

The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence

FRANCIS A. LONGSTAFF, PEDRO SANTA-CLARA,
and EDUARDO S. SCHWARTZ*

ABSTRACT

Although traded as distinct products, caps and swaptions are linked by no-arbitrage relations through the correlation structure of interest rates. Using a string market model, we solve for the correlation matrix implied by swaptions and examine the relative valuation of caps and swaptions. We find that swaption prices are generated by four factors and that implied correlations are lower than historical correlations. Long-dated swaptions appear mispriced and there were major pricing distortions during the 1998 hedge-fund crisis. Cap prices periodically deviate significantly from the no-arbitrage values implied by the swaptions market.

THE GROWTH IN INTEREST-RATE SWAPS during the past decade has led to the creation and rapid expansion of markets for two important types of swap-related derivatives: interest-rate caps and swaptions. These over-the-counter derivatives are widely used by many firms to manage their interest-rate risk exposure and collectively represent the largest class of fixed-income options in the financial markets. The International Swaps and Derivatives Association (ISDA) estimates that the total notional amount of caps and swaptions outstanding at the end of 1997 was over \$4.9 trillion, which was more than 300 times the \$15 billion notional of all Chicago Board of Trade Treasury note and bond futures options combined.

Caps and swaptions are generally traded as separate products in the financial markets, and the models used to value caps are typically different from those used to value swaptions. Furthermore, most Wall Street firms use a piecemeal approach in calibrating their models for caps and swaptions, making it difficult to evaluate whether these derivatives are fairly priced

* All authors are from the Anderson School at UCLA. We acknowledge the capable research assistance of Martin Dierker and Bing Han. We are grateful for the comments of seminar participants at the University of British Columbia, the University of California at Riverside, Capital Management Sciences, Chase Manhattan Bank, Countrywide, Credit Suisse First Boston, Daiwa Securities, M.I.T, the Nippon Finance Association, the Norinchukin Bank, the Portuguese Finance Network meetings, Risk Magazine Conferences in Boston, London, and New York, Salomon Smith Barney in London and New York, Simplex Capital, Sumitomo Bank, the University of Texas at Austin, and the University of Washington. We are particularly grateful for the comments and suggestions of Alan Brace, Qiang Dai, Yoshihiro Mikami, João Pedro Nunes, Soetojo Tanudjaja, Toshiki Yotsuzuka, an anonymous referee, and the editor, René Stulz. All errors are our responsibility.

relative to each other. Financial theory, however, implies no-arbitrage relations that must be satisfied by cap and swaption prices. Specifically, a cap can be represented as a portfolio of options on individual forward rates. In contrast, a swaption can be viewed as an option on a portfolio of individual forward rates. Because of this, standard option pricing theory such as Merton (1973) implies that the relation between cap and swaption prices, or between different swaption prices, is driven primarily by the correlation structure of the forward rates. Given a unified valuation framework capturing these correlations, the no-arbitrage relations among cap and swaption prices can be tested directly.

This paper conducts an empirical analysis of the relative valuation of caps and swaptions using an extensive data set of interest-rate option prices. For the valuation framework, we use a string market model of the term structure of interest rates which blends the market-model framework of Brace, Gatarek, and Musiela (1997) and Jamshidian (1997) with the string-shock framework of Santa-Clara and Sornette (2001), Goldstein (2000), and Longstaff and Schwartz (2001). This approach has the important advantages of incorporating correlations directly into the model in a simple way and providing a unified framework for valuing fixed-income derivatives. The empirical approach taken in the paper consists of first solving for the covariance matrix implied by the market prices of all traded swaptions. This is the matrix equivalent of the familiar technique of solving for the implied volatility of an option. Once the implied covariance matrix has been identified, we can directly examine the implications for the relative values of caps and swaptions.

The empirical results provide a number of interesting insights into the fixed-income derivatives market. We find evidence of four statistically significant factors in the covariance matrix implied from market swaption prices. This contrasts with results based on historical covariance matrices which typically find only two to three factors, but is consistent with more recent evidence by Knez, Litterman, and Scheinkman (1994). Our results indicate that the market considers factors that contribute little to the unconditional volatility of term-structure movements, but represent a major source of conditional volatility during periods of market stress. Our results also indicate that the correlations among forward rates implied from swaption prices tend to be lower than those observed historically.

We then examine the relative valuation of swaptions and find that most swaptions tend to be valued fairly relative to each other. The major exception is during the 12-week period immediately following the announcement in September 1998 of massive trading losses at Long Term Capital Management. During this turbulent period, there is strong evidence of significant distortions in the quoted prices of many swaptions, a finding independently corroborated by interviews with many fixed-income derivatives traders. We also find that long-dated swaptions generally tend to be undervalued relative to other swaptions throughout the sample period.

Turning to the relative valuation of caps and swaptions, we find that the median differences between model and market cap prices are close to zero.

The distribution of differences, however, is skewed towards the right and all of the mean differences are positive and significant. This suggests that caps are typically valued fairly relative to swaptions, but that there are periodically large discrepancies between the two markets. This is particularly true during the hedge-fund crisis during late 1998. Alternatively, these results may imply that a more general model, such as one that allows a time-varying covariance structure, might be needed to capture fully the relative pricing of caps and swaptions.¹

Finally, we contrast the hedging performance of the string market model with that of the standard Black model often used in practice. Despite using only four hedging portfolios to hedge all of the swaptions in the sample, the string market model performs slightly better than the Black model, which uses a different hedge portfolio for each of the 34 swaptions in our sample.

The remainder of this paper is organized as follows. Section II provides a brief introduction to cap and swaption markets. Section III describes the string market model framework used to value interest-rate derivatives. Section IV discusses the data. Section V presents the empirical results. Section VI compares the implications of the string model for fixed-income derivatives with those of the Black model. Section VII summarizes the results and makes concluding remarks.

I. The Caps and Swaptions Markets

This section provides a brief introduction to the caps and swaptions markets. We first describe the characteristics of caps and explain how they are used in the financial markets. We then discuss the features of swaptions and their uses.

A. The Caps Market

Many financial market participants enter into financial contracts in which they pay or receive cash flows tied to some floating rate such as Libor. To hedge the risk created by the variability of the floating rate, firms often enter into derivative contracts that are essentially calls or puts on the level of the Libor rate. These types of derivatives are known as interest-rate caps and floors.

Specifically, a cap gives its holder a series of European call options or caplets on the Libor rate, where each caplet has the same strike price as the others, but a different expiration date.² Typically, the expiration dates for the caplets are on the same cycle as the frequency of the underlying Libor rate. For example, a five-year cap on three-month Libor struck at six percent represents a portfolio of 19 separately exercisable caplets with quar-

¹ One example of this type of model is Collin-Dufresne and Goldstein (2000). We are grateful to the referee for this insight.

² For many currencies, the market convention is for the cap to be on the three-month Libor rate. In some markets, however, caps may be on the six-month Libor rate. For example, Yen caps with maturities greater than one year are usually on the six-month Libor rate.

terly maturities ranging from one-half to five years, where each caplet has a strike price of 0.06.³ The cash flow associated with a caplet expiring at time T is $(a/360)\max(0, L(\tau, T) - K)$ where a is the actual number of days during the period from τ to T , $L(\tau, T)$ is the value at time τ of the Libor rate applicable from time τ to T , and K is the strike price. Note that while the cash flow on this caplet is received at time T , the Libor rate is determined at time τ , which means that there is no uncertainty about the cash flow from the caplet after Libor is set at time τ . The series of cash flows from the cap provides a hedge for an investor who is paying Libor on a quarterly or semiannual floating-rate note, where each quarterly or semiannual caplet hedges an individual floating coupon payment. In addition to caps, market participants often use interest-rate floors. These are similar to caps, except that the cash flow from an individual floorlet with expiration date T is $(a/360)\max(0, K - L(\tau, T))$. Thus, floors are essentially a series of European put options on the Libor rate. The market for interest-rate caps and floors is generically termed the caps market.

Market prices for caps and floors are universally quoted relative to the Black (1976) model. Specifically, let $D(t, T)$ denote the value at time t of a discount bond maturing at time T , and let $F(t, \tau, T)$ denote the value at time t for the Libor forward rate applicable to the period from time τ to T . Since $L(\tau, T) = F(\tau, \tau, T)$, a caplet can be viewed as an option on an individual Libor forward rate. Applying the Black model to this forward rate results in the following closed-form expression for the time-zero value of a caplet with expiration date T :

$$D(0, T) \frac{a}{360} [F(0, \tau, T)N(d) - KN(d - \sqrt{\sigma^2\tau}/2)], \quad (1)$$

where

$$d = \frac{\ln(F(0, \tau, T)/K) + \sqrt{\sigma^2\tau}/2}{\sqrt{\sigma^2\tau}}$$

and

$$F(0, \tau, T) = \frac{360}{a} \left(\frac{D(0, \tau)}{D(0, T)} - 1 \right)$$

and where σ is the volatility of changes in the logarithm of the forward rate. With this closed-form solution, the price of a cap is given by summing the values of the constituent caplets. Thus, a cap is simply a portfolio of indi-

³ The standard market convention is to omit the first caplet since the cash flow from this caplet is set at time $t = 0$ and is not stochastic.

vidual options, each on a different forward Libor rate. The market convention is to quote cap prices in terms of the implied value of σ , which sets the Black model price equal to the market price. Note that the convention of quoting cap prices in terms of the implied volatility from the Black model does not necessarily mean that market participants view the Black model as the most appropriate model for caps. Rather, implied volatilities from the Black model are simply a more convenient way of quoting prices, because implied volatilities tend to be more stable over time than the actual dollar price at which a cap would be traded.

B. The Swaptions Market

The underlying instrument for a swaption is an interest rate swap. In a standard swap, two counterparties agree to exchange a stream of cash flows over some specified period of time. One counterparty receives a fixed annuity and pays the other a stream of floating cash flows tied to the three-month Libor rate. Counterparties are identified as either receiving fixed or paying fixed in the swap. Although principal is not exchanged at the end of a swap, it is often more intuitive to think of a swap as involving a mutual exchange of \$1 at the end of the swap. From this perspective, the cash flows from the fixed leg are identical to those from a bond with coupon rate equal to the swap rate, whereas the cash flows from the floating leg are identical to those from a floating rate note. Thus, a swap can be viewed as exchanging a fixed rate coupon bond for a floating rate note.⁴

At the time a swap is initiated, the coupon rate on the fixed leg of the swap is specified. Intuitively, this rate is chosen to make the present value of the fixed leg equal to the present value of the floating leg. To illustrate how the fixed rate is determined, designate the current date as time zero and the final maturity date of the swap as time T . The fixed rate at which a new swap with maturity T can be executed is known as the constant maturity swap rate and we denote it by $FSR(0,0,T)$, where the first argument refers to time zero, the second argument denotes the start date of the swap which is time zero for a standard swap, and T is the final maturity date of the swap. Once a swap is executed, then fixed payments of $FSR(0,0,T)/2$ are made semiannually at times 0.50, 1.00, 1.50, ..., $T - 0.50$, and T . Floating payments are made quarterly at times 0.25, 0.50, 0.75, ..., $T - 0.25$, and T and are equal to $a/360$ times the three-month Libor rate at the beginning of the quarter, where a is the actual number of days during the quarter. This feature is termed setting in advance and paying in arrears. Abstracting from credit issues, a floating rate note paying three-month Libor quarterly must be worth par at each quarterly Libor reset date. Because the initial value of

⁴ For discussions about the economic role that interest-rate swaps play in financial markets, see Bicksler and Chen (1986), Turnbull (1987), Smith, Smithson, and Wakeman (1988), Wall and Pringle (1989), Macfarlane, Ross, and Showers (1991), Sundaresan (1991), Litzenberger (1992), Sun, Sundaresan, and Wang (1993), and Gupta and Subrahmanyam (2000).

a swap is zero, the initial value of the fixed leg must also be worth par. Setting the time-zero values of the two legs equal to each other and solving for the swap rate gives

$$FSR(0,0,T) = 2 \left[\frac{1 - D(0,T)}{A(0,0,T)} \right], \quad (2)$$

where $A(0,0,T) = \sum_{i=1}^{2T} D(0, i/2)$ is the present value of an annuity with first payment six months after the start date and final payment at time T . Swap rates are continuously available from a wide variety of sources for standard swap maturities such as 2, 3, 4, 5, 7, 10, 12, 15, 20, 25, and 30 years.

For many swaptions, the underlying swap has a forward start date. In a forward swap with a start date of τ , fixed payments are made at time $\tau + 0.50$, $\tau + 1.00$, $\tau + 1.50$, ..., $T - 0.50$, and T and floating rate payments are made at times $\tau + 0.25$, $\tau + 0.50$, $\tau + 0.75$, ..., $T - 0.25$, and T . At the start date τ , the value of the floating leg equals par. Discounting this time- τ value back to time zero implies that the time-zero value of the floating cash flows is $D(0, \tau)$. Because the forward swap has a time-zero value of zero, the time-zero value of the fixed leg must also equal $D(0, \tau)$. This implies that the forward swap rate $FSR(0, \tau, T)$ must satisfy

$$FSR(0, \tau, T) = 2 \left[\frac{D(0, \tau) - D(0, T)}{A(0, \tau, T)} \right]. \quad (3)$$

After a swap is executed, the coupon rate on the fixed leg may no longer equal the current market swap rate and the value of the swap can deviate from zero. Let $V(t, \tau, T, c)$ be the value at time t to the counterparty receiving fixed in a swap with forward start date $\tau \geq t$ and final maturity date T , where the coupon rate on the fixed leg is c . The value of this forward swap is given by

$$V(t, \tau, T, c) = \frac{c}{2} \sum_{i=1}^{2(T-\tau)} D(t, \tau + i/2) + D(t, T) - D(t, \tau), \quad (4)$$

where the first two terms in this expression represent the value of the fixed leg of the swap, and the third term is the present value of the floating leg, which will be worth par at time τ . For $t > \tau$, the swap no longer has a forward start date and the value of the swap on semiannual fixed coupon payment dates is given by the expression

$$V(t, \tau, T, c) = \frac{c}{2} \sum_{i=1}^{2(T-t)} D(t, \tau + i/2) + D(t, T) - 1. \quad (5)$$

Note that in either case, the value of the swap is just a linear combination of zero-coupon bond prices.

Swaptions or swap options allow their holder to enter into a swap with a prespecified fixed coupon rate, or to cancel an existing swap. Intuitively, swaptions can also be viewed as calls or puts on coupon bonds. Natural end users of swaptions are government agencies and firms coming to the capital markets to borrow funds. These entities use swaptions for the same reasons many firms issue callable or puttable debt—to cancel a swap with an above-market coupon rate or to enter into a new swap at a below-market coupon rate.

There are two basic types of European swaptions.⁵ The first is the option to enter a swap and receive fixed payments. For example, let τ be the expiration date of the swaption, c be the coupon rate on the swap, and T be the final maturity date on the swap. The holder of this option has the right at time τ to enter into a swap with a remaining term of $T - \tau$, and receive the fixed annuity of c . Because the value of the floating leg will be par at time τ , this option is equivalent to a call option on a bond with a coupon rate of c and a remaining maturity of $T - \tau$ where the strike price of the call is \$1. This option is generally called a τ into $T - \tau$ receivers swaption, where τ is the maturity of the option and $T - \tau$ is the tenor of the underlying swap. This swaption is also known as a τ by T receivers swaption. Note that if the option holder is paying fixed at rate c in a swap with a final maturity date of T , then exercising this option has the effect of canceling the original swap at time τ since the two fixed and two floating legs cancel each other out. Observe, however, that when the option is used to cancel the swap at time τ , the current fixed for floating coupon exchange is made first.

The second type of swaption is the option to enter a swap and pay a fixed rate, and the cash flows associated with this option parallel those described above. An option that gives the option holder the right to enter into a swap at time τ with final maturity date at time T and pay fixed is generally termed a τ into $T - \tau$ or a τ by T payers swaption. Again, this option is equivalent to a put option on a coupon bond where the strike price is the value of the floating leg at time τ of \$1. A τ by T payers swaption can be used to cancel an existing swap with final maturity date at time T where the option holder is receiving fixed at rate c .

From the symmetry of the European payoff functions, it is easily shown that a long position in a τ by T receivers swaption and a short position in a τ by T payers swaption with the same coupon has the same payoff as receiving fixed in a forward swap with start date τ and coupon rate c . A standard no-arbitrage argument gives the receivers/payers parity result that at time t , $0 \leq t \leq \tau$, the value of the forward swap must equal the value of the

⁵ For a discussion of the characteristics of American-style swaptions, see Longstaff, Santa-Clara, and Schwartz (2001). Callable bonds are also very similar to swaptions. For a discussion of callable bonds, see Bliss and Ronn (1998).

receivers swaption minus the value of the payers swaption. When the coupon rate c equals the forward swap rate $FSR(t, \tau, T)$, the forward swap is worth zero and the receivers and payers swaptions have identical values. In this case, the swaptions are said to be at the money forward.

As in the caps markets, the convention in the swaptions market is to quote prices in terms of their implied volatility relative to a standard pricing model. In swaption markets, prices are quoted as implied volatilities relative to the Black (1976) model as applied to the forward swap rate. Again, this does not mean that the market views this model as the most accurate model for swaptions. To illustrate how prices are quoted in the swaptions market, consider a τ by T European payers swaption where the fixed coupon rate equals c . Under the assumption that the forward swap rate follows a lognormal process under the annuity measure (the measure where the value of the annuity $A(t, \tau, T)$ is used as the numeraire), the Black model implies that the value of this swaption at time zero is

$$\frac{1}{2} A(0, \tau, T) [FSR(0, \tau, T) N(d) - cN(d - \sigma\sqrt{\tau})], \quad (6)$$

where

$$d = \frac{\ln(FSR(0, \tau, T)/c) + \sigma^2\tau/2}{\sigma\sqrt{\tau}},$$

where $N(\cdot)$ is again the cumulative standard normal distribution function and σ is the volatility of the logarithm of the forward swap rate. The value of the corresponding receivers swaption is given from the receivers/payers parity result. In the special case where the swaption is at-the-money forward, $c = FSR(0, \tau, T)$ and equation (6) reduces to

$$(D(0, \tau) - D(0, T)) [2N(\sigma\sqrt{\tau}/2) - 1]. \quad (7)$$

Because this receivers swaption is at the money forward, the value of the corresponding payers swaption is identical. When an at-the-money-forward swaption is quoted at an implied volatility of σ , the actual price that is paid by the purchaser of the swaption is given by substituting σ into equation (7).⁶

In the previous section, we showed that caps are simple portfolios of options on individual forward rates. In contrast, swaptions can be viewed as

⁶ Smith (1991) describes the application of the Black (1976) model to European swaptions. Brace et al. (1997), Jamshidian (1997), and others demonstrate that the Black model for swaptions can be derived within an internally consistent no-arbitrage model of the term structure in which the numeraire is the value of an annuity.

options on portfolios of forward rates. To see this, recall that a swaption is an option on the forward swap rate in the Black (1976) model. Furthermore, forward swap rates can be expressed as nearly linear functions of individual forward rates, where the weights are related to the durations of the cash flows from the fixed leg of the swap.⁷ From this, it follows that the swaption can be thought of as an option on a linear combination or portfolio of forward rates. Merton (1973) presents a number of no-arbitrage propositions including the well-known result that the value of an option on a portfolio must be less than or equal to that of a corresponding portfolio of options. This inequality is strict if the assets underlying the individual options are not perfectly correlated. Although the forward swap rate is only approximately linear in the individual forward rates, the key implication of the Merton result, namely that the relative value of a portfolio of options and an option on a portfolio is determined by the correlations between the underlying assets, is directly applicable to caps and swaptions. This key implication motivates many of the empirical tests later in the paper. In particular, we solve for the correlation matrix among forwards implied by a set of swaption prices, and then examine the extent to which other fixed-income options satisfy the no-arbitrage restrictions imposed by the correlation structure of forwards.

Finally, while both caps and swaptions are quoted in terms of the Black (1976) model, it should be recognized that the Black model is being actually used in different ways in these markets. In particular, the caps market uses the forward short-term Libor rate as the underlying state variable in the Black model, whereas the swaptions market uses longer-term forward swap rates. Because forward swap rates are nearly linear in individual forward rates, the lognormality assumption implicit in the Black model cannot hold simultaneously for both individual forward rates and forward swap rates, since a linear combination of lognormal variates is not lognormal. This is the sense in which the two markets use different models; the inputs used in the Black model differ across the two markets. In addition, since the volatilities used in the Black model are for fundamentally different rates, direct comparisons between the quoted implied volatilities of caps and swaptions are invalid. This has important implications for the risk management of portfolios of caps and swaptions.

II. The Valuation Framework

In this section, we develop a general string market model for valuing fixed-income derivatives such as caps and swaptions. We then describe how to invert the model to solve for the implied covariance matrix that best fits observed market prices.

⁷ This well-known rule of thumb or approximation can be obtained by differentiating the expression for the forward swap rate in equation (3) with respect to either spot or forward rates. For example, see Fabozzi (1997, Chapter 5).

A. The String Market Model

In a series of recent papers, Brace et al. (1997), Jamshidian (1997), and others develop term-structure models in which either Libor forward rates or forward swap rates are taken to be fundamental and their dynamics modeled directly using a Heath, Jarrow, and Morton (1992) framework. This class of models is often referred to as market models since they are based on the forwards of observable term rates in the market rather than on instantaneous forward rates. This approach has the advantage of solving some technical problems associated with continuously compounded lognormal rates as well as paralleling the standard practitioner approach of basing models on term rates. Libor-based and swap-based market models have been applied to a variety of interest-rate derivative valuation problems. Because the structure of these models is closely related to that of the Heath et al. framework, they share many of the same calibration issues and have typically only been implemented with a small number of factors.

In another recent literature, Kennedy (1994, 1997), Goldstein (2000), Longstaff and Schwartz (2001), and Santa-Clara and Sornette (2001) model the evolution of the term structure as a stochastic string. In this approach, each point along the term structure is a distinct random variable with its own dynamics, but which may be correlated with the other points along the term structure. Thus, string models are inherently high-dimensional models. Surprisingly, however, string models can actually be much easier to calibrate than models with fewer factors. The reason for this is that string models are directly parameterized by the correlation function for the points along the string. This direct approach is generally much more parsimonious than the standard approach of parameterizing the elements of a matrix of diffusion coefficients. The advantages of the string model approach to parameterization become increasingly important as the number of factors driving the term structure increases. Santa-Clara and Sornette show that the string model approach generalizes the Heath et al. (1992) framework for instantaneous forward rates while preserving its intuitive structure and appeal.

In this paper, we blend the market model setup with the string model approach of calibration to develop a valuation framework for fixed-income derivatives. This approach has the advantage of allowing us to develop the model in terms of the forward Libor rates that underlie the prices of caps and swaptions. At the same time, this approach makes it possible to directly model the correlation structure among Libor forwards in a simple way even when there are a large number of factors. Capturing the correlation structure is particularly important in this study; recall from earlier discussion that the correlation structure among forwards plays a central role in determining the relative valuation of caps and swaptions. We designate this valuation framework the string market model (SMM).

In this model, we take the Libor forward rates out to 10 years $F_i \equiv F(t, T_i, T_i + 1/2)$, $T_i = i/2$, $i = 1, 2, \dots, 19$, to be the fundamental variables

driving the term structure. Similarly to Black (1976), we assume that the risk-neutral dynamics for each forward rate are given by

$$dF_i = \alpha_i F_i dt + \sigma_i F_i dZ_i, \quad (8)$$

where α_i is an unspecified drift function, σ_i is a deterministic volatility function, dZ_i is a standard Brownian motion specific to this particular forward rate, and $t \leq T_i$.⁸ Note that although each forward rate has its own dZ_i term, these dZ_i terms are correlated across forwards. The correlation of the Brownian motions together with the volatility functions determine the covariance matrix of forwards, Σ . This is different from traditional implementations of multifactor models that use several uncorrelated Brownian motions to shock each forward rate. This seemingly minor distinction actually has a number of important implications for the estimation of model parameters from market data.

To model the covariance structure among forwards in a parsimonious but economically sensible way, we make the assumption that the covariance between dF_i/F_i and dF_j/F_j is time homogeneous in the sense that it depends only on $T_i - t$ and $T_j - t$.⁹ Furthermore, since our objective is to apply the model to swaps that make fixed payments semiannually, we make the simplifying assumption that these covariances are constant over six-month intervals. With these assumptions, the problem of capturing the covariance structure among forwards reduces to specifying a 19 by 19 time-homogenous covariance matrix Σ .

One of the key differences between this string market model and traditional multifactor models is that our approach allows the parameters of the model to be uniquely identified from market data. For example, if there are N forward rates, the covariance matrix Σ has only $N(N + 1)/2$ distinct parameters. Thus, market prices of fixed-income derivatives contain information on at most $N(N + 1)/2$ covariances, and no more than $N(N + 1)/2$ parameters can be uniquely identified from the market data. Since the string market model is parameterized by Σ , the parameters of the model are econometrically identified. In contrast, a typical implementation with constant coefficients of a traditional N -factor model of the form

$$dF_i = \alpha_i F_i dt + \sigma_{i1} F_i dZ_1 + \sigma_{i2} F_i dZ_2 + \dots + \sigma_{iN} F_i dZ_N, \quad (9)$$

⁸ We assume that the initial value of F_i is positive and that the unspecified α_i terms are such that standard conditions guaranteeing the existence and uniqueness of a strong solution to equation (8) are satisfied. These conditions are described in Karatzas and Shreve (1988, Chapter 5). In addition, we assume that α_i is such that F_i is nonnegative for all $t \leq T_i$.

⁹ Although the assumption of time homogeneity imposes additional structure on the model, it has the advantage of being more consistent with traditional dynamic term-structure models in which interest rates are determined by the fundamental state of the economy. In addition, time homogeneity facilitates econometric estimation because of the stationarity of the model's specification. For discussions of the advantages of time-homogeneous models, see Andersen and Andreasen (2000) and Longstaff et al. (2001).

would require N parameters for each of the N forwards, resulting in a total of N^2 parameters. Given that there are only $N(N + 1)/2 < N^2$ separate covariances among the forwards, the general specification in equation (9) cannot be identified using market information unless additional structure is placed on the model. Similar problems also occur when there are fewer factors than forwards. By specifying the covariance or correlation matrix among forwards directly, the string market model avoids these identification problems. String models also have the advantage of being more parsimonious. For example, up to $N \times K$ parameters would be needed to specify a traditional K -factor model. In contrast, only $K(K + 1)/2$ parameters would be needed to specify a string market model with rank K .¹⁰

Although the string is specified in terms of the forward Libor rates, it is much more efficient to implement the model using discount bond prices. By definition,

$$F_i = \frac{360}{a} \left[\frac{D(t, T_i)}{D(t, T_i + 1/2)} - 1 \right]. \quad (10)$$

Thus, the forward rates F_i can all be expressed as functions of the vector of discount bond prices with maturities 0.50, 1, ..., 10. Conversely, these discount bond prices can be expressed as functions of the string of forward rates, assuming that standard invertibility conditions are satisfied.¹¹ Applying Itô's Lemma to the vector D of discount bond prices gives

$$dD = r D dt + J^{-1} \sigma F dZ, \quad (11)$$

where r is the spot rate, $\sigma F dZ$ is the vector formed by stacking the individual terms $\sigma_i(t, T_i) F_i dZ_i$ in the forward rate dynamics in equation (8), and J^{-1} is the inverse of the Jacobian matrix for the mapping from discount bond prices to forward rates. Since each forward depends only on two discount bond prices, this Jacobian matrix has the following simple banded diagonal form.¹²

¹⁰ These types of identification problems parallel those which occur in general affine term-structure models. The specification and identification issues associated with affine term-structure models are discussed in an important recent paper by Dai and Singleton (2000).

¹¹ The primary condition is that the determinant of the Jacobian matrix for the mapping from discount bond prices to forward swap rates be nonzero. If this condition is satisfied, local invertibility is implied by the Inverse Function Theorem.

¹² For notational simplicity, discount bonds are expressed as functions of their maturity date in the Jacobian matrix. The Jacobian matrix represents the derivative of the 19 forwards $F_{0.50}, F_{1.00}, F_{1.50}, \dots, F_{9.50}$ with respect to the discount bond prices $D(1.00), D(1.50), D(2.00), \dots, D(10.00)$. Since $\sigma(T_i - t) = 0$ for $T_i \leq 0.50$, $D(0.50)$ is not stochastic and does not affect the diffusion term in equation (11).

$$J = \begin{bmatrix} -\frac{D(0.50)}{D^2(1.00)} & 0 & 0 & \dots & 0 & 0 & 0 \\ \frac{1}{D(1.50)} & -\frac{D(1.00)}{D^2(1.50)} & 0 & \dots & 0 & 0 & 0 \\ 0 & \frac{1}{D(2.00)} & -\frac{D(1.50)}{D^2(2.00)} & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & \frac{1}{D(9.50)} & -\frac{D(9.00)}{D^2(9.50)} & 0 \\ 0 & 0 & 0 & \dots & 0 & \frac{1}{D(10.00)} & -\frac{D(9.50)}{D^2(10.00)} \end{bmatrix}$$

It is important to observe that the drift term rD in equation (11) does not depend on the drift term α_i in equation (8). The reason for this is that discount bonds are traded assets in this complete markets setting and their instantaneous expected return is equal to the spot rate under the risk-neutral measure.¹³ Thus, this string market model formulation has the advantage of allowing us to avoid specifying the complicated drift term α_i , making the model numerically easier to work with than formulations based entirely on forward rates. Again, since our objective is a discrete-time implementation of this model, we make the simplifying assumption that r equals the yield on the shortest maturity bond at each time period.¹⁴

The dynamics for D in equation (11) provide a complete specification of the evolution of the term structure. This string market model is arbitrage free in the sense that it fits the initial term structure exactly and the expected rate of return on all discount bonds equals the spot rate under the

¹³ The bond market is complete in the sense that there are as many traded bonds as there are sources of risk. Thus, although no discount bond is a redundant asset, the market is complete and all fixed income derivatives can be priced under a risk-neutral measure in which the expected returns on all bonds equals the riskless rate. For a discussion of this point, see Santa-Clara and Sornette (2001).

¹⁴ Extensive numerical tests indicate that this discretization assumption has little effect on the results; we find that this approach gives values for European swaptions that are virtually identical to those implied by their closed-form solutions.

risk-neutral pricing measure. Furthermore, the model allows each point along the curve to be a separate factor, but also allows for a general correlation structure through Σ . To complete the parameterization of the model, we need only specify Σ in a way that matches the market or the historical behavior of forward rates.

B. Implied Covariance Matrices

Rather than specifying the covariance matrix Σ exogenously, our approach is to solve for the implied matrix Σ that best fits the observed market prices of some set of market data. Specifically, we imply the covariance matrix from the set of all observed European swaption prices.

In solving for the implied covariance matrix, it is important to note that a covariance matrix must be positive definite (or at least positive semidefinite) to be well defined. This means that care must be taken in designing the algorithm by which the covariance matrix is implied from the data to insure that this condition is satisfied. Standard results in linear algebra imply that a matrix is positive definite if, and only if, the eigenvalues of the matrix are all positive.¹⁵

Motivated by this necessary and sufficient condition, we use the following procedure to specify the implied covariance matrix. First, we estimate the historical correlation matrix of percentage changes in forward rates H from a time series of forward rates taken from a five-year period prior to the beginning of the sample period used in our study.¹⁶ We then decompose the historical correlation matrix into its spectral representation $H = U\Lambda U'$, where U is the matrix of eigenvectors and Λ is a diagonal matrix of eigenvalues. Finally, we make the identifying assumption that the implied covariance matrix is of the form $\Sigma = U\Psi U'$, where Ψ is a diagonal matrix with non-negative elements. This assumption places an intuitive structure on the space of admissible implied covariance matrices.¹⁷ Specifically, if the eigenvectors are viewed as factors, then this assumption is equivalent to assuming that the factors that generate the historical correlation matrix also generate the implied covariance matrix, but that the implied variances of these factors may differ from their historical values. Viewed this way, the identification assumption is simply the economically intuitive requirement that the market prices swaptions based on the factors that drive term-structure move-

¹⁵ For example, see Noble and Daniel (1977).

¹⁶ We implement this procedure using the historical correlation matrix rather than the covariance matrix to simplify the scaling of implied eigenvalues. We have also implemented this procedure using the historical covariance matrix. Not surprisingly, the eigenvectors from the historical covariance matrix are very similar to those obtained from the historical correlation matrix.

¹⁷ This assumption is equivalent to requiring that the historical correlation matrix H and the implied covariance matrix Σ commute, that is, $H\Sigma = \Sigma H$. We are grateful to Bing Han for this observation.

ments. Extensive numerical tests suggest that virtually any realistic implied correlation matrix can be closely approximated by this representation.¹⁸

Given this specification, the problem of finding the implied covariance matrix reduces to solving for the implied eigenvalues along the main diagonal of Ψ that best fit the market data. Since there are typically far more swaptions than eigenvalues, we solve for the implied eigenvalues by standard numerical optimization where the objective function is the root mean squared error (RMSE) of the percentage differences between the market price and the model price, taken over all swaptions. Specifically, for a given choice of the elements of the diagonal matrix Ψ , we form the estimated covariance matrix $U\Psi U'$ and then simulate 2,000 paths of the vector of discount bond prices using the string market model dynamics in equation (11). In simulating correlated Brownian motions, we use antithetic variates to reduce simulation noise. The time homogeneity of the model is implemented in the following way. During the first six-month simulation interval, the full 19 by 19 versions of the matrices Σ and J are used to simulate the dynamics of the 19 forward rates. After six months, however, the first forward becomes the spot rate, leaving only 18 forward rates to simulate during the second six-month period. Because of the time homogeneity of the model, the relevant 18 by 18 covariance matrix is given by taking the first 18 rows and columns of Σ ; the last row and column is dropped from the covariance matrix Σ . Similarly, the first row and column are dropped from the Jacobian since they involve derivatives with respect to the first forward, which has now become the spot rate. This process is repeated until the last six-month period, when only the final forward rate remains to be simulated.

Using the paths generated, we then value the individual at-the-money-forward European swaptions by simulation and evaluate the RMSE. In simulating the prices of swaptions, we use the following procedure. First, recall that since we simulate the evolution of the full vector of discount bond prices of all maturities ranging up to 10 years, these bond values are available at the expiration date τ of the swaption for each of the simulated paths of the term structure. From these discount bond prices at time τ , we can calculate the value of the underlying swap for each path. Specifically, the value of the swap $V(\tau, \tau, T, c)$ at time τ is given by the expression

$$V(\tau, \tau, T, c) = \frac{c}{2} \sum_{i=1}^{2(T-\tau)} D(\tau, \tau + i/2) + D(\tau, T) - 1, \quad (12)$$

¹⁸ We note that there are alternative ways of specifying the correlation matrix. For example, Rebonato (1999) independently offers a method to construct correlation matrices among forward rates. In our framework, however, Rebonato's approach requires optimizing over a large set of parameters and is computationally infeasible. Additionally, we examined a variety of specifications where the covariance between the i th and j th forwards is of the form $e^{a+bT_i}e^{a+bT_j}e^{c|T_i-T_j|}$, where a , b , and c are calibrated to fit swaption prices based on the RMSE criterion. These types of specifications generally performed poorly relative to the specification used in this paper.

where c is the fixed coupon rate of the swap which is equal to the forward swap rate $FSR(0, \tau, T)$ defined by equation (3). Thus, the value of the underlying swap at the expiration date τ of the swaption is easily calculated using the vector of discount bond prices. Once the value of the underlying swap at time τ is determined, the cash flow from the swaption at time τ is simply $\max(0, V(\tau, \tau, T, c))$ for a receivers swaption and $\max(0, -V(\tau, \tau, T, c))$ for a payers swaption. For each path, we then discount the cash flow from the option by multiplying by the compounded money-market factor $\prod_{i=0}^{2\tau-1} D(i, i + 1/2)$. Finally, we average the discounted cash flows over all paths. Since at-the-money-forward receivers and payers swaptions have the same value, we use the average of the simulated receivers and payers swaptions as the simulated value of the swaption.

We iterate this entire process over different choices of the eigenvalues until convergence is obtained, using the same seed for the random number generator at each iteration to preserve the differentiability of the objective function with respect to the eigenvalue. Although 19 implied eigenvalues are required for the covariance matrix Σ to be of full rank, implied covariance matrices of lower rank can easily be nested in this specification by solving for the first N eigenvalues and then setting the remaining $19 - N$ equal to zero.¹⁹

III. The Data

In conducting this study, we use three types of data: Libor and swap data defining the term structure of interest rates, market-implied volatilities for European swaptions, and market-implied volatilities for Libor interest-rate caps. Together with the term-structure data, these implied volatilities define the market prices of swaptions and caps. The source of all data is the Bloomberg system, which collects and aggregates market quotations from a number of brokers and dealers in the OTC swap and fixed-income derivatives market.

The term-structure data consists of weekly observations (Friday closing) for the 6-month and 1-year Libor rates as well as midmarket 2-year, 3-year, 4-year, 5-year, 7-year, and 10-year par swap rates for the period from January 17, 1992, to July 2, 1999. These maturities are the standard maturities

¹⁹ Although the numerical optimization is conceptually straightforward, there are a number of ways in which the search algorithm can be accelerated. For example, a least squares algorithm similar to Longstaff and Schwartz (2001) can be used to approximate forward swap rates as linear functions of the individual forward rates. Given a covariance matrix, this linear approximation then implies closed-form expressions for the variance of individual forward swap rates at the expiration dates of the swaptions, which can then be used to provide a closed-form approximation to the value of the swaption. This closed-form approximation can then be corrected for bias by an iterative process of comparing the simulated values given by the string market model to those implied by this approximation, and then adjusting the approximation. The implied eigenvalues can then be determined by optimizing the closed-form approximation rather than having to resimulate paths of the term structure at each iteration. With this type of algorithm, solving for the implied eigenvalues typically takes less than 10 seconds using a 750 MHz Pentium III processor.

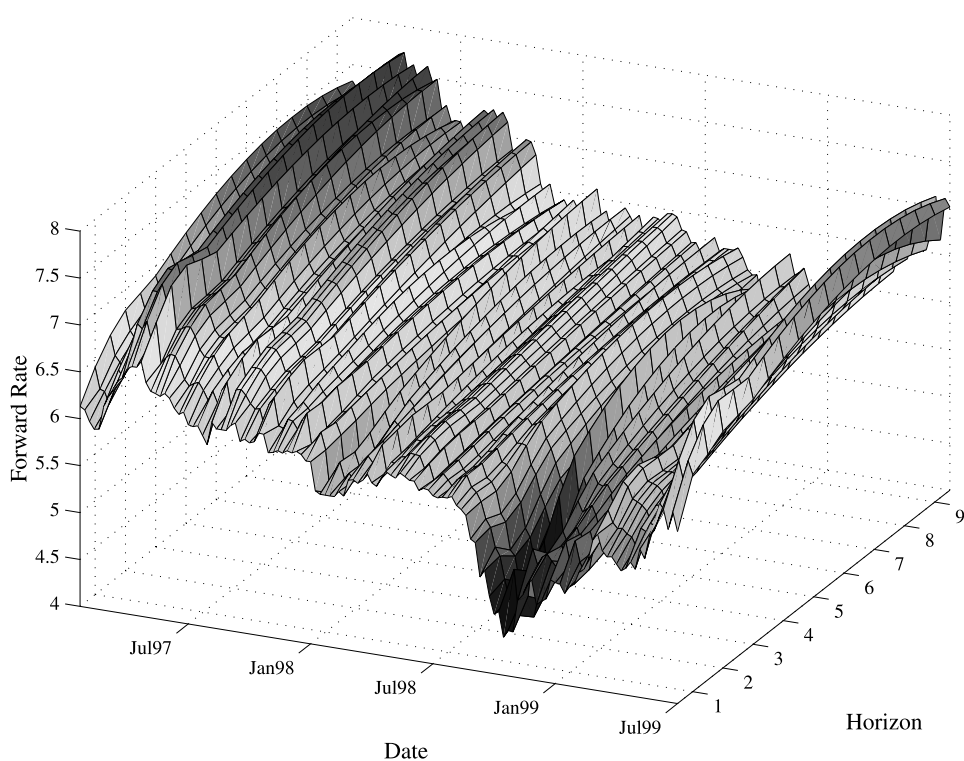


Figure 1. Time series of six-month Libor forward rates. The data set consists of weekly observations for six-month Libor forward rates starting at 0.5 to 9.5 years, for the period from January 24, 1997, to July 2, 1999. The forward rates are computed from the 6-month and 1-year Libor rates as well as the 2-year, 3-year, 4-year, 5-year, 7-year, and 10-year midmarket swap rates using a cubic spline to interpolate the par curve and then bootstrapping the forward curve. All data are obtained from the Bloomberg system. The weekly data for interest rates represent Friday closing rates. The forwards are denoted by the number of years until the beginning of the period covered by the 6-month forward rate. The total number of observations in the sample is 128.

for which swap rates are quoted in the market. From these rates, we solve for the term structure of six-month Libor forward rates out to 10 years in the following way. We first use the 6-month and 1-year Libor rates to solve for the 6-month and 1-year par rates. We then use a standard cubic spline algorithm to interpolate the par curve at semiannual intervals. Finally, we solve for 6-month forward rates by bootstrapping the interpolated par curve.²⁰ The term structures of Libor forward rates for the in-sample period from January 24, 1997, to July 2, 1999, are graphed in Figure 1. The term-

²⁰ Following the market convention, we discount cash flows using the swap curve as if it were the riskless term structure. Since the cash flows from both legs of a swap are discounted using this curve, however, this convention has little or no effect on valuation results.

structure data for the 5-year ex ante period from January 17, 1992, to January 17, 1997, is used to estimate the historical correlation matrix H from which the eigenvectors used in solving for the implied covariance matrices are determined. This ex ante correlation matrix is shown in Table I; all of the in-sample results are based on this ex ante correlation matrix. Note that the correlations are generally smooth monotonically decreasing functions of the distance between forward rates. One interesting exception is the correlation between the first and second forwards; the first two forwards display a significant amount of independent variation, hinting at money-market factors not present in longer-term forward rates.

The swaption data consists of weekly midmarket implied volatilities for 34 at-the-money-forward European swaptions for the in-sample period from January 24, 1997, to July 2, 1999. These 34 swaptions represent all of the standard quoted τ by T European swaption structures where the final maturity date of the underlying swap is less than or equal to 10 years, $T \leq 10$. As described earlier, the market convention is to quote swaption prices in terms of their implied volatility relative to the Black (1976) model for at-the-money-forward European swaptions given in equation (7); the market prices of these swaptions are given by substituting the implied volatilities into the Black model. Table II provides summary statistics for the implied volatilities. Figure 2 graphs the implied volatilities over time; Figure 3 shows a number of examples of the shape of the swaption implied volatility surface at different points in time during the sample period.

Observe that there is a significant spike in these implied volatilities during the fall of 1998. This spike coincides with the hedge-fund crisis precipitated by the announcement in early September 1998 of massive trading losses by Long Term Capital Management (LTCM). The sudden threat to the solvency of LTCM, which had been widely viewed as a premier client by many Wall Street firms, created a near panic in the financial markets. In the subsequent weeks, a number of other highly leveraged hedge funds also announced that they had experienced large trading losses on positions similar to those held by LTCM. Examples of these funds included Convergence Capital Management, Ellington Capital Management, D. E. Shaw & Co., and MKP Capital Management. In an effort to stabilize the market, the Federal Reserve Bank of New York persuaded a consortium of 16 investment and commercial banks to inject \$3.6 billion into LTCM in exchange for virtually all of the remaining equity in the fund. The prompt action by the Federal Reserve, announced to the markets on September 24, 1998, allowed LTCM to avoid insolvency and reduced the pressure on the fund to unwind trading positions at illiquid fire-sale prices, which would have exacerbated the problems at other hedge funds to which the consortium members had considerable risk exposure.

The interest-rate cap data consists of weekly midmarket implied volatilities for 2-year, 3-year, 4-year, 5-year, 7-year, and 10-year caps for the same period as for the swaptions data, January 24, 1997, to July 2, 1999. By market convention, the strike price of a T -year cap is simply the T -year

Table I
Correlation Matrix of Log Changes in Six-month Libor Forward Rates

The correlation matrix is based on weekly changes in the logarithm of individual six-month Libor forward rates for the ex ante period from January 17, 1992, to January 17, 1997. The forward rates are computed from the 6-month and 1-year Libor rates as well as the 2-year, 3-year, 4-year, 5-year, 7-year and 10-year midmarket swap rates using a cubic spline to interpolate the par curve and then bootstrapping the forward curve. All data are obtained from the Bloomberg system. The weekly data for interest rates represent Friday closing rates. The horizons of the 6-month forward rates used to compute the correlation matrix range from 0.50 years to 9.50 years forward, giving a total of 19 time series of forward rates. The total number of observations is 262.

	0.50	1.00	1.50	2.00	2.50	3.00	3.50	4.00	4.50	5.00	5.50	6.00	6.50	7.00	7.50	8.00	8.50	9.00	9.50	
0.50	1.000																			
1.00	0.340	1.000																		
1.50	0.579	0.950	1.000																	
2.00	0.547	0.946	0.995	1.000																
2.50	0.544	0.904	0.964	0.983	1.000															
3.00	0.519	0.812	0.883	0.918	0.975	1.000														
3.50	0.476	0.771	0.841	0.884	0.954	0.996	1.000													
4.00	0.440	0.731	0.800	0.849	0.929	0.984	0.996	1.000												
4.50	0.410	0.692	0.762	0.815	0.901	0.967	0.986	0.997	1.000											
5.00	0.388	0.656	0.727	0.781	0.871	0.944	0.969	0.986	0.996	1.000										
5.50	0.373	0.621	0.693	0.747	0.838	0.914	0.943	0.966	0.983	0.995	1.000									
6.00	0.365	0.585	0.659	0.711	0.798	0.873	0.905	0.933	0.957	0.978	0.994	1.000								
6.50	0.360	0.546	0.621	0.668	0.747	0.816	0.850	0.882	0.912	0.942	0.970	0.991	1.000							
7.00	0.358	0.501	0.576	0.615	0.681	0.739	0.773	0.807	0.844	0.882	0.922	0.960	0.988	1.000						
7.50	0.355	0.448	0.522	0.551	0.600	0.643	0.674	0.709	0.749	0.796	0.847	0.901	0.950	0.986	1.000					
8.00	0.349	0.391	0.462	0.479	0.507	0.531	0.557	0.591	0.634	0.686	0.748	0.816	0.884	0.944	0.986	1.000				
8.50	0.340	0.333	0.400	0.404	0.411	0.414	0.434	0.465	0.508	0.564	0.633	0.712	0.797	0.879	0.945	0.987	1.000			
9.00	0.329	0.279	0.341	0.333	0.320	0.302	0.315	0.343	0.385	0.442	0.515	0.603	0.700	0.799	0.886	0.952	0.989	1.000		
9.50	0.317	0.233	0.290	0.271	0.239	0.202	0.209	0.232	0.271	0.329	0.404	0.497	0.603	0.714	0.819	0.903	0.961	0.991	1.000	

Table II
Summary Statistics for OTC Market At-the-Money-Forward
European Swaption Volatilities

The data set consists of 128 weekly observations from January 24, 1997, to July 2, 1999, of midmarket implied Black-model volatilities for the indicated N into M at-the-money-forward European swaption structure, where N denotes years until option expiration and M denotes the length of the underlying swap in years.

N	M	Mean	Standard Deviation	Minimum	Median	Maximum	Serial Correlation
0.50	1.00	14.60	3.37	9.70	13.50	27.00	0.926
1.00	1.00	16.10	2.98	11.80	15.35	11.80	0.945
2.00	1.00	16.66	2.27	13.10	16.20	13.10	0.940
3.00	1.00	16.42	1.99	13.00	16.10	13.00	0.932
4.00	1.00	16.08	1.72	13.00	15.95	13.00	0.941
5.00	1.00	15.73	1.48	12.90	15.70	12.90	0.933
0.50	2.00	15.35	3.10	10.40	14.45	10.40	0.919
1.00	2.00	16.08	2.59	12.20	15.50	12.20	0.943
2.00	2.00	16.31	2.05	13.00	16.00	13.00	0.936
3.00	2.00	16.02	1.74	12.90	15.85	12.90	0.939
4.00	2.00	15.70	1.51	12.80	15.70	12.80	0.931
5.00	2.00	15.38	1.34	12.70	15.50	12.70	0.925
0.50	3.00	15.34	2.93	10.60	14.65	10.60	0.913
1.00	3.00	15.89	2.35	12.20	15.30	12.20	0.937
2.00	3.00	16.00	1.84	12.90	15.70	12.90	0.933
3.00	3.00	15.72	1.57	12.80	15.65	12.80	0.927
4.00	3.00	15.42	1.37	12.70	15.50	12.70	0.929
5.00	3.00	15.09	1.23	12.60	15.20	12.60	0.920
0.50	4.00	15.32	2.73	10.80	14.70	10.80	0.906
1.00	4.00	15.67	2.14	12.10	15.30	12.10	0.930
2.00	4.00	15.72	1.64	12.80	15.65	12.80	0.933
3.00	4.00	15.45	1.42	12.70	15.50	12.70	0.930
4.00	4.00	15.14	1.27	12.60	15.20	12.60	0.918
5.00	4.00	14.80	1.14	12.50	15.00	12.50	0.917
0.50	5.00	15.30	2.57	11.00	14.70	11.00	0.897
1.00	5.00	15.42	1.93	12.00	15.20	12.00	0.926
2.00	5.00	15.45	1.49	12.70	15.40	12.70	0.924
3.00	5.00	15.18	1.32	12.60	15.25	12.60	0.932
4.00	5.00	14.85	1.18	12.50	15.00	12.50	0.919
5.00	5.00	14.48	1.04	12.50	14.70	12.50	0.909
0.50	7.00	15.09	2.41	11.00	14.60	11.00	0.887
1.00	7.00	15.10	1.71	12.00	15.00	12.00	0.914
2.00	7.00	15.03	1.38	12.40	15.05	12.40	0.921
3.00	7.00	14.77	1.24	12.30	14.90	12.30	0.913

swap rate. To parallel the features of swaptions and to simplify the analysis, we assume that caps are on the 6-month Libor rate rather than the 3-month rate.²¹ The market prices of caps are then given by substituting the implied volatility into the Black model (1976) given in equation (1), where $T - \tau = 1/2$. Table III presents summary statistics for the market cap volatilities

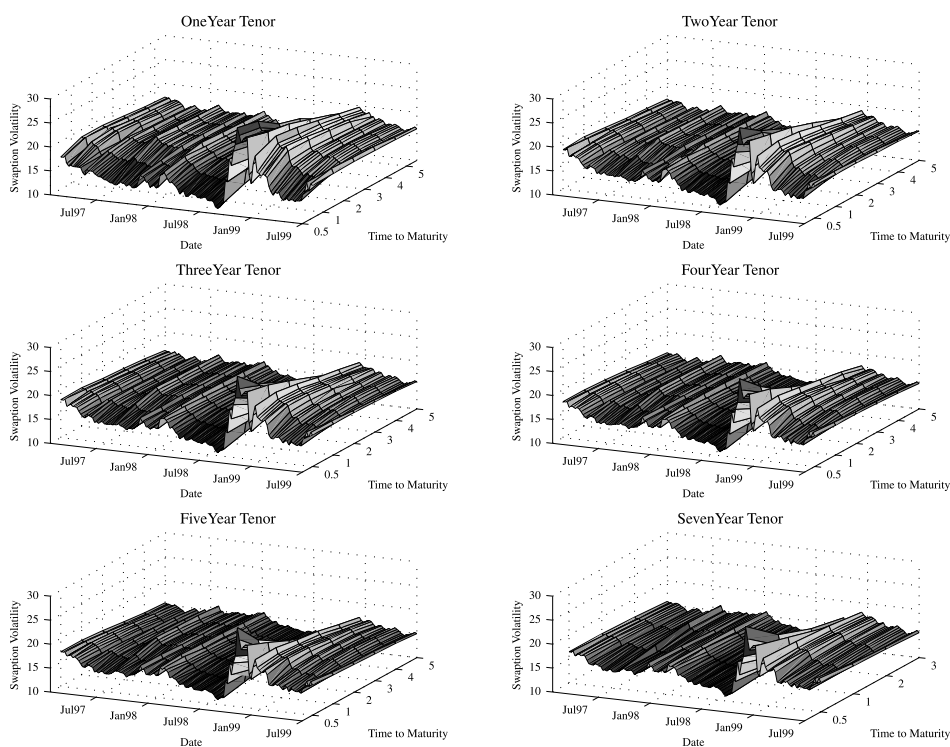


Figure 2. Time series of swaption volatilities. The data set consists of 128 weekly observations from January 24, 1997, to July 2, 1999, of midmarket implied Black-model volatilities for the indicated N into M at-the-money-forward European swaption structure, where N denotes years until option expiration (time to maturity) and M denotes the length of the underlying swap in years (the tenor). The subplots show, for each tenor, the implied volatilities of options with different times to maturity.

during the sample period. The implied volatilities display a time-series pattern similar to those observed for swaptions. Figure 4 also graphs the time series of cap volatilities.

IV. The Empirical Results

In this section, we report the empirical results from the study. First, we examine how many implied factors are required to explain the market prices of swaptions. We then study the relative valuation of swaptions in the string

²¹ This assumption is relatively innocuous. We have spoken with several caps dealers who indicated that the implied volatilities for caps on six-month Libor would typically be equal to or an eighth to a quarter below the implied volatility for a cap on three-month Libor. Diagnostic tests presented later in the paper indicate that this assumption has virtually no effect on the empirical results.

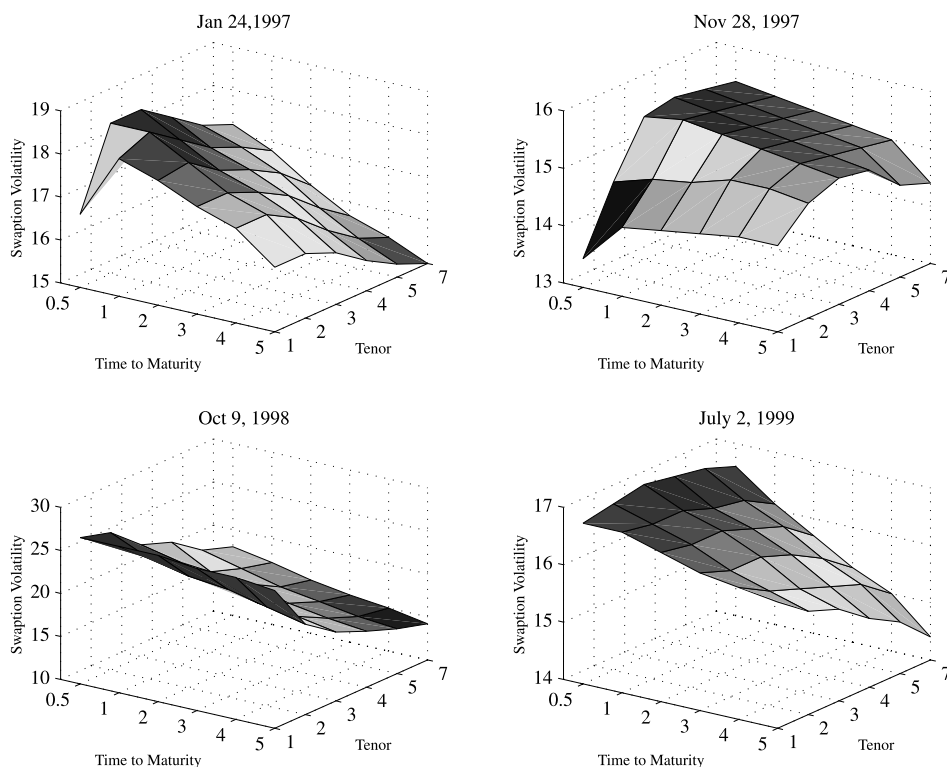


Figure 3. Examples of swaption volatility surfaces. This figure plots the quoted volatilities of swaptions on four different dates of the sample. Each figure shows quotes for swaptions with maturities between six months and five years on underlying swaps with horizons at the maturity of the options between one and seven years. Note that we do not use the four-into-seven and the five-into-seven swaptions in the empirical study since they are less liquid.

market model. Finally, we examine the relative valuation of both caps and swaptions in the string market model.

A. How Many Implied Factors?

Many researchers have studied the question of how many factors or principal components are needed to capture the historical variation in the term structure. For example, recent papers by Litterman and Scheinkman (1991) and others find that most of the variation in term-structure movements is explained by two or three factors. One important recent exception is Knez et al. (1994), who find evidence of a significant fourth factor affecting short-term interest rates.

An important advantage of our approach is that it offers a completely different perspective on this issue. Rather than focusing on the number of factors in historical term-structure data, we infer from swaption prices the

Table III
Summary Statistics for OTC Market Libor
Interest-rate Cap Volatilities

The data set consists of 128 weekly observations from January 24, 1997, to July 2, 1999, of midmarket implied Black-model volatilities for the indicated cap maturities. All data are obtained from Bloomberg.

Cap Maturity	Mean	Standard Deviation	Minimum	Median	Maximum	Serial Correlation
2 Year	15.34	3.19	10.60	14.25	28.00	0.945
3 Year	16.43	2.75	12.10	15.60	25.20	0.940
4 Year	16.75	2.41	12.50	16.10	23.50	0.932
5 Year	16.84	2.21	12.90	16.30	22.75	0.941
7 Year	16.46	1.86	12.75	16.15	20.87	0.933
10 Year	15.97	1.55	12.60	15.90	19.75	0.919

actual number of factors that market participants view as important influences on the term structure. Since the implied factor structure is forward looking, the number of implied factors need not be the same as those obtained historically. Intuitively, this approach is analogous to the familiar technique of solving for the implied volatility in option prices; implied volatilities typically do not equal estimates of volatility based on historical data, and often provide more accurate forecasts of future volatility.²²

We estimate the implied number of factors using an incremental likelihood ratio test based on all 128 weekly observations for each of the 34 European swaptions in the data set. Recall that when all but the first N eigenvalues in the diagonal matrix Ψ are equal to zero, the implied covariance matrix is of rank N , or equivalently, the implied covariance matrix is generated by N factors. For a given value of N , and for the i th week $i = 1, 2, \dots, 128$, we use the procedure described in Section III.B to solve for the N -implied eigenvalues that minimize the sum of squared percentage swaption pricing errors, where the percentage errors are defined as the differences between the simulated and market values of each swaption, expressed as a percentage of the market value of the swaption. Note that these pricing errors arise because we are trying to fit 34 swaption prices with only $N < 34$ parameters. Thus, these errors have an interpretation very similar to that of the residuals from a nonlinear least squares regression. We repeat the process of solving for the N eigenvalues that minimize the sum of squared per-

²² We note that other researchers have also used the approach of backing out factors from asset prices such as bonds. Important recent examples of this approach include Longstaff and Schwartz (1992), Chen and Scott (1993), Pearson and Sun (1994), Duffie and Singleton (1997), de Jong and Santa-Clara (1999), Dai and Singleton (2000), Duffee (2000), and many others. Our approach differs in that we use the information in swaption prices to address the question of the number of factors. Intuitively, it is clear that since swaptions have nonlinear payoffs, their prices may contain more information about market estimates of the conditional volatility of factors than can be recovered from bond prices alone.

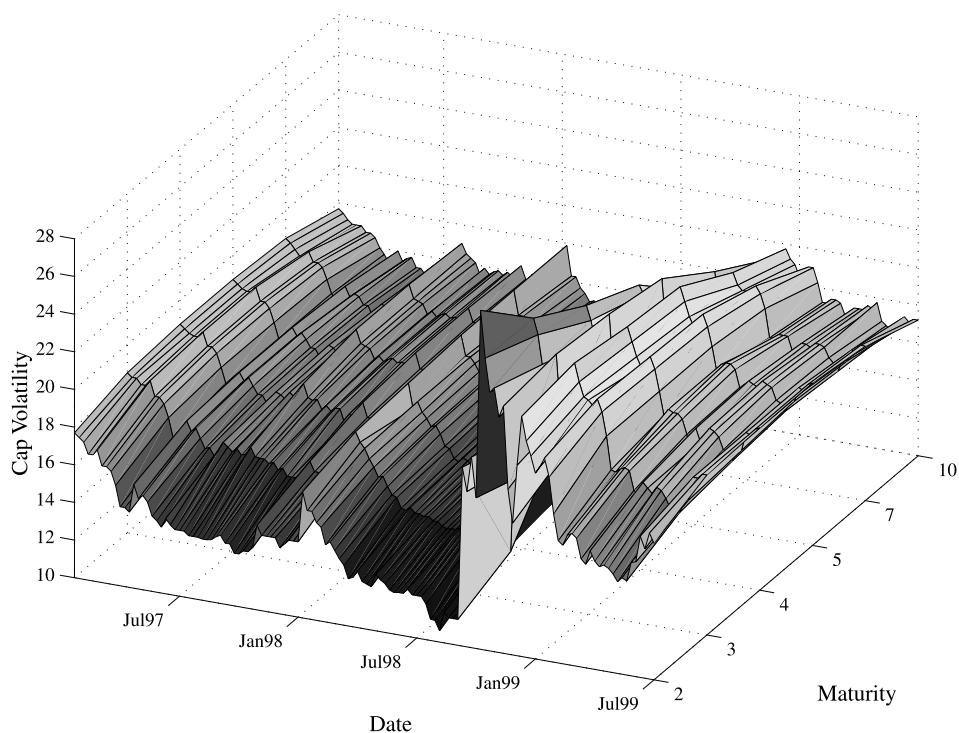


Figure 4. Time series of cap volatilities. The data set consists of 128 weekly observations from January 24, 1997, to July 2, 1999, of midmarket implied Black-model volatilities for the indicated cap maturities.

centage pricing errors for each of the 128 weeks in the sample period. Adding the sum of squared errors over all 128 weeks gives the total sum of squared errors. We then repeat this entire procedure for the case of $N + 1$ eigenvalues, where the same seed for the random number generator is used for all values of N to insure comparability in the results. Under the null hypothesis of equality, $128 \times 34 = 4,352$ times the difference between the logarithms of the sum of squared errors for N and $N + 1$ factors is asymptotically distributed as a chi-square variate with 128 degrees of freedom.

Table IV reports the results from the incremental pairwise comparisons as N ranges from one to seven. As shown, the pairwise comparisons are statistically significant for two versus one, three versus two, and four versus three factors, and are insignificant for all of the other comparisons. These results imply that there are four significant factors underlying the covariance matrix of forwards used by the market in the pricing of European swaptions. These results contrast with the earlier empirical work mentioned above, which finds only two to three factors in historical term-structure movements. It is important to mention, however, that most of these earlier studies focus on Treasury bonds whereas our results apply to the swap curve. Thus, it is

Table IV
Likelihood Ratio Tests for the Number of Implied Factors
in European Swaption Prices

This table reports the likelihood ratio test statistics from pairwise incremental comparisons of the number of factors. In each case, we solve for the N -implied eigenvalues that minimize the sum of squared errors for the swaption prices and then compare with the sum of squared errors obtained by solving for the $N + 1$ implied eigenvalues that best fit the data. The difference between the sum of squared errors is asymptotically χ^2_{128} under the null hypothesis of equality for the full sample, and χ^2_{64} for the two half samples. For a given vector of eigenvalues, the sum of squared errors is given by first forming the implied covariance matrix from the diagonal matrix of eigenvalues and the historical eigenvectors, simulating 2,000 paths of evolution of the term structure using the string model, and then solving for the individual swaption values by simulation. In generating simulated paths, the same seed for the random number generator is used to insure comparability across the number of factors. The data set consists of 128 weekly observations of 34 swaption values for the period from January 24, 1997, to July 2, 1999, giving a total of 4,352 observations. The critical value of χ^2_{128} is 168.1332 at the 99 percent level. The critical value of χ^2_{64} is 93.2169 at the 99 percent level.

N Factors	$N + 1$ Factors	Test Statistic	p -Value
A. Full Sample Period			
1	2	3352.183	0.000
2	3	2220.194	0.000
3	4	284.847	0.000
4	5	86.211	0.998
5	6	13.204	1.000
6	7	21.991	1.000
7	8	3.959	1.000
B. First Half of the Sample Period			
1	2	4481.614	0.000
2	3	1699.481	0.000
3	4	76.549	0.135
4	5	137.540	0.000
5	6	24.084	1.000
6	7	24.516	1.000
7	8	7.324	1.000
C. Second Half of the Sample Period			
1	2	2377.247	0.000
2	3	922.156	0.000
3	4	162.201	0.000
4	5	17.689	1.000
5	6	2.048	1.000
6	7	7.505	1.000
7	8	0.607	1.000

possible that the existence of a credit factor influencing swap rates but not Treasury rates could reconcile our results with those obtained by earlier researchers. Because of these results, all of our subsequent analysis is based on implied covariance matrices generated by four eigenvalues, resulting in four-factor or rank-four implied covariance matrices.

As a robustness check, we also conduct the incremental likelihood ratio tests using only the first half of the sample period (64 weeks) and also using only the second half of the sample period (64 weeks). Since the hedge-fund crisis of fall 1998 occurred entirely during the second half of the sample period, this diagnostic addresses whether the results about the number of factors are specific to this volatile period. As shown, however, the subperiod results are similar to those for the entire period. In both the first and second subperiods, the likelihood ratio tests find evidence of four statistically significant factors. Thus, the results about the number of factors are not artifacts of the hedge-fund crisis of fall 1998.²³

To provide some insight into the four implied factors that market participants view as driving the term structure, Figure 5 graphs the first four eigenvectors, which define the weights of the first four factors, from the historical correlation matrix in Table I. As illustrated, these factors closely resemble those found in earlier papers. The first factor essentially generates parallel shifts in the term structure. The second factor generates shifts in the slope of the term structure. The third factor is a curvature factor that generates movements in the term structure where short-term and long-term rates move in opposite directions from the midterm rates. Finally, the fourth factor primarily affects the shape of the very short end of the term structure, possibly reflecting the influence of the Federal Reserve or other monetary authorities. Thus, this fourth factor has an interpretation very similar to the fourth factor found by Knez et al. (1994) in their study of short-term rates.

Since the eigenvectors used in solving for the implied covariance matrix have the interpretation of term-structure factors, the fitted eigenvalues can be viewed as the implied variances of the factors. To illustrate this, Figure 6 graphs the time series of fitted values for each of the four eigenvalues used to define the implied covariance matrix. The first eigenvalue shows the relative volatility over time of the parallel shift factor. The volatility of this factor was very stable during much of 1997, decreased somewhat during the early part of 1998, and then increased significantly during the fall of 1998 when the financial stability of a number of highly visible hedge funds was threatened by severe trading losses. The volatility of the term-structure slope factor decreased significantly during 1997, and was quite low during most of

²³ It is interesting to note that the four significant factors during the first half of the sample are the first, second, third, and fifth, while the four significant factors during the entire sample period and during the second half of the sample period are the first, second, third, and fourth. Thus, one could argue that as many as five factors could occasionally be needed to describe swaption prices. We take the more parsimonious view that there are only four significant factors based on the results for the full sample period.

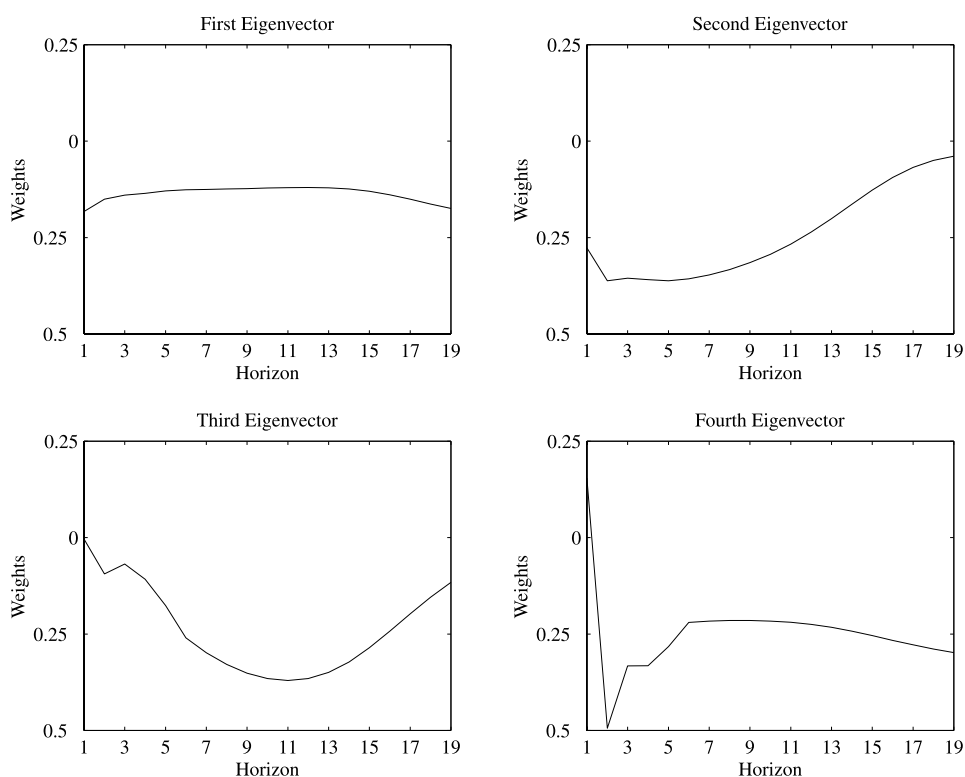


Figure 5. Eigenvector weights. The four subplots show the weights of the first four eigenvectors of the historical correlation matrix. The correlation matrix is based on weekly changes in the logarithm of individual six-month Libor forward rates for the ex ante period from January 17, 1992, to January 17, 1997. The forward rates are computed from the six-month and one-year Libor rates as well as the two-year, three-year, four-year, five-year, seven-year and 10-year midmarket swap rates using a cubic spline to interpolate the par curve and then bootstrapping the forward curve. The weekly data for interest rates represents Friday closing rates. The horizons of the six-month forward rates used to compute the correlation matrix range from 0.50 years to 9.50 years forward, giving a total of 19 time series of forward rates. The total number of observations is 262.

1998. In the fall of 1998, however, the volatility of this factor suddenly increased by a factor of nearly 10, but then quickly returned to levels near those at the beginning of the sample period. The volatility of the curvature factor shows a pattern similar to that of the slope factor; the volatility decreases significantly during 1997, is generally low during most of 1998, and then spikes dramatically during the fall of 1998. The behavior of the volatility of the short-term or fourth factor suggests one possible way of reconciling these results with the historical evidence on the number of factors. The implied volatility of this fourth factor is often quite small and can actually be zero. During periods of market stress such as the fall of 1998,

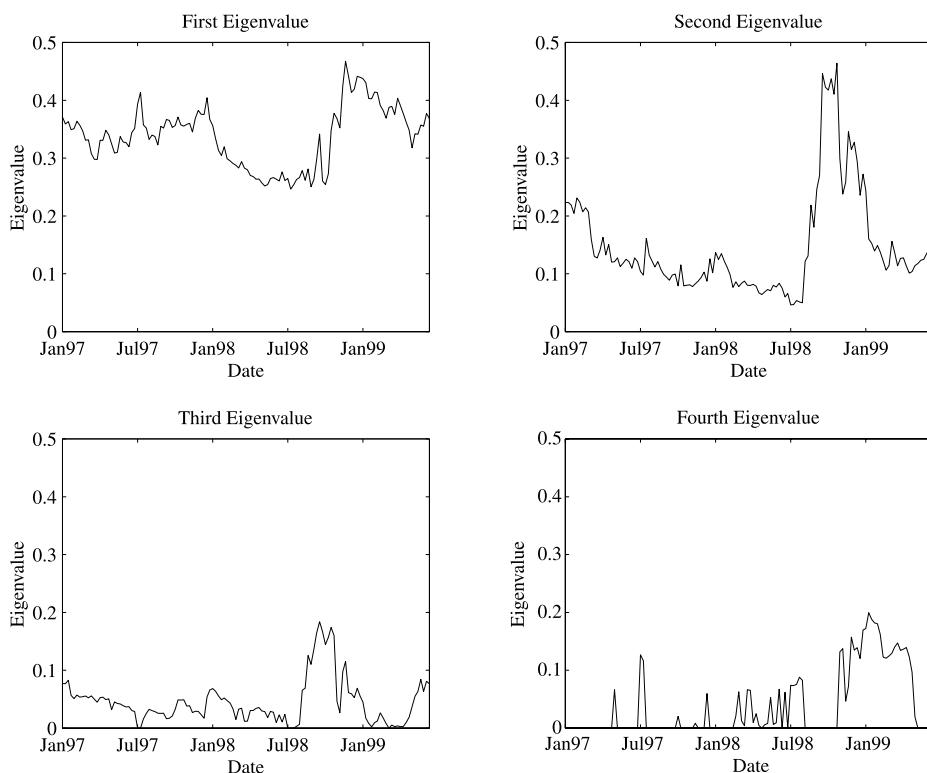


Figure 6. Time series of eigenvalues. The four subplots show the eigenvalues computed from the 128 weekly implied correlation matrices from January 24, 1997, to July 2, 1999, obtained by fitting the model to the swaption data, keeping fixed the eigenvectors of the historical correlation matrix.

however, the volatility of this factor can suddenly increase and become a major source of term-structure movements. Thus, the time-series pattern of the volatility of the fourth factor suggests that this may be more of an event-related factor that only becomes important in periods of extreme market stress. Since historical analysis of the number of factors is typically based on unconditional tests, factors that have time-varying volatilities that are usually small or zero may not show up in these types of standard tests. Despite this, these factors could represent a serious source of conditional volatility risk to market participants who would appropriately incorporate their effects into the market prices of swaptions. Recent papers by Hull and White (1999) and Jagannathan and Sun (1999) independently confirm that three factors are not sufficient to fully capture the pricing of interest rate caps and swaptions. Peterson, Stapleton, and Subrahmanyam (2000) find that going from one to two term-structure factors has a significant effect on the valuation of swaptions.

B. The Implied Correlation Matrix

As discussed, the implied eigenvalues uniquely determine the implied covariance matrix. In this sense, our approach is simply the matrix version of the familiar technique of inverting option prices to solve for the implied volatility of the underlying asset. One natural question that arises is how closely the implied correlation matrix matches the historical correlation matrix. To compare the two, we do the following. Based on the results of the likelihood ratio tests in the previous section, we set $N = 4$ and use the corresponding four implied eigenvalues for each week to define a diagonal matrix Ψ for each week. This diagonal matrix Ψ has the four implied eigenvalues as the first four elements along the diagonal, and zeros as the remaining diagonal elements. From Ψ and the historical matrix of eigenvectors U , the implied covariance matrix for that week is defined by $\Sigma = U\Psi U'$. Standardizing the covariance matrix gives the implied correlation matrix for that week. We repeat this process for all 128 weeks in the sample, resulting in a series of 128 implied correlation matrices.

To obtain summary measures of implied correlations, we then compute the matrix of average implied correlations by simply taking the time series average of each element in the implied correlation matrix over all 128 weeks. We then take the difference between the matrix of average implied correlations and the historical correlation matrix in Table I and report these differences in Table V. To provide some sense of the time-series variation in these differences, Table VI reports the matrix of standard deviations of the implied correlations.

As shown in Table V, there are clearly systematic differences between the historical and implied correlations. The differences along the main diagonal are all zero, of course, since the main diagonals of both the implied and historical correlation matrices consist of ones. As we move away from the main diagonal, however, the differences are almost all negative, which means that the implied correlations tend to be lower than the historical correlations. Most of the differences are on the order of 0.05 to 0.10, but a few are as large as 0.20. The largest differences are typically for the correlation of two-to-three-year forwards with seven-to-nine-year forwards. The only notable positive difference is for the correlation between the first and second forwards.

Table VI shows that there is a fair amount of time-series variation in the implied correlations, indicating that the implied correlation matrix is not constant over time. In general, however, the standard deviations do not appear to be excessively variable; most of the standard deviations range from 0.05 to 0.20. The largest standard deviation is for the correlation between the first and second forwards. Intuitively, however, it is this correlation that is likely to be the hardest to estimate since it only affects one of the swaptions; all of the other correlations affect multiple swaption values.

Table V
**Differences Between the Average Implied Correlations and the Historical Correlations
of Six-month Libor Forward Rates**

These differences are calculated by averaging the 128 weekly implied correlation matrices from January 24, 1997, to July 2, 1999, obtained by fitting the model to the swaption data, and then subtracting the historical correlations shown in Table I.

	0.50	1.00	1.50	2.00	2.50	3.00	3.50	4.00	4.50	5.00	5.50	6.00	6.50	7.00	7.50	8.00	8.50	9.00	9.50	
0.50	0.000																			
1.00	0.054	0.000																		
1.50	0.002	0.007	0.000																	
2.00	-0.000	0.008	0.000	0.000																
2.50	-0.012	-0.011	-0.006	-0.004	0.000															
3.00	-0.030	-0.047	-0.027	-0.021	-0.007	0.000														
3.50	-0.044	-0.058	-0.038	-0.032	-0.015	-0.002	0.000													
4.00	-0.057	-0.070	-0.049	-0.043	-0.023	-0.005	-0.001	0.000												
4.50	-0.068	-0.082	-0.061	-0.055	-0.033	-0.012	-0.006	-0.002	0.000											
5.00	-0.076	-0.095	-0.074	-0.068	-0.046	-0.022	-0.014	-0.007	-0.002	0.000										
5.50	-0.083	-0.112	-0.091	-0.085	-0.064	-0.039	-0.028	-0.018	-0.009	-0.003	0.000									
6.00	-0.088	-0.130	-0.111	-0.107	-0.088	-0.064	-0.051	-0.039	-0.026	-0.014	-0.005	0.000								
6.50	-0.090	-0.150	-0.133	-0.132	-0.118	-0.098	-0.085	-0.070	-0.053	-0.037	-0.020	-0.006	0.000							
7.00	-0.090	-0.168	-0.154	-0.156	-0.150	-0.136	-0.124	-0.109	-0.090	-0.069	-0.046	-0.023	-0.006	0.000						
7.50	-0.085	-0.177	-0.167	-0.172	-0.175	-0.170	-0.160	-0.145	-0.125	-0.103	-0.075	-0.047	-0.021	-0.004	0.000					
8.00	-0.080	-0.178	-0.174	-0.181	-0.189	-0.192	-0.183	-0.170	-0.152	-0.129	-0.101	-0.068	-0.036	-0.013	-0.003	0.000				
8.50	-0.073	-0.172	-0.172	-0.179	-0.193	-0.200	-0.193	-0.181	-0.164	-0.142	-0.114	-0.080	-0.046	-0.020	-0.005	-0.001	0.000			
9.00	-0.067	-0.159	-0.164	-0.171	-0.187	-0.196	-0.191	-0.181	-0.166	-0.145	-0.117	-0.084	-0.049	-0.020	-0.003	0.001	0.001	0.000		
9.50	-0.057	-0.142	-0.150	-0.156	-0.172	-0.183	-0.179	-0.170	-0.155	-0.137	-0.110	-0.078	-0.043	-0.013	0.003	0.008	0.005	0.002	0.000	

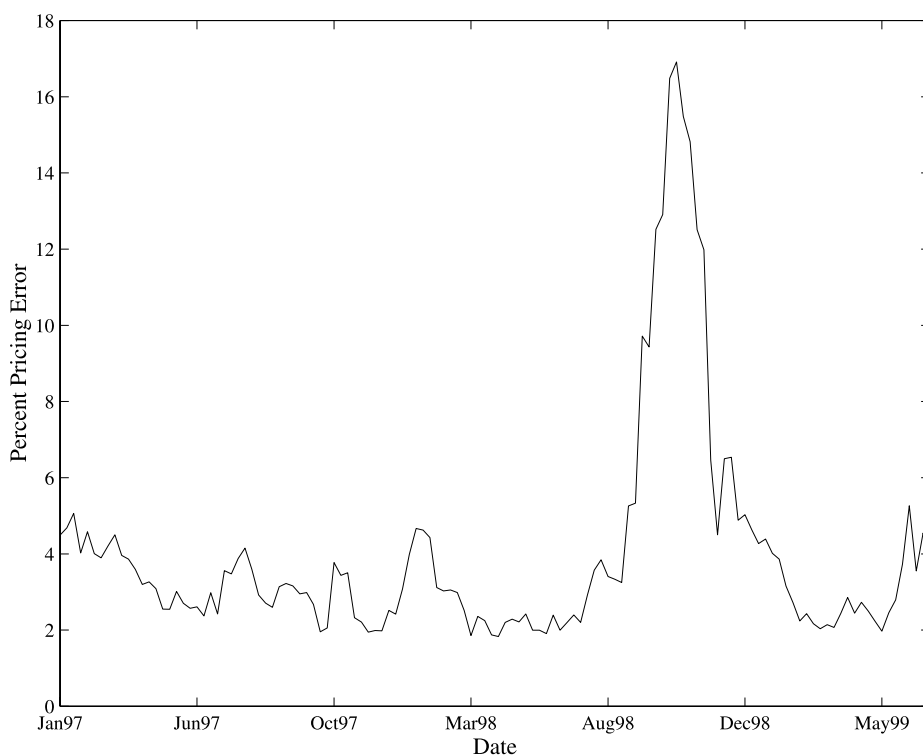


Figure 7. Time series of swaptions RMSE of four-factor string market model. The graph shows the RMSE of the differences between the fitted model swaptions prices and the market prices expressed as a percentage of the market prices. The data set consists of 128 weekly observations from January 24, 1997, to July 2, 1999, for 34 at-the-money-forward European swaption structures.

C. The Relative Valuation of Swaptions

The structure of the string market model imposes a number of constraints on the dynamic evolution of the term structure. Because of this, it is important to examine how well the model is able to describe the underlying structure of market swaption prices. Recall that the string market model is attempting to explain the cross section of 34 swaption prices using only four parameters each period. Thus, the model places a number of overidentifying restrictions of swaption prices and the pricing errors from fitting the string market model provide insights into how well these overidentifying restrictions are satisfied by the data.

To address this issue, Figure 7 graphs the RMSEs for the 34 swaptions in the sample for each of the 128 weeks in the sample period. Recall that these RMSEs are computed by first estimating the four eigenvalues that best fit the market swaption prices for that week, pricing the swaptions by simulating paths of the string market model, and then taking the percentage

differences between the market and model prices. As illustrated, the RMSEs from this time-homogeneous string market model are generally very small; the model typically captures the shape of the swaption volatility surface quite closely. Leaving out the exceptional period in the fall of 1998, the RMSEs are generally between two to three percent. These RMSEs are roughly about one-third to one-half of the size of the bid-offer spread.²⁴ The median RMSE is 3.10 percent and the standard deviation of the RMSEs is 2.98 percent.

Although the string market model fits the swaptions market well during most of the sample period, the fall of 1998 is clearly a major outlier. During this period, the RMSE spikes up to as high as 16 percent. The period during which the RMSE exceeds five percent begins with the week of September 11, 1998. Interestingly, this is just a few days after the well-publicized letter from John Meriwether to the investors of LTCM informing them that the fund had lost 52 percent of its capital through the end of August due to major trading losses in a number of markets. The RMSEs remain consistently above five percent for the 10-week period from September 11, 1998, to November 13, 1998, which closely aligns with the period during which most of the uncertainty about the survival of many of the hedge funds involved in the crisis was being resolved.

The failure of the string market model to capture the shape of the swaptions volatility surface during this period raises two possibilities: either the assumption of time-homogeneity is too restrictive, or quoted prices in the swaptions market were inconsistent with the absence of arbitrage. Although we cannot completely resolve this classical “joint-hypothesis” problem, we have conducted extensive interviews with many swaptions traders who experienced this period. These traders generally made two points. First, because of the turbulence in the market, the liquidity in the swaptions market was less than typical, and the quality of the market quotations collected by Bloomberg could be questioned. Second, there was an almost uniform belief among traders that there were in fact arbitrage opportunities in the markets. Many traders during this period felt that the fear of a complete market meltdown prevented them from executing trades that otherwise would have been viewed as highly profitable during ordinary circumstances.²⁵

Going beyond the overall RMSEs, it is also useful to examine the valuation errors for individual swaptions. Although the overall RMSEs are generally small and the fitting procedure requires pricing errors to have a mean close to zero, individual swaptions could still potentially display systematic

²⁴ Bloomberg reports that the typical bid-offer spread for these swaptions is about one unit of Black-model implied volatility; for a typical implied volatility of 16 percent, a 1 percent volatility bid-offer spread represents about 6 percent of the value of an at-the-money-forward swaption.

²⁵ Liu and Longstaff (2000) demonstrate that rational investors facing realistic margin constraints may actually choose to underinvest in arbitrages, or avoid investing in an arbitrage altogether, because of the risk that the arbitrage opportunity may widen further before it ultimately converges.

Table VII
Summary Statistics for Percentage Swaption Valuation Errors

The summary statistics reported are for the differences between the fitted model price and the market price expressed as a percentage of the market price. The data set consists of 128 weekly observations from January 24, 1997, to July 2, 1999, for the indicated N into M at-the-money-forward European swaption structures, where N denotes years until option expiration and M denotes the tenor or length of the underlying swap in years. The t -statistic for the mean and the standard deviation reported are adjusted for first-order serial correlation. The overall statistics are simple averages of the corresponding columns with the exception of the overall minimum, median, and maximum, which are computed from all observations.

N	M	Mean	t -Stat. Mean	Standard Deviation	Minimum	Median	Maximum	Serial Correlation
0.50	1.00	-2.456	-1.26	1.954	-19.623	-1.032	6.406	0.906
1.00	1.00	-1.123	-1.45	0.775	-12.922	-1.560	5.423	0.727
2.00	1.00	-2.124	-3.57	0.595	-8.583	-2.006	4.383	0.776
3.00	1.00	-0.924	-1.30	0.712	-5.146	-1.127	6.726	0.845
4.00	1.00	0.501	0.38	1.304	-4.030	-0.177	11.406	0.917
5.00	1.00	4.815	1.95	2.475	-1.209	3.470	24.683	0.935
0.50	2.00	-1.312	-1.06	1.238	-16.379	-1.033	5.057	0.853
1.00	2.00	-0.538	-1.43	0.375	-9.588	-0.273	3.654	0.575
2.00	2.00	-1.542	-3.24	0.476	-5.757	-1.577	2.163	0.809
3.00	2.00	0.289	0.42	0.685	-3.487	-0.131	9.523	0.855
4.00	2.00	1.906	1.35	1.411	-3.041	1.005	14.468	0.916
5.00	2.00	5.612	2.51	2.239	-0.995	4.289	24.027	0.924
0.50	3.00	-0.030	-0.01	2.332	-22.139	1.786	8.077	0.907
1.00	3.00	-1.696	-2.01	0.843	-13.825	-1.294	2.517	0.871
2.00	3.00	-1.920	-7.08	0.271	-4.701	-1.907	0.794	0.745
3.00	3.00	0.328	0.71	0.461	-6.406	0.112	6.868	0.793
4.00	3.00	2.033	1.69	1.201	-2.552	1.323	14.843	0.895
5.00	3.00	5.446	2.51	2.168	-0.193	3.976	22.971	0.934
0.50	4.00	-0.547	-0.19	2.815	-27.444	1.191	9.925	0.899
1.00	4.00	-2.523	-1.68	1.503	-16.314	-2.057	3.209	0.923
2.00	4.00	-2.209	-8.81	0.251	-5.382	-2.141	-0.034	0.763
3.00	4.00	0.151	0.40	0.373	-2.371	0.014	5.346	0.804
4.00	4.00	1.775	1.56	1.137	-2.494	1.124	12.965	0.912
5.00	4.00	5.409	2.64	2.045	-0.317	3.839	19.778	0.939
0.50	5.00	-1.529	-0.50	3.079	-32.517	0.387	9.693	0.897
1.00	5.00	-3.013	-2.62	1.149	-18.546	-2.631	6.900	0.814
2.00	5.00	-2.402	-6.42	0.374	-7.244	-2.148	-0.087	0.817
3.00	5.00	-0.446	-1.70	0.263	-3.433	-0.396	3.546	0.743
4.00	5.00	1.239	1.21	1.026	-1.963	0.383	10.447	0.915
5.00	5.00	5.615	2.79	2.013	-0.444	4.609	18.406	0.944
0.50	7.00	-3.604	-0.91	3.941	-40.981	-1.438	9.483	0.909
1.00	7.00	-5.365	-1.93	2.779	-26.164	-3.978	4.643	0.937
2.00	7.00	-5.096	-4.36	1.170	-16.876	-4.053	-2.015	0.909
3.00	7.00	-3.382	-15.35	0.220	-7.276	-3.154	-1.265	0.669
Overall		-0.255	-1.38	1.343	-40.981	-0.416	24.683	0.852

patterns of mispricing. To investigate this possibility, Table VII reports summary statistics for the pricing errors of individual swaption structures.

As shown, there are some clear patterns in the valuation errors. First, many of the valuation errors are highly serially correlated, implying that deviations between the model and market prices are persistent. Generally,

the most persistent errors are for the swaptions with five years to maturity, while the least persistent errors occur for the swaptions with one, two, or three years to maturity.²⁶

Table VII shows that although many of the means for the individual swaption valuation errors are significantly different from zero (after correcting the standard errors for serial correlation), the largest valuation errors occur for the swaptions with five years to maturity. In addition, the means for these five-year swaptions are all positive and greater than four percent. Note that the large positive means for these swaptions results in most of the other means being negative since there is an implicit adding-up-to-zero constraint imposed by the fitting procedure. Although smaller in magnitude, the mean differences for the swaptions with two years to maturity are also generally significantly different from zero. Another interesting feature of the valuation errors is that they tend to be skewed. This can easily be seen by comparing the mean valuation errors with the median errors. Note that the distribution of valuation errors tends to be skewed towards smaller values for short-maturity swaptions and towards larger values for long-maturity swaptions. Taken together, these results strongly suggest that there are significant and predictable valuation errors.

D. The Relative Valuation of Caps and Swaptions

In the string market model, cash flows from fixed-income derivatives can be expressed in terms of the fundamental forward rates defining the term structure. Thus, once the covariance matrix Σ has been estimated from the market prices of swaptions, the values of other fixed-income derivatives such as caps are uniquely determined by the string market model. In this sense, by parameterizing the model with swaption prices, which are essentially options on baskets of forwards, the model implies prices for caps, which can be viewed as baskets of options on individual forward rates. As in Merton (1973), the covariance matrix Σ determines the relation between the prices of options on portfolios and portfolios of options. It is important to note that the relation between swaption and cap prices implied by the model is a contemporaneous one; the prices of caps at time t in the model are implied from the prices of swaptions at time t . In this sense, the relative value relation implied by the model between caps and swaptions is similar to the put/call parity formula for options, which also places restrictions on the relative values of simultaneously observed call and put prices.

The main diagonal of the implied covariance matrix represents the implied variance of the individual forward rates as they roll down in maturity and become the spot rate. In particular, the implied variance of each forward

²⁶ It is important to note, however, that some of the persistence in these pricing errors may arise because the data consists of weekly observations of swaption prices where the maturities are typically multiple years. Thus, the overlapping nature of the data may induce serial correlation in the estimated pricing errors. We are grateful to the referee for pointing out this potential source of serial correlation in the pricing errors.

rate during the last period before it becomes the spot rate is the first element on the diagonal, the implied variance of each forward rate during the next-to-last period before it becomes the spot rate is given by the second element on the diagonal, and so forth. Since this provides a complete specification of the volatility of all forwards, the main diagonal uniquely determines the values of individual caplets (conditional on the number of eigenvalues fitted), which then determine the values of caps. Thus, once the model is fitted to the swaptions market, we can directly examine the implications for the valuation of caps. In the absence of arbitrage, the values of caps implied by the swaptions market should match the actual market prices of caps.

To examine the relative valuation of caps and swaptions, we use the main diagonal of the implied covariance matrix and solve for the implied values of 2-year, 3-year, 4-year, 5-year, 7-year, and 10-year at-the-money caps, using the Black model given in equation (1) and the initial term structure to value individual caplets. Since the Black model gives a closed-form expression for caplet prices, we do not need to solve for caps prices by simulation. The use of the Black model for pricing caplets is appropriate here since the lognormal dynamics for forward rates given in equation (8) imply that the Black model holds for individual caplets, since caplets are simple European options on individual forward rates. Note that the variance used in the Black model for valuing a caplet is simply the average variance for the corresponding forward rate from the present until the forward becomes the spot rate. Thus, the variance for the caplet maturing in 6 months is the first diagonal element, the variance for the caplet maturing in 12 months is the average of the first and second diagonal elements, and so forth. We repeat this procedure for each of the 128 weeks in the sample period and report summary statistics for the differences between the market and implied prices in Table VIII.

As illustrated, the hypothesis that market cap prices match the values implied by the swaption market is rejected for all of the maturities. The mean percentage pricing errors range from a high of 23.326 for the two-year caps down to 5.665 for the five-year caps. The positive means imply that the market cap prices are undervalued relative to swaptions. Note that these percentage pricing errors also tend to display a significant amount of persistence as evidenced by their first-order serial correlation coefficients.

A different perspective is obtained by focusing on the median values of the pricing differences. The median pricing errors are all within three percent of zero, and the overall median is only 0.862, which suggests that the caps and swaptions markets are usually consistent; the significant mean percentage pricing are primarily due to periodic large positive errors, resulting in a skewed, somewhat bimodal distribution of errors.

As an additional diagnostic, we also recompute the pricing errors under the assumption that the Black volatilities for caps on six-month Libor are 0.25 volatility points below those for caps on three-month Libor. Recall from the earlier discussion that there could be a slight difference in the quoted

Table VIII
Summary Statistics for Percentage Cap Valuation Errors

The summary statistics reported are for the percentage difference between the model price implied from fitting the string market model to the swaptions market and the market price expressed as a percentage of the market price. The data set consists of 128 weekly observations from January 24, 1997, to July 2, 1999, for the indicated maturities. The t -statistic for the mean and the standard deviation of the mean reported are adjusted for first-order serial correlation. The overall statistics are simple averages of the corresponding columns with the exception of the overall minimum, median, and maximum, which are computed from all observations.

Maturity	Mean	t -Stat. Mean	Standard Deviation	Minimum	Median	Maximum	Serial Correlation
2.00	23.326	3.50	6.669	-19.484	2.589	113.307	0.663
3.00	11.782	2.73	4.319	-11.359	-0.705	63.361	0.664
4.00	7.666	2.36	3.245	-9.739	-2.120	47.687	0.646
5.00	5.665	2.14	2.644	-9.191	-2.385	38.071	0.641
7.00	6.021	2.58	2.332	-8.941	-0.054	32.379	0.685
10.00	7.864	3.57	2.200	-4.592	2.883	32.578	0.734
Overall	10.387	2.81	3.568	-19.484	0.862	113.307	0.672

volatilities for caps on six-month Libor rather than on three-month Libor. The results, however, are virtually the same as those reported in Table VIII.

As another test of the Merton (1973) no-arbitrage bounds, we recompute the percentage pricing differences under the assumption that the correlations between all forwards equals one. This is done by fitting only a single implied eigenvalue to the market prices of the swaptions; all of the remaining eigenvalues are set equal to zero. This specification results in a rank-one covariance matrix, which in turn, implies perfect correlation among all forward rates. Following Merton, it is easily shown that the model price from the one-factor model should provide a lower bound for the value of a cap. This is directly an implication of the fact that the value of a portfolio of options should be greater than or equal to the value of an option on a portfolio. Thus, no-arbitrage considerations imply that the percent pricing differences from the one-factor model should all be positive. Virtually all of the cap prices satisfy this no-arbitrage bound. The mean and median values of the percentage pricing differences are now all negative. Of the $128 \times 34 = 4,352$ observations, only 8 or 0.18 percent are positive.

In summary, the evidence suggests that while caps and swaptions almost always satisfy the strictest no-arbitrage restriction of Merton (1973), the values of caps and swaptions are frequently inconsistent with each other. This is consistent with Hull and White (1999) who independently find that a set of cap and swaptions prices for a single day in August 1999 cannot be reconciled within the context of a three-factor model. Similarly, Jagannathan and Sun (1999) find that caps and swaptions appear significantly mispriced in a three-factor Cox, Ingersoll, and Ross (1985) framework. Our results suggest the possibility that whereas buy-and-hold arbitrages may not be feasible, dynamic trading strategies exploiting inconsistencies in the

relative valuation of caps and swaptions may be profitable. Of course, however, these results can also be interpreted as evidence that the model may be too restrictive. For example, a time-varying covariance structure may be necessary to capture fully the relative values of caps and swaptions.

V. A Comparison to the Black Model

In this paper, we have examined the relative valuation of caps and swaptions using a multifactor string market model of the term structure. As an additional issue, it is also useful to contrast the performance of the multifactor string market model with the standard Black model often applied to caps and swaptions in practice.

Before making any comparisons, however, it is important to first understand the key differences between the two modeling approaches. The string market model is a unified multifactor framework in which the same calibration is used, for example, in pricing and hedging all of the swaptions in the sample. In contrast, 34 separate specifications of the Black model are needed to price and hedge the 34 swaptions in the sample, each specification with a different forward swap rate as the underlying factor, and each with a distinct volatility calibration. In this context, the Black model is more appropriately viewed as a collection of different univariate models, where the relationship between the underlying factors is left unspecified.²⁷ In contrast, the string market model provides a complete unified description of the multivariate relationships among all points along the term structure.

The piecemeal way in which the Black model is typically used results in many limitations to its applicability. Because the Black model requires a different calibration for each swaption, it does not place any overidentifying restrictions on swaption prices. For example, even if 34 different versions of the Black model are fitted exactly to the prices of the 34 swaptions in the sample, these calibrations tell us nothing about what the price of a 35th swaption would be if it were introduced into the sample. Thus, the Black model cannot be extended to other swaptions with different maturities, expiration dates, or strike prices. In practice, the volatilities or prices of the 34 original swaptions would be interpolated or extrapolated to price a 35th swaption. It is important to note, however, that it is the assumptions about interpolation or extrapolation that determine the pricing of the 35th swaption in this situation, not the Black model. In contrast, once calibrated, the string market model can be used to price any other fixed-income derivative such as a swaption with a different maturity, exercise date, or strike price,

²⁷ Because the 34 forward swaps underlying the 34 swaptions in our sample can be expressed in terms of just 19 distinct forward rates, it is tempting to argue that the dimensionality of the Black model cannot be higher than the number of forward rates. Since Black model volatilities can be specified arbitrarily, however, these volatilities may not be linked by the dependence of the forward swap rates on a common set of forward rates. Hence, the Black model is best viewed as a collection of models that may not be strictly compatible with each other.

interest-rate caps, exotic interest rate options with payoffs that depend on multiple forward rates, and even American-style swaptions.

Another problem with the way that the Black model is applied in practice is that it cannot be used to hedge portfolios of options. Since each swaption has its own underlying asset in the Black model, a swaption can only be hedged with its own associated forward swap. Thus, hedging the 34 swaptions in the sample requires 34 distinct hedging instruments. Since the relationships between different forward rates are left unspecified in the Black framework, there is no clear way in which the risks of different swaptions can be aggregated without making ad hoc auxiliary assumptions unrelated to the Black model itself. In a strict sense, it is not appropriate to aggregate the hedge ratios that are computed using different calibrations of the Black model since there is no guarantee that the different calibrations will be internally consistent. Thus, the Black model provides no guidance on how one swaption can be hedged with another. This implies that cross-hedging is not possible using only the Black model.²⁸ In contrast, the string market model implies that all risks can be hedged using four factors and that the risks of different types of fixed-income derivatives can be directly aggregated at a portfolio level.

Although the Black model places no overidentifying restrictions on prices, it is possible to contrast the two models in terms of their abilities to hedge fixed-income derivatives. The Black model implies that changes in swaption values are driven entirely by changes in the corresponding underlying forward swap rate.²⁹ This means that changes in 34 distinct forward swap rates would be needed to explain the variation in the prices of the 34 swaptions in the sample. In contrast, the string market model implies that the variation in the prices of the 34 swaptions can be explained in terms of the changes of only four factors. Thus, the string model attempts to explain the variation in swaption prices using far fewer state variables than this interpretation of the Black model.

To compare the two models, we do the following. Using only information available at time t , we solve for the hedge ratios for each swaption with respect to the state variables. In the Black model, the underlying state variable is the forward swap rate and the hedge ratio is given analytically by differentiating the pricing expression. In the string market model, the hedge ratios for the four factors are computed by simulation by varying the initial curve at time t and recomputing swaption prices. In doing this, we use the eigenvalues from time t and the eigenvectors determined from the ex ante period prior to the beginning of the sample period; all of the information used in computing hedge ratios at time t in the string market model is ob-

²⁸ In practice, the risk of fixed income derivative portfolios is often calculated by computing the sensitivity of Black model prices to changes in individual forward rates. Note, however, that this approach is much more consistent with the string market model than with the Black model.

²⁹ It is easily shown from the Black model expression in equation (6) that changes in the forward swap rate are spanned by the returns on two separate portfolios of zero-coupon bonds with values $D(t, \tau) - D(t, T)$ and $A(t, \tau, T)$, respectively.

servable at time t . We then solve for the pricing errors for both models by taking the change in swaption prices from time t to $t + 1$ and subtracting the change in the hedging portfolio over the same period, where the change in the hedging portfolio is given from the hedge ratios and the changes in the individual state variables. Since we only have data on at-the-money-forward swaptions rather than repeated observations on the prices of a specific swaption, we make the identifying assumption that the volatility for the swaption at time $t + 1$ is the same as the volatility for the at-the-money-forward swaption at time $t + 1$ in computing price changes.³⁰ These differences directly measure the hedging errors resulting from using the hedge ratios and hedging instruments implied by the two models over a one-week horizon.

Intuitively, one might suspect that the Black model would perform much better in this hedging analysis since it hedges with the 34 specific underlying forward swaps whereas the string market model uses only four hedges for all 34 swaptions. Surprisingly, however, the string market model actually performs slightly better than the Black model. For the case of receivers swaptions, the Black model explains 89.28 percent of the variability in the price changes for the 34 swaptions over the 127-week period. In contrast, the string market model explains 89.35 percent of the variability in the price changes.³¹ For the case of payers swaptions, the Black model explains 92.46 percent of the variability in the price changes whereas the string market model explains 92.48 percent. We acknowledge, of course, that the differences in the explanatory power between the two models are economically quite small and only affect the fourth decimal place. Nevertheless, they are still impressive when one considers that the string market model is able to result in a slightly better hedge while using 30 fewer hedging instruments.

Intuitively, the reason why the common factors driving term-structure movements in the string market model have incremental power in explaining swaption price changes is not hard to understand. In the Black model, the volatility of forward swap rates is assumed to be constant. In the string market model, however, forward swap rates are essentially baskets of individual forward rates. As the composition of the basket changes over time, either as the option approaches maturity or as the swaption moves away from being at-the-money forward, the string market model captures the fact that the revised basket should have a different volatility. Thus, the string market model is better able to capture the variation in swaption-implied volatilities over time.³² Several recent papers also confirm that the hedging performance of single-factor term-structure models is inferior to that of multi-

³⁰ This assumption is not very restrictive. We have also repeated the tests using only the subsample for which the swaption is at the money at both times t and $t + 1$. The results are very similar to those reported.

³¹ The percent of variability is computed by simply taking one minus the ratio of the variance of the hedging errors divided by the variance of the actual price changes. This measure is essentially the R^2 for the hedge.

³² We acknowledge, of course, that some of the variation in implied volatilities is probably due to stochastic volatility, which neither the Black nor the string market model incorporate.

factor models; for example, see Gupta (1999) and Driessen, Klaassen, and Melenberg (2000).

VI. Conclusion

Using a string market model framework calibrated to swaptions market data, we study the relative pricing of swaptions and interest-rate caps. We find evidence that the market considers the risk of four factors in the valuation of swaptions. This contrasts with earlier work documenting that two to three factors captures the historical behavior of term-structure movements. Our results suggest that the market may consider factors that may contribute little to the unconditional variance of term-structure movements, but can periodically affect the conditional variance of term-structure movements significantly.

Focusing on the valuation of swaptions, we find that swaption prices are generally well described by the time-homogeneous string market model with the exception of the short period during the hedge-fund crisis of late 1998. We also find evidence that long-dated swaptions in particular appear to be slightly undervalued by the market. Although we stop short of claiming that there are arbitrage opportunities in this market, the results clearly suggest the need for additional research.

Finally, we examine the relative valuation of caps and swaptions using the time-homogeneous string market model. Once the covariance structure among forwards has been implied from the market prices of swaptions, cap prices are determined from the main diagonal of the implied covariance matrix. This simple restriction is tested directly by comparing the prices of caps implied by the fitted string market model to their market prices. We find that although the median differences between the two markets are close to zero, there can be large differences between the two, particularly during periods of market stress. Again, since our results are based on market quotations rather than actual transactions, we cannot definitively conclude that there are arbitrage opportunities across the caps and swaptions market. These results, however, clearly indicate the possibility that differences in the way that models are calibrated and used in the caps market and the swaptions market may introduce a wedge between the relative prices of these important fixed-income derivatives.

REFERENCES

- Andersen, Leif, and Jesper Andreasen, 2000, Factor dependence of Bermudan swaption prices: Fact or fiction? Working paper, General Re Financial Products.
- Bicksler, James L., and Andrew H. Chen, 1986, An economic analysis of interest rate swaps, *Journal of Finance* 41, 645–655.
- Black, Fischer, 1976, The pricing of commodity contracts, *Journal of Financial Economics* 3, 167–179.
- Bliss, Robert R., and Ehud I. Ronn, 1998, Callable U.S. Treasury bonds: Optimal calls, anomalies, and implied volatilities, *Journal of Business* 71, 211–252.

- Brace, Alan, Darius Gatarek, and Marek Musiela, 1997, The market model of interest rate dynamics, *Mathematical Finance* 7, 127–155.
- Chen, Ren-Raw, and Louis O. Scott, 1993, Maximum likelihood estimation for a multifactor equilibrium model of the term structure of interest rates, *Journal of Fixed Income* 3, 14–31.
- Collin-Dufresne, Pierre, and Robert S. Goldstein, 2000, “True” stochastic volatility and generalized affine models of the term structure, Working paper, Carnegie Mellon University.
- Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.
- Dai, Qiang, and Kenneth J. Singleton, 2000, Specification analysis of affine term structure models, *The Journal of Finance* 55, 1943–1978.
- De Jong, Frank, and Pedro Santa-Clara, 1999, The dynamics of the forward interest rate curve: A formulation with state variables, *Journal of Financial and Quantitative Analysis* 34, 131–157.
- Driessen, Joost, Pieter Klaassen, and Bertrand Melenberg, 2000, The performance of multifactor term structure models for pricing and hedging caps and swaptions, Working paper, Tilburg University.
- Duffee, Gregory R., 2000, Term premia and interest rate forecasts in affine models, Working paper, the University of California at Berkeley.
- Duffie, Darrell, and Kenneth J. Singleton, 1997, An econometric model of the term structure of interest rate swap yields, *Journal of Finance* 52, 1287–1321.
- Fabozzi, J. Frank, 1997, *Fixed Income Mathematics*, 3rd ed. (McGraw-Hill, New York).
- Goldstein, Robert S., 2000, The term structure of interest rates as a random field, *Review of Financial Studies* 13, 365–384.
- Gupta, Anurag, 1999, An examination of the static and dynamic accuracy of interest rate option pricing models in the cap-floor market, Working paper, Case Western Reserve University.
- Gupta, Anurag, and Marti G. Subrahmanyam, 2000, An empirical investigation of the convexity bias in the pricing of interest rate swaps, *Journal of Financial Economics* 55, 239–279.
- Heath, David, Robert Jarrow, and Andrew Morton, 1992, Bond pricing and the term structure of interest rates, *Econometrica* 60, 77–106.
- Hull, John, and Alan White, 1999, Forward rate volatilities, swap rate volatilities, and the implementation of the Libor market model, Working paper, University of Toronto.
- Jagannathan, Ravi, and Guoqiang Sun, 1999, Valuation of swaps, caps, and swaptions, Working paper, Northwestern University.
- Jamshidian, Farshid, 1997, Libor and swap market models and measures, *Finance and Stochastics* 1, 290–330.
- Karatzas, Ioannis, and Steven Shreve, 1988, *Brownian Motion and Stochastic Calculus*, (Springer Verlag Inc., New York).
- Kennedy, Douglas P., 1994, The term structure of interest rates as a Gaussian random field, *Mathematical Finance* 4, 247–258.
- Kennedy, Douglas P., 1997, Characterizing Gaussian models of the term structure of interest rates, *Mathematical Finance* 7, 107–118.
- Knez, Peter J., Robert Litterman, and Jose Scheinkman, 1994, Explorations into factors explaining money market returns, *Journal of Finance* 49, 1861–1882.
- Litterman, Robert, and Jose Scheinkman, 1991, Common factors affecting bond returns, *Journal of Fixed Income* 1, 54–61.
- Litzenberger, Robert H., 1992, Swaps: Plain and fanciful, *Journal of Finance* 42, 831–850.
- Liu, Jun, and Francis A. Longstaff, 2000, Losing money on arbitrages: Optimal dynamic portfolio choice in markets with arbitrage opportunities, Working paper, University of California at Los Angeles.
- Longstaff, Francis A., Pedro Santa-Clara, and Eduardo S. Schwartz, 2001, Throwing away a billion dollars: The cost of suboptimal exercise strategies in the swaptions market, Working paper, UCLA.
- Longstaff, Francis A. and Eduardo S. Schwartz, 1992, Interest rate volatility and the term structure: A two-factor general equilibrium model, *Journal of Finance* 47, 1259–1282.
- Longstaff, Francis A., and Eduardo S. Schwartz, 2001, Pricing American options by simulation: A simple least squares approach, *Review of Financial Studies* 14, 113–147.

- Macfarlane, John, Daniel Ross, and Janet Showers, 1991, The interest rate swap market: Yield mathematics, terminology, and conventions, in Carl R. Beidleman, ed.: *Interest Rate Swaps* (Business One Irwin, Homewood, IL).
- Merton, Robert C., 1973, The theory of rational option pricing, *Bell Journal of Economics and Management Science* 4, 141–183.
- Noble, Ben, and James Daniel, 1977, *Applied Linear Algebra*, 2nd ed. (Prentice-Hall, Inc., Englewood Cliffs, NJ).
- Pearson, Neil D., and Tong-Sheng Sun, 1994, An empirical examination of the Cox, Ingersoll, and Ross model of the term structure of interest rates using the method of maximum likelihood, *Journal of Finance* 54, 929–959.
- Peterson, Sandra, Richard C. Stapleton, and Marti G. Subrahmanyam, 2000, The valuation of American-style swaptions in a two-factor spot-futures model, Working paper, New York University.
- Rebonato, Riccardo, 1999, On the simultaneous calibration of multi-factor log-normal interest-rate models to Black volatilities and to the correlation matrix, *Journal of Computational Finance* 2, 5–27.
- Santa-Clara, Pedro, and Didier Sornette, 2001, The dynamics of the forward interest rate curve with stochastic string shocks, *Review of Financial Studies* 14, 149–185.
- Smith, Clifford W., Charles W. Smithson, and Lee Wakeman, 1988, The market for interest rate swaps, *Financial Management* 17, 34–44.
- Smith, David, 1991, A simple method for pricing interest rate swaptions, *Financial Analysts Journal*, May–June, 72–76.
- Sun, Tong-Sheng, Suresh M. Sundaresan, and Ching Wang, 1993, Interest rate swaps: An empirical investigation, *Journal of Financial Economics* 36, 77–99.
- Sundaresan, Suresh M., 1991, Valuation of swaps, in Sarkis J. Khoury, ed.: *Recent Developments in International Banking and Finance* vol. 5 (Elsevier Science Publishers, New York).
- Turnbull, Stuart M., 1987, Swaps: Zero-sum game?, *Financial Management* 16, 15–21.
- Wall, Larry D., and John Pringle, 1989, Alternative explanations of interest rate swaps: A theoretical and empirical analysis, *Financial Management* 16, 15–21.