

On Pricing and Hedging in the Swaption Market: *How Many Factors, Really?*

RONG FAN, ANURAG GUPTA, AND PETER RITCHKEN

RONG FAN

is an assistant vice president, Portfolio Credit Management at Countrywide Bank in Thousand Oaks, CA.
rong_fan@countrywide.com

ANURAG GUPTA

is an associate professor of finance at the Weatherhead School of Management at Case Western Reserve University in Cleveland, OH.
anurag@case.edu

PETER RITCHKEN

is Kenneth Walter Haber Professor and chair, Department of Banking and Finance, Weatherhead School of Management at Case Western Reserve University in Cleveland, OH.
peter.ritchken@case.edu

We examine how the number of factors and their associated volatility structures affect pricing accuracy and hedging performance in the swaption market. In spite of the fact that low dimensional one- and two-factor models do not reflect historical correlations that exist among forward rates, they are capable of accurately pricing swaptions as well as higher order multi-factor models. However, regarding hedging, there are significant benefits in using multi-factor models. This is true even if one uses the same number of hedging instruments in both single and multi-factor models. Our empirical findings have strong implications for the modeling and risk management of an array of actively traded derivatives that closely relate to swaptions.

According to the Option Clearing Corporation, the notional principal of derivatives held by U.S. commercial banks is about \$50 trillion, with interest rate derivatives contracts accounting for 80% of this total. In this large market, the primary option contracts are caps and floors on interest rates and swaptions, which are options to enter or cancel swaps. The notional amount of these over-the-counter contracts exceeds seven trillion dollars, making them among the most important interest rate claims that trade. Based on the size of this market, it is not surprising that significant effort has been placed on developing pricing models for these claims. However, the

empirical testing of these models, especially in the swaption market, has lagged behind the theoretical advances made in this area.

It is widely believed that since the term structure of interest rates is driven by multiple factors, interest rate claims should be valued using multifactor models. Standard arbitrage arguments of Heath, Jarrow and Morton [1992] have shown that for pricing purposes, models differ according to the assumptions imposed on the volatility and correlation structures of forward rates. The exact specification of the volatility structures for forward rates, and the appropriate number of factors to be used, are considered to depend on the particular application. For example, Rebonato [1999] argues that while a one-factor model, in which forward rates are instantaneously perfectly correlated, might suffice for the pricing of caps, it is very unlikely to be useful for pricing swaptions, since they depend heavily on the correlation among forward rates.¹

In this article, we investigate a broad range of one, two, three, and four-factor models with different volatility structures, that incorporate varying degrees of level and maturity dependence to pick up skewness and volatility hump effects. We closely examine how these models perform in pricing and hedging swaptions. Typical data sets in this market consist of at-the-money contracts, with an array of expiry dates and maturity dates of the underlying swaps. Characterizing the biases in prices produced by

the models along these two dimensions is important, and has not been well documented by empiricists to date.² A good model should properly explain all expiry and maturity effects. Of course, if we had data on the prices of away-from-the-money swaptions, then skew effects could also be examined. While such data are not available, we have price data on caps with multiple strikes, which we use to explore skew effects.

In the first part of the article, we investigate the impact of adding factors on pricing swaptions. To measure pricing effectiveness, we calibrate a model using known swaption and term structure data. Once the parameters are estimated, we can price claims in the future *contingent* on the future term structure. If the "out-of-sample" residuals are "small" and have no biases, then the model is viewed positively. In our analysis, the "out-of-sample" tests are conducted one, two, three, and four weeks after calibration.

We replicate some of the results of Longstaff, Santa-Clara, and Schwartz [2001], who price swaptions using models with up to eight factors. We first confirm their results that increasing the number of factors up to four provides improved precision in estimating out-of-sample swaption prices. Longstaff, Santa-Clara, and Schwartz [2001], argue that the large improvements obtained by adding factors is due to the fact that low dimensional models are unable to produce realistic correlations among forward rates and getting these correlation effects right is crucial for pricing swaptions. We show that similar results are obtained for caps, which may be less sensitive to correlation effects. We then test whether the improvement in these types of models is due to increasing the number of factors, or increasing the number of free parameters. We do this by considering several one- and two-factor models that have the same number of free parameters as our four-factor model, and compare their pricing performance. In this case, we find that for pricing swaptions, the benefits of increasing the number of factors beyond one are insignificant. Indeed, we conclude that from a pricing perspective only, there appears to be little advantage in moving from a one- to a four-factor model. Our pricing results hold true even when swaption prices are generated up to four weeks after the parameters were estimated. For pricing purposes, the importance of models that better reflect the correlation structures among forward rates appears to be minor.

We also address the importance of level dependence in the volatility structures. Since volatility skews are hard

to assess in the swaption market, because the prices of contracts that are available are restricted to at-the-money contracts, we turn to the cap market to get evidence for the pricing performance over the strike price domain.³ Our results show that incorporating level dependence in the volatility structure is extremely important for away-from-the-money caps, and that proportional dependent structures are better than both square root or level independent structures. For at-the-money swaptions, the effect of level dependence is minor. However, the evidence from cap prices suggests that away-from-the-money swaptions would be better priced using proportional models rather than level independent structures.

The removal of systematic expiry date and underlying swap maturity biases in out-of-sample pricing is a necessary condition for a model to be useful for hedging, but it is not sufficient. The second goal of our research is to investigate the ability of alternative models in *hedging* swaptions. In particular, we carefully quantify the benefits, if any, of using higher order factor models over lower order models. We investigate the effectiveness of alternative hedging strategies using different models and differing numbers of hedging instruments, and produce convincing evidence that multifactor models are essential for reducing the risk in the hedged positions for swaptions. We also demonstrate that allowing additional hedging instruments in a one- and two-factor model does not improve their hedging performance. Our main conclusion is that while accurate swaption prices can be obtained from a one-factor model, one- and even two-factor models cannot hedge swaptions well, and the benefits of multifactor models are significant.

The article proceeds as follows. We first review the empirical literature. We then discuss the set of 18 different models that we evaluate. In the following two sections, we discuss our data and model implementation. In the next section, we examine pricing accuracy by closely examining a nested set of principal component based models. In the section after that, we compare the performance of our best principal component based model with alternative lower order models. In the following section, we examine the hedging effectiveness of our models. We end with our conclusions.

LITERATURE REVIEW

There is much controversy regarding the selection of models for pricing and hedging caps, swaptions, and other interest rate derivatives. Models differ according to

the number of stochastic drivers, forward rate volatility structures, and the use of time varying parameters. Empirical tests of models have, for the most part, focused on models with time stationary volatility structures. Amin and Morton [1994] present one of the early tests of alternative forward rate volatility structures. They find that the single factor generalized Vasicek model provides the best out-of-sample pricing performance. Ritchken and Chuang [1999] show that a generalization of this model, that captures the hump in the volatility of forward rates, leads to significant improvements. Gupta and Subrahmanyam [2005] find that a one-factor lognormal forward rate model outperforms other competing one-factor models in pricing caps, with two-factor models improving pricing performance only marginally. However, for hedging, they find a significant advantage in moving from one-to two-factor models. Bühler, Uhrig, Walter, and Weber [1999] reject the deterministic volatility structure for forward rates in favor of a model where volatility is proportional to the level of rates. Unlike Gupta and Subrahmanyam [2005], they find no advantage in moving beyond a one-factor model. More recently, Jarrow, Li, and Zhao [2007] show that even a three-factor model with stochastic volatility and jumps cannot completely capture the smile/skew patterns observed in this market.

In contrast to the cap/floor market, few empirical studies have been conducted on swaptions. Longstaff, Santa-Clara, and Schwartz [2001] use a string model framework to test the relative valuation of caps and swaptions using at-the-money cap and swaptions data, and find evidence for using at least a four-factor model for swaptions. Their criterion for evaluating models is based on the sum of squared percentage pricing errors. In other words, their criterion is based on pricing accuracy, not on hedging precision.⁴ Peterson, Stapleton, and Subrahmanyam [2003] develop an extension of the lognormal model of Black and Karasinski [1991] to multiple factors and provide evidence that the addition of a third factor is helpful in pricing swaptions. In this regard, these studies provide support for Rebonato's claim of the importance of appropriately modeling the correlation structure for swaptions. In a recent study, Duarte [2007] shows that including information about mortgage-backed securities hedging activity improves the performance of models for pricing swaptions.

Not all studies, however, indicate that multiple factors are necessary for improving pricing performance for swaptions. For example, Driessen, Klaassen, and

Melenberg [2003] investigate the performance of several Gaussian models, where volatility structures are deterministic functions of their maturities. They find that one of their one-factor models prices swaptions no worse than their multi-factor models, and to the same degree of accuracy as Longstaff, Santa-Clara, and Schwartz's [2001] multi-factor model. Jagannathan, Kaplin, and Sun [2003] investigate the pricing of swaptions using multifactor Cox, Ingersoll, and Ross models, and conclude that increasing the number of factors does not necessarily improve pricing performance.

Very few studies have compared the abilities of different models for hedging swaptions. Longstaff, Santa-Clara, and Schwartz [2001] briefly consider hedging, in the context of their four-factor model, relative to the Black model, but they do not evaluate the benefits of hedging using an increasing number of factors.⁵ The most related studies on swaption hedging are by Driessen, Klaassen, and Melenberg [2003] and Fan, Gupta, and Ritchken [2003]. Driessen, Klaassen, and Melenberg use their Gaussian models to demonstrate that if the number of hedge instruments is equal to the number of factors, multi-factor models outperform one-factor models in hedging caps and swaptions. However, they claim that by using a large set of hedge instruments, their one-factor models perform as well as multi-factor models. This last finding is the opposite of what Gupta and Subrahmanyam [2005] find in the cap market. In the context of unspanned stochastic volatility, Fan, Gupta, and Ritchken [2003] show that even swaption straddles can be well hedged using LIBOR bonds alone, if at least a three-factor model is used. However, in their research, the focus is on examining whether bonds span the swaption market, not on investigating the role of the number of factors in a model for pricing and hedging swaptions.⁶

In summary, some studies find that four factors are necessary for pricing swaptions, while others find that one and two-factor models are satisfactory. The few studies on hedging swaptions have also produced mixed results.

THE BASIC MODELS

Caps and swaptions are actively traded, and, according to market convention, their prices are quoted in volatility form using the standard Black [1976] model, with instruments at different maturities and strikes trading at different implied volatilities. Since the volatilities used in the Black model are for forward rates, for the case of

caps, and swap rates, for the case of swaptions, direct comparisons of cap and swaption volatilities are not meaningful. Indeed, the Black formula should be viewed only as a nonlinear transformation from prices into volatilities and vice versa. This market convention provides a convenient way of communicating prices because volatilities tend to be more stable over time than actual dollar prices. The market convention does not imply that participants in this market view the Black model as being appropriate.

Let $f(t, T)$ denote the forward interest rate at time t for instantaneous riskless borrowing or lending at date T . The dynamics of forward rates are given by

$$df(t, T) = \mu_f(t, T)dt + \sum_{n=1}^N \sigma_{f_n}(t, T)dw_n(t), \quad (1)$$

with $f(0, s)$ given for $s \geq 0$

where $\{dw_n(t) | n = 1, 2, \dots, N\}$ are standard independent Wiener increments. The volatility structures, $\{\sigma_{f_n}(t, T) | n = 1, 2, \dots, N\}$, could, in general, be functions of all path information up to date t .

Heath, Jarrow, and Morton [1992] show that to avoid riskless arbitrage, the drift term, under the equivalent martingale measure, is completely determined by the volatility functions in the above equation. Specifically:

$$\mu_f(t, T) = \sum_{n=1}^N \sigma_{f_n}(t, T) \int_t^T \sigma_{f_n}(t, u) du$$

This implies that for pricing purposes, only the volatility structures need to be specified and estimated.

In performing empirical tests on pricing and hedging effectiveness, there are some important features that we have to consider. First, all the models that we study have time stationary volatility structures. Our volatility structures are chosen to accommodate a large number of different skew patterns in implied volatilities. This is accomplished by introducing varying degrees of level dependence into the volatility structures. Both level dependence and time-stationarity assumptions come at some cost. Level dependence results in increasing the number of state variables that are necessary to characterize the dynamics of the term structure, thereby increasing the computational complexity. Without time dependence in the volatility structures, it is not possible to construct

models that can match an array of swaption prices exactly. All the models that we examine have, at most, four free parameters in the volatility structures, so at best we could only match four swaptions at any point in time. In the models considered by Longstaff, Santa-Clara, and Schwartz [2001], the number of free parameters equals the number of stochastic drivers. In contrast, some of our lower order one- and two-factor models contain as many free parameters as our four-factor models. Let

$$\sigma_{f_1}(t, T) = \left[(a + b(T-t)e^{-\kappa(T-t)} + c) \right] f(t, T)^\gamma \quad (2)$$

This parametric volatility structure nests many well known models. First, when $\gamma = 0$, the volatility structure is a deterministic function of maturity. With $c = b = 0$, the model reduces to the generalized Vasicek model, commonly referred to as the single factor Hull and White [1993] model, in which forward rate volatilities dampen with their maturity. With b and c released from 0, the model can accommodate hump shapes in forward rate volatilities. Such models have been considered by Moraleda and Vorst [1997] and Ritchken and Chuang [1999]. They have the attractive feature that analytical solutions can be set up for pricing many derivative contracts. However, they have also been criticized since interest rate volatilities do not depend on their levels and can therefore become negative. When γ is positive, the volatility structure does depend on the level of rates. In this case, however, computational problems emerge. In particular, the term structure is no longer Markovian in a finite number of state variables.⁷ Fortunately, Monte Carlo simulation provides a powerful tool for pricing European claims, and, recently, Longstaff and Schwartz [2001], among others, have shown that American claims can also be priced using rather simple and efficient methods.⁸

We consider three parametric one-factor models, where the volatility structure is of the form in Equation (2). The models differ according to the level dependence parameter, γ . In particular, we consider models where γ is zero, one half, and one. All these models have four parameters.

In our two-factor models, the first volatility structure is of the form in Equation (2), with $c = 0$. The second volatility structure is of the form:

$$\sigma_{f_2}(t, T) = d \cdot [r(t)]^\gamma \quad (3)$$

For this two-factor model, the shocks to forward rates consist of two types. The first type of shock depends on

the level of the forward rate and on the maturity. Over the short end, the structure permits an increasing volatility, but eventually the shock dampens with maturity. The second shock has a "parallel" effect over maturity. The absolute magnitude of this effect is driven by the level of the short rate. A structure similar to this model has been empirically examined by Inui and Kijima [1998]. For $\gamma = 0$, this model nests the two-factor generalized Vasicek model of Hull and White [1993]. For $\gamma = 0.5$ and 1, the model is similar to a generalized Cox, Ingersoll, and Ross [1985] model and a proportional model, respectively. Note that all three of our two-factor models also have four free parameters.

The final set of models is based on modifying the loadings provided by the principal components of the historical correlation matrix of forward rates along the lines of the string models of Longstaff, Santa-Clara, and Schwartz [2001]. In these models, we consider a discrete set of M maturities, say $\{\tau_1, \tau_2, \dots, \tau_M\}$ with $\tau_1 < \tau_2, \dots, \tau_M$. Then:

$$\sigma_{f_j}(t, t + \tau_j) = g(\tau_j) f(t, t + \tau_j)^\gamma \quad (4)$$

where $g(\cdot)$ is a deterministic function of the maturity of the forward rate, that is estimated primarily using principal component analysis on historical data. In particular, take the case where $\gamma = 0$. Using historical data on weekly forward rates, a correlation matrix of forward rate changes, separated by three months for maturities less than a year, and six months thereafter (up to 10 years maturity), is obtained. In particular, 22 forward rate maturities are used and a 22×22 correlation matrix is established. The matrix of eigenvectors (principal components) is computed, and the first four eigenvectors are retained. Let

$$T^* = \{\tau_1 = 0.25, \tau_2 = 0.5, \tau_3 = 0.75, \tau_4 = 1, \tau_5 = 1.5, \\ \tau_6 = 2, \tau_7 = 2.5, \dots, \tau_{21} = 9.5, \tau_{22} = 10\}$$

represent the set of 22 forward rate maturities, and let h_i be a 22×1 vector representing the i th eigenvector for $i = 1, 2, 3$, and 4. Then, define:

$$g_i(\tau_j) = \lambda_i h_{ij} \text{ where } i=1,2,3,4 \text{ and } j=1,2,\dots,22$$

where h_{ij} is the j th element of the i th eigenvector, and the λ_i values are the free parameters, the i th one representing

the scaling factor for all the elements of the i th principal component and is implied out at any date t using date t swaption data.

The principle behind such a procedure is simple. As shown by several researchers, including Litterman and Scheinkman [1991], the first four historical principal components identify the four most important types of orthogonal shocks to the forward rate curve. Since the exact contribution of each of these shocks may vary over time, the eigenvalues for the future period may be different from the eigenvalues over the historical period. Since, in an efficient market, the swaption data reflects all available information on the set of forward looking correlations among forward rates, this data should be used to establish the eigenvalues.

When γ differs from 0, the same analysis is done, except that the correlation structure for the principal component analysis is estimated over transformed values of forward rates. For example, when $\gamma = 1$, the correlation is estimated over percentage changes in forward rates.

The above method, which we term an adapted Principal Component Analysis (PCA) method, has been used by Longstaff, Santa-Clara, and Schwartz [2001] for proportional models ($\gamma = 1$) and by Driessen, Klaassen, and Melenberg [2003] for absolute models ($\gamma = 0$). In addition to these, we also permit $\gamma = 0.5$. Notice that, like all our other models, our four-factor model has four free parameters that can be implied out using option data. Notice too, that if a three-factor model is used, then only the first three principal components are retained and the number of free parameters drops by one.

In summary, we consider 12 PCA models (three one-factor, three two-factor, three three-factor and three four-factor models) and six parametric models (three one-factor and three two-factor models). All the parametric models and the four-factor PCA models have four free parameters. The other models have as many free parameters as stochastic drivers.

DATA

The data for this study consist of U.S. dollar swaption and cap prices. The swaptions data set comprises volatilities of swaptions of six-month, one-, two-, three-, four-, and five-year maturities, with the underlying swap maturities of one, two, three, four, and five years each (in all, there are 30 swaption contracts). As per market convention, a swaption is considered at-the-money when the

strike rate equals the forward swap rate for an equal maturity swap. The cap prices are for a 10-month period (March 1–December 31, 1998), across four different strikes (6.5%, 7%, 7.5%, and 8%) and four maturities (two-, three-, four-, and five-year), obtained from Bloomberg Financial Markets. For swaptions, the data consists of at-the-money volatilities for a 32-month period (March 1, 1998–October 31, 2000), obtained from DataStream.

For constructing the yield curve, we use futures and swap data. For the short end of the curve (up to one-year maturity), we use the five nearest futures contracts on any given date. These futures rates are interpolated, and then convexity corrected, using the Vasicek [1977] model to obtain the forward rates for three-, six-, nine-, and 12-month maturities. The rest of the yield curve out to five years is estimated using the forward rates bootstrapped at six month intervals from market swap rates. The futures and swap data are obtained from DataStream. Eventually, we obtain weekly forward rate curves that start one year before our cap and swaption data begin, and extend to the end of our swaption data period.

For the principal component analysis, we follow the procedure used in Longstaff, Santa-Clara, and Schwartz [2001]. Specifically, we use the one-year history of forward rates prior to the beginning of our swaption data, to estimate the correlation structure of forward rates. For example, for $\gamma = 1$, we estimate the percentage changes in forward rates from the historical time series of forward rates. We then decompose the correlation matrix, R , into $U \Lambda U'$, where U is the matrix of eigenvectors and Λ^* is a diagonal matrix of eigenvalues. Finally, we retain the first four eigenvectors and assume a covariance structure for forward rates, Σ , given by $\Sigma = U \Lambda U'$, where Λ is a diagonal matrix with the first four diagonal elements positive, the others zero. A similar analysis is done for the models with $\gamma = 0$ and $\gamma = 0.5$.

MODEL IMPLEMENTATION

We consider a discrete implementation of the multifactor Heath-Jarrow-Morton model. Towards this goal, we divide the time interval into trading intervals of length Δt , and label the periods with consecutive integers. Let $f^{\Delta t}(t, j)$ be the forward rate at period t , for the time interval $[j\Delta t, (j+1)\Delta t]$. Let $\Delta f^{\Delta t}(t, j)$ represent the change in the forward rate over a time increment Δt . That is

$$\Delta f^{\Delta t}(t, j) = f^{\Delta t}(t+1, j) - f^{\Delta t}(t, j) \text{ for } j \geq t$$

The actual magnitude of this change could depend on the forward rate itself, its maturity date, and other factors.

We start with an initial forward rate curve, $\{f^{\Delta t}(0, j), j = 0, 1, \dots, m\}$, that is chosen to match the observed term structure at date 0 for all maturities up to the date $m\Delta t$. Notice that $f^{\Delta t}(0, 0)$ is just the spot rate for the immediate period, $[0, \Delta t]$. Over each time increment, the forward rates change as follows:

$$\Delta f^{\Delta t}(t, j) = \mu_f^{\Delta t}(t, j)\Delta t + \sum_{n=1}^N \sigma_{f_n}^{\Delta t}(t, j)\sqrt{\Delta t} Z_{t+1}^{(n)} \quad (5)$$

where $Z_{t+1}^{(n)}$ is a standard normal random variable, j is an integer larger than the current date, t , $\mu_f^{\Delta t}(t, j)$ is the drift term, and $\sigma_{f_n}^{\Delta t}(t, j)$ is the volatility term associated with the n th factor, $n = 1, 2, \dots, N$, where the N standard normal random variables are independent. The discrete time equivalent of the Heath-Jarrow-Morton restriction is given by

$$\mu_f^{\Delta t}(t, j) = \sum_{n=1}^N \mu_{f_n}^{\Delta t}(t, j)$$

where

$$\mu_{f_n}^{\Delta t}(t, j) = \left[\sigma_{f_n}^{\Delta t}(t, j) \right]^2 \frac{\Delta t}{2} + \sigma_{f_n}^{\Delta t}(t, j) \sigma_{p_n}(t, j)$$

and

$$\sigma_{p_n}(t, j) = \sum_{i=t+1}^{j-1} \sigma_{f_n}^{\Delta t}(t, i)\Delta t$$

for $n = 1, 2, \dots, N$.

Prices of European interest rate claims can be computed using Monte Carlo simulation. Specifically, we simulate $K = 2000$ different paths, each path initiated at date 0 where the initial term structure is given. Consider the k th simulation. Given the date 0 term structure, forward rates are updated recursively using Equation (5). This gives the k th path of the term structure of forward rates. At date 0, \$1 is placed in a fund that rolls over at the short rate. At period t , the value of the money fund, $M(t; k)$, is given by:

$$M(t; k) = \prod_{i=0}^{t-1} e^{f^{\Delta t}(i, i) \Delta t}$$

Consider a claim that pays out in period T_E . Using simulation, a set of forward rates at this date can be computed, and hence all bond prices and swap rates can be recovered. In addition, the accumulated money fund, $M(T_E; k)$, along this path is known. The terminal value of this claim for this path can then be computed. Let $C(T_E; k)$ be this value. The date 0 value of the claim, for this path, is approximated by:

$$C(0; k) = \frac{C(T_E; k)}{M(T_E; k)}$$

The value of the claim at date 0 is then given by the average of all these values obtained over the K paths. Specifically:

$$C(0) = \frac{\sum_{k=1}^K C(0; k)}{K}$$

Since repeated calls are used to estimate the parameters of the process, it is important that the pricing algorithms be as efficient as possible. Hence, we use $\Delta t = 0.125$ years. Rather than price all the contracts separately, we simulate the money fund and forward rates along paths for a 10-year period, and at each relevant maturity date along the path, all the appropriate caplet and swaption prices are computed. We repeat this procedure $K = 2000$ times, and use antithetic variate variance reduction techniques, to establish the fair prices of all our contracts. We ran extensive robustness checks to ensure that the benefits of increasing the sample size and decreasing the time partition were negligible. Further, to the extent possible, we use the same stream of random numbers to price the same contracts with different volatility structures. This ensures that the difference in prices of the contracts is more tightly attributable to the different volatility structures, rather than to sampling error.

PRICING PERFORMANCE OF PCA MODELS

Like Amin and Morton [1994], Driessen, Klaassen, and Melenberg [2003], Longstaff, Santa-Clara, and

Schwartz [2001], and Moraleda and Pelsser [2000], we estimate model parameters from cross sectional options data. At any date, we fit models to the prices of swaptions for different expiry dates and underlying swap maturities. Our objective function is to minimize the sum of squared percentage errors between theoretical and actual prices, using a non-linear least squares procedure. An alternative objective would be to minimize the sum of squared errors in prices. However, since prices of swaptions can range from a couple of basis points to a thousand basis points, which is almost four orders of magnitude apart, such a minimization would place more weight on the expensive contracts.

The first set of experiments examines the pricing issues for at-the-money swaptions. Using the swaption data, for each odd week, using mid-week data, we establish the best fit for the prices of all swaptions. Given these parameter estimates, we forecast one-, two-, three- and four-week out-of-sample residuals, which are all labeled according to their in-sample time period, model, and contract.

A similar set of experiments is conducted on cap prices. For each of our models, we establish the best fit for the prices of caps for the four strikes and four maturities. We do this using mid-week data, in separate optimizations, for every odd week. We then use the parameter estimates, together with the term structure the following week, to generate one week "out-of-sample" residuals. In addition, we also compute two-, three- and four-week out-of-sample residuals. These residuals are stored for each model, for each contract, and for each date.

Exhibit 1 presents the average absolute percentage errors for one-week out-of-sample swaption and cap prices produced by the PCA models.

The results show that for any given value of γ , increasing the number of factors improves the pricing performance for swaptions. This confirms the results obtained by Longstaff, Santa-Clara, and Schwartz [2001], who conclude that four factors are necessary for pricing swaptions. Exhibit 1 also shows that, for a fixed number of factors, increasing γ has no significant effect on reducing the average absolute percentage errors for swaptions. Since the level dependence parameter, γ , controls the skewness of the risk neutral distribution, and since all the swaptions are at-the-money contracts, these results are not surprising.⁹

Exhibit 1 also presents the results for caps. For any given number of factors, the average absolute percentage errors for models with $\gamma = 1$ are consistently lower than

EXHIBIT 1

Absolute Average Pricing Errors—PCA Models

This exhibit presents the average absolute percentage errors, one week out-of-sample, for swaptions and caps, for the PCA models tested in this article. The swaption data are biweekly from March 1, 1998–October 31, 2000, consisting of 70 data sets. Hence, each error reported for swaptions is an average across 30 contracts over 70 dates (hence an average of 2100 individual errors). The cap data are biweekly from March 1–December 31, 1998, consisting of 22 data sets. So each error reported for caps is an average across 16 contracts over 22 dates (hence an average of 352 individual errors). The standard error of the mean is reported in parenthesis. The options are priced using Monte Carlo simulation with 4000 paths for the evolution of the term structure. In generating the paths, the same seeds for the random number generator were used to ensure consistency across the models.

	Swaptions			Caps		
	$\gamma = 0$	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 0$	$\gamma = 0.5$	$\gamma = 1$
One-factor	4.22 (0.08)	4.22 (0.08)	4.52 (0.09)	18.9 (0.79)	16.4 (0.75)	15.8 (0.76)
Two-factor	3.81 (0.07)	3.57 (0.07)	3.64 (0.08)	18.1 (0.75)	16.0 (0.72)	14.9 (0.80)
Three-factor	3.39 (0.07)	3.15 (0.07)	3.18 (0.08)	18.1 (0.75)	16.1 (0.75)	14.6 (0.79)
Four-factor	2.99 (0.06)	3.03 (0.07)	3.01 (0.07)	18.1 (0.75)	16.5 (0.78)	14.7 (0.79)

those for models with $\gamma = 0$, with the errors for models with $\gamma = 0.5$ being in between the two. In contrast to swaptions, the advantage of increasing the number of factors appears to be small. Also, in contrast to swaptions, the cap errors are much higher. However, these errors are one-week out-of-sample errors across a range of strikes and are comparable to those reported for at-the-money contracts in other studies, for example, Driessen, Klaassen, and Melenberg [2003], and Longstaff, Santa-Clara, and Schwartz [2001].

In order to further investigate the effects of level dependence, Exhibit 2 presents the proportion of times that a PCA model with a particular γ value produces one week out-of-sample cap prices closer to actual market prices, than the same model with a different γ value. These proportion tests are based on a total of 352 residuals each, since there are 16 cap contracts on each of the 22 dates. For example, for the one-factor PCA model, 66% of the residuals from a $\gamma = 0.5$ model are closer to zero than the residuals produced by the $\gamma = 0$ model. For each of the PCA models, the precision increases with an increase in the γ value.

These results provide conclusive evidence for the use of proportional models over square root and absolute models for explaining the skew effects in cap prices, since they price away-from-the-money contracts better. This

suggests that away-from-the-money swaptions would also be better priced by proportional models.

In order to further examine the effects of increasing the number of factors, Exhibit 3 reports the proportion of times a higher order PCA model outperforms a lower order PCA model in pricing caps. Specifically, for each of the four PCA models and the three γ values, we estimate the one-week out-of-sample prices, and compute the residuals. Again, since there are 16 contracts at each date and 22 dates, we have 352 residuals for each model-gamma combination.

For each pair of models, Exhibit 3 reports the proportion of times (out of 352) that one model produces smaller residuals than the other. For example, for $\gamma = 1$, the two-factor PCA model produces prices closer to the actual market prices 57% of the time, when compared to the one-factor PCA model. The results show that, at the 5% level of significance, for all levels of γ , the two-factor, three-factor and four-factor PCA models outperform the one-factor PCA model. The three-factor model outperforms the two-factor model. However, the four-factor model does not produce results significantly different from the three-factor model.

Exhibit 4 presents similar tests for swaptions. In each of the 70 time periods, there are 30 contracts, for a total of 2100 residuals. The table presents the percentage of

EXHIBIT 2

Level Dependence Comparisons for Caps using PCA Models

This exhibit compares the one week out-of-sample predictions of each PCA model for each of the 16 contracts over all 22 dates, for different γ values. For example, for the one-factor model, based on all 352 pairwise comparisons of residuals, the model with $\gamma = 0.5$ beat the model with $\gamma = 0$ 66% of the time. A starred cell indicates that the proportion is significantly different from 50% at the 5% level of significance.

No. of Factors	$\gamma = 0.5$ vs $\gamma = 0$	$\gamma = 1$ vs $\gamma = 0$	$\gamma = 1$ vs $\gamma = 0.5$
1	0.66*	0.59*	0.53*
2	0.72*	0.69*	0.73*
3	0.78*	0.76*	0.70*
4	0.78*	0.79*	0.75*

times the higher order model produces smaller absolute residuals than the lower order model. The results are presented for the in-sample residuals, as well as for one-, two-, and four-week out-of-sample residuals, only for the proportional model.

These results show that within a PCA framework, the two-factor model outperforms the one-factor model, the three-factor model outperforms the two-factor model, and the four-factor model outperforms the three-factor model. These results are significant at the 1% level of significance, and are consistent with those reported in other studies. When the results are broken down and analyzed by expirations (or underlying swap maturities), the same trend is observed. Overall, the marginal improvement from adding a factor decreases as the number of factors increases, and the contribution of the fourth factor, while statistically significant, is small. The relative performance of PCA models for swaptions remains the same, even as the out-of-sample period is increased to four weeks. Although not reported in the article, we also observe that the inference drawn from caps remains unchanged as the out-of-sample period increases from one to four weeks.

In summary, we show that for at-the-money swaptions, level dependence is not important, but increasing the number of factors (within PCA models) to four improves pricing performance. Since level dependence is important for away-from-the-money caps, we postulate that it is important for away-from-the-money swaptions as well. Finally, for caps, increasing the number of factors to three or four significantly improves pricing performance.

Longstaff, Santa-Clara, and Schwartz [2001] attribute the better performance of the multi-factor PCA models for swaptions to the fact that they are capable of producing more realistic correlation structures for the forward rates. On the other hand, since caps are less sensitive to correlation effects, the improvement in their pricing performance, as the number of factors increases, might be surprising, and raises the possibility that there are other explanations. For example, the improvements in performance might be due to the fact that higher order models have more free parameters than lower order models, and the improvements arise because of these extra degrees of freedom.

PRICING PERFORMANCE OF PARAMETRIC AND PCA MODELS

In this section, we establish whether the improvement in pricing swaptions is due to increasing the number of free parameters in the model, or due to increasing the number of stochastic drivers. In order to do this, we focus on the two parametric models with one and two stochastic drivers, and on the four-factor PCA model (which is the best model among the PCA models), all of which have four free parameters. The four-factor PCA model clearly comes closest to matching the historical correlations among

EXHIBIT 3

Comparison of Out-of-Sample Performance of PCA Models for Caps

This exhibit compares the one week out-of-sample predictions for the PCA models for each of the 16 contracts over all 22 dates. For example, the first row of the table shows the proportion of the times that the two-factor model produces a residual smaller than the one-factor model. All four PCA models are compared pairwise. A starred cell indicates that the proportion is statistically significantly different from 50% at the 5% level of significance. The data are biweekly from March 1–December 31, 1998.

Models	$\gamma = 0$	$\gamma = 0.5$	$\gamma = 1$
2 vs 1	0.56*	0.57*	0.57*
3 vs 1	0.57*	0.56*	0.62*
4 vs 1	0.54*	0.57*	0.59*
3 vs 2	0.63*	0.53*	0.58*
4 vs 2	0.55*	0.50	0.54*
4 vs 3	0.51	0.46	0.51

EXHIBIT 4

Comparison of Out-of-Sample Performance of PCA Models for Swaptions

This exhibit compares the one-, two-, and four-week out-of-sample predictions of each PCA model (for $\gamma = 1$), for each of the 30 contracts over all 70 dates. The results are aggregated over all expirations and underlying maturities for the swaptions. Results for in-sample tests are also reported for comparison. Each entry in the exhibit is based on the number of times a particular model, represented by the first number in the row, produces a residual closer to zero than the model it is competing with. Each entry is based on 2100 residuals. All four PCA models are compared pairwise. A starred cell indicates that the proportion is statistically significantly different from 50% at the 1% level of significance. The data are biweekly from March 1, 1998–October 31, 2000.

Models	Weeks Out-of-Sample			
	0	1	2	4
2 vs 1	0.66*	0.66*	0.64*	0.63*
3 vs 1	0.75*	0.72*	0.69*	0.67*
4 vs 1	0.71*	0.69*	0.66*	0.65*
3 vs 2	0.68*	0.66*	0.65*	0.62*
4 vs 2	0.63*	0.63*	0.59*	0.58*
4 vs 3	0.53*	0.52*	0.53*	0.53*

forward rates. If this, indeed, is an important feature of the swaption market, then this model should outperform one- and two-factor models with the same degrees of freedom. On the other hand, if correlations are less important, then the one- and two-factor models should be comparable to the four-factor model. In addition, the previous results for PCA models were based on aggregate statistics. In this section, we examine the potential biases in the models when the contracts are broken down by expiry and underlying swap maturity dates.

Exhibit 5 compares the average absolute swaption errors in the out-of-sample prices by each model.

The one-week out-of-sample results are similar to the results obtained for the PCA models, where the importance of γ was found to be minor. In light of the importance of γ in pricing away-from-the-money caps, we therefore focus on models with $\gamma = 1$. Exhibit 5 shows the two-, three- and four-week out-of-sample performance of the models with $\gamma = 1$.

What immediately stands out is the performance of the one-factor parametric model. In particular, the one-week out-of-sample model prices swaptions with average absolute errors less than three percent, which is typically within the bid-ask spreads in the swaptions market. Indeed,

for the one-factor parametric model, of the 2100 (70×30) one-week out-of-sample residuals, 84% (56%) were within one (one-half) Black vol, and almost one-third were within one-quarter of a Black vol.

Exhibit 5 also shows that, in aggregate, the two-factor parametric model outperforms the one-factor parametric and the four-factor PCA models. Therefore, of all the models, there is a slight preference for the two-factor model. The important point here, however, is that the four-factor model, at the aggregate level, does not dominate the lower order models. However, since the above conclusions are based on aggregate results, they may mask any expiration date and swap maturity biases.

Effects of Level Dependence on Pricing

To more closely examine the effects of level dependence on the pricing of swaptions, we compute the average percentage errors, one week out-of-sample, for each swaption expiration and underlying swap maturity, for each model. Exhibit 6 presents the average percentage errors for each swaption expiration, plotted against the underlying swap maturity, for all three models and for all three γ values.

In general, the one-factor model underprices swaptions with lower underlying swap maturities, and overprices swaptions with higher underlying swap maturities. However, the figures clearly show that the inclusion of level dependence does not eliminate this bias. Indeed, the pricing patterns are strikingly similar for all three values of γ , across all swaption expirations and underlying swap maturities. There appears to be no benefit of incorporating level dependence in the one-factor model, which confirms the results that were obtained at the aggregate level in Exhibit 5.

The last two columns of Exhibit 6 show that the two-factor parametric model and the four-factor PCA model produce similar results. Interestingly, in the PCA models, it is not just the lower underlying swap maturity swaptions that are underpriced. For some swaption expirations, even the longer underlying swap maturity swaptions are underpriced, with the medium swap maturity swaptions being overpriced. Exhibit 6 indicates that incorporating level dependence has little effect on pricing at-the-money swaptions.

As discussed earlier for the PCA models, we turn to our cap data to assess the effects of level dependence in forward rate volatility structures. Exhibit 7A presents the

EXHIBIT 5

Average Absolute Pricing Errors for Swaptions—Parametric and Four-Factor PCA Models

This exhibit presents the average absolute percentage errors, one-, two-, three-, and four-week out-of-sample, for swaptions across expirations and maturities (of underlying swaps), for the one-factor and two-factor parametric models and the four-factor PCA model. The swaption data are biweekly from March 1, 1998–October 31, 2000, consisting of 70 data sets. Each error reported in this table is an average across 30 contracts over 70 dates (hence an average of 2100 individual errors). The standard error of the mean is reported in parentheses. The standard error three- and four-weeks out-of-sample may be understated, since they are not corrected for potential serial correlation. The swaptions are priced using Monte Carlo simulation, with 4000 paths for the evolution of the term structure. In generating the paths, the same seeds for the random number generator were used to ensure consistency across the models.

	Out-of-Sample Period					
	1 week			2 weeks	3 weeks	4 weeks
	$\gamma = 0$	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 1$	$\gamma = 1$	$\gamma = 1$
One-factor Parametric	2.43 (0.05)	2.57 (0.05)	2.96 (0.06)	3.23 (0.06)	3.48 (0.07)	3.70 (0.07)
Two-factor Parametric	2.57 (0.05)	2.36 (0.05)	2.48 (0.05)	2.74 (0.06)	3.03 (0.06)	3.16 (0.06)
Four-factor PCA	2.99 (0.06)	3.03 (0.07)	3.01 (0.07)	3.34 (0.07)	3.56 (0.08)	3.80 (0.08)

average absolute pricing errors for caps, one week out-of-sample.

Exhibit 7A shows that models with $\gamma = 1$ are significantly better for pricing away-from-the-money caps. In Exhibit 7B, we provide pairwise comparisons across models with different values. For all cases, the models with $\gamma = 1$ are the best, and the models with $\gamma = 0$ are the worst. Exhibit 8 presents the breakdown of the biases by moneyness for differing maturity dates and models.

Exhibit 8 clearly identifies skew effects, especially for the models with $\gamma = 0$ and $\gamma = 0.5$. The models with $\gamma = 1$ explain more of the skew effects than the other models. In light of the evidence on skews in the cap market, and the indifference of level dependence for at-the-money swaptions, we only consider models with $\gamma = 1$.

Effects of the Number of Factors on Pricing

So far, we have presented results for swaptions at the aggregate level. We now examine individual contract pricing errors in more detail, across the models, to determine if there is any significant benefit to moving beyond a one-factor model for pricing at-the-money swaptions. Exhibit 9 presents the box and whiskers plots of one-week out-of-sample errors for the one- and two-factor models, and the four-factor PCA model. The plots are presented separately for each swaption expiration, with the

errors for the three models plotted across underlying swap maturities.

For the six-month expiry contracts, the one-factor model produces fairly unbiased prices with small variances for all five underlying swap maturities. Indeed, with few exceptions, over all 30 expiry-swap maturity combinations, the bias and the interquartile range for the one-factor model appear to be no worse than those for the four-factor model. Further, for any expiry date, the trends of the biases across swap maturities appear to be random. The plots suggest that the largest differences between the models arise for short term expiration dates.

The percentage errors for almost all of the 30 expiration-swap maturity buckets are within three percent. In fact, over 80% of all the contracts have one-week out-of-sample residuals within one Black vol, with almost 60% of the contracts being within half a Black vol. This represents a reasonable bid-ask spread during the time period analyzed in this study. The precision of the three models is similar.

Exhibit 10 provides pairwise comparisons among the three models for each expiry-underlying swap maturity combination. While there are some expiry date-swap maturity combinations where the four-factor model performs the best, overall, there appears to be a preference for the lower order parametric models. Indeed, the differences in the performance of the models is small, with

EXHIBIT 6

Average Percentage Errors for Swaptions

This exhibit presents plots of the average percentage errors, one week out-of-sample, for swaptions across underlying swap maturities, for each expiration, for the one-factor and two-factor parametric and the four-factor PCA models. The plots are presented for three values of γ : $\gamma = 0$ (dotted line), $\gamma = 0.5$ (hashed line), and $\gamma = 1$ (solid line). The swaption data are biweekly from March 1, 1998–October 31, 2000, consisting of 70 data sets. The swaptions are priced using Monte Carlo simulation with 4000 paths for the evolution of the term structure. In generating the paths, the same seeds for the random number generator were used to ensure consistency across the models.




EXHIBIT 7A

Absolute Average Pricing Errors for Caps—Parametric and Four-Factor PCA Models

This exhibit presents the average absolute percentage errors, one week out-of-sample, for caps across strikes and maturities, for the one- and two-factor parametric models and the four-factor PCA model. The cap data are biweekly from March 1–December 31, 1998, consisting of 22 data sets. So each error reported in this exhibit is an average across 16 contracts over 22 dates (hence an average of 352 individual errors). The standard error of the mean is reported in parentheses. The caps are priced using Monte Carlo simulation with 4000 paths for the evolution of the term structure. In generating the paths, the same seeds for the random number generator were used to ensure consistency across the models.

	$\gamma = 0$	$\gamma = 0.5$	$\gamma = 1$
One-factor Parametric	17.9 (0.73)	15.0 (0.70)	14.4 (0.76)
Two-factor Parametric	18.4 (0.75)	15.6 (0.68)	15.1 (0.72)
Four-factor PCA	18.1 (0.75)	16.5 (0.78)	14.7 (0.79)

EXHIBIT 7B

Level Dependence Comparisons for Caps—Parametric and Four-Factor PCA Models

This exhibit presents the results of the proportion tests for the one- and two-factor parametric models and the four-factor PCA model, for different values of γ . For example, the first column of this exhibit shows the proportion of times a model with $\gamma = 0.5$ outperforms a model with $\gamma = 0$. The total number of contracts used for each proportion was 352 (22×16). The second and third columns show the proportion of times a $\gamma = 1$ model beats a $\gamma = 0.5$ and a $\gamma = 0$ model, respectively. The starred cells indicate the cases where the null hypothesis that the proportion of wins is 50% is rejected at the 5% level of significance.

	$\gamma = 0.5$ vs $\gamma = 0$	$\gamma = 1$ vs $\gamma = 0.5$	$\gamma = 1$ vs $\gamma = 0$
One-factor Parametric	0.61*	0.51	0.67*
Two-factor Parametric	0.70*	0.59*	0.63*
Four-factor PCA	0.72*	0.65*	0.71*

all models producing a high percentage of prices within bid-ask spreads.

Exhibit 10 confirms this inference, with more formal pairwise proportion tests across the models. The exhibit presents the proportion of times that one model produces more precise one-week out-of-sample prices than another, for contracts broken down by expiration and swap maturity. We first examine the summary statistics. In comparing the residuals produced by the different models, contract by contract, the two-factor parametric model produces smaller residuals than the one-factor model 55% of the time, and smaller residuals than the four-factor PCA model 54% of the time. The four-factor PCA model does not outperform the one-factor parametric model. Therefore, we conclude that on an aggregate basis at the 1% level of significance, the two-factor parametric model is the best, with no significant differences between the one-factor parametric and

the four-factor PCA models. However, the exhibit allows us to identify where some models fail. Specifically, the reason for the poor performance of the four-factor PCA model is its inability to price both very short- and very long-dated swaptions.

The results in Exhibit 10 reaffirm the conclusion that was derived from the aggregate results, that all three models produce results that are satisfactory, with the two-factor parametric model being marginally better. Unlike Longstaff, Santa-Clara, and Schwartz [2001], we do not conclude that higher order models are necessary for effectively pricing swaptions, one-week out-of-sample and even beyond, up to four weeks. Therefore, for pricing swaptions, it may not be necessary to require accurate calibration of correlations among forward rates, and high dimensional models such as three- or four-factor models may not be necessary.

EXHIBIT 8

Average Percentage Errors for Caps

This exhibit presents plots of the average percentage errors, one week out-of-sample, for caps across strikes for each maturity, for the one-factor and two-factor parametric and the four-factor PCA models. The plots are presented for three values of γ : $\gamma = 0$ (dotted line), $\gamma = 0.5$ (hashed line), and $\gamma = 1$ (solid line). The cap data are biweekly from March 1–December 31, 1998, consisting of 22 data sets. The in-sample optimizations for the models were conducted over four unknown parameters, with caps being priced using Monte Carlo simulation with 4000 paths for the evolution of the term structure. In generating the paths, the same seeds for the random number generator were used to ensure consistency across the models.




EXHIBIT 9

Box and Whiskers Plots of Swaption Pricing Errors

This exhibit presents the box and whiskers plots of the one week out-of-sample pricing errors for the one-factor and two-factor parametric and the four-factor PCA models with $\gamma = 1$, for swaptions across expirations and underlying swap maturities. The swaption data are biweekly from March 1, 1998–October 31, 2000, consisting of 70 data sets. The swaptions are priced using Monte Carlo simulation with 4000 paths for the evolution of the term structure. In generating the paths, the same seeds for the random number generator were used to ensure consistency across the models. In each exhibit, the first box corresponds to the one-factor parametric model, the second box to the two-factor parametric model, and the third box to the four-factor PCA model.









EXHIBIT 10

Comparison of Parametric and Four-Factor PCA Models for Swaptions

This exhibit compares the one week out-of-sample predictions for the one-factor and two-factor parametric models and the four-factor PCA model, with $\gamma = 1$, for each of the 30 contracts over all 70 dates. For example, the first entry in the top row shows that in 39% of the 70 cases, the two-factor model outperforms the one-factor model for one week ahead pricing of the 0.5×1 swaptions. The starred (dagger) cells indicate that the proportion is significantly different from 50% and the second (first) model in the comparison wins. All tests are conducted at the 1% level of significance.

Expiration	Swap Maturity	2P vs 1P	4PCA vs 2P	4PCA vs 1P
0.5	1	0.39	0.22*	0.28*
0.5	2	0.53	0.37	0.50
0.5	3	0.32*	0.25*	0.24*
0.5	4	0.31*	0.49	0.28*
0.5	5	0.57	0.34*	0.28*
1	1	0.53	0.70†	0.56
1	2	0.44	0.59	0.37
1	3	0.46	0.60	0.50
1	4	0.60	0.50	0.51
1	5	0.61	0.46	0.57
2	1	0.31	0.67†	0.46
2	2	0.49	0.59	0.57
2	3	0.50	0.61	0.63
2	4	0.78†	0.44	0.72†
2	5	0.82†	0.31*	0.84†
3	1	0.74†	0.22*	0.36
3	2	0.41	0.53	0.37
3	3	0.31*	0.70†	0.44
3	4	0.36	0.65†	0.57
3	5	0.67†	0.57	0.70
4	1	0.88†	0.31*	0.90†
4	2	0.71†	0.51	0.75†
4	3	0.65†	0.53	0.71†
4	4	0.63	0.37	0.40
4	5	0.53	0.27*	0.37
5	1	0.81†	0.53	0.61
5	2	0.64	0.27*	0.36
5	3	0.44	0.40	0.37
5	4	0.50	0.31*	0.46
5	5	0.64	0.46	0.61
Average		0.55†	0.46*	0.51

HEDGING PERFORMANCE OF PARAMETRIC AND PCA MODELS

Our analysis shows that *conditional* on a future term structure, all three models are capable of producing fairly precise estimates of swaption prices one to four weeks after the parameters for the volatility structure have been estimated. The removal of systematic expiry date and swap maturity biases from a pricing model is a necessary condition for any viable model for swaptions. However, a good model should also be able to hedge effectively. If the model is correct, the risk of carrying a hedged position over a time increment is entirely due to the fact that the volatility parameters are not known with certainty and continuous revisions were not accomplished over the time increment. If one model consistently produces hedges that are more effective than another model, then it must be the case that the first model, with its volatility structure, better captures the true dynamics of the term structure and the true sensitivity of options to movements in the underlying term structure.

Our hedging experiments are conducted as follows. Given any calibrated n -factor model, we can establish a hedge position for a particular swaption using n different LIBOR discount bonds. For example, for a four-factor model, four price changes for each swaption are recorded, each price change arising after a small shock is applied to a single factor. In addition, the four price changes to a set of discount bonds are computed. The unique portfolio of the four bonds is then established that hedges the swaption against instantaneous shocks consistent with the model. The construction of the hedged position, at any date t , only uses information available at date t . This analysis is repeated for all contracts and for all models. The hedge position is maintained unchanged for one week, and the hedged and unhedged residuals are obtained and stored. The analysis is repeated for holding periods of two, three and four weeks as well. Finally, this entire procedure is repeated every second week, over all 70 weeks, for which data was available.¹⁰

If a one-factor model is correct, then only one hedging instrument is necessary. In this case, any zero coupon bond could be used to hedge a particular swaption. Similarly, for a four-factor model, four hedging instruments are required and, theoretically, any four zero coupon bonds could be used to establish a unique hedge. Since the contract underlying the swaption is a forward starting swap, we restrict the maturities of the hedging

instruments to the time period where cash flows are received.

In our first set of analyses, we evaluate hedging strategies where the number of hedging instruments used equals the number of factors in the model. For the one-factor model, we take the discount bond corresponding to the maturity date of the underlying swap. For the two-factor model, the two hedging instruments are taken as the discount bonds corresponding to the expiration date of the swaption and the maturity date of the underlying swap. Finally, for the four-factor model, the hedging instruments are taken to be these two discount bonds, plus two additional bonds that have maturities equally spaced between the expiry and underlying swap maturity date.

It could be argued that these tests may be biased in favor of the four-factor model, since that model allows more hedging instruments. In our second set of analyses, we examine alternative hedging strategies using multiple zero coupon bonds as hedging instruments for the one- and two-factor models. In particular, we examine six additional strategies using the one-factor model, and two additional strategies for the two-factor model, while allowing no such flexibility to the four-factor model. One such strategy involves using the same four bonds to hedge, regardless of how many factors the models have. For example, for the two into three swaption, we hedge using bonds with maturities of three, four, five, and six years to maturity.

Effects of the Number of Factors on Hedging


The benchmark hedging period is one week, i.e., the hedge is set up and then evaluated after one week. For each swaption contract, a time series of weekly unhedged and hedged residuals is obtained. Exhibit 11 compares the box and whisker plots of hedging errors for each contract type across the three models. It also presents the unhedged pricing errors.

Exhibit 11 immediately shows that all hedges are doing their job! However, it also shows that the one-factor model has larger errors. While the spread of the errors, as indicated by the width of the inner box, are of similar sizes for the three models, the bias in the results are greatest for the one-factor model. As an example, consider the six-month expiry contracts. The biases in the hedging errors, as indicated by the difference between the median error and zero, are large and positive for the one- and two-factor models. In contrast, the interquartile

EXHIBIT 11

Box and Whiskers Plots of Swaption Hedging Errors

This exhibit presents the box and whiskers plots for the one week out-of-sample hedging errors for the one-factor and two-factor parametric and the four-factor PCA models (with $\gamma = 1$), for swaptions across all expirations and underlying swap maturities. The corresponding plots for the unhedged swaptions are presented for comparison purposes. In each exhibit, the first box corresponds to the unhedged swaption, the second box to the one-factor parametric model, the third box to the two-factor parametric model, and the fourth box to the four-factor PCA model.



ranges are somewhat similar. This phenomenon holds true for almost all expiry dates, with the exception of the long term contracts. Here, the bias is small, but the variance of the residuals produced by the one-factor model is larger.

Exhibit 11 clearly shows that the ratio of the standard deviation of the hedged position versus the standard deviation of the unhedged swaption is not an appropriate measure of the hedging effectiveness of a model when

the model is biased. The appropriate metric is based on the root mean squared error of the hedging errors, since it accounts for any bias in the model. Exhibit 12 presents the root mean squared error (multiplied by 10,000) for each contract type. As a result, each entry can be interpreted in basis points.

As an example, consider the six-month maturity contract on a swap of two years. The unhedged root mean squared error is 12.1 basis points, while the one-factor

EXHIBIT 12

Absolute Hedging Errors for Swaptions

This exhibit presents the root mean squared errors (in basis points) of the hedged and unhedged portfolios for the one-factor and two-factor parametric and the four-factor PCA models, one week out-of-sample. In the hedge portfolios, the number of hedging instruments used equals the number of factors in the model. The swaption data are biweekly from March 1, 1998–October 31, 2000, consisting of 70 data sets. The root mean square of the hedging errors for a contract, across all dates, is multiplied by 10,000 so that it can be interpreted as a basis point error. The corresponding root mean squared errors for the unhedged swaptions are also presented, for comparison.

Expiration	Swap Maturity	Unhedged Swaption	One-Factor	Two-factor	Four-Factor
0.5	2	12.1	6.2	3.6	3.0
0.5	3	18.1	7.0	5.9	4.2
0.5	4	23.3	7.7	6.3	5.2
0.5	5	28.2	8.2	6.5	5.9
1	2	13.1	5.5	3.6	3.2
1	3	18.9	6.6	5.0	4.6
1	4	24.1	7.7	8.8	5.7
1	5	28.9	8.2	8.7	6.2
2	2	13.1	5.6	3.6	4.3
2	3	18.7	7.4	5.0	8.4
2	4	23.0	8.4	6.6	5.8
2	5	27.5	9.5	8.4	7.0
3	2	12.6	5.8	3.8	5.4
3	3	17.3	7.4	8.5	5.0
3	4	21.9	8.7	7.2	6.3
3	5	26.9	10.6	8.1	8.4
4	2	11.4	5.7	4.3	3.7
4	3	15.8	7.0	5.5	5.3
4	4	20.9	8.5	6.3	6.3
4	5	25.9	10.2	7.7	7.6
5	2	10.8	5.4	4.1	4.5
5	3	16.4	7.8	5.6	5.7
5	4	21.5	9.9	7.2	7.1
5	5	26.3	11.7	8.8	9.5

hedged position has a root mean squared error of 6.2 basis points. The four-factor PCA model, however, has a root mean squared error of 3.0 basis points, indicating it is almost twice as effective. In comparing the root mean squared errors, contract by contract, the benefits of the four-factor model become apparent.

As a more formal test of hedging effectiveness, we conduct pairwise comparisons of the hedging residuals produced by each model. Specifically, for each contract and for each week, the hedging residual is computed and the model with the smallest absolute value of hedging error is identified. The results are presented in Exhibit 13. For example, for the six-month contract, with swap maturity of two years, out of 70 hedging experiments, the two- and four-factor models produce smaller absolute hedging residuals than the one-factor model on 87% of occasions. In 66% of occasions, the four-factor model outperforms the two-factor model. If the hedges are maintained unchanged for two weeks, the performance of the multifactor models improves even further. Indeed, over the four-week period, the four-factor model outperforms the one-factor model in all but one week.

Exhibit 13 clearly shows that over all contracts, and for all weeks, the multifactor model is dominant. For hedging purposes, multifactor models are necessary!

The Effects of Increasing the Number of Hedging Instruments

Driessen, Klaassen, and Melenberg [2003] show that when bucket strategies are used for hedging, the performance of the one-factor models improves significantly. In light of their result, we set up experiments where the benefits of using additional hedging instruments in a one- and two-factor model could be assessed.

The hedging instruments for these tests are selected as follows. Let the swaption expiry date be labeled date 1, and the underlying swap maturity date be labeled date 4. So far, for the one-factor model, we only considered hedging using the date 4 maturity discount bond. In contrast, the four-factor model uses the date 1 and 4 bonds, together with two equally spaced bonds in between. Let these two corresponding dates be 2 and 3, respectively. For the one-factor model, we now consider hedging the swaption with two discount bonds, with three bonds, and with four bonds. In particular, with two bonds, we choose bonds 3 and 4; with three bonds, we choose bonds 2, 3

and 4; and with 4 bonds, we use the same set of bonds as the four-factor model.

If more than one hedging instrument is used in a one-factor model, the hedge ratios are not unique, and a somewhat arbitrary rule must be made to construct the specific bond hedges. We use equal number of bonds in each hedge. Of course, other allocations can be considered, as well as other rules for obtaining unique hedge ratios, but our goal was only to assess if there were any strategies using lower order models that could lead to significant improvements.

For the two-factor model, we use two hedging instruments corresponding to dates 1 and 4. When using three hedging instruments, we choose the third bond to be the middle maturity between 1 and 4, and with four hedging instruments, we again use bonds 1, 2, 3, and 4. With three hedging instruments, we assume that the numbers in the short bond and the middle bond are equal; for the four hedging instruments, unique hedges are constructed by assuming the number of bonds in the first two maturities to be equal, as well as the number of bonds in the last two maturities to be equal.

The one-factor model with one hedging instrument, corresponding to the swap maturity (labeled bond 4), is used as the benchmark. The number of times the benchmark model produces smaller absolute residuals than each of the challenging models is recorded, across all weeks for each contract. Exhibit 14 summarizes the results aggregated over all contracts.

For the one-factor model, none of the potential enhancements produce better results. Similarly, no improvements are obtained for the two-factor model. Other bond allocations were tested, but we could not identify systematic ways of improving the performance of lower order models using multiple hedging instruments.¹¹

CONCLUSION

This article carefully examines the role of volatility structures and factors in the pricing of swaptions. Among the PCA models, we show that increasing the number of factors up to four improves pricing performance. Similar results are obtained for caps. Since away-from-the-money caps provide information about skew effects, we use that data to establish that level dependence is important, and that proportional models are better than absolute or square root models. Our results reconfirm the results of Longstaff, Santa-Clara, and Schwartz [2001] and suggest that the

EXHIBIT 13

Comparison of Parametric and PCA Models in Hedging Swaptions

This exhibit presents the fraction of times one model outperforms the other model in hedging forecasts, for the one-factor and two-factor parametric and the four-factor PCA models, one week out-of-sample. For example, the first proportion in the first row indicates that the two-factor model outperforms the one-factor model 87% of the time. In the hedge portfolios, the number of hedging instruments used equals the number of factors in the model. The swaption data are biweekly from March 1, 1998–October 31, 2000, consisting of 70 data sets. Therefore, for each contract, the proportions are computed from a comparison of 70 hedging errors.

Expiration	Swap Maturity	1 week out-of-sample			2 weeks out-of-sample			4 weeks out-of-sample		
		2 vs 1	4 vs 1	4 vs 2	2 vs 1	4 vs 1	4 vs 2	2 vs 1	4 vs 1	4 vs 2
0.5	2	0.87	0.87	0.66	0.93	0.99	0.68	0.94	0.99	0.75
0.5	3	0.81	0.86	0.67	0.93	0.90	0.71	0.93	0.96	0.84
0.5	4	0.67	0.87	0.76	0.77	0.88	0.75	0.82	0.94	0.79
0.5	5	0.59	0.81	0.73	0.70	0.84	0.72	0.78	0.91	0.78
1	2	0.81	0.77	0.61	0.91	0.90	0.64	0.97	0.97	0.65
1	3	0.71	0.73	0.54	0.86	0.88	0.58	0.94	0.94	0.65
1	4	0.63	0.70	0.57	0.74	0.86	0.57	0.79	0.93	0.62
1	5	0.56	0.69	0.61	0.62	0.78	0.74	0.71	0.84	0.72
2	2	0.73	0.66	0.50	0.78	0.75	0.57	0.90	0.85	0.53
2	3	0.67	0.53	0.31	0.72	0.65	0.30	0.94	0.74	0.31
2	4	0.59	0.69	0.59	0.65	0.68	0.51	0.82	0.84	0.62
2	5	0.61	0.64	0.63	0.62	0.67	0.57	0.68	0.78	0.53
3	2	0.63	0.61	0.46	0.70	0.67	0.52	0.76	0.68	0.49
3	3	0.66	0.67	0.60	0.64	0.70	0.55	0.74	0.74	0.53
3	4	0.66	0.69	0.57	0.68	0.67	0.57	0.74	0.72	0.56
3	5	0.69	0.66	0.56	0.65	0.65	0.67	0.75	0.74	0.62
4	2	0.64	0.63	0.53	0.71	0.72	0.54	0.79	0.74	0.51
4	3	0.70	0.69	0.60	0.65	0.67	0.61	0.76	0.68	0.51
4	4	0.77	0.73	0.61	0.70	0.65	0.62	0.79	0.75	0.59
4	5	0.74	0.71	0.60	0.70	0.67	0.59	0.72	0.66	0.53
5	2	0.73	0.66	0.50	0.75	0.62	0.45	0.74	0.68	0.49
5	3	0.81	0.76	0.54	0.72	0.74	0.62	0.76	0.71	0.60
5	4	0.80	0.77	0.66	0.72	0.68	0.62	0.71	0.69	0.62
5	5	0.73	0.66	0.47	0.72	0.70	0.58	0.69	0.66	0.50
Average (1 week out)		0.70	0.71	0.58	0.73	0.75	0.59	0.80	0.80	0.60

Gaussian models of Driessen, Klaassen, and Melenberg [2003] could be improved upon by incorporating level dependence in the volatility structures.

One limitation of the PCA models is that the number of free parameters implied out from derivatives data equals the number of stochastic drivers. Therefore, it is unclear whether the four-factor PCA model outperforms lower

order models due to the number of factors or due to the number of free parameters. If it is due to more factors, it supports the importance of forward rate correlations for modeling swaption products. However, if it is due to the fact that higher order models have more degrees of freedom which permit better fits to swaptions, then the need to model correlations well is questionable. By analyzing lower

EXHIBIT 14

Comparison of Models using Different Hedging Instruments

This exhibit presents the fraction of times one model specification outperforms the same model using different hedging instruments, one week out-of-sample, for the one-factor and two-factor parametric models. For example, the first row indicates that within the one-factor model, the hedging strategy using one hedging instrument outperforms that using two hedging instruments 52% of the time. The swaption data are biweekly from March 1, 1998–October 31, 2000, consisting of 70 data sets. Underlying discount bonds are used as hedging instruments, and are labeled as follows. The bond expiring on the swaption expiration date is labeled “1,” while the bond expiring on the underlying swap maturity date is labeled “4.” Bonds “2” and “3” correspond to bonds with maturities equally spaced between the swaption expiration date (bond 1) and the underlying swap maturity date (bond 4). For the one-factor model, the hedge using bond 4 is used as the benchmark hedge. For the two-factor model, the hedge using bonds 1 and 4 is the benchmark. When using three hedging instruments within the two-factor model, the third bond is chosen to be the middle maturity between bonds 1 and 4. The four hedging instruments within the two-factor model are bonds 1, 2, 3, and 4.

Model Comparison	Fraction of wins
One-Factor Parametric Model	
One vs two instruments (4 vs 4, 3)	0.52
One vs three instruments (4 vs 4, 3, 2)	0.54
One vs four instruments (4 vs 4, 3, 2, 1)	0.55
One instrument (4 vs 3)	0.53
One instrument (4 vs 2)	0.57
One instrument (4 vs 1)	0.66
Two-Factor Parametric Model	
Two vs three instruments	0.50
Two vs four instruments	0.53

order models with the same degrees of freedom, we address this issue and conclude that, from a pricing perspective, one-factor models are sufficient.

Regarding hedging effectiveness, however, the one-factor models are significantly outperformed by multifactor models. We show that the higher order models perform consistently better for all contracts, regardless of their expiration date and underlying swap maturity. We show that the popular criterion of assessing hedging performance based on the percentage of unhedged *variance* explained by the hedging instrument could lead to flawed conclusions, since some models have systematic biases in their hedges. Using *root mean squared error* as the criterion, which accounts for variance and bias, leads to more precise conclusions. In the hedging tests, the performance of multifactor models is not only superior for hedges maintained over one week, but also for hedges maintained unchanged for as long as four weeks.

While, at a superficial level, our results echo the final conclusions reached by Longstaff, Santa-Clara, and Schwartz [2001], our conclusions are arrived at for different reasons. In particular, for pricing purposes, they find it necessary to use four-factor models. In contrast, we

find a one-factor model that is satisfactory for pricing, and only recommend multifactor models for hedging. Our results are consistent with Andersen and Andreasen [2001], who identify useful one-factor models for pricing Bermudan swaptions. With regard to both pricing and hedging, our results align with those obtained by Gupta and Subrahmanyam [2005], who show that in the cap market, one-factor models are good enough for pricing, but two-factor models are required for hedging. Our results also align with Fan, Gupta, and Ritchken [2003], who show that multifactor models that do not exhibit unspanned stochastic volatility are capable of accurately hedging swaption straddles. In addition, our hedging results for swaptions show that using multiple instruments within a lower order model does not improve hedging performance. These results differ from Driessen, Klaassen, and Melenberg [2003], possibly due to their incorrect use of variance reduction as a criterion of hedging effectiveness. Bias effects are important when a misspecified model is used.

It would be interesting to study how well the models we examine perform on swaptions with strikes away-from-the-money. Unfortunately, we do not have prices of such contracts. For at-the-money swaptions, the effects of level

dependence in the forward rate volatility structure are not apparent. Indeed, absolute and square root models work as well as proportional models. For the Gaussian models, extremely accurate analytical approximations are available for swaption prices. However, based on our results for caps, it is likely that away-from-the-money swaptions might be more sensitive to the level dependence parameter, γ , hence models with $\gamma = 1$ are preferable.

While our study indicates that there is a need for multi-factor models for hedging exotics like Bermudan swaptions, much work remains to be done here. For example, Andersen and Andreasen [2001] explore the common market practice of pricing swaption based products by continuously refitting a model with time varying parameters to swaption prices. The common market practice for hedging, as far as we can assess, is to perturb selected forward rates, one at a time, and to reprice under each shocked curve. The resulting price changes of claims can then be hedged using Eurodollar futures of the appropriate maturities. It remains for future research to assess if the hedges set up, based on this popular market approach, are competitive with hedges set up using our higher order multi-factor models.

We obtain small out-of-sample residuals for swaptions and larger residuals for caps. Collin-Dufresne and Goldstein [2003] have argued that the high residuals for caps, found in many studies, are due to the absence of state variables driving correlation and volatility risk, and that caps might be more sensitive to these factors than swaptions. It remains for future research to evaluate the performance of these models in the cap market.

Continuing to extract information from interest rate derivative products will, no doubt, lead us to an increased understanding of the nature of the volatility structure of forward rates. A deeper understanding of these structures is crucial for interest rate risk management and value-at-risk.

ENDNOTES

¹A cap consists of a portfolio of caplets, each caplet representing an option on an individual forward rate. In contrast, a swaption can be viewed as an option on a portfolio of forward rates. As a result, the relationship between caps and swaptions is determined largely by the correlation structure among the forward rates. Rebonato claims that since forward rates are perfectly correlated in a one-factor model, such models will tend to overprice swaptions.

²Exceptions are Jaganathan, Kaplin and Sun [2003] and Fan, Gupta, and Ritchken [2003], who characterize the pricing performance of their models according to the expiry dates of the swaption.

³Incorporating level dependence in the volatility structure results in distributions of forward rates that are no longer Gaussian. This has large implications for pricing claims that derive value from the tails of the distribution.

⁴Hull and White [2000] develop very similar models to the Longstaff, Santa-Clara, and Schwartz models using the LIBOR market based model of Brace, Gatarek, and Musiela [1997]. A big motivation for extending the existing models was to permit an analysis of volatility structures that were not proportional to their levels. Hull and White [2000] test their extended model using data for a single day, and provide preliminary support for multi-factor models where volatilities are nonlinear in forward rates.

⁵In addition, their hedging tests are not based on the construction of portfolios of traded instruments.

⁶The empirical evidence on the importance of unspanned stochastic volatility for pricing and hedging fixed income derivatives is also discussed by Li and Zhao [2006] and Casassus, Collin-Dufresne, and Goldstein [2005].

⁷For a discussion of Markovian Heath, Jarrow, and Morton models, see Ritchken and Sankarasubramanian [1995a] and Bhar and Chiarella [1995].

⁸For a review of simulation approaches for pricing claims, see Boyle, Broadie, and Glasserman [1997].

⁹Prices of away-from-the-money contracts convey more information about the tails of the conditional distribution of forward rates and are more sensitive to γ . For example, see Ritchken and Sankarasubramanian [1995b].

¹⁰For a discussion of how best to use simulation models for establishing hedge ratios, see Jäckel [2001].

¹¹For example, in the one-factor model, it is true that choosing a different maturity bond as the single bond to hedge with can make a difference. For example, using the expiry date bond as the only hedging instrument gave very poor results relative to using the swap maturity dated bond. Since the choice of hedging instruments should not really matter, if the model is correct, this only provides more evidence that the one-factor model is a poor model for hedging.

REFERENCES

- Amin, K., and A. Morton. "Implied Volatility Functions in Arbitrage-free Term Structure Models." *Journal of Financial Economics*, 35 (1994), pp. 141–180.
- Andersen, L., and J. Andreasen. "Factor Dependence of Bermudan Swaptions: Fact or Fiction?" *Journal of Financial Economics*, 62 (2001), pp. 3–37.

- Bhar, R., and C. Chiarella. "Transformation of Heath-Jarrow-Morton Models to Markovian Systems." *The European Journal of Finance*, 3 (1995), pp. 1-26.
- Black, F. "The Pricing of Commodity Contracts." *Journal of Financial Economics*, 3 (1976), pp. 167-179.
- Black, F., and P. Karasinski. "Bond and Option Pricing when Short Rates are Lognormal." *Financial Analysts Journal*, 47 (1991), pp. 52-59.
- Boyle, P., M. Broadie, and P. Glasserman. "Monte Carlo Methods for Security Pricing." *Journal of Economic Dynamics and Control*, 21 (1997), pp. 1267-1322.
- Bühler, W., M. Uhrig, U. Walter, and T. Weber. "An Empirical Comparison of Forward and Spot-rate Models for Valuing Interest-rate Options." *Journal of Finance*, 54 (1999), pp. 269-305.
- Carr, P., and G. Yang. "Simulating Bermudan Interest Rate Derivatives." Working paper, Bank of America Securities, 1999.
- Casassus, J., P. Collin-Dufresne, and R.S. Goldstein. "Unspanned Stochastic Volatility and Fixed Income Derivatives Pricing." *Journal of Banking and Finance*, 29 (2005), pp. 2723-2749.
- Collin-Dufresne, P., and R.S. Goldstein. "Stochastic Correlation and the Relative Pricing of Caps and Swaptions in a Generalized-Affine Framework." Working paper, Carnegie Mellon University, 2003.
- Cox, J., J. Ingersoll, and S. Ross. "A Theory of the Term Structure of Interest Rates." *Econometrica*, 53 (1985), pp. 363-384.
- Driessen, J., P. Klaassen, and B. Melenberg. "The Performance of Multi-factor Term Structure Models for Pricing and Hedging Caps and Swaptions." *Journal of Financial and Quantitative Analysis*, 38 (2003), pp. 635-672.
- Duarte, J. "The Causal Effect of Mortgage Refinancing on Interest Rate Volatility: Empirical Evidence and Theoretical Implications." Forthcoming, *The Review of Financial Studies*, 2007.
- Fan, R., A. Gupta, and P. Ritchken. "Hedging in the Possible Presence of Unspanned Stochastic Volatility: Evidence from Swaption Markets." *Journal of Finance*, 58 (2003), pp. 2219-2248.
- Gupta, A., and M. Subrahmanyam. "Pricing and Hedging Interest Rate Options: Evidence from Cap-Floor Markets." *Journal of Banking and Finance*, 29 (2005), pp. 701-733.
- Heath, D., R. Jarrow, and A. Morton. "Bond Pricing and the Term Structure of Interest Rates." *Econometrica*, 60 (1992), pp. 77-105.
- Hull, J., and A. White. "One-factor Interest Rate Models and the Valuation of Interest Rate Derivative Securities." *Journal of Financial and Quantitative Analysis*, 28 (1993), pp. 235-254.
- . "Forward Rate Volatilities, Swap Rate Volatilities, and the Implementation of the LIBOR Market Model." *Journal of Fixed Income*, 10 (2000), pp. 46-62.
- Inui, K., and M. Kijima. "A Markovian Framework in Multi-Factor Heath-Jarrow-Morton Models." *Journal of Financial and Quantitative Analysis*, 33 (1998), pp. 423-440.
- Jäckel, P. "Greeks With Monte Carlo." *Wilmott* (July 2007), pp. 55-63.
- Jagannathan, R., A. Kaplin, and S. Sun. "An Evaluation of Multi-Factor CIR Models Using LIBOR, Swap Rates, and Cap and Swaption Prices." *Journal of Econometrics*, 116 (2003), pp. 113-146.
- Jarrow, R., H. Li, and F. Zhao. "Interest Rate Caps Smile Too! But Can the LIBOR Market Models Capture It?" *Journal of Finance*, 62 (2007), pp. 345-382.
- Li, H., and F. Zhao. "Unspanned Stochastic Volatility: Evidence from Hedging Interest Rate Derivatives." *Journal of Finance*, 61 (2006), pp. 341-378.
- Litterman, T., and J. Scheinkman. "Common Factors Affecting Bond Returns." *Journal of Fixed Income* (June 1991), pp. 54-61.
- Longstaff, F., P. Santa-Clara, and E. Schwartz. "The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence." *Journal of Finance*, 56 (2001), pp. 2067-2109.
- Longstaff, F., and E. Schwartz. "Valuing American Options by Simulation: A Simple Least-Squares Approach." *The Review of Financial Studies*, 14 (2001), pp. 113-147.
- Moraleda, J., and A. Pelsser. "Forward versus Spot Interest Rate Models of the Term Structure: An Empirical Comparison." *The Journal of Derivatives* (Spring 2000), pp. 9-21.
- Moraleda J., and T. Vorst. "Pricing American Interest Rate Claims with Humped Volatility Models." *Journal of Banking and Finance*, 21 (1997), pp. 1131-1137.

Peterson, S., R. Stapleton, and M. Subrahmanyam. "A Multi-Factor Spot-Rate Model for the Pricing of Interest Rate Derivatives." *Journal of Financial and Quantitative Analysis*, 38 (2003), pp. 847-880.

Ritchken, P., and I. Chuang. "Interest Rate Option Pricing with Volatility Humps." *Review of Derivatives Research*, 3 (1999), pp. 237-262.

Ritchken, P., and L. Sankarasubramanian. "Volatility Structures of Forward Rates and the Dynamics of the Term Structure." *Mathematical Finance*, 5 (1995a), pp. 55-72.

———. "The Importance of Forward Rate Volatility Structures in Pricing Interest Rate Sensitive Claims." *Journal of Derivatives*, (1995b), pp. 25-41.

Rebonato, R. "On the Simultaneous Calibration of Multi-factor Lognormal Interest Rate Models to Black Volatilities and to the Correlation Matrix." *Journal of Computational Finance*, 4 (1999), pp. 5-27.

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

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Institutional Investor PLC.

The most recent two editions of this title are only ever available at <http://www.euromoneyplc.com>