correlations between forward Libor rate changes of different maturities, as implied by option-based estimation of the three-factor model in equation (9).

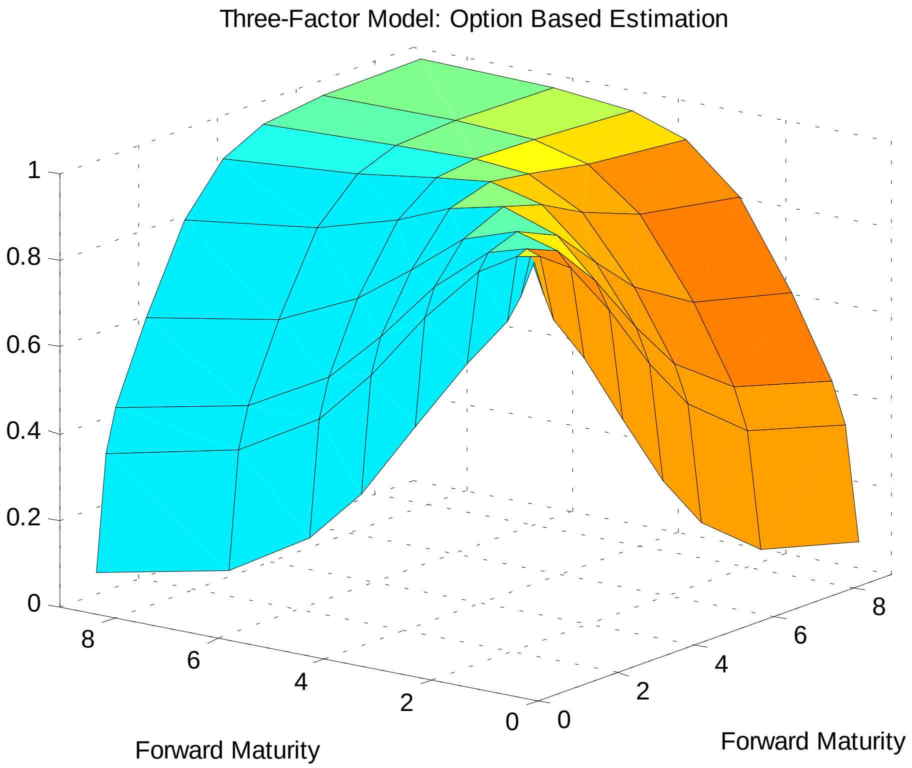

Figure 11. Option Based Correlation Matrix: Subset of Moments. The figure presents results for option-based estimation with a subset of option moment conditions. Only caps and swaptions with 3-month and 1-year option maturities are included. The figure graphs the correlations between forward Libor rate changes of different maturities, as implied by the correlation parameterization in equation (8) and option-based estimation for the described subset.