BOUNDING BERMUDAN SWAPTIONS IN A SWAP-RATE MARKET MODEL

MARK S. JOSHI AND JOCHEN THEIS

Abstract. We develop a new method for finding upper bounds for Bermudan swaptions in a swap-rate market model. By comparing with lower bounds found by exercise boundary parametrization, we find that the bounds are well within bid-offer spread. As an application, we study the dependence of Bermudan swaption prices on the number of instantaneous factors used in the model. We also establish an equivalence with LIBOR market models and show that virtually identical lower bounds for Bermudan swaptions are obtained.

1. Introduction

The pricing of Bermudan swaptions under market models is a long-standing tricky problem. As the drifts of the rates are state-dependent and the volatilities are typically time-dependent, the only feasible pricing method is by Monte Carlo simulation. However, to price an option with early exercise opportunities by Monte Carlo one needs to know the exercise strategy which is tightly bound up with knowing the price one wishes to compute.

As the price of a Bermudan swaption is the supremum of the prices over all exercise strategies (stopping times), a lower bound can always be found by picking some exercise strategy. More generally, one can optimize over a class of exercise strategies to find a good lower bound. Such approaches have been developed by Anderson, [1], and Jäckel, [6] in the context of LIBOR market models. Jäckel shows that in certain cases where comparison with a non-recombining tree is possible that his method is very effective but for the general case the comparison is not feasible. However, in the absence of a good upper bound, one can never be sure how good these lower bounds are in general. Here we develop a method for upper bounds in the context of a swap-rate market model which gives upper bounds within a fraction of a vega

Date: February 7, 2002.

of the lower bound found by adapting Jäckel's method to swap-rate models. Thus we can be sure that both the lower bound and upper bound are tight.

We proceed by adapting a method proposed by Rogers, [15], as well as Haugh and Kogan, [5], in the context of equity and FX options. Taking B_t as numeraire, they make the observation that as the price of an option, $D(0)$, is equal to

$$B_0 \sup_{\tau} \mathbb{E}(B_\tau^{-1} D_\tau^{'}), \tag{1.1}$$

where $D_\tau^{'}$ indicates the exercise value of the Bermudan at time τ, and the supremum is taken over all stopping times, τ, the price can only be increased by taking a supremum over all random times. However, if we allow all random times then there is a clear winner: exercise with maximal foresight. Thus we have the upper bound

$$B_0 \mathbb{E}(\max_t B_t^{-1} D_t), \tag{1.2}$$

where the max is taken over the exercise dates of the Bermudan.

This upper bound is, of course, too crude to be useful. However, the same argument holds if a martingale of initial value zero is subtracted from the portfolio. Thus if we take any portfolio, P, of derivatives of initial value zero and consider

$$B_0 \mathbb{E}(\max_t B_t^{-1}(D_t - P_t)), \tag{1.3}$$

we also have an upper bound.

Thus we can optimize over the possible portfolios, P_t, to obtain upper bounds. Rogers shows the existence of a portfolio, P_t, which attains the price of the option, but the proof is non-constructive. His argument is not easily adapted to practical pricing, but it does suggest a way to proceed. Consider a class of portfolios, P_t^α, indexed by α and then optimize over possible values of α to obtain a best upper bound.

It is important to realize that the portfolio can be a dynamic trading strategy. Our solution is to use a weighted sum of European swaptions each one associated to an exercise date of the Bermudan swaption, together with a short position in zero coupon bonds to ensure that the initial value is zero. At the time of expiry of the European swaption, we assume that it is cash settled and the money is used to buy a European swaption of the next shortest maturity. The parameters we optimize over are the notionals of the European swaptions in the initial portfolio.

A crucial part of the procedure is that we need to know the values of the European swaptions at all exercise times, this means that we need to be able to price them before their expiries. It is here that the use of a swap-rate market model is preferable to a LIBOR market model, as the European swaptions can be priced instantly via the Black formula. Whilst one could approximate the prices in a LIBOR market model, using an equivalent swaption volatility formula, see for example [12], [13] or [7], this results in additional approximation errors and is more time consuming.

An advantage of optimizing over notionals, rather than other parameters is that the dependence the value of the portfolio at a given time and yield curve as a function of the notionals is purely linear and therefore easily recomputed as the values change during an optimization.

Our procedure is therefore as follows.

(1) Generate a set of training paths.
(2) For this set of training paths, optimize over the notionals to obtain a best upper bound.
(3) Generate a second set of paths with independent variates and use this to estimate the expectation, (1.3).

We discuss the implementation details in greater depth in Section 2. We stress that we use an independent set of paths for the estimation of the expectation in order to avoid the possibility that biasing may arise from the optimization procedure exploiting the precise structure of the sample drawn. We can therefore be confident that the upper bound found from step 3 is accurate up to the convergence of that Monte Carlo simulation.

We found that in all cases the upper bound could be refined to be within a vega of the lower bound. For particularly humped yield curves, the upper bound was least effective. However, we found that in such cases the upper bound could be greatly improved by using larger portfolios of European swaptions. In particular, we included additional European swaptions with strikes ten percent above and below the strike of the Bermudan. We present results in Section 6.

One consequence of our results is therefore that we can be confident that the Jäckel exercise strategy is sufficiently accurate to be used for pricing without the worry that it is failing to capture a lot of value.

As one application of our results, we study the dependence of Bermudan prices on the number of instantaneous factors driving the swap-rates. We find in general a small but marked price increase as the number of factors increases from one to two, but only slight increases thereafter. The difference between the lower bound for one factor and the upper bound for a full factor model was in general less than half a vega and therefore would lie well within bid-offer spread. We present these results in Section 6. We stress that these results reflect changes in the number of instantaneous factors and therefore only reflect changes in instantaneous correlations, and do not therefore affect decorrelation which arises from the differing shapes of the volatility curves for swaptions. We discuss the issues of calibration and the meaning of factors further in Section 3. We also discuss the differences between our results and those of [2] and [10].

The final issue we consider is the similarity in prices of Bermudan swaptions under LIBOR and swap-rate based market models. We show that the lower bound estimation procedure leads to very similar prices under the two models provided equivalent covariance structures are used. We discuss how to make the covariance structures equivalent in Section 3 and present the numerical results in Section 6.

Whilst we believe our results are interesting, we still feel that there remains the issue of how to construct the comparison martingale in an intuitive fashion rather than just plugging the notionals into an optimizer and seeing what comes out.

An approach to the estimation of Bermudan swaption prices with similar theorerical underpinnings but quite different practicalities has been previously introduced by Anderson and Broadie, [3]; their approach involves running additional Monte Carlo simulations within the main Monte Carlo simulation but does not require an optimization procedure. For a full discussion of the history of this problem we refer the reader to [3].

2. Implementation and notation

Let us fix some notation. We study a Bermudan swaption with strike K associated to times

$$t_0 < t_1 < \cdots < t_n.$$

We denote by f_j the forward rate for a deposit running from t_j to t_{j+1}. We denote by SR_j the swap rate for the times $t_j, \dots, t_n$ and by B_j the price of a zero-coupon bond expiring at time t_j.

In a swap-rate market model, we take the rates f_j to be log-normally distributed in the real-world measure with a possibly time-dependent volatility. When we take a bond, B_j, as numeraire and pass to the pricing measure, the swap rate has a state-dependent drift and the volatility does not change, see [8]. The drift involves the swap rates and the instantaneous covariances between rates – i.e. the instantaneous correlation times the product of volatilities.

The fact that the drifts are state-dependent complicates the implementation of either model as a Monte Carlo simulation. We employ the techniques of [4] to allow us to step the rates over several years at once.

Let g_j denote either f_j or SR_j . We suppose that we are given the covariance matrix, C^l, of $\log g_j$ over each period of evolution from t_{l-1} to t_l. (Take $t_{-1} = 0$.) We discuss the provenance of this covariance matrix in section 3.

Let A^l denote a pseudo-square root of C^l. If the rates $\log g_j$ had constant drift, μ_j, then we could simulate their evolution precisely via

$$\log g_j(t_l) = \log g_j(t_{l-1}) + \mu_j(t_l - t_{l-1}) + \sum_{i=0}^{n-1} A^l_{ji} Z_j$$

with Z_j a vector of independent $N(0,1)$ draws. We approximate the evolution of g_j via the use of a predictor-corrector method. We first compute an approximation to the drift across the time step by substituting the covariance elements across the time step for the instantaneous covariance elements, and using the values of the rates at the start of the time step.

This gives an initial guess for the terminal values of the rates. Using this initial guess, we then recompute the drift at the end of the step using the same covariance elements. The average of the two drifts is then our best guess for the drift and we re-evolve using the original $N(0,1)$ draws. This predictor-corrector method is shown to be highly accurate for the LIBOR market model in [4] and works equally well for the swap-rate model.

In order to carry out our optimization, we therefore generate a set of training paths. We take B_{n-1} as numeraire. For this set of training paths, we store all the necessary information for computing the maximum value along each path for any set of weights. In particular, we stored the ratios of the values of each swaption at each time with the numeraire and the ratio of the exercised value of the Bermudan with the numeraire.

For any set of weights, it is then possible to rapidly get an estimate of the upper bound with no new path generation, simply by computing using the existing set of paths and stored values. Note that this would not be possible if we optimized over strikes or other parameters. For example, we could use trigger swaps and optimize over the trigger level but we would then need to reprice the trigger swap over every path, which would require repeated calls to the Black formula and cause a great decrease in speed. We typically used around 16384 paths.

Once we have the upper bound as a function of the notionals, we optimize to get the lowest upper bound. We did so by employing the simplex method as detailed in [11].

Once the optimal parameters for the set of training paths had been found, a second simulation was run to evaluate the expectation. The second simulation was running using different variates in order to avoid the possibility of biasing arising from the optimization exploiting any inaccuracies in the approximate upper bound function. We typically ran 2 to the power 18 paths to be sure that the simulation was well converged.

Note that any inaccuracies in the upper bound approximation used for optimization would not affect the status of the final upper bound as an upper bound. Any such inaccuracies may lead to a worse upper bound however. Any inaccuracy in the final upper bound only arises to the extent that the final simulation has not converged.

In order, to maximize the convergence rate of our simulations, we worked with high-dimensional Sobol numbers combined with Brownian bridging techniques.

3. Calibration

One of the trickiest aspects of working with market models is their calibration. Whilst calibrating the LIBOR market model to caplet prices is trivial and immediate, and for the swap-rate market model calibration to swaption prices is immediate, there is the basic problem that calibration is too easy in that one can find many calibrations to the same prices and needs extra information to fix the calibration.

We first recall some standard techniques for calibrating the LIBOR market model, see for example [13]. If we allow the caplet prices to be a function of time, then we have for a forward rate expiring at time T that

$$df_T = \mu_T dt + \sigma_T(t) f_T dW_t.$$

In order to calibrate to the caplet market we need

$$\int_0^T \sigma_T(t)dt = \hat{\sigma}_T^2 T \tag{3.1}$$

where $\hat{\sigma}_T$ is the caplet implied volatility.

There are clearly many such choices of σ_T. One solution, which we use, is to require that

$$\sigma_T(t) = \sigma(T-t) \tag{3.2}$$

with σ a function independent of t. This means that every forward rate has the same volatility as a function of the amount of time to its own maturity.

Following Rebonato, [12, 13], we use a functional form

$$\sigma(\tau) = ((a+b\tau)e^{-c\tau} + d)H(\tau), \tag{3.3}$$

where $H(\tau)$ is one for $\tau \geq 0$ and zero otherwise. The volatility function for f_j can then be adjusted by a constant multiplicative factor K_j to ensure that (3.1) is satisfied exactly. The cut-off H ensures that each forward rate stops moving after its own expiry.

To run a simulation, we also need the instantaneous correlations between forward rates, ρ_{ij}. We take the instantaneous correlation matrix to be of the form

$$\rho_{ij} = e^{-\beta|t_i - t_j|}.$$

The covariance between the logs of f_i and f_j over the period $[s,t]$ is then

$$\int_s^t \rho_{ij}\sigma_{t_i}(r)\sigma_{t_j}(r)dr.$$

We can thus compute the covariance matrices and run our simulation.

However, we wish to calibrate to the prices of the underlying European swaptions which are essentially options on the rates SR_j. We also wish to calibrate our swap-rate based model. We proceed by using an equivalence between swap-rate models and forward-rate models.

By writing a swap-rate as a function of the underlying forward rates, we can write

$$d\,\mathrm{SR}_i = \mu_i dt + \sum_{j=i}^{n-1} \frac{\partial\,\mathrm{SR}_i}{\partial f_j} df_j. \tag{3.4}$$

It therefore follows that

$$d\log\mathrm{SR}_i = \tilde{\mu}_i dt + \sum_{j=i}^{n-1} \frac{\partial\,\mathrm{SR}_i}{\partial f_j}\frac{f_j}{\mathrm{SR}_i} d\log f_j. \tag{3.5}$$

Thus if we let

$$Z_{ij} = \frac{\partial\,\mathrm{SR}_i}{\partial f_j}\frac{f_j}{\mathrm{SR}_i}, \tag{3.6}$$

for $i \leq j$ and zero otherwise, we can write in vector terms, ignoring drifts,

$$d\log\mathrm{SR} = Z d\log f. \tag{3.7}$$

This means that if $C_f(s,t)$ is the forward-rate covariance matrix across a period $[s,t]$, we can approximate the swap-rate covariance matrix across this interval by

$$C_{\mathrm{SR}} = Z(0)C_f(s,t)Z(0)^t. \tag{3.8}$$

Alternatively, if we wish to prescribe the swap-rate covariance matrix we can invert (3.8).

When pricing a Bermudan swaption, we generally wish to prescribe the swap-rate variances so as to ensure the exact pricing of the underlying European swaptions. However, it is difficult to get a handle on the time-dependence of the European swaptions and their covariances. We therefore adopt a compromise in which we allow the caplets to determine the correlation structure and the swaptions to determine the variances.

We therefore obtain a first guess, $C_{f,1}$ for C_f by calibration to the caplets. This implies a first guess, $C_{\mathrm{SR},1}$, for the swap-rate covariance matrix.

Let the desired variance for SR_j over $[0,t_j]$ be V_j. We want $C_{\mathrm{SR}\,1}(0,t_j)$ to have V_j in the jj entry. Let

$$\lambda_j = \sqrt{\frac{V_j}{C_{\mathrm{SR}\,1}(0,t_j)_{jj}}}. \tag{3.9}$$

Now let Λ be the diagonal matrix with Λ_{jj} equal to λ. The matrix

$$C_{\text{SR}}(0, t_j) = \Lambda C_{\text{SR}\,1}(0, t_j)\Lambda$$

now has the required variance on the diagonal for each value of j.

We therefore define

$$C_{\text{SR}}(s, t) = \Lambda Z(0) C_{f1}(s, t) Z(0)^t \Lambda, \tag{3.10}$$

and

$$C_f(s, t) = Z(0)^{-1} \Lambda Z(0) C_{f1}(s, t) Z(0)^t \Lambda (Z(0)^t)^{-1}. \tag{3.11}$$

We have to modify (3.10) when $t > t_0$, to ensure that the swap-rates do not change after their own expiry. If

$$t_j \leq s \leq t \leq t_{j+1}, \tag{3.12}$$

then we use (3.10) but zero the rows and columns pertaining to rates that have already reset, that is columns 0 to j. For the general case, we break the covariance matrix into a sum of individual matrices for which (3.12) hold.

The method we have given for calibrating the LIBOR market model is essentially that of [7], see [13] for further discussion. In [7] it is shown to be highly effective for calibrating to swaption prices. To use this technique for calibrating swap-rate market models would appear to be new. One of our principal results numerically demonstrated in Section 6, is that a LIBOR market model and swap-rate market model with these calibrations yield the same results when developing lower bounds for Bermudan swaptions.

4. Factor reduction

In Section 3, we used an instantaneous correlation matrix for the forward rates of the form

$$\rho_{ij} = e^{-\beta|t_i - t_j|}. \tag{4.1}$$

This leads to a full factor model for the rates evolution. However, many front offices use a two or three factor model. It is important to realize that these models are short-stepped so only decorrelation coming from changes in the instantaneous correlation matrix are affected by this factor reduction; the terminal decorrelation effects arising from the shape of the instantaneous volatility curves will not be affected.

We therefore study how changing the rank of the instantaneous correlation matrix affects the price of a Bermudan swaption. Note that if one carried out factor reduction by trimming the long stepped covariance matrices, then the results would be affected by the length of the steps, and introducing extra steps would change the prices.

Our technique for reducing the factors is to cut-off the lower eigenvalues. In particular, we diagonalize the correlation matrix to get the eigenvectors λ_j and associated eigenvectors e_j. To get a rank r matrix we take the matrix A_r such that column j is $\sqrt{\lambda_j}$ for $j < r$ and zero otherwise. We then form $B = A_r A_r^t$ which is a covariance matrix but not a correlation matrix as the diagonal elements are not equal to one. We therefore take the correlation matrix, C, to be given by

$$C_{ij} = \frac{B_{ij}}{\sqrt{B_{ii} B_{jj}}}. \tag{4.2}$$

In Section 6, we give numerical results for the upper and lower bounds as a function of the number of factors. Qualitatively, we find that the transition from one to two factors gives a small but clear increase in price. Our lower bound for the two-factor model is generally around the same level as the upper bound for the one-factor model. For increasing factors beyond that we see slight but not insignificant improvements.

It is interesting to note that in all our tests the difference between the lower bound for the one factor model and the upper bound for the ten factor model is generally less than a vega, and the impact of factor dependence is therefore insubstantial compared to the size of bid-offer spreads.

The issue of factor dependence has previously been studied by Andersen and Andreason, [2], and by Longstaff, Santa-Clara and Schwartz, [10]. The thesis of the latter paper was that banks are throwing away large sums of money by pricing and hedging Bermudan swaptions using low-factor models. The former paper argued that the latter paper was mistaken and that, in fact, a two-factor model gives lower prices than a one factor model for Bermudan swaptions.

Ultimately, the answer to the question lies in what one calibrates to, and in how one defines a factor. If one takes a model in which all volatiltiies are flat and the forward rates in a LIBOR market model are driven by a single factor then the model is effectively a BDT type model and one obtains a lower price, see [9] or [13]. The lower number of factors in this case is not just smaller in the instantaneous sense but also in the sense that the covariance matrix across long time steps is of rank one and in this case, the model is certainly failing to capture certain features of the market.

Andersen and Andreasen calibrate their models to both swaptions and caplets simultaneously and achieve the result that the price decreases as the number of factors increases, whereas we use the caplets only to infer the correlation structure and achieve the result that the price increases. The main issue is therefore whether it is appropriate to calibrate the model both caps and swaptions or whether one should place more emphasis on having the correct correlation structure. The issue is really more financial than mathematical. Rebonato has argued in [14] and [13] that the simultaneous calibration is not appropriate and we refer the reader to his work.

5. The lower bound

In this section, we recall the results of Jäckel, [6], and discuss their implementation in the context of a swap-rate market model. Jäckel suggested using an exercise strategy based on the next forward rate and the swap-rate running from the end of the forward rate to the final time. Thus at time t_j, we examine the levels of f_j and SR_{j+1} in order to determine an exercise strategy.

In particular, Jäckel suggested exercising at time t_i according to whether

$$f_i(t_i) - \left(p_{i1} \frac{\mathrm{SR}_{i+1}(0)}{\mathrm{SR}_{i+1}(t_i) + p_{i2}} + p_{i3} \right),$$

is positive for payer's swaptions and negative for receiver's, and p_{ij} are the parameters to be optimized over. The additional constraint of never exercising out of the money is also added.

As with the upper bound, one first develops a set of training paths and then optimizes over the parameters to obtain a best lower bound. Once this is done one reprices using a second set of paths to avoid biasing. One can then be sure that the lower bound is accurate up to the level of convergence of the second Monte Carlo. We refer the reader to [6] for further details. We used a simplex type algorithm for the optimization, [11].

6. Numerical Results

In this section, we present numerical results and graphs. For each test, we present the values of a, b, c, and d used to generate the instantaneous volatility structure, the European swaption volatilities calibrated to, the value of β and the yield curve. The "K" factors are always taken to be one. We also given the strike and type of the Bermudan swaption and the reset times.

The yield curve is given via a functional form so that the forward rate from year k to year $k+1$ is equal to

$$(A + Bk)e^{-Ck} + D.$$

This is used to fix the discount factor after all integer numbers of years. Other discount factors are found by log-linear interpolation. We adopt this approach as it allows easy specification and communication of a large class of plausible yield curves.

In each case, we present the lower bounds for swap-rate and LIBOR market models, the upper bound for the swap-rate model and the lower bound for the swap-rate model with volatilities increased by one percent. In all cases, we used 16384 training paths and 131072 pricing paths using low-discrepancy numbers and Brownian bridging techniques. This ensured that the final Monte Carlo simulations were converged to within a fraction of a basis point.

For our first test, we took ten yearly rates starting in five years with the following parameters

				Swaption Vols
A	-1%	a	5%	11.83%
B	0%	b	10%	11.50%
C	30%	c	50%	11.13%
D	6%	d	10%	10.80%
		beta	0.1	10.55%
PayReceive	pay			10.39%
Strike	0.059077858			10.30%
				10.28%
				10.32%
				10.45%

Factors	Swap lower bound	Upper bound	Bumped vol	LIBOR lower bound
1	0.0461	0.0467	0.0499	0.0462
2	0.0475	0.0485	0.0514	0.0476
3	0.0477	0.0490	0.0516	0.0479
4	0.0479	0.0492	0.0518	0.0480
5	0.0479	0.0493	0.0518	0.0480
6	0.0479	0.0494	0.0519	0.0481
7	0.0480	0.0494	0.0519	0.0481
8	0.0480	0.0495	0.0519	0.0481
9	0.0480	0.0495	0.0520	0.0481
10	0.0480	0.0495	0.0519	0.0481

For our second test, we took ten yearly rates starting in five years with the following parameters

				Swaption Vols
A	0%	a	5%	11.88%
B	0%	b	10%	11.53%
C	30%	c	50%	11.15%
D	6%	d	10%	10.81%
		beta	0.1	10.56%
PayReceive	pay			10.39%
Strike	0.06			10.30%
				10.28%
				10.32%
				10.45%

Factors	Swap lower bound	Upper bound	Bumped vol	LIBOR lower bound
1	0.0448	0.0454	0.0485	0.0450
2	0.0462	0.0471	0.0500	0.0462
3	0.0463	0.0476	0.0502	0.0465
4	0.0465	0.0478	0.0504	0.0466
5	0.0465	0.0479	0.0504	0.0467
6	0.0466	0.0480	0.0505	0.0467
7	0.0466	0.0480	0.0505	0.0467
8	0.0466	0.0480	0.0505	0.0467
9	0.0466	0.0481	0.0505	0.0467
10	0.0466	0.0481	0.0505	0.0468

For our third test, we took ten yearly rates starting in five years with the following parameters

				Swaption Vols
A	0%	a	0%	11.88%
B	0%	b	0%	11.53%
C	30%	c	50%	11.15%
D	6%	d	10%	10.81%
		beta	0.1	10.56%
PayReceive	pay			10.39%
Strike	6.00%			10.30%
				10.28%
				10.32%
				10.45%

Factors	Swap lower bound	Upper bound	Bumped vol	LIBOR lower bound
1	0.0309	0.0311	0.0344	0.0309
2	0.0319	0.0328	0.0356	0.0319
3	0.0321	0.0332	0.0359	0.0321
4	0.0322	0.0333	0.0360	0.0322
5	0.0323	0.0334	0.0361	0.0322
6	0.0323	0.0335	0.0361	0.0323
7	0.0323	0.0335	0.0361	0.0323
8	0.0324	0.0335	0.0362	0.0323
9	0.0324	0.0335	0.0362	0.0323
10	0.0324	0.0335	0.0362	0.0323

For our fourth test, we took the same parameters as for the third test but took the rates to be half yearly.

Swaption Vols
9.24%
9.31%
9.39%
9.46%
9.54%
9.62%
9.70%
9.78%
9.88%
10.00

Factors	Swap lower bound	Upper bound	Bumped vol	LIBOR lower bound
1	0.0162	0.0163	0.0181	0.0162
2	0.0165	0.0168	0.0184	0.0165
3	0.0166	0.0169	0.0185	0.0166
4	0.0166	0.0169	0.0185	0.0166
5	0.0166	0.0170	0.0185	0.0166
6	0.0166	0.0170	0.0184	0.0166
7	0.0166	0.0170	0.0185	0.0166
8	0.0166	0.0170	0.0186	0.0166
9	0.0166	0.0170	0.0185	0.0166
10	0.0166	0.0170	0.0186	0.0166

For our fifth test, we took ten half year rates starting in five years

				Swaption Vols
A	-1%	a	-6%	15.62%
B	2%	b	4%	15.82%
C	30%	c	50%	15.97%
D	6%	d	16%	16.09%
		beta	0.1	16.19%
PayReceive	pay			16.28%
Strike	7.22%			16.36%
				16.45%
				16.55%
				16.71%

Factors	Swap lower bound	Upper bound	Bumped vol	LIBOR lower bound
1	0.0357	0.0360	0.0378	0.0357
2	0.0363	0.0368	0.0385	0.0363
3	0.0362	0.0370	0.0383	0.0364
4	0.0365	0.0371	0.0386	0.0364
5	0.0365	0.0371	0.0386	0.0364
6	0.0365	0.0371	0.0387	0.0364
7	0.0363	0.0371	0.0384	0.0364
8	0.0365	0.0372	0.0387	0.0363
9	0.0366	0.0372	0.0387	0.0365
10	0.0366	0.0372	0.0387	0.0364

Note the slight noisiness in the lower bounds here. As the lower bounds results from an optimization procedure, it will sometimes stop at a local rather than global minimum.

For our sixth test, we took ten half year rates starting in half a year.

				Swaption Vols
A	-1%	a	-6%	14.83%
B	2%	b	4%	15.37%
C	30%	c	50%	15.83%
D	6%	d	16%	16.18%
		beta	0.1	16.44%
PayReceive	pay			16.65%
Strike	7.22%			16.80%
				16.93%
				17.05%
				17.21%

We include results on the upper bound with three swaptions for each expiry as well the results for one swaption.

Factors	Swap lower bound	Upper bound	Bumped vol	LIBOR lower bound	Improved Upper
1	0.0305	0.0311	0.0320	0.0305	0.0306
2	0.0308	0.0316	0.0323	0.0308	0.0312
3	0.0309	0.0317	0.0324	0.0309	0.0314
4	0.0310	0.0318	0.0325	0.0309	0.0315
5	0.0310	0.0319	0.0324	0.0309	0.0315
6	0.0310	0.0319	0.0325	0.0309	0.0316
7	0.0311	0.0319	0.0326	0.0310	0.0316
8	0.0310	0.0319	0.0325	0.0310	0.0316
9	0.0311	0.0319	0.0326	0.0310	0.0316
10	0.0311	0.0319	0.0326	0.0310	0.0316

7. Control variates

The upper bound was obtained by constructing a hedging portfolio for the portfolio, if this portfolio is a good hedge then we can expect the portfolio consisting of the difference to have much lower variance than the original product. This means that if we price the difference portfolio according to Monte Carlo then we can expect the simulation to converge much faster than for the original product. In other words, we can use the trading strategy in European swaptions as a control variate for the Monte Carlo simulation.

Whilst the fact that it is generally slower to find the upper bound than the lower bound means that this will not buy us much when trying to price the product, it does mean that if one wishes to run many different Monte Carlo simulations, for example in order to compute Greeks, then the technique becomes worthwhile. As we are pricing using low-discrepancy numbers and using Brownian bridge techniques we present a convergence table rather than standard error numbers as it is not clear what a standard error estimate means when using low discrepancy numbers. We do not address the issue of whether the use of a control variate could speed up estimation of the exercise boundary, but only look at the convergence of the price given an exercise strategy.

We present data in the same format as the previous section. We give one example. There are ten half yearly rates starting in half a year.

				Swaption Vols
A	-1%	a	0%	11.34%
B	0%	b	4%	11.53%
C	30%	c	50%	11.69%
D	6%	d	10%	11.83%
		beta	0.1	11.95%
PayReceive	pay			12.05%
Strike	5.91%			12.13%
				12.20%
				12.24%
				12.30%

Paths	Unhedged price	Hedged Price
1048576	0.00919	0.00920
524288	0.00920	0.00920
262144	0.00920	0.00919
131072	0.00920	0.00919
65536	0.00921	0.00919
32768	0.00925	0.00921
16384	0.00924	0.00922
8192	0.00919	0.00922
4096	0.00922	0.00925
2048	0.00931	0.00926
1024	0.00927	0.00920
512	0.00927	0.00921
256	0.00889	0.00910

The price of the Bermudan is correct to within a basis point after only 256 paths.

References

[1] L. Anderson, A simple approach to the pricing of Bermudan swaptions in the multi-factor LIBOR market model, *Journal of Computational Finance,* 3(2):5-32, Winter 1999/2000

[2] L. Andersen, J. Andreasen, Factor dependence of Bermudan swaptions: Fact or fiction, Gen Re working paper October 2000

[3] L. Anderson, M. Broadie, A primal-dual simulation algorithm for pricing multi-dimensional American options, preprint July 2001

[4] C. Hunter, P. Jäckel, M. Joshi, Getting the drift, *Risk* July 2001.

[5] M.B. Haugh, L. Kogan, Pricing American Options: A Duality Approach. Preprint 2001.

[6] P. Jäckel, Using a non-recombining tree to design a new pricing method for Bermudan swaptions, QUARC working paper 2000.

[7] P. Jäckel, R. Rebonato, Accurate and optimal calibration to co-terminal swaptions European swaptions in a FRA-based BGM framework, QUARC working paper

[8] F. Jamishidian, LIBOR and swap market models and measures, *Finance and Stochastics,* 1, 293-330, (1997)

[9] D. S. Kainth, R. Rebonato, A comparison of the Black Derman Toy model and the BGM model with flat term structure and perfect correlation in the pricing of Bermudan swaptions, Royal Bank of Scotland QUARC working paper

[10] F. Longstaff, E. Santa-Clara, E. Schwartz, Throwing away a billion dollars: the cost of suboptimal exercise in the swaptions market, UCLA working paper 2000

[11] W.H. Press, S.A. Teutolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in C, Cambridge University Press 1992

[12] R. Rebonato, Volatility and correlation, Wiley 1999

[13] R. Rebonato, The modern approach to pricing interest rate derivatives, to appear

[14] R. Rebonato, Review of Interest Rate Models - Theory and Practice by Brigo and Mercurio, Risk November 2001, 96-96

[15] L.C.G. Rogers, Monte Carlo valuation of American options, preprint, University of Bath 2001.

QUARC, Royal Bank of Scotland Group Risk, 2nd Floor, Waterhouse Square, 138-142 High Holborn, London EC1N 2TH