Factor Dependence and Estimation Risk for Cap-Related Interest Rate Exotics

JEROEN KERKHOF

JEROEN KERKHOF
is with Quantitative Fixed Income Research, Lehman Brothers Int. in London, UK.
Jeroen.Kerkhof@Lehman.com

During the 1990s, we witnessed a spectacular growth of trading derivative instruments that continues into the new millennium. For example, the turnover of exchange-traded financial derivatives is estimated at about $192 trillion (see Jeanneau [2002]). The notional amount of outstanding OTC is estimated at $128 trillion (see Jameson and Gadanecz [2003]) consisting for the larger part of (short-term) interest rate products.

The two main liquid (over-the-counter) interest rate derivatives markets are the caps (floors) and swaptions market. Though the largest part of derivatives traded in these markets are the plain vanilla caps, floors, and swaptions, there exists a reasonably sizable market in more exotic products. These products can be traded separately or, as is often the case, as part of a more complex structured deal. In either case these deals are over-the-counter and not very liquid. Therefore, in these markets, besides hedging model risk considerable pricing model risk can also exist (see, e.g., Hull and Sno [2002]). Therefore, it is important for both traders and risk management divisions to get an idea about the price range of the value of the exotic derivative. In this article we investigate the estimation risk involved in pricing exotic interest rate derivatives. These price ranges due to estimation risk can provide risk management divisions guidelines in setting model reserves for the various products and traders about bid-ask spreads.

We adopt the popular Libor market model (see Brace et al. [1997]; Miltersen et al. [1997], and Jamshidian [1997]) for analyzing several cap-related interest rate exotics: deferred caps, autocaps, sticky caps, ratchet caps, and discrete barrier caps.[1] Contrary to ordinary caps, these products are sensitive to the joint distribution of the forward rates term structure. In the case of the Libor market model, this means the specification of the correlation matrix of the forward rates. In practice, this correlation matrix is either estimated historically or calibrated to the swaption market as swaptions are correlation-sensitive liquid products. Unfortunately, the relation of caps and swaptions markets remains somewhat unclear (see, e.g., Longstaff et al. [2001a] and de Jong et al. [2001]). Calibration of the Libor market model to swaption prices often results in unrealistically unstable correlation matrices and consequently prices and especially sensitivities. Therefore, we adopt historical correlation in this article.

The Libor market model can be specified in numerous ways. Among other things, the number of factors needs to be specified. In this article we investigate to what extent the cap-related exotics are sensitive to the number of factors used in estimating the correlation matrix. We find that the autocap, ratchet cap, and sticky caps are very sensitive to the number of factors used. Price ranges due to estimation risk are investigated using the stationary bootstrap of Politis and Romano [1994] and the window length selection method of Politis and White [2003]. As one would expect, the prices of the correlation-sensitive products contain a substantial degree of estimation risk.

The remainder of this article is structured as follows. Section II introduces notation and reviews the Libor market model and discrete string model. Section III describes the exotic derivatives investigated. Section VI discusses the bootstrap technique for determining estimation risk. The results are presented in Section V. Finally, Section IV concludes.

NOTATION AND MODELS

Notation

In this section we introduce the notation necessary for the remainder and give a description of the Libor market model/discrete string model. First, we define a finite set of dates, the so-called *tenor structure*

$$t = T_0 < T_1 < T_2 < \ldots < T_{N+1}. \tag{1}$$

We indicate the current time by t, and $T_1, \ldots, T_{N+1}$ denote the forward tenor dates. This gives a spot LIBOR rate (for $[T_0, T_1]$) and N forward LIBOR rates from (for $[T_i, T_{i+1}]$, $i = 1, \ldots, N$). We define $\delta_i = \delta(T_i, T_{i+1})$ as the so-called daycount fractions (for an extensive treatment on day-count fractions, see Miron and Swannell [1992]), which are determined by the maturity of the LIBOR rate and are most often approximately equal to 3 or 6 months. Let the forward LIBOR rate from T_i to T_{i+1} at time $t \le T_i$ be denoted by $F(t, T_i, T_{i+1})$ which is defined as

$$F(t,T_i,T_{i+1}) \equiv \frac{1}{\delta_i}\left(\frac{D(t,T_i)-D(t,T_{i+1})}{D(t,T_{i+1})}\right), \quad t \le T_i \text{ and } i = 1, \ldots, \tag{2}$$

where $D(t, T)$ denotes the value of a discount bond at time t with maturity T. For notational convenience, we define $F_i(t) \equiv F(t, T_i, T_{i+1})$.

Libor Market Model

In this section, we briefly describe the Libor market model (introduced by Brace et al. [1997], Miltersen et al. [1997], and Jamshidian [1997]). In Kerkhof and Pelsser [2002] it was shown that this model is observationally equivalent to the discrete string model (introduced by Longstaff et al. [2001] and Longstaff et al. [2001a]). The models are characterized by the following dynamics for the N individual forward rates:

$$\frac{dF_i(t)}{F_i(t)} = \alpha_i^M(t)dt + \sigma_i(t)dZ_i^M(t), \tag{3}$$

$$= \alpha_i^M(t)dt + \Gamma_i(t)'dW^M(t), \quad i = 1,\ldots,N \tag{4}$$

where $\left\{Z_i^M(t)\right\}_{i=1}^N$ are (correlated) Wiener processes under probability measure M with

$$d\left[Z_i^M, Z_j^M\right](t) = \rho_{i,j}\ dt, \quad i,j = 1,\ldots,N,$$

W^M denotes a K-dimensional ($K \le N$) standard Wiener process and $\rho_{i,j}$ denotes the instantaneous correlation between Z_i and Z_j. Equation (3) corresponds to the discrete string model (DSM) setting whereas Equation (4) corresponds to the Libor market model (LMM) setting. Due to absence-of-arbitrage restrictions, $\alpha_i^M = 0$ if $M = \mathbb{Q}^{i+1}$, the T_{i+1}-forward martingale measure associated with numeraire $D(\bullet, T_{i+1})$.

The covariance matrix (in DSM notation, and constant for notational convenience) of the log forward rate changes is given by

$$\Sigma = \begin{bmatrix} \sigma_1^2 & \cdots & \rho_{1,N}\sigma_1\sigma_N \\ \vdots & \ddots & \vdots \\ \rho_{1,N}\sigma_1\sigma_N & \cdots & \sigma_N^2 \end{bmatrix}. \tag{5}$$

Monte-Carlo Pricing

The forward rate dynamics of F_i given in Equation (3) are easily generated using Monte-Carlo methods under the T_i-forward martingale measure using an Euler method, for all i. In the sequel we are interested in pricing derivatives that depend on several forward rate settings, which complicates the pricing. Therefore, we want to model the dynamics of all forward rates under one measure. Though in principle the choice of measure is arbitrary, it is convenient to use the T_{N+1}-forward measure, also referred to as the terminal measure. Changing the measure from the T_i-forward martingale measure to the terminal measure gives the following drifts for (3) and (4):

$$\alpha_i^{N+1}(t) = -\sigma_i(t)\sum_{k=i+1}^{N}\frac{\delta_k F_k(t)}{1+\delta_k F_k(t)}\sigma_k(t)\rho_{i,k}, \quad i=1,\dots,N. \quad (6)$$

Note that the dynamics now have a stochastic drift term, and therefore in an Euler scheme, the dynamics of the forward rates are no longer exact but for the last forward rate. More advanced methods exist to approximate the stochastic drift term in (6), such as the predictor-corrector method proposed in Hunter et al. [2001] (see Kloeden and Platen [1999] for more general predictor-corrector methods), which is employed in this article.

For the exotic derivatives treated in the sequel no analytical formula is available in the Libor market model. In computing the prices, several variance reduction techniques can be applied. Since analytical prices are available for caps and floors, these are used as control variates. Furthermore, antithetic variables are used.

EXOTIC INTEREST RATE DERIVATIVES

In this section we describe a number of commonly used cap-related exotic interest rate derivatives. For a more extensive list of interest rate exotics, see, for example, Brigo and Mercurio [2001], Rebonato [2002],Pelsser [2000], and Hunt and Kennedy [2000] For all products a unit notional amount is assumed.

Cap/Floor

We start with the most basic instruments: caps and floors. A cap (floor) is a portfolio of options on LIBOR rates called caplets (floorlets). Caps (floors) provide protection against high (low) interest rates. The payoff of a caplet (floorlet) with strike κ fixing at T_i and paying at T_{i+1} is given by

$$\mathrm{Cpl}_i(T_{i+1}) = \delta_i\left(F_i(T_i)-\kappa\right)^+ \quad (7)$$

$$\mathrm{Frl}_i(T_{i+1}) = \delta_i\left(\kappa-F_i(T_i)\right)^+. \quad (8)$$

Choosing the discount bond maturing at T_{i+1}, $D(T_i, T_{i+1})$, as numeraire and working with the associated martingale measure (usually denoted as the T_{i+1}*-forward measure*) we find convenient pricing equations for caplets and floorlets:

$$\mathrm{Cpl}_i(T_0) = \delta_i D(T_0,T_{i+1})\mathbb{E}_0^{i+1}\left[\left(F_i(T_i)-\kappa\right)^+\right], \quad (9)$$

$$\mathrm{Frl}_i(T_0) = \delta_i D(T_0,T_{i+1})\mathbb{E}_0^{i+1}\left[\left(\kappa-F_i(T_i)\right)^+\right], \quad (10)$$

where $\mathbb{E}_0^{i+1}$ denotes the conditional expectation with respect to the measure $\mathbb{Q}^{i+1}$ at time T_0.

In the LMM, where the forward LIBOR rates follow a lognormal distribution, caplets and floorlets can be priced explicitly using the Black formula (see Black [1976]):

$$\mathrm{Cpl}_i(T_0) = \delta_i D(T_0,T_{i+1})\left(F_i(T_i)\Phi(d_+)-\kappa\Phi(d_-)\right), \quad (11)$$

$$\mathrm{Frl}_i(T_0) = \delta_i D(T_0,T_{i+1})\left(\kappa\Phi(-d_-)-F_i(T_i)\Phi(-d_+)\right), \quad (12)$$

where Φ denotes the cumulative distribution of the standard Gaussian distribution and

$$d_\pm = \frac{\log(F_i(T_i)/\kappa)\pm\frac{1}{2}\Sigma_i^2}{\Sigma_i}, \quad \text{and} \quad \Sigma_i^2 = \int_0^{T_i}\sigma_i^2(s)ds. \quad (13)$$

Deferred Cap

The standard cap/floor caplet/floorlet payments occur at their "natural" times, that is, one period after their setting. There also exist caps/floors in the market for which all caplet/floorlet payments occur at the final time, T_{N+1}. We term this a *deferred cap/floor*. The value of a deferred cap is given by

$$\mathrm{DC}(T_0) = D(T_0,T_{N+1})\sum_{i=1}^{N}\delta_i\mathbb{E}_0^{N+1}\left[(F_i(T_i)-\kappa)^+\right], \quad (14)$$

Similar valuation formulas can be derived for deferred floors. We expect the deferred cap/floor to be sensitive to the shape of the term structure. It should be slightly sensitive to correlation (see correlation in the drift term in (6)), but probably not too much.

Autocap

An *autocap* (*autofloor*) is similar to an ordinary cap (floor), but the holder is only allowed to exercise $l \le N$ instead of the normal N caplets (floorlets). Furthermore, the caplets (floorlets) must be exercised when they are in the money at a settlement date. The value of an autocap can be written as

$$\mathrm{AC}_{1:N}^{l}(T_0) = \sum_{i=1}^{N}D(T_0,T_{i+1})\delta_i\mathbb{E}_0^{i+1}\left[(F_i(T_i)-\kappa)^+\right]\mathbf{I}\{E_i\}, \quad (15)$$

where

$$E_i = \left\{ \sum_{j=1}^{i-1} \mathbf{I}(F_j(T_j) > \kappa) \leq l - 1 \right\}.$$

A similar formula can be derived for autofloors. One should note that the $\left\{E_i\right\}_{i=1}^{N}$ not only depends on the current rate, but also on the previously observed spot Libor rates $\{F_n(T_n)\}_{n=1}^{i}$.

Discrete Barrier Cap

A *discrete barrier cap (floor)* is similar to a standard cap (floor) with the difference that the payoff of its underlying caplets (floorlets) are conditional on the event that previous spot Libor rates have or have not (depending on the type of barrier option) hit a certain level, known as the barrier. Though more complex variants exist, we consider discrete barrier caps where the barrier condition is only reviewed at fixing dates. The discrete barrier cap (floor) is characterized by a series of strikes $\{\kappa_i\}_{i=1}^{N}$ and barriers $\{H_i\}_{i=1}^{N}$. The value of an up-and-out discrete barrier cap is given by

$$\mathrm{DBC}^{\mathrm{uo}}(T_0) = \sum_{i=1}^{N} D(T_0, T_i)\delta_i \mathrm{E}_0^{i+1} \times \left[\left(F_i(T_i) - \kappa_i\right)^+ \prod_{n=1}^{i} \mathrm{I}(F_n(T_n) < H_n) \right] \quad (16)$$

From (16) we see that caplet i only pays out in case all previous spot Libor rates, $\{F_n(T_n)\}_{n=1}^{i}$, are below their barriers. The discounted payoff of a down-and-in barrier cap is given by

$$\mathrm{DBC}^{\mathrm{di}}(T_0) = \sum_{i=1}^{N} D(T_0, T_i)\delta_i \mathrm{E}_0^{i+1} \times \left[\left(F_i(T_i) - \kappa_i\right)^+ \left(1 - \prod_{n=1}^{i} \mathrm{I}(F_n(T_n) > H_n)\right) \right] \quad (17)$$

Thus, a down-and-in discrete barrier caplet i pays out if at least one of the previous spot Libor rates, $\{F_n(T_n)\}_{n=1}^{i}$,was set below its barrier. A portfolio of an up-and-in and an up-and-out barrier caplet equals an ordinary caplet. Combining the following features in/out, up/down, standard/digital, and cap/floor results in 16 different types of discrete (digital) barrier caps and floors.

Ratchet Cap

A *ratchet option* (also denoted as a one-way floater) is similar to a standard cap/floor with the difference that the strikes are variable over time and depend on settings of previous spot Libor rates. The value of a ratchet option is given by

$$\mathrm{Ra}(T_0) = \sum_{i=1}^{N} D(T_0, T_{i+1})\delta_i \mathbb{E}_0^{i+1}\left[\left(F_i(T_i) - \kappa_i\right)^+\right], \quad (18)$$

where κ_1 is given and for $i > 1$

$$\kappa_i = \kappa\left(\kappa_{i-1}, \ldots, \kappa_1; F_{i-1}(T_{i-1}), \ldots, F_1(T_1)\right). \quad (19)$$

In the sequel we consider a ratchet option with as strike the spot Libor at the previous setting

$$\kappa_i = F_{i-1}(T_{i-1}). \quad (20)$$

Furthermore, we consider a *sticky cap* where the strike equals the minimum of the spot Libor at the previous setting and the previous strike

$$\kappa_i = F_{i-1}(T_{i-1}) \wedge \kappa_{i-1}. \quad (21)$$

Both the ratchet and the sticky cap are expected to be very sensitive to the correlation matrix.

ESTIMATION RISK IN THE LMM

From the dynamics of the forward rates given in (3) and (4) combined with (6), it is clear that the prices of the exotic derivatives treated in the previous section are dependent on the volatility functions, $\{\sigma_i(t)\}_{i=1}^{N}$, and the correlation matrix, which we denote by ρ. If we have an analytical pricing formula, say f, which is a function of $\theta \equiv (\sigma_1, \ldots, \sigma_N, \rho_{1,2}, \ldots, \rho_{N-1,N})$, an application of the delta method (see, e.g., Van der Vaart [1998]) gives the standard errors of the option prices

$$\sqrt{T}\left(f(\hat{\theta}) - f(\theta)\right) \underset{\rightarrow}{d} \mathrm{N}\left(0, \frac{\partial f}{\partial \theta'} \Sigma_\theta \frac{\partial f}{\partial \theta}\right), \quad (22)$$

where Σ_θ denotes the covariance matrix of θ. However, for the exotic derivatives discussed in Section III, no analytical formula is available and prices need to be determined using numerical methods. In principle one could try to determine the standard errors by numerically computing the partial derivatives of f, but this does not look very attractive and it contains more information than needed.

A second and more promising method to get standard errors of the option prices is using the bootstrap (see Efron [1979] for the original work). Loosely speaking, the bootstrap theorem (under some regularity conditions) says that

$$\sqrt{T}\left(f(\hat{\theta}^{*})-f(\hat{\theta})\right)\underset{\rightarrow}{d} \mathrm{N}\left(0,\frac{\partial f}{\partial\theta'}\Sigma_{\theta}\frac{\partial f}{\partial\theta}\right), \tag{23}$$

where is $\hat{\theta}^{*}$ estimated on a bootstrap sample of the original data. For the ordinary bootstrap, a bootstrap sample is a sample drawn with replacement from the original sample, $\{X_1, \ldots, X_T\}$, say $\left\{X_1^*,\ldots,X_T^*\right\}$ with the original sample length.[2]

The problem with an ordinary bootstrap is that it does not take time dependence of the observations into account. A method that can take time dependence into account is the moving block bootstrap, where one draws blocks of observations instead of single observations. This method can generate time dependence in the data; however, the bootstrap samples resulting from a moving block bootstrap are not necessarily stationary even if the original data are. In case of LMM, the (return) data are generated from a stationary process.[3] An alternative method, capable of taking into account time dependence and stationarity is the stationary bootstrap (see Politis and Romano [1994]). This method can be described as follows. First, we randomly select a bootstrap observation X_1^* from the original T observations. Suppose that $X_1^* = X_j$. The second bootstrap observation is then equal to X_{j+1} with probability $1 - p$; otherwise, it is picked at random from the original T observations, where p is a given positive constant smaller than or equal to 1.[4] Thus, having the ith bootstrap observation $X_i^* = X_{i*}$, we take the next bootstrap observation according to

$$X_{i+1}^{*}=\begin{cases}X_{i^*+1} & \text{with prob. } 1-p,\\ \text{a random draw from } \left\{X_1,\ldots,X_T\right\}, & \text{with prob. } p,\end{cases} \tag{24}$$

for $i = 1, \ldots, T$.

By applying the stationary bootstrap to our data of forward rates, we can get B bootstrap estimators $\hat{\theta}^*$, where B denotes the number of bootstraps. For each bootstrap estimator $\hat{\theta}^*$ we can compute the prices of the exotic options using the Monte-Carlo simulation method resulting in an estimator for standard errors of the option prices. Using Equation (22) confidence intervals are easily constructed.

RESULTS

We use the U.S. forward curve from March 19, 1997 to May 14, 2003 on a weekly basis to estimate correlations. We investigate quarterly compounding cap products, as these are the most liquid for the U.S. market. First, we investigate which products are most sensitive to the number of factors used. In case of a K-factor model, we perform principal components analysis (PCA) on the historical covariance matrix and set all but the K largest eigenvalues equal to zero. From the resulting covariance matrix, we compute the correlation matrix and transform this into the LMM covariance matrix using the volatility term structure that can be calibrated to market prices of caps.[5] PCA can also be done on the correlation matrix (see, e.g., Kerkhof and Pelsser [2002] for an application to Libor market models and Basilevsky [1995] in general).

In Exhibit 1 we present the prices in basis points of the various products using a one-factor LMM with a

EXHIBIT 1
One-Factor Model Prices

maturity	Cap	Deferred cap	Autocap	Sticky cap	Ratchet cap	UI barrier cap
1Year	16.01	15.87	1.32	19.01	11.71	10.98
	(0.00)	(0.07)	(0.13)	(0.03)	(0.08)	(0.12)
2Years	50.79	49.28	4.29	67.96	26.70	20.67
	(0.00)	(0.31)	(0.42)	(0.12)	(0.35)	(0.44)
3Years	94.82	89.85	7.89	133.45	40.81	26.04
	(0.00)	(0.64)	(0.81)	(0.24)	(0.70)	(0.84)
4Years	145.22	134.15	12.11	210.32	54.52	29.49
	(0.00)	(1.00)	(1.27)	(0.37)	(1.11)	(1.30)
5Years	200.31	180.14	17.23	295.33	68.05	32.21
	(0.00)	(1.37)	(1.76)	(0.51)	(1.57)	(1.79)

This table gives the prices for the cap, deferred cap, autocap, sticky cap, ratchet cap, and up-and-in barrier cap for a one-factor Libor market model. The initial term structure is flat at 4% and the volatility term structure is flat at 20%. The strike for the deferred cap, autocap, sticky cap (initial strike), and ratchet cap (initial strike) equals 4%. The barrier for the up-and-out barrier equals 5%. The number of simulations equals 10,000 and the simulation standard errors are within brackets.

Exhibit 2
Factor Dependence

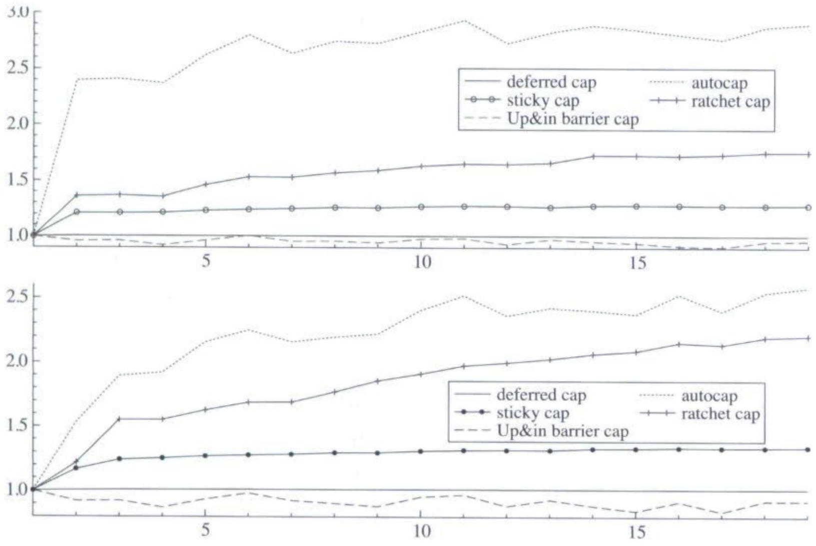

This figure gives the prices of the *K*-factor model relative to the one-factor model for the deferred cap, autocap, sticky cap, ratchet cap, and up-and-in discrete barrier cap. The term structure is assumed to be flat at 4% and the volatility term structure is assumed to be flat at 20%. All products are ATM, the barrier for the up-and-in barrier cap equals 5%, the start strike for the sticky cap and ratchet cap equals 4%, and the number of caps for the autocap equals half the number of settings. The upper panel gives a 3-year deal and the lower panel a 5-year deal. The number of simulations equals 10,000.

flat term structure at 4% and a flat volatility term structure at 20%. Besides prices simulation standard errors are also given. Cap prices are computed using the analytical Black formula and therefore do not contain simulation error.

In Exhibit 2 we investigate the influence of the number of factors used in fitting the correlation matrix. We see that the deferred cap is, as expected, hardly sensitive to correlation. The up-and-in barrier cap is somewhat sensitive to the number of factors in the correlation matrix; prices computed using more than one factor are between 90% and 100% of the one-factor price. We find that the sticky and ratchet cap, are rather sensitive to the number of factors, although in case of the sticky cap, the prices flatten out after factor 4 at about 1.3 times the one-factor price. The autocap is very sensitive to correlation matrix. Even for a small number of factors, prices are already up twice the one-factor price. However, for all products, one should be careful when looking at high factor models as part of the correlation in the correlation matrix is due to noise. It is quite likely that high factor models are overfitting noise instead of "true" correlation.

In Exhibit 3 we also investigate the factor dependence of the various products, this time with a hump-shaped volatility term structure given in Exhibit 4. We find that in this case all products except the ratchet are less sensitive to the number of factors used. Besides the ratchet, the autocap is most sensitive. However, in this situation the autocap is valued higher with the one-factor model than the more than one-factor models.

In order to compute the standard errors of the computed prices, we use the stationary bootstrap with $p = 0.1$ (expected block size equals 10), where p was determined using the automatic block size determination method of Politis and White [2003] on spot Libor rates. For each bootstrap sample we compute the LMM covariance matrix using PCA as described above. In Exhibits 5 and 6 we present the relative prices of the autocap, sticky cap, ratchet cap, and up-and-in barrier cap with their 95% confidence intervals. Results of the deferred cap are not shown as it has hardly any estimation risk. We clearly see that, in case one uses a model with more than one factor, estimation risk can be considerable. Furthermore, we see that for all but the up-and-in barrier cap, the confidence regions grow with the number of factors, but as expected the growth rate decreases with the number of factors. We see that for the up-and-in barrier option the price range computed using two factors more or less covers all prices computed by higher factor models and one might be inclined to conclude that two factors suffice in pricing the up-and-in barrier. However, risk management divisions are advised to set a model reserve larger than or equal to the worst case limit of the price range (depending on whether the bank is on the buy or sell side, this is the upper or lower confidence limit). For autocaps and sticky caps, the price range given by the five-factor model more or less covers the prices computed by the higher factor models. However, for the ratchet cap, it is not so clear and risk management divisions are advised to set considerable model reserves.

EXHIBIT 3

Factor Dependence Hump-Shaped Volatility

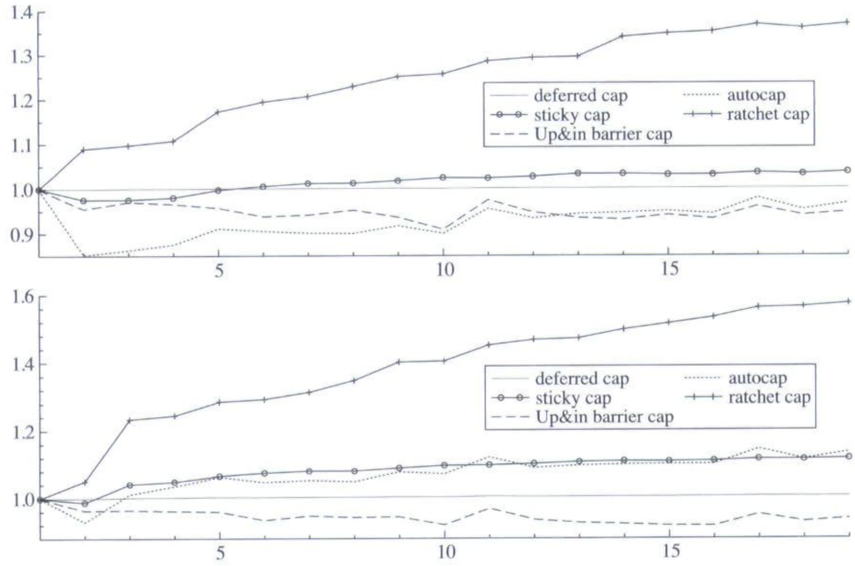

This figure gives the prices of the K-factor model (number of factors on the horizontal axis) relative to the one-factor model for the deferred cap, autocap, sticky cap, ratchet cap, and up-and-in discrete barrier cap. The term structure is assumed to be flat at 4% and the volatility term structure is assumed to be hump shaped. All products are ATM, the barrier for the up-and-in barrier cap equals 5%, the start strike for the sticky cap and ratchet cap equals 4%, and the number of caps for the autocap equals half the number of settings. The upper panel gives a 3-year deal and the lower panel a 5-year deal. The number of simulations equals 10,000.

EXHIBIT 4

Hump-Shaped Volatility Term Structure

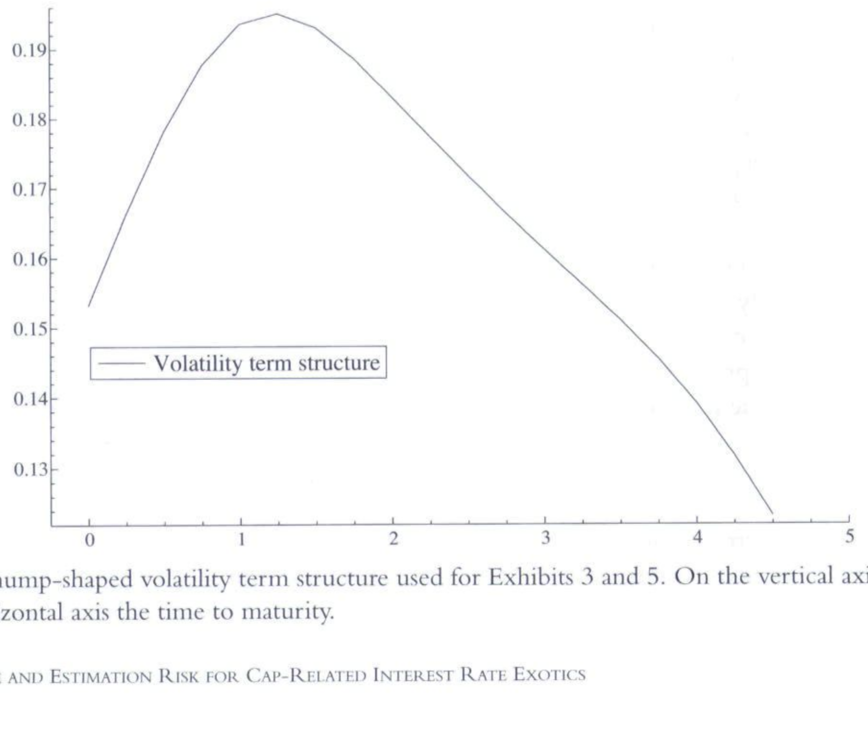

This figure presents the hump-shaped volatility term structure used for Exhibits 3 and 5. On the vertical axis is the level of caplet volatility and on the horizontal axis the time to maturity.

EXHIBIT 5
Estimation Risk

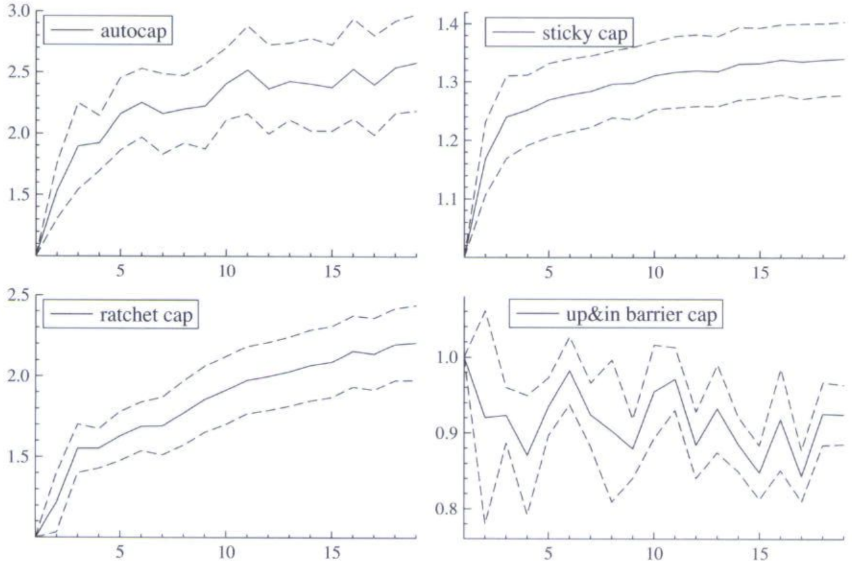

This figure gives the prices of the K-factor model (number of factors on the horizontal axis) relative to the one-factor model including the 95% confidence regions. The term structure is assumed to be flat at 4% and the volatility term structure is assumed to be flat at 20%. All products are ATM, the barrier for the up-and-in barrier cap equals 5%, the start strike for the sticky cap and ratchet cap equals 4%, and the number of caps for the autocap equals half the number of settings. The upper panels present the results of a 3-year deal and the lower panels present the results of a 5-year deal. The number of simulations equals 10,000.

CONCLUSIONS

In this article we presented the (stationary) bootstrap as a method to take estimation risk into account when computing exotic interest rate derivatives prices. This provides traders and risk managers with a price range of values for the exotic product contrary to a single estimate. Based on these price ranges risk managers can set model reserves for trading desks. Traders themselves get an idea of the size of bid-ask spread they should provide. We find that this estimation risk is very much product dependent. Autocaps, sticky caps, and especially ratchet caps are very sensitive to correlation resulting in considerable standard errors for these products. For some products risk managers are advised to set reserves up to the price of the product.

ACKNOWLEDGMENT

I would like to thank Bertrand Melenberg and Hans Schumacher for their valuable comments. Any remaining errors are mine.

REFERENCES

Basilevsky, A. *Statistical Factor Analysis and Related Methods: Theory and Applications*. New York: John Wiley & Sons, Inc., 1995.

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

Brace, A., Gatarek, D., and Musiela, M. "The Market Model of Interest Rate Dynamics." *Mathematical Finance*, 7 (1997), pp. 127-155.

Brigo, D., and Mercurio, F. *Interest Rate Models: Theory and Practice*. Berlin, Heidelberg, New York: Springer-Verlag, 2001.

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

Efron, B. "Bootstrap Methods: Another Look at the Jackknife." *Annals of Statistics*, 7 (1979), pp. 1-26.

EXHIBIT 6
Estimation Risk Hump Shaped Volatility

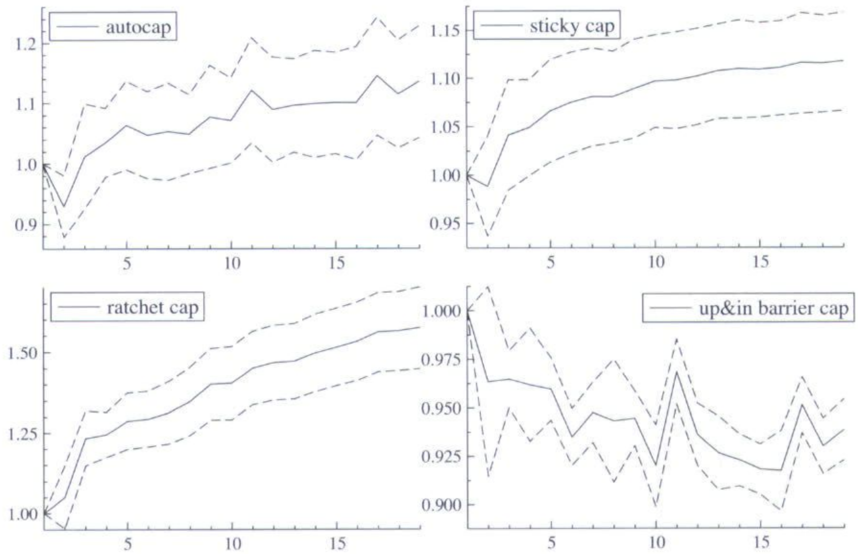

This figure gives the prices of the *K*-factor model (number of factors on the horizontal axis) relative to the one-factor model including the 95% confidence regions. The term structure is assumed to be flat at 4% and the volatility term structure is assumed to be hump shaped. All products are ATM, the barrier for the up-and-in barrier cap equals 5%, the start strike for the sticky cap and ratchet cap equals 4%, and the number of caps for the autocap equals half the number of settings. The upper panels present the results of a 3-year deal and the lower panels present the results of a 5-year deal. The number of simulations equals 10,000.

Hull, J., and Suo, W. "A Methodology for Assessing Model Risk and its Application to the Implied Volatility Function Model." *Journal of Financial and Quantitative Analysis*, 37 (2002), pp. 297-318.

Hunt, P., and Kennedy, J. *Financial Derivatives in Theory and Practice*. Chichester: John Wiley & Sons, 2000.

Hunter, C., Jäckel, P., and Joshi, M.S. "Getting the Drift." *Risk*, (2001), pp. 81-84.

Jameson, A., and Gadanecz, B. OTC Derivatives Market Activity in the Second Half of 2002, Bank for International Settlements, Basel, *Technical report.*, 2003

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

Jeanneau, S. Derivative Markets. *BIS Quarterly Review*, (2002) pp. 35-42.

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

Kloeden, P.E., and Platen, E. *Numerical Solution of Stochastic Differential Equations*, 3rd ed. Berlin: Springer Verlag, 1999.

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

Longstaff, F., Santa-Clara, P., and Schwartz, E. "Throwing Away a Billion Dollars: The Cost of Suboptimal Exercise Strategies in the Swaptions Market." *Journal of Financial Economics*, 62 (2001b), pp. 39-66.

Miltersen, K. R., Sandmann, K., and Sondermann, D. "Closed Form Solutions for Term Structure Derivatives with Lognormal Interest Rates." *Journal of Finance*, 52 (1997), pp. 409-430.

Miron, P., and Swannell, P. *Pricing and Hedging Swaps*. London: Euromoney Books, 1992.

Pelsser, A. *Efficient Methods for Valuing Interest Rate Derivatives*, Berlin: Springer Verlag, 2000.

Politis, D.N., and Romano, J.P. "The Stationary Bootstrap." *Journal of the American Statistical Association*, 89 (1994), pp. 1303-1313.

Politis, D.N., and White, H. "Automatic Block-Length Selection for the Dependent Bootstrap." Forthcoming in *Econometric Reviews* (2003).

Politis, D., Romano, J., and Wolf, M. *Subsampling*. New York: Springer-Verlag, 1999.

Rebonato, R. *Modern Pricing of Interest-Rate Derivatives: The Libor Market Model and Beyond*, NJ: Princeton University Press, 2002.

Van der Vaart, A.W. *Asymptotic Statistics*. Cambridge University Press, 1998.

ENDNOTES

[1]The same methodology can also be applied to the more general class of HJM models, but we have selected the Libor market model for convenience and its popularity.

[2]One can also draw samples smaller than the original sample, in which case we speak of subsampling (see Politis et al. [1999] for an excellent treatment).

[3]Strictly speaking this only holds for a forward rate under its own forward martingale measure due to the stochastic drift under a different measure, but this effect is minor.

[4]Of course, $p = 1$ gives the ordinary bootstrap.

[5]Historical and implied volatility functions can therefore differ.