JUGGLING SNOWBALLS

CHRISTOPHER BEVERIDGE AND MARK JOSHI

1. Introduction

Whilst many advances have been made in recent years on the pricing of early exercisable derivatives using Monte Carlo simulation, it is still a hard problem. The fundamental problem being that whilst existing techniques work, they generally require a degree of specialist handcrafting or sub-Monte Carlo simulations. The former requiring an undesirable amount of research for new pay-offs and the latter resulting in models that are slower than satisfactory.

We study the pricing of contracts which give the holder a sequence of cash flows up until the time of cancellation, often referred to as cancellable products, in the LIBOR market model. As an example, we focus on the pricing of "snowball" contracts. These contracts present special difficulties since the coupons are path dependent. In addition, the high-dimensionality of the LIBOR market model makes the problem intractable on lattices.

Our approach will be to combine the least-squares method of Carrière, [5], (see also [12],) with various enhancements including a judicious choice of basis functions, a trigger-type optimization and the identification of sub-optimal points of exercise.

In particular, our approach involves

- working directly with the cancellable product, as in [10],
- a smarter choice of basis functions which are in a certain sense generic,
- showing that certain points are provably sub-optimal and so can be excluded from regressions,
- combining the Andersen approach with the least-squares approach.

This will enable us to obtain very tight bounds, that is bounds that are small compared to one vega, whilst not requiring sub-Monte Carlo simulations for the lower bound. This improves on the work of Bender, Kolodko and Schoenmakers [2], [3], who found that existing techniques were not sufficient and improved the lower bound by using sub-Monte Carlo simulations to estimate continuation values; this technique is known as *policy iteration.*

We consider the same example as in [2], obtaining very similar lower and upper bounds for cancellable snowball swaps in a full-factor LIBOR market model. Our upper bounds, as well as our lower bounds, take considerably less time to compute. In addition, we obtain similar results in what we consider a tougher test, pricing a cancellable snowball swap in a full-factor-displaced-diffusion LIBOR market model with an upward sloping yield curve and time-dependent volatility.

2. What is a snowball?

Under a snowball swap, the issuer receives the floating LIBOR rate and pays a "snowball" coupon. We have a set of tenor dates $0 = T_0 < T_1 < \ldots < T_n$. Let f_j denote the prevailing

LIBOR rate from T_j to T_{j+1} and let δ_j be the accrual factor. Snowball coupons typically have the following structure, where K_i denotes the i^{th} coupon before multiplication by the accrual factor,

$$\begin{aligned} K_i &= I, \text{ for } i = 0, \ldots, h, \\ K_i &= (K_{i-1} + A_i - f_i(T_i))_+, \text{ for } i = h+1, \ldots, n-1. \end{aligned}$$

Generally, the snowball coupon is fixed for the first year of the contract and after this period the issuer has the right to cancel the swap. In addition, A_i usually increases with i, contributing to the 'snowball' effect.

We considered the same example as in [2]. In particular, we considered pricing a semi-annual snowball swap over 10 years, where $h = 1$ and the issuer had the right to cancel the swap at half yearly intervals from $T_2 = 1$ onwards. For our parameters, we took $I = 7\%$, and had that A_i increased on an annual basis by 0.25%, starting from $A_2 = 3\%$. For example, we had $A_2 = 3\%$, $A_3 = 3\%$, $A_4 = 3.25\%$, $A_5 = 3.25\%$, ..., $A_{19} = 5\%$. In addition, the i^{th} coupon was settled at T_{i+1}, and so had a natural time-lag, and we assumed the swap had a \$1 notional.

3. The LIBOR Market Model

Let $P(t, T)$ denote the price at time t of a zero-coupon bond paying one at its maturity, T. Using no-arbitrage arguments,

$$f_j(t) = \frac{\frac{P(t,T_j)}{P(t,T_{j+1})} - 1}{\delta_j}.$$

Here, we work solely in the *spot LIBOR measure*, which corresponds to using the discretely-compounded money market account as numeraire. The value of this numeraire at time t will be denoted $N(t)$.

Under the displaced-diffusion LIBOR market model, the forward rates f_i are assumed to have the following evolution

$$df_i(t) = \mu_i(f, t)(f_i(t) + \alpha_i)dt + \sigma_i(t)(f_i(t) + \alpha_i)dW(t), \tag{3.1}$$

where the σ_i's are deterministic F $(n - 1 \geq F \geq 1)$ dimensional row vectors, the α_i's are constant displacement coefficients, W is a standard F-dimensional Brownian motion process under the relevant martingale measure, and the μ_i's are uniquely determined by no-arbitrage requirements.

4. Improvements to Standard Methods

Most exotic interest rate products are cancellable and many practioners work with such products by considering the non-cancellable product and the right to enter into the opposite product. However, there are numerous simplifications if one instead works with the cancellable product directly; see [10].

It is standard not to regress basis functions over out-of-the-money points when estimating continuation values; however, it is not clear what this means for cancellable products so we give a new method of identifying sub-optimal points. This identification also allows for acceleration to upper bounds as in [11].

4.1. **Excluding Sub-optimal Points.** Consider a breakable contract where the issuer receives a sequence of cash flows Y_i at each tenor date, until the time of exercise, where a deterministic rebate, $R \in (-\infty, \infty)$, is received. The time of exercise is decided by the issuer. Assume that the product can be exercised on a subset of the tenor dates $t_1, \ldots, t_k$, so that

$$\{t_1, \ldots, t_k\} \subset \{T_1, \ldots, T_n\}.$$

Let Z_j denote all cash flows generated by the product between t_{j-1} and t_j measured in units of the numeraire. As described in [2], at each exercise time t_i, the continuation value of our breakable note, $V_{i+1}(t_i)$, is given by

$$\frac{V_{i+1}(t_i)}{N(t_i)} = \sup_{\tau \in \Gamma_{i+1}} E_{t_i}\left(\sum_{j=i+1}^{\tau} Z_j + \frac{R}{N(t_\tau)}\right), \tag{4.1}$$

where Γ_j denotes the set of stopping times taking values in the set $\{j, j+1, \ldots, k\}$ representing the set of possible exercise times, $E_t(.)$ denotes the conditional expectation taken in the equivalent martingale measure associated with using $N(t)$ as numeraire, given the information available at time t.

From (4.1),

$$\begin{aligned} \frac{V_{i+1}(t_i)}{N(t_i)} &\geq \max_{h \in \{i+1, \ldots, k\}} E_{t_i}\left(\sum_{j=i+1}^{h} Z_j + \frac{R}{N(t_h)}\right), \\ &:= L_i. \end{aligned}$$

Optimal exercise should occur if and only if the payoff received upon exercise is greater than or equal to the continuation value. Therefore, if $R < L_i$, exercise should not occur.

In many cases it will not be possible to evaluate the conditional expectations on the right hand side for all $h \in \{i+1, \ldots, k\}$; for products with non-analytic underlying bonds it is not. A simple, yet effective alternative is to use a special case of the above. In particular,

$$\begin{aligned} \frac{V_{i+1}(t_i)}{N(t_i)} &\geq E_{t_i}\left(\sum_{j=i+1}^{h} Z_j + \frac{R}{N(t_h)}\right)\Bigg|_{h=i+1}, \\ &:= \tilde{L}_i, \end{aligned} \tag{4.2}$$

where exercise should never occur at t_i if

$$R < \tilde{L}_i. \tag{4.3}$$

In the common case that coupons are paid with a natural time-lag, rebates are deterministic, and exercise can occur at each subsequent coupon date, all cash flows in the above conditional expectation are known as of t_i, making the conditional expectation in (4.2) trivial to evaluate. Even when the value of Z_{i+1} is not known as of t_i, provided it only depends on a single outstanding rate it should lend itself to easy evaluation in most models via a Black-type formula. For more complicated coupon structures (for example, if Z_{i+1} contains non-linear terms) when the conditions given above are not met, certain approximations may be needed to apply the result, which would result in approximately sub-optimal points being identified, rather than provably sub-optimal points. However, this result has a wide range of applicability, and as we shall see, produced good results.

4.2. **Choosing Basis Functions.** When pricing complicated products, the success of the least-squares method depends heavily on choosing an appropriate set of basis functions. However, coming up with suitable basis functions can be tedious, and at present it is as much art as science. Here we present a reasonably generic set which has performed well in all examples so far considered. In choosing our basis functions, we follow the advice given in [13], but differ slightly in the exact choice of functions. We pick a set of core variables and then use all quadratic polynomials in them.

In particular, at exercise time t_i corresponding to tenor date T_j, these core variables are

- $f_j(T_j)$,
- $\mathrm{SR}_{j+1}(T_j)$ (the swap-rate starting at the next tenor date and running to final tenor date, evaluated at T_j),
- $\mathrm{FL}_j(T_j)$ (the value of the floating leg of the swap running from T_j to the final tenor date, also evaluated at T_j), so

$$\begin{aligned} \mathrm{FL}_j(T_j) &= P(T_j, T_j) - P(T_j, T_n), \\ &= 1 - P(T_j, T_n). \end{aligned}$$

- the current snowball coupon, K_j.

This final variable encapsulates the path dependence, and would not be useful for non-path dependent products.

5. Results

5.1. **LIBOR Market Model calibrations.** Two different calibrations for the LIBOR market model were considered. The first is the same as the one used in [2], and the second was used to provide a harder example with which to test the various methods. In each case, we considered the pricing of deals where the underlying swap lasted 10 years and had half-yearly payments, the first payment being due in six months. As such, $\delta_i = 0.5$ for all i and $n = 20$. We used a full-factor LIBOR market model, because, as noted in [2], this is where existing methods struggled most.

In the first set-up, we assumed a flat initial forward rate curve of 3.5% and zero displacements (i.e. $\alpha_i = 0$ for all i). In addition, each forward rate was assumed to have a constant volatility of 20%, and the instantaneous correlation between forward rates i and j was assumed to be equal to

$$\rho_{i,j} = \exp\left[\frac{\log(0.3)}{n-2}|i-j|\right].$$

As such, for $t < T_i$,

$$\sigma_i(t) = 0.2 e_i,$$

where e_i denotes the i^{th} row of a pseudo-square root of ρ. We will refer to this set-up as LIBOR market model set-up one.

In the second set-up, we assumed an initially increasing forward rate curve, such that $f_i(0) = 0.02 + 0.002i + x$, where x was varied accross examples. In addition, for $t < T_i$, $(f_i(t) + \alpha_i)$ was assumed to have time-dependent volatility of the form

$$p_i(t) = (0.05 + 0.09(t_i - t)) \exp(-0.44(t_i - t)) + 0.2, \tag{5.1}$$

and we assumed the same correlation structure as above. We used a spectral decomposition to determine the pseudo-square root of the covariance matrix across each time step in the examples considered. Displacements for all forward rates were assumed to be 1.5%. Due to (5.1) and the presence of non-zero displacements, this set-up has significantly larger effective volatilities than LIBOR market model set-up one. This set-up will be referred to as LIBOR market model set-up two.

In evolving the forward rates, we used the predictor-corrector drift approximation from [7], see [4] for further discussion. We also refer the reader to that paper and [9] for further details on how the evolution was carried out. Due to the accuracy of the predictor-corrector method, we evolved the forward rates using one step per tenor date. Mersenne Twister pseudo-random numbers were used throughout.

5.2. **Methods Used and Notation.**

5.2.1. *Least-squares Method.* We used the extension to the original least-squares method suggested in [1] (see also [10] for a detailed description of the method used), which allows for cancellable products to be priced directly, to obtain lower bounds. We considered two sets of basis functions. Specifically, as suggested in [13] and [2], at tenor date T_j, we used all quadratic polynomials in $f_j(T_j)$, $\mathrm{SR}_j(T_j)$, and K_j. We secondly used those from Subsection 4.2.

We also considered excluding sub-optimal points according to (4.3) in the least-squares methods used. We denote estimated lower bounds obtained using the least-squares like method as

$$L_b^{\mathrm{LS}},$$

and those obtained by also excluding sub-optimal points as

$$L_b^{\mathrm{LS},+},$$

where in each case b is used to denote the set of basis functions used ($b = \mathrm{BKS}$ indicates the first set described above is used and $b = \mathrm{New}$ the second). In order to investigate the exclusion of sub-optimal points, we also considered only excluding sub-optimal points in the second independent simulation, and therefore including all points in the least squares regressions. Estimated lower bounds using this approach are denoted as

$$L_b^{\mathrm{LS},1/2+}.$$

5.2.2. *Least-squares and Andersen Method.* We also considered improving the estimated continuation value obtained with the least-squares method by using it in conjunction with the Andersen method, as described in [2]. Lower bounds obtained using this combined method will be denoted as

$$L_b^{\mathrm{LSA}},\ L_b^{\mathrm{LSA},+},\ \text{or}\ L_b^{\mathrm{LSA},1/2+},$$

in an obvious extension to the notation above. We briefly describe this method. At each time frame, one first performs a regression of the continuation value against the basis functions, yielding an estimate of the continuation value, $\hat{C}$. One then adds a parameter α and instead of exercising according to whether $\hat{C} \leq E$ where E is the exercise value, one requires

$$\hat{C} + \alpha \leq E.$$

An optimization is then carried out over α to maximize the value of the Bermudan starting at that time frame. The optimization, being one-dimensional, is rapid to perform.

5.2.3. *Policy Iteration Method.* We also compared with the policy iteration method mentioned above. Prices obtained using this method will be denoted as

$$L_b^{\text{PI, method}},$$

where method is used to denote one of the methods discussed previously.

5.2.4. *Andersen–Broadie Method.* To obtain upper bounds, we used an extension to the Andersen–Broadie method described in [10]. As input, this method uses an exercise strategy developed by one of the lower bound methods mentioned above. Upper bounds will be denoted

$$U_b^{\text{method}}.$$

In estimating upper bounds, we identified and excluded sub-optimal points according to (4.3) when using exercise strategies that would not exercise at these points (i.e. the exercise strategies denoted with a +).

5.3. **Numerical Results.** We now present numerical results. We used a two-pass approach. With the first pass, 10^5 paths used to estimate the exercise strategy, and the second pass used an independent simulation for the pricing; this is standard practice to remove foresight bias.

In Table 5.1, we compare the accuracy of the various methods for the first calibration. To see the values without any of the improvements introduced here look at the top quarter of the table. See the bottom line for all the improvements used together. The number of paths used in the second pass for each simulation is the same as in [2], and is such that the relative standard error was less than 0.5% for all examples except those for L^{PI}, where the relative standard error was less than 1%. In particular, 10^7 paths were used for L_b for each method; $2.5 * 10^4$ outer and 500 inner paths were used for U_b^{LS} and $U_b^{\text{LS},+}$; 10^4 outer and 500 inner paths were used for U_b^{LSA} and $U_b^{\text{LSA},+}$; and $5 * 10^4$ outer and 500 inner paths were used for L_b^{PI} for each method. In order to give an idea of scale and hence assess whether the various bounds were sufficiently accurate, we compared the duality gaps obtained with vegas calculated using finite differencing and a 1% change in volatility. A duality gap less than the vega indicates it is adequate from a practical point of view. We see that very tight bounds can be obtained using

$$L_{\text{New}}^{\text{LSA},+}.$$

	L_{BKS}	U_{BKS}	Dual. Gap	$L^{\text{PI}}_{\text{BKS}}$	Vega
LS	77.37	119.88	42.51	103.47	26.10
LSA	92.65	111.11	18.44	106.08	20.42
LS,+	97.64	111.18	14.12	106.26	22.91
LSA,+	100.19	110.30	10.11	106.73	21.05
	L_{New}	U_{New}	Dual. Gap	$L^{\text{PI}}_{\text{New}}$	Vega
LS	91.04	115.00	23.96	104.30	24.51
LSA	101.15	109.36	8.21	106.82	20.28
LS,+	103.80	109.39	5.59	106.96	22.58
LSA,+	105.67	109.19	3.52	106.77	20.66

Table 5.1. Results for LIBOR market model set-up one. All numbers are in basis points.

Given similar levels of accuracy can be obtained, an important test of the various methods is to compare the time taken to obtain a given level of accuracy. In Figure 5.1, we compare the convergence of

$$L^{\text{PI,LSA,+}}_{\text{New}} \text{ and } L^{\text{LSA,+}}_{\text{New}}$$

as a function of time. Calculations involving the policy iteration method were carried out in conjunction with the variance reduction technique introduced in [2]. It is clear from Figure 5.1 that the new techniques are significantly faster than policy iteration; for a given amount of time, the standard error obtained with policy iteration was more than six times that obtained when not using it.

Note that simulation times do not include the time taken to build the approximate exercise strategy, which was essentially equal in each case. We used the same number of first pass paths as that used in Table 5.1. In addition, the example considered is particularly favorable to the policy iteration method, in that exercise occurred at the first exercise time in over 40% of paths (see [2]). Clearly, due to the use of sub-Monte Carlo simulations in assessing exercise strategies, we would expect significantly greater differences in the times taken for examples where exercise generally occurs at later exercise dates or examples with more exercise dates. We also note that increased convergence could be obtained through the use of low-discrepancy numbers in combination with Brownian bridging (see [8]), the use of control variates (see [6]), and so on.

By excluding sub-optimal points from the Andersen-Broadie style method, we also obtained significant time savings. In particular, it took on average over 40% longer to compute upper bounds for a given exercise strategy when sub-optimal points were not excluded. However, excluding sub-optimal points produced no noticeable improvement in the accuracy of the upper bound.

While these results are appealing, LIBOR market model set-up one provides a simple example. In order to provide a tougher test for the various methods, we also considered LIBOR market model set-up two. Table 5.2 compares the accuracy of the different methods with $x = 0$, where the number of paths used in each simulation was the same as that used in Table 5.1. This ensured that the relative standard errors were less than 1% except for

	L_{BKS}	U_{BKS}	Dual. Gap	$L^{\text{PI}}_{\text{BKS}}$	Vega
LS	73.55	144.67	71.12	116.57	25.24
LSA	104.15	131.68	27.53	121.22	19.56
LS,+	104.03	131.49	27.46	121.09	23.03
LSA,+	110.64	129.85	19.21	121.68	20.69
	L_{New}	U_{New}	Dual. Gap	$L^{\text{PI}}_{\text{New}}$	Vega
LS	102.22	135.11	32.89	118.79	23.78
LSA	114.35	127.97	13.62	122.56	20.07
LS,+	117.64	127.71	10.07	122.38	22.48
LSA,+	121.61	126.99	5.38	122.84	20.37

TABLE 5.2. Results for LIBOR market model set-up two with $x = 0$. All numbers are in basis points.

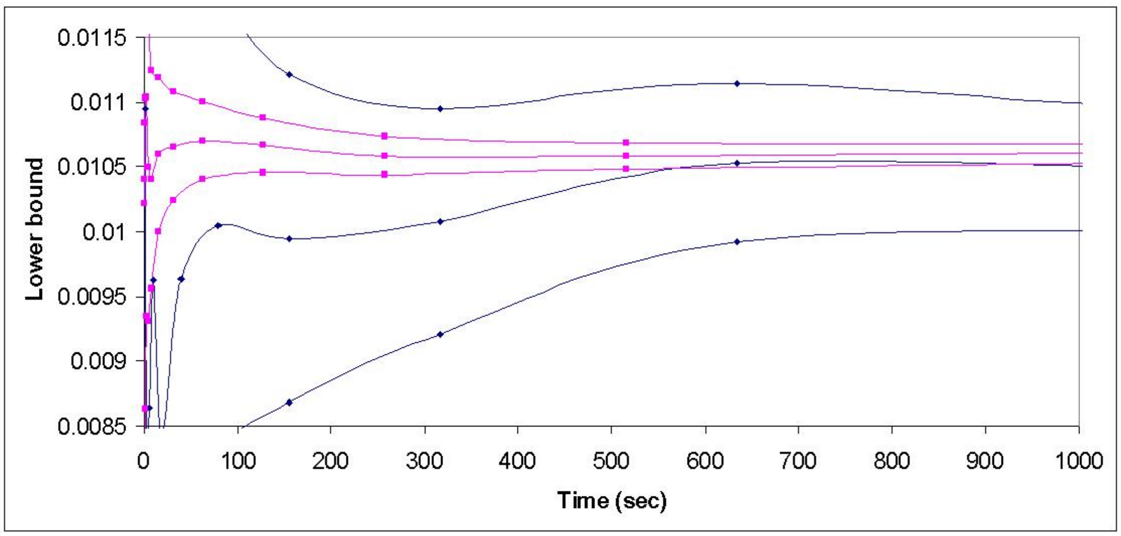

FIGURE 5.1. Plot of

$$L^{\text{PI,LSA,+}}_{\text{New}} \pm 2\text{SE (blue)}$$

and

$$L^{\text{LSA,+}}_{\text{New}} \pm 2\text{SE (pink)}$$

versus time using LIBOR market model set-up 1, where SE denotes standard error. For each colour, the centre curve gives a point estimate of the lower bound, and the outer curves make up a corresponding approximate 95% confidence interval. Lower bounds, standard errors and times were obtained by averaging the results found from five independent simulations.

V^{PI}, where they were less than 2%. We obtained similar results to Table 5.1, except now the difference between the various sets of basis functions is more pronounced. Further, the new set of basis functions introduced in Section 4 was generally required to produce sufficient

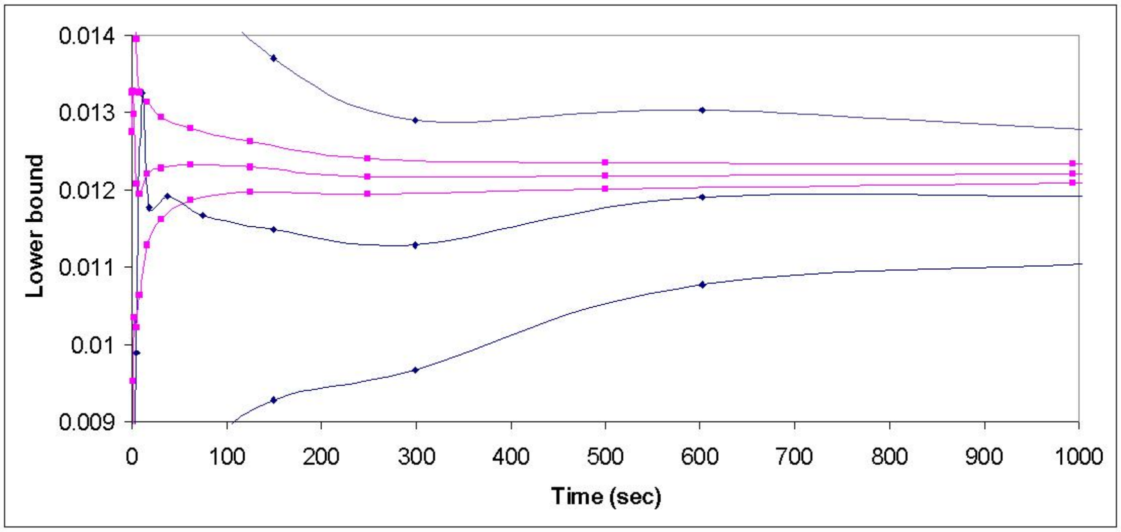

FIGURE 5.2. Plot of

$$L_{\text{New}}^{\text{PI,LSA},+} \pm 2\text{SE (blue)}$$

and

$$L_{\text{New}}^{\text{LSA},+} \pm 2\text{SE (pink)}$$

versus time using LIBOR market model set-up 2 with $x = 0$, where SE denotes standard error. For each colour, the centre curve gives a point estimate of the lower bound, and the outer curves make up a corresponding approximate 95% confidence interval. Lower bounds, standard errors and times were obtained by averaging the results found from five independent simulations.

bounds. Importantly, it was again possible to obtain sufficient (and in the case of $L_{\text{New}}^{\text{LSA},+}$, very tight) bounds without policy iteration. We also look at the convergence of the most accurate methods in Figure 5.2, and again see that the policy iteration method is significantly slower.

It is worth noting that in each example, policy iteration produced excellent lower bounds, even when given a not particularly great exercise strategy as input.

We now examine the exclusion of sub-optimal points in more detail. In Figure 5.3, the proportion of exercises which occur at sub-optimal points for each exercise time is plotted for various exercise strategies. This is done using LIBOR market model set-up two with $x = 0.5\%$. However, very similar graphs were also obtained when $x = 0$ and $x = 1\%$. Corresponding lower bounds are given in Table 5.3. It is interesting to note that using the Andersen adjustment significantly increases the proportion of sub-optimal exercises at most exercise times, yet still brings about significant improvements in the lower bound. Also, without the Andersen adjustment, the proportion of sub-optimal exercises at each exercise time is very low.

Table 5.3 compares the lower bounds obtained under LIBOR market model set-up two for $x = 0$, $x = 0.5\%$, and $x = 1\%$, where the larger the value of x, the greater the value of the underlying swap is to the issuer, and the less likely it is to be cancelled. The number of paths

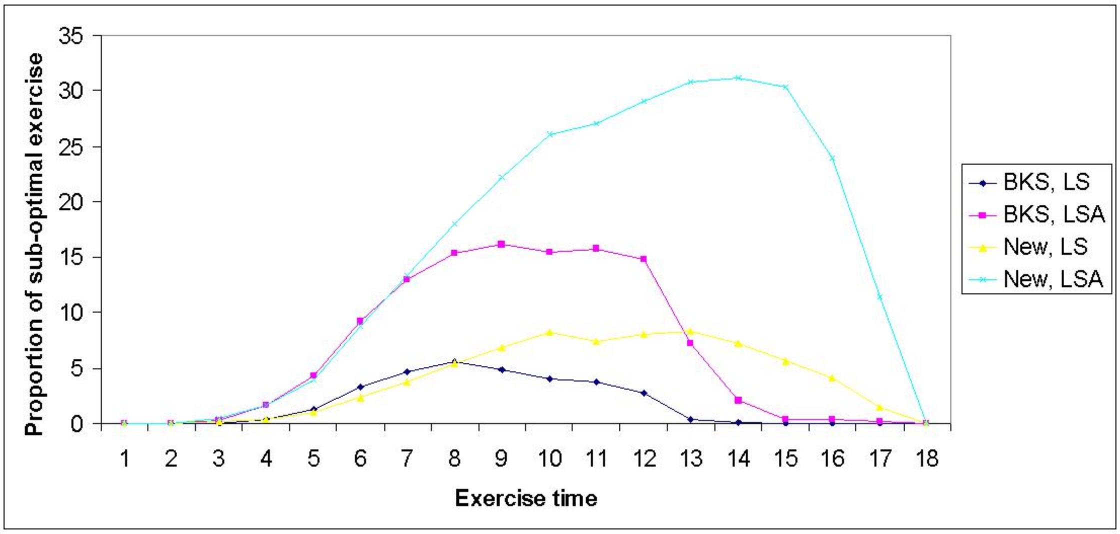

FIGURE 5.3. Plot of proportion of sub-optimal exercises at each exercise time using LIBOR market model set-up two with $x = 0.5\%$.

used were the same as above. From Table 5.3, we see that excluding sub-optimal points brings about improvements in two main ways; by improving continuation value estimates and by ensuring that sub-optimal exercise does not occur in the second independent simulation. In particular, by comparing Table 5.3 with Figure 5.3, when, as was the case with no Andersen improvement, very little sub-optimal exercise occurs, most of the improvement is obtained through improved continuation value estimates. However, when significant proportions of exercises occur at sub-optimal points, as when the Andersen improvement was applied, improvements are also obtained by simply ensuring that sub-optimal exercise does not occur in the second pass.

Table 5.3 also demonstrates that it is possible to obtain very tight bounds for varying values of the underlying swap to the issuer using

$$L_{\text{New}}^{\text{LSA},+}.$$

In particular, the duality gap obtained using this strategy was never over six basis points, regardless of the level of the initial forward rate curve.

In Figure 5.4, the distribution of exercise times is plotted for various strategies. Similar patterns were obtained for all other cases considered. It is clear that excluding sub-optimal points did not alter the distribution of exercise times much at all, but rather reduced the effect of incorrect decisions by their reducing their frequency and/or impact.

6. Conclusion

From the results presented, a clear conclusion emerges: it is possible to obtain tight bounds on snowball prices by using the hybrid Andersen/least-squares method, picking the right basis functions and excluding sub-optimal points. The lower bound does not require sub-Monte Carlo simulations and therefore is quick to run. Once this strategy has been obtained, sub-simulations can then be applied to obtain ultra-tight upper and lower bounds, but these

L_{BKS}	$x = 0$	$x = 0.5\%$	$x = 1\%$
LS	73.55	503.17	998.75
LS,1/2+	74.89	504.84	1000.70
LS,+	104.03	541.02	1038.19
LSA	104.15	539.87	1035.37
LSA,1/2+	107.21	543.84	1039.87
LSA,+	110.64	548.77	1046.04
L_{New}			
LS	102.30	539.73	1036.55
LS,1/2+	102.81	540.42	1037.60
LS,+	117.64	555.95	1052.95
LSA	114.53	550.73	1046.83
LSA,1/2+	118.44	555.87	1053.32
LSA,+	121.61	560.37	1058.40
$U_{\text{New}}^{\text{LSA},+}$	126.99	566.02	1063.26

TABLE 5.3. Results for LIBOR market model set-up two with $x = 0$, $x = 0.5\%$ and $x = 1\%$. All numbers are in basis points.

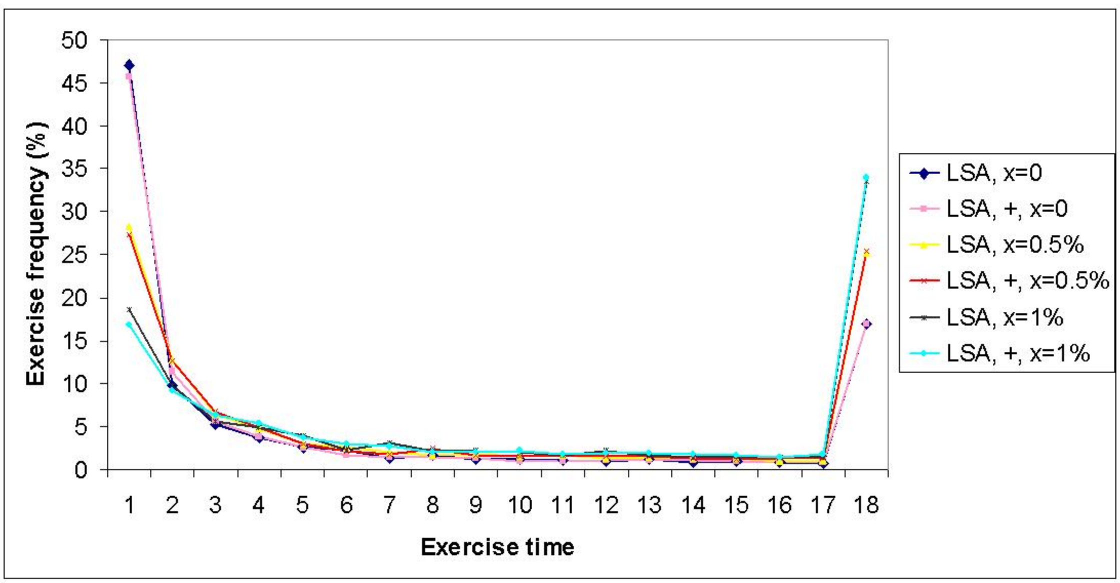

FIGURE 5.4. Plot of exercise frequency for each exercise time using LIBOR market model set-up two with $x = 0$, $x = 0.5\%$, and $x = 1\%$. All graphs were obtained using New basis functions.

bounds are mainly useful for testing purposes, since our improved strategy is more than good enough.

References

[1] A. Amin. Multi-factor cross currency LIBOR market model: implementation, calibration and examples, 2007. http://www.geocities.com/anan2999/.

[2] C. Bender, A. Kolodko, and J. Schoenmakers. Iterating cancellable snowballs and related exotics. *Risk*, pages 126–130, 2006.

[3] C. Bender, A. Kolodko, and J. Schoenmakers. Enhanced policy iteration for American options via scenario selection. *Quantitative Finance*, 8:135–146, 2008.

[4] C. J. Beveridge, N. A. Denson, and M. S. Joshi. Comparing discretizations of the LIBOR market model in the spot measure, 2008. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1207482.

[5] J. F. Carrière. Valuation of the early-exercise price for options using simulation and nonparametric regression. *Insurance: Mathematics and Economics*, 19:19–30, 1996.

[6] P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, New York, 2004.

[7] C. Hunter, P. Jäckel, and M. S. Joshi. Getting the drift. *Risk*, 14:81–84, 2001.

[8] P. Jäckel. *Monte Carlo Methods in Finance*. John Wiley and Sons Ltd., New York, 2001.

[9] M. S. Joshi. *The Concepts and Practice of Mathematical Finance*. Cambridge University Press, London, 2003.

[10] M. S. Joshi. Monte Carlo bounds for callable products with non-analytic break costs, 2006. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=907407.

[11] M. S. Joshi. A simple derivation of and improvements to Jamshidian's and Rogers' upper bound methods for Bermudan options. *Applied Mathematical Finance*, 14:197–205, 2007.

[12] F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: a simple least squares approach. *The Review of Financial Studies*, 14:113–147, 2001.

[13] V. Piterbarg. A practitioner's guide to pricing and hedging callable LIBOR exotics in forward LIBOR models. *Journal of Computational Finance*, 8:65–119, 2004.

Centre for Actuarial Studies, Dept of Economics, University of Melbourne, Victoria 3010, Australia