PRACTICAL POLICY ITERATION: GENERIC METHODS FOR OBTAINING RAPID AND TIGHT BOUNDS FOR BERMUDAN EXOTIC DERIVATIVES USING MONTE CARLO SIMULATION

CHRISTOPHER BEVERIDGE AND MARK JOSHI

ABSTRACT. We introduce a set of improvements which allow the calculation of very tight lower bounds for Bermudan derivatives using Monte Carlo simulation. These tight lower bounds can be computed quickly, and with minimal hand-crafting. Our focus is on accelerating policy iteration to the point where it can be used in similar computation times to the basic least-squares approach, but in doing so introduce a number of improvements which can be applied to both the least-squares approach and the calculation of upper bounds using the Andersen–Broadie method. The enhancements to the least-squares method improve both accuracy and efficiency.

Results are provided for the displaced-diffusion LIBOR market model, demonstrating that our practical policy iteration algorithm can be used to obtain tight lower bounds for cancellable CMS steepener, snowball and vanilla swaps in similar times to the basic least-squares method.

1. INTRODUCTION

The problem of pricing derivative contracts with early exercise features using Monte Carlo simulation has until recently been considered a very difficult problem. The difficulty arises because Monte Carlo simulation is a method which naturally moves forward in time, whereas exercise decisions require comparisons between the value received upon exercise and that of the unexercised product, the latter not being easily attainable in a simulation. However, Monte Carlo simulation is generally required to calculate prices in models with high-dimensional state spaces, or to handle path-dependent products, where alternative methods can become troublesome. An important example where Monte Carlo simulation is required is the LIBOR market model; it is the benchmark model in a market where early exercise features are very popular.

In recent years a number of important advances have been made in the pricing of early exercisable derivatives using Monte Carlo simulation, but the problem is not fully solved. The least-squares approach of Longstaff and Schwartz, [29], and Carrière, [14], is widely used for calculating lower bounds, but is very sensitive to the initial choice of basis functions. What to use for basis functions is generally not obvious, and a significant amount of time is often required for each new product to investigate different sets of basis functions before this approach can be used confidently to obtain accurate lower bounds. In particular, according to Brace, [10], "the choice of regression variables is partly an art and partly a science, and should be tailored to the callable instrument being valued". Kolodko and Schoenmakers, [28], developed a generic improvement to the least-squares (or any other) approach, called the policy iteration method, which reduced the importance of the initial choice of basis functions. However, their method involved using sub-Monte Carlo simulations to assess the improved exercise strategy in most cases, which in its basic form made it too slow to be of any real practical use; see [9]. Broadie and Cao, [13], introduced an efficiency improvement to the policy iteration method and applied it to equity derivatives. While their improvement produces useful efficiency improvements, they still note that policy iteration "may be too expensive to apply to some Bermudan options".

While lower bounds are often used in practice for pricing, upper bounds are useful for assessing the accuracy of lower bound methodologies. The predominant approach to upper bounds is the dual approach introduced by Rogers, [33], and Haugh and Kogan, [20], and improved upon by Andersen and Broadie, [3].

In addition, a number of extensions have been suggested to apply the methods mentioned above to cancellable contracts, that is contracts which give the holder a sequence of cash flows up until the time of cancellation. The most notable are those of Amin, [2], who focuses on lower bounds and Joshi, [27], who works with upper bounds. These are important because most exotic interest rate products are cancellable and there are numerous simplifications if one works with the cancellable product directly, rather than decomposing the product into the non-cancellable product and the complementary callable product; see [27]. As such, this paper focuses on cancellable contracts, but all techniques can also be used for callables. However, these extensions do not remove the significant hand-crafting required for the least-squares approach, or the speed issues with policy iteration.

In this paper, we introduce a set of improvements which allow the calculation of very tight lower bounds quickly, and with minimal handcrafting. Our ultimate goal is to accelerate policy iteration to the point where it is as fast as the basic least-squares approach, but in doing so use a number of improvements which can be applied to both the least-squares approach and the calculation of upper bounds using the Andersen–Broadie method. In particular, the results of this paper can be broken into four main categories:

- a number of improvements to the policy iteration algorithm which make it fast,
- a new adaptive approach to choosing basis functions in the least-squares method,
- a double regression approach to the approximation of continuation values using the least-squares method,
- a new control variate for Bermudan options to be used in conjunction with the least-squares method. The control variate suggested is based on the idea of Delta hedging.

These improvements will be discussed in the context of the LIBOR market model, but with minor modifications they can just as easily be applied to other models where Monte Carlo simulation is used to price Bermudan options. We will refer to the improved policy iteration algorithm suggested in this paper as practical policy iteration, and show that it can be run in similar simulation times to the standard least-squares method, yet produce significant improvements in accuracy.

Our work bears some similarity to that of Jensen and Svenstrup, [24], who look at improving the efficiency of both lower bound methods and the Andersen–Broadie method for upper bounds. In particular, Jensen and Svenstrup focus on the particular case of Bermudan swaptions in the LIBOR market model. However, the techniques introduced in this paper are significantly more generic, and can be applied to a wide range of exotic interest rate derivatives as well as Bermudan options in other markets. In addition, while in [24] only efficiency is addressed, our improvements can also significantly improve accuracy.

It is worth mentioning the potential impact of recent advances in parallel processing technology; see [17]. Since Monte Carlo simulation is extremely parallel, it can fully exploit these advances, with increases in speed of factors in the hundreds appearing achievable. As the methods in this paper are based on Monte Carlo simulation, they stand to gain significantly from these improvements, and we believe they will become even more attractive from a practical point of view in the near future.

The paper is organised as follows. In Section 2, the displaced-diffusion LIBOR market model, which is to be used in later examples, is briefly introduced, with additional notation introduced in Section 3. Section 4 contains the main body of the paper, focusing on making policy iteration fast. In particular, different variance reduction techniques and other efficiency improvements are discussed in relation to policy iteration, and these will be shown to be particularly effective in the setting of the LIBOR market model. Section 5 introduces a simple yet effective way of choosing basis functions for the least-squares method requiring minimal changes for different products, and in Section 6 a generic improvement to least-squares continuation value estimates based on a double regression approach is introduced. Section 7 is used to describe a new control variate based on the least-squares continuation value estimates and the concept of Delta hedging. Section 8 presents results for the displaced-diffusion LIBOR market model, showing that very tight bounds can be obtained using policy iteration in similar computation times to the standard least-squares method. In particular, two challenging products are considered, cancellable snowball and CMS steepener swaps, as well as Bermudan swaptions.

2. The Displaced-Diffusion LIBOR Market Model

Since it was given a firm theoretical base in the fundamental papers by Brace, Gatarek and Musiela, [11], Musiela and Rutkowski, [30], and Jamshidian, [23], the LIBOR market model has become a very popular method for pricing interest rate derivatives. It is based on the idea of evolving the yield curve directly through a set of discrete market observable forward rates, rather than indirectly through use of a single non-observable quantity which is assumed to drive the yield curve. A distinct advantage of this approach is the ability to easily calibrate to a large number of simpler financial contracts, often used in the hedging process for exotics; see [1].

Suppose we have a set of tenor dates, $0 = T_0 < T_1 < \ldots < T_{n+1}$, with corresponding forward rates $f_0, \ldots, f_n$. Let $\delta_j = T_{j+1} - T_j$, and 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,

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

where $f_j(t)$ is said to reset at time T_j, after which point it is assumed that it does not change in value. We work solely in the *spot LIBOR measure*, which corresponds to using the discretely-compounded money market account as numeraire, because this has certain practical advantages; see [25]. This numeraire is made up of an initial portfolio of one zero-coupon bond expiring at time T_1, with the proceeds received when each bond expires being reinvested in bonds expiring at the next tenor date, up until T_n. More formally, the value of the numeraire portfolio at time t will be,

$$N(t) = P(t, T_{\eta(t)}) \prod_{i=1}^{\eta(t)-1} (1 + \delta_i f_i(T_i)),$$

where $\eta(t)$ is the unique integer satisfying

$$T_{\eta(t)-1} \leq t < T_{\eta(t)},$$

and thus gives the index of the next forward rate to reset.

Under the displaced-diffusion LIBOR market model, the forward rates that make up the state variables of the model are assumed to be driven by the following process

$$df_i(t) = \mu_i(f,t)(f_i(t) + \alpha_i)dt + \sigma_i(t)(f_i(t) + \alpha_i)dW_i(t), \tag{2.1}$$

where the $\sigma_i(t)$'s are deterministic functions of time, the α_i's are constant displacement coefficients, the W_i's are standard Brownian motions under the spot LIBOR martingale measure, and the $\mu_i's$ are uniquely determined by no-arbitrage requirements. It is assumed that W_i and W_j have correlation $\rho_{i,j}$ and throughout $\{\mathcal{F}_t\}_{t\geq 0}$ will be used to denote the filtration generated by the driving Brownian motions. In addition, all expectations will be taken in the spot LIBOR probability measure. The requirement that the discounted price processes of the fundamental tradeable assets, that is the zero-coupon bonds associated to each tenor date, be martingales in the pricing measure, dictates that the drift term is uniquely given by

$$\mu_i(f,t) = \sum_{j=\eta(t)}^{i} \frac{(f_j(t) + \alpha_j)\delta_j}{1 + f_j(t)\delta_j} \sigma_i(t)\sigma_j(t)\rho_{i,j};$$

see [12]. The existence of a unique equivalent martingale measure means that the displaced-diffusion LIBOR market model forms a complete market; see [19].

Displaced-diffusion is used as a simple way to allow for the skews seen in implied caplet volatilities that have long persisted in interest rate markets; see [25]. In particular, the use of displaced-diffusion allows for the wealth of results concerning calibrating and evolving rates in the standard LIBOR market model to be carried over with only minor changes. The model presented collapses to the standard LIBOR market model when $\alpha_i = 0$ for all values of i.

All methods introduced below will use the displaced-diffusion LIBOR market model as an example, as will the numerical examples considered in Section 8.

3. Notation

Here we introduce the notation for a generic contract, which we will use to discuss the methods introduced in this paper.

Consider a contract where the issuer continues to receive net cash flows S_i at each tenor date T_i, until the time of exercise. The coupons paid by the issuer at each tenor date will be denoted by s_i. The time of exercise is decided by the issuer. Assume that the product can be exercised on a subset of the tenor dates $t_1, \ldots, t_h$, so that

$$\{t_1, \ldots, t_h\} \subset \{T_1, \ldots, T_{n+1}\}.$$

Upon exercise at t_j, the issuer receives a discounted exercise value denoted by H_j. Let C_j denote all cash flows generated by the product between t_{j-1} and t_j measured in units of the numeraire. Note that this case reduces to that of a Bermudan callable option if all cash flows, C_j, are zero.

When working with breakables, the case that exercise does not occur needs to be allowed for, both in terms of cash flows and in terms of ensuring that a finite stopping time is used as an exercise strategy so that we can apply the Optional Sampling theorem when calculating upper bounds and using control variates. To do this, we assume an additional exercise time, t_{h+1}, occurring after all cash flows are received, where it is assumed that exercise has to occur and zero rebate is received upon exercise.

Let $\bar{\tau}$ denote a given exercise strategy taking values in the set

$$\{1, 2, \ldots, h, h+1\},$$

representing the set of possible exercise times.

Finally, for ease of exposition, let $\mathbb{E}^i(.) = \mathbb{E}(.|\mathcal{F}_{t_i})$.

4. Making Policy Iteration Fast

A general improvement to lower bound methods, called the policy iteration method, was introduced by Kolodko and Schoenmakers, [28], and further studied in [7], [5], [6], and [13]. In practical terms, an exercise strategy developed using a given lower bound methodology is improved upon by determining each exercise decision through comparing the estimated continuation value obtained using a sub-Monte Carlo simulation and this exercise strategy with the value received upon exercise. In [5], it was demonstrated that even after performing a single iteration on a reasonable exercise strategy, very tight lower bounds could be obtained for cancellable snowball swaps in a full-factor LIBOR market model. However, the general need for sub-simulations makes using even one iteration quite expensive if policy iteration is implemented naively. Thus while capable of significantly improving the accuracy of lower bounds obtained by basic exercise strategies, policy iteration suffers from its computational cost.

In this section, a number of improvements to the policy iteration algorithm will be described which make it significantly faster. In Section 4.1, we give a financial interpretation for the variance reduction technique introduced in [5], and describe how to significantly improve it. In Section 4.2, an adaptive approach to terminating the sub-simulations is introduced. A collection of additional improvements, including the use of sub-optimal points (Section 4.3), and other details about the practical implementation of policy iteration (Section 4.4), are also introduced.

4.1. **Andersen–Broadie Hedge Control Variate.** We briefly review the idea behind variance reduction technique introduced in [5] to be used with policy iteration. Let Y_i denote the discounted value process for the Bermudan at t_i using $\bar{\tau}$. Then,

$$Y_i = H_i \mathrm{I}_{\bar{A}_i} + \mathbb{E}^i \left((C_{i+1} + Y_{i+1}) \right) \mathrm{I}_{\bar{A}_i^c},$$

where $\bar{A}_i$ denotes the event that $\bar{\tau}$ exercises at t_i.

Consider a hedge to this Bermudan that involves buying (or dynamically replicating, which is the same in a complete market) the underlying product exercised according to $\bar{\tau}$. The underlying product is held until $\bar{\tau}$ says exercise. At this point, the hedge is exercised and the underlying Bermudan exercised according to $\bar{\tau}$ starting at the next time frame is purchased. The discounted value of this hedge portfolio is given by

$$M_i = \sum_{j=1}^{i} (Y_j + C_j - \mathbb{E}^{j-1}(Y_j + C_j)).$$

Then, trivially,

$$M_{\bar{\tau}} + \mathbb{E}^0(Y_1 + C_1) = H_{\bar{\tau}} + \sum_{j=1}^{\bar{\tau}} C_j,$$

and therefore

$$\mathbb{E}^0(Y_1 + C_1) = H_{\bar{\tau}} + \sum_{j=1}^{\bar{\tau}} C_j - M_{\bar{\tau}}.$$

This suggests that an approximation to the optimal hedge, M, would be a good control variate when pricing Bermudan options. Since in the policy iteration method, the information required to use the approximation to the optimal hedge introduced by Andersen and Broadie, [3], is naturally available, this can be used at minimal additional computational cost. In [3], the approximation to the optimal hedge is used to compute upper bounds, where as here we use it as a control variate to improve the policy iteration method. However, the idea is the same. In particular, rather than using the underlying product to hedge, the product exercised according to the non-iterated strategy is used. This works because the non-iterated product can be valued quickly by sub-Monte Carlo simulations while maintaining the martingale property for the discounted value of the approximate hedge; see [18].

We briefly review the Andersen–Broadie approximate hedge. The hedge portfolio starts with nothing and sells the numeraire asset to purchase the equivalent of the underlying Bermudan option exercised according to the non-iterated strategy. When the non-iterated strategy says exercise, the hedge returns the exercise value and we re-purchase the non-iterated product starting at the next time frame. Any additional cash received/needed is dealt with by buying/selling the numeraire asset. Nested Monte Carlo simulations using the non-iterated strategy are used to estimate the continuation value of the hedge when needed.

Naturally, one would want to carry out the approximate hedge until the iterated strategy says exercise when using it as a control variate. However, by doing this an upward bias is introduced into the lower bound, which is now proved.

4.1.1. *Bias in Andersen–Broadie Hedge Control Variate.* We look at the Andersen–Broadie hedge as a control variate in the context of the method suggested in [5], and show that an upward bias is introduced if the control variate is stopped when the iterated exercise strategy says exercise. The intuition is simple. Since the same sub-simulations determine both the value of the hedge and when the hedging is stopped, this restricts the possible values of the hedge across each step, introducing a bias.

Let $\hat{M}_i$ denote the value of the approximate hedge described in [3], so

$$\hat{M}_i = \sum_{j=1}^{i} (\hat{Y}_j + C_j - Z_{j-1}), \tag{4.1}$$

where

$$\begin{aligned} Z_{j-1} &= \mathbb{E}^{j-1}\widehat{(\tilde{Y}_j + C_j)}, \\ \tilde{Y}_j &= H_j \mathrm{I}_{A_j} + (\tilde{Y}_{j+1} + C_{j+1})\mathrm{I}_{A_j^{\mathrm{C}}}, \\ \hat{Y}_j &= H_j \mathrm{I}_{A_j} + Z_j \mathrm{I}_{A_j^{\mathrm{C}}}, \end{aligned}$$

with A_j denoting the event that exercise occurs according to the non-iterated strategy at t_j and Z_j denoting the Monte Carlo estimate of the t_j continuation value using the non-iterated strategy.

Now, suppose $\hat{M}_{\tau^*}$ is used as a control variate, where τ^* denotes the exercise time of the iterated exercise strategy. Unfortunately, as we now show, $\mathbb{E}\left(\hat{M}_{\tau^*}\right) < 0$, which causes an upward bias when using $\hat{M}_{\tau^*}$ as a control variate.

Re-writing $\hat{M}_{\tau^*}$ gives

$$\begin{aligned} \hat{M}_{\tau^*} &= \sum_{j=1}^{\tau^*} \left(\hat{Y}_j + C_j - Z_{j-1}\right), \\ &= \sum_{j=1}^{h+1} \left(\hat{Y}_j + C_j - Z_{j-1}\right) \mathrm{I}_{j \le \tau^*}, \\ &= \sum_{j=1}^{h+1} \left(\hat{Y}_j + C_j - Z_{j-1}\right) \mathrm{I}_{j-1 < \tau^*}. \end{aligned}$$

Now consider the expectation of each term in the sum separately; consider the j^{th} term. We introduce two new sigma-algebras. Let $\mathcal{H}$ contain the information generated by the driving Brownian motions up until time t_{j-1} and the sub-simulations $Z_0, \ldots, Z_{j-1}$. In addition, let $\mathcal{G}$ contain the information generated by the driving Brownian motions up until t_j and the sub-simulations $Z_0, \ldots, Z_{j-1}$. This notation will be used throughout this section. Trivially, $\mathcal{H} \subset \mathcal{G}$. Throughout we will assume that the sub-simulations at different exercise times are independent, meaning

$$\mathbb{E}(Z_j|\mathcal{G}) = \mathbb{E}(\tilde{Y}_{j+1} + C_{j+1}|\mathcal{G}).$$

Now, by the Tower property,

$$\begin{aligned}\mathbb{E}\left[\left(\hat{Y}_j + C_j - Z_{j-1}\right)\mathrm{I}_{j-1<\tau^*}\right] &= \mathbb{E}\left[\mathbb{E}\left[\left(\hat{Y}_j + C_j - Z_{j-1}\right)\mathrm{I}_{j-1<\tau^*}|\mathcal{H}\right]\right], \\ &= \mathbb{E}\left[\mathrm{I}_{j-1<\tau^*}\mathbb{E}\left[\left(\hat{Y}_j + C_j - Z_{j-1}\right)|\mathcal{H}\right]\right], \\ &= \mathbb{E}\left[\mathrm{I}_{j-1<\tau^*}\left(\mathbb{E}\left[\hat{Y}_j + C_j|\mathcal{H}\right] - Z_{j-1}\right)\right],\end{aligned}$$

using that both $\mathrm{I}_{j-1<\tau^*}$ and Z_{j-1} are $\mathcal{H}$-measurable.

It is useful to introduce the following lemma.

Lemma 1. $\mathbb{E}\left[\hat{Y}_j + C_j|\mathcal{H}\right] = \mathbb{E}\left[\tilde{Y}_j + C_j|\mathcal{H}\right]$.

Proof. Writing out $\hat{Y}_j$,

$$\mathbb{E}\left[\hat{Y}_j + C_j|\mathcal{H}\right] = \mathbb{E}\left[H_j\mathrm{I}_{A_j} + Z_j\mathrm{I}_{A_j^{\mathrm{C}}} + C_j|\mathcal{H}\right]. \tag{4.2}$$

Using the Tower property on the middle term on the right hand side of (4.2),

$$\begin{aligned}\mathbb{E}\left[Z_j\mathrm{I}_{A_j^{\mathrm{C}}}|\mathcal{H}\right] &= \mathbb{E}\left[\mathbb{E}\left[Z_j\mathrm{I}_{A_j^{\mathrm{C}}}|\mathcal{G}\right]|\mathcal{H}\right], \\ &= \mathbb{E}\left[\mathrm{I}_{A_j^{\mathrm{C}}}\mathbb{E}\left[Z_j|\mathcal{G}\right]|\mathcal{H}\right], \qquad (4.3)\\ &= \mathbb{E}\left[\mathrm{I}_{A_j^{\mathrm{C}}}\mathbb{E}\left[\tilde{Y}_{j+1} + C_{j+1}|\mathcal{G}\right]|\mathcal{H}\right], \qquad (4.4)\\ &= \mathbb{E}\left[\mathrm{I}_{A_j^{\mathrm{C}}}\left(\tilde{Y}_{j+1} + C_{j+1}\right)|\mathcal{H}\right],\end{aligned}$$

where we have used that $\mathrm{I}_{A_j^{\mathrm{C}}}$ is $\mathcal{G}$-measurable to move it in and out of expectations conditional on $\mathcal{G}$, and the Tower property again for the final step.

As a result,

$$\begin{aligned}\mathbb{E}\left[\hat{Y}_j + C_j|\mathcal{H}\right] &= \mathbb{E}\left[H_j\mathrm{I}_{A_j} + \left(\tilde{Y}_{j+1} + C_{j+1}\right)\mathrm{I}_{A_j^{\mathrm{C}}} + C_j|\mathcal{H}\right], \\ &= \mathbb{E}\left[\tilde{Y}_j + C_j|\mathcal{H}\right].\end{aligned}$$

□

Putting this together,

$$\begin{aligned}\mathbb{E}\left[\left(\hat{Y}_j + C_j - Z_{j-1}\right)\mathrm{I}_{j-1<\tau^*}\right] &= \mathbb{E}\left[\mathrm{I}_{j-1<\tau^*}\left(\mathbb{E}\left[\hat{Y}_j + C_j|\mathcal{H}\right] - Z_{j-1}\right)\right], \\ &= \mathbb{E}\left[\mathrm{I}_{j-1<\tau^*}\left(\mathbb{E}\left[\tilde{Y}_j + C_j|\mathcal{H}\right] - Z_{j-1}\right)\right], \\ &= -\mathbb{E}\left[\mathrm{I}_{j-1<\tau^*}\epsilon_{j-1}\right], \\ &< 0,\end{aligned}$$

where ϵ_{j-1} is used to denote the Monte Carlo error for the $(j-1)^{\text{th}}$ sub-simulation. The final inequality easily follows from financial intuition: the larger the Monte-Carlo error on the $(j-1)^{\text{th}}$ sub-simulation, the greater the estimate of continuation value at t_{j-1}, and the more likely that $j-1 < \tau^*$. Thus, $\mathrm{I}_{j-1<\tau^*}$ and ϵ_{j-1} are positively correlated.

4.1.2. *Removing the Bias.* There are a number of different approaches one can take to remove the bias introduced by using the Andersen–Broadie approximate hedge as a control variate. In [5], a modified control variate is used. The Andersen–Broadie approximate hedge is used as a control, but the hedging is instead carried out until one period after the iterated strategy says exercise. However, by taking this approach a significant portion of the variance reduction is lost. In particular, in [32], Rasmussen proved that it was never optimal to stop a control variate after the exercise time of the underlying option. We now introduce two more effective alternatives.

Semi-Analytic Approach. We estimate the bias and subtract it. Our approach bears some similarity to Fries' technique for removing the bias in the first pass of the least-squares method; see [16]. In particular, the bias is given by

$$\sum_{j=1}^{h+1} \mathbb{E}\left[\mathrm{I}_{j-1<\tau^*}\epsilon_{j-1}\right]. \tag{4.5}$$

Consider each term in the sum separately. Using the Tower property,

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j\right] = \mathbb{E}\left[\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j|\mathcal{G}\right]\right]. \tag{4.6}$$

By the Central Limit Theorem, conditional on $\mathcal{G}$,

$$\epsilon_j \overset{\mathrm{d}}{\approx} \mathrm{N}(0, \sigma^2/m),$$

where m is the number of paths used for the t_j sub-simulation. As such, it is possible to approximately calculate the inner expectation of (4.6) for each outer path of the policy iteration simulation. In particular,

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j|\mathcal{G}\right] \approx \begin{cases} 0 & \text{if } \tau^* \leq j-1; \\ \mathbb{E}\left[\mathrm{I}_{\{a+\frac{\sigma}{\sqrt{m}}X>H_j\}}\frac{\sigma}{\sqrt{m}}X\right] & \text{otherwise,} \end{cases} \tag{4.7}$$

where $X \overset{\mathrm{d}}{=} \mathrm{N}(0,1)$, and

$$a := \mathbb{E}^j\left[\tilde{Y}_j + C_j\right]$$

denotes the continuation value at t_j using the non-iterated strategy. Of course a and σ will not be known exactly (we are trying to estimate

a with the sub-simulation), however they can be estimated using the results of the sub-simulation at t_j. The estimates will be denoted by $\hat{a}$ and $\hat{\sigma}$.

Ignoring the randomness of $\hat{a}$ and $\hat{\sigma}$, it is then straight-forward to compute the expectation on the right hand side of (4.7), obtaining

$$\mathbb{E}\left[\mathrm{I}_{\{\hat{a}+\frac{\hat{\sigma}}{\sqrt{m}}X>H_j\}}\frac{\hat{\sigma}}{\sqrt{m}}X\right] = \frac{\hat{\sigma}}{\sqrt{m}}\phi\left(\frac{H_j-\hat{a}}{\hat{\sigma}/\sqrt{m}}\right),$$

where $\phi(x)$ denotes the standard Normal pdf; see [16].

Given that it is possible to sample

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j|\mathcal{G}\right] \tag{4.8}$$

approximately on each path, at each exercise time with very little additional effort, we can easily use Monte Carlo to estimate each term in (4.5). In particular, if D_i denotes (4.8) sampled for the i^{th} outer path, where p outer paths are used,

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j\right] \approx \frac{1}{p}\sum_{i=1}^{p} D_i.$$

Adding these individual estimates together gives an estimate of the bias introduced by using the Andersen–Broadie hedge control variate, and this can be subtracted to obtain an approximately unbiased estimate of the lower bound using the iterated strategy. While this approach is very effective, it relies on the Central Limit Theorem and could potentially break down if very few sub-simulation paths are used. However, at a small additional cost in computation time, there is an alternative that is guaranteed to be unbiased by construction.

Numerical Approach. If the sub-simulations used to determine the iterated exercise strategy and the value of the Andersen–Broadie hedge at each exercise time are independent, that is performed using independent random numbers, then ϵ_j and $\mathrm{I}_{j<\tau^*}$ are uncorrelated. To see this, consider

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j\right] = \mathbb{E}\left[\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j|\mathcal{G}\right]\right]. \tag{4.9}$$

Since independent sub-simulations are used, it is possible to write

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j|\mathcal{G}\right] = \begin{cases} 0 & \text{if } \tau^* \leq j-1; \\ \mathbb{E}\left[\mathrm{I}_{\{X_1>H_j\}}X_2\right] & \text{otherwise,} \end{cases} \tag{4.10}$$

where X_1 and X_2 are independent and $\mathbb{E}\left[X_2\right]=0$. As a result

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j|\mathcal{G}\right] = 0,$$

and therefore,

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j\right] = 0,$$

for all values of j. As such, each term in (4.5) is zero and therefore no bias is present.

At first sight this approach may appear to be expensive computationally in that double the number of sub-simulations must be performed. However, this is not the case. At each exercise time, a sub-simulation to determine the iterated exercise strategy must be performed. However, in terms of the Andersen–Broadie hedge, an additional independent sub-simulation is only needed when the iterated and non-iterated strategies disagree. Provided the numerical method reduces the variance by a factor of two or more over the method in [5], it will be worthwhile since the increase in time should be less than two-fold for any reasonable input strategy. Similarly, for a good non-iterated exercise strategy, the numerical removal of the bias does not take much additional time over the semi-analytical approach. However, there is the benefit that we are guaranteed to be free of bias.

4.2. **Adaptive Termination for Sub-Simulations.** In [5], a fixed number of paths is used for the sub-simulations. However, this is inefficient.

When choosing the number of paths for the sub-simulations to determine the iterated strategy, the primary concern is the accuracy of the iterated exercise strategy. When using the Andersen–Broadie approximate hedge as a control variate, the number of inner paths used to value this control variate also has an impact on the variance of the iterated lower bound, but this is of secondary importance. The convergence of the iterated strategy to the case of no Monte Carlo error is controlled by the probability of making the wrong decision due to the Monte Carlo error at each exercise time. In particular, the level of convergence depends on

$$\Pr\left((Z_j - H_j)(a - H_j) < 0\right), \tag{4.11}$$

for each value of j, where, as above, a denotes the true continuation value at t_j using the non-iterated strategy.

By re-writing Z_j as

$$Z_j \approx a + \frac{\sigma}{\sqrt{m}} X_j,$$

where $X_j \overset{\mathrm{d}}{=} \mathrm{N}(0,1)$, it is possible to write (4.11) as

$$\begin{aligned}
\Pr\left((Z_j - H_j)(a - H_j) < 0\right) &\approx \Pr\left((a + \frac{\sigma}{\sqrt{m}} X_j - H_j)(a - H_j) < 0\right), \\
&= \Pr\left((a - H_j)^2 < -\frac{\sigma}{\sqrt{m}}(a - H_j) X_j\right), \\
&= \Pr\left(|X_j| > \frac{|H_j - a|}{\sigma/\sqrt{m}}\right). \qquad (4.12)
\end{aligned}$$

From (4.12), the convergence of the iterated lower bound is controlled by the size of

$$\frac{|H_j - a|}{\sigma/\sqrt{m}}. \qquad (4.13)$$

The greater the size of (4.13), the greater the convergence to the case of no Monte Carlo error, and the tighter the lower bound should be.

This suggests that the sub-simulations should continue until a certain level is achieved in (4.13), subject to a given minimum and maximum number of paths. One issue in doing this is what to use for a and σ, since these will obviously not be known. However, it is possible to use the least-squares continuation value estimate for a and an estimate based on the results of the sub-simulation up to the current path for σ. By using this estimate of σ, a slight bias will enter the estimates obtained using the sub-Monte Carlo simulations. However, numerical tests indicate this has essentially no impact on results.

By terminating sub-simulations based on the method described here, time is not wasted simulating paths when the exercise decision is already clear. This can result in significant time savings.

4.3. **Excluding Sub-Optimal Points.** Significant time savings can be obtained by considering provably sub-optimal points, that is points at which it can be shown that the optimal exercise strategy would not exercise. Bender, Kolodko and Schoenmakers, [6], introduced the use of sub-optimal points in the policy iteration method and demonstrated that significant reductions in computation time could be achieved since sub-simulations to assess the iterated exercise strategy are not required at sub-optimal points.

When using the Andersen–Broadie hedge as a control variate together with the numerical bias removal, the exclusion of sub-optimal points is trivial to apply. Sub-simulations for the exercise strategy are not required at sub-optimal points, and sub-simulations used to value the hedge are only needed when the iterated and non-iterated strategies disagree.

However, if the semi-analytical bias removal technique is used, there are two additional changes to be made. First, if the non-iterated strategy does not exclude sub-optimal points, then a sub-simulation will still be required to update the Andersen–Broadie hedge control variate if the iterated and non-iterated strategies disagree. Second, the exclusion of sub-optimal points effects the bias removal approximation. In particular, (4.7) becomes

$$\mathbb{E}\left[\mathrm{I}_{j<\tau^*}\epsilon_j|\mathcal{G}\right] \approx \begin{cases} 0 & \text{if } \tau^* \leq j-1; \\ 0 & \text{if } \tau^* > j-1 \text{ and sub-optimal at } t_j; \\ \frac{\hat{\sigma}}{\sqrt{m}}\phi\left(\frac{H_j-\hat{a}}{\hat{\sigma}/\sqrt{m}}\right) & \text{otherwise.} \end{cases}$$

The middle line above follows since at provably sub-optimal points where exercise has not occurred previously the indicator is clearly one, and the expectation of the Monte Carlo error is zero.

All other improvements introduced are not affected by the exclusion of sub-optimal points.

4.4. **Other Implementation Issues.** The use of control variates can be applied to the sub-simulations when using policy iteration. This allows the number of paths used in the sub-simulations to be reduced without affecting accuracy and can bring about significant reductions in computation times. Note that when applying the adaptive termination method of Section 4.2, the maximum and minimum number of paths can also be reduced accordingly with no loss of accuracy. A new control variate suitable for this purpose will be introduced in Section 7.

Another implementation issue arises when the Andersen–Broadie hedge control variate is used with policy iteration. We can easily write,

$$\mathbb{E}\left[\sum_{j=1}^{\tau^*} C_j + H_{\tau^*} - \hat{M}_{\tau^*}\right] \tag{4.14}$$

$$= \mathbb{E}\left[\sum_{j=1}^{\tau^*} C_j + H_{\tau^*} - \left(\hat{M}_{\tau^*} + Z_0\right)\right] + \mathbb{E}\left[Z_0\right],$$

$$= \mathbb{E}\left[\sum_{j=1}^{\tau^*} C_j + H_{\tau^*} - \left(\hat{M}_{\tau^*} + Z_0\right)\right] \tag{4.15}$$

$$+ \mathbb{E}\left[Y_1 + C_0\right]. \tag{4.16}$$

Rather than estimate (4.14) directly, it is beneficial to break-up the estimation procedure. This is done by using a separate simulation to estimate the non-iterated lower bound, which is equivalent to calculating (4.16). In a second independent simulation the policy iteration improvement is estimated, where the policy iteration improvement is given by (4.15). The policy iteration improvement is equal to

$$\mathbb{E}\left[\sum_{j=1}^{\tau^*} C_j + H_{\tau^*}\right] - \mathbb{E}\left[\sum_{j=1}^{\tau} C_j + H_{\tau}\right],$$

with τ denoting the non-iterated exercise strategy, and thus gives the improvement in accuracy due to policy iteration.

By achieving this separation, it is easier to apply different improvements to the two simulations. For example, the use of quasi-Monte Carlo together with Brownian bridging as described in [22] can easily be applied to the non-iterated lower bound only, but not elsewhere. In addition, the number of paths used to estimate the non-iterated lower-bound is no longer tied to that needed to estimate the policy iteration improvement. Note that provided independent simulations are used, breaking up the lower bound procedure does not affect variance.

Finally, in the context of the displaced-diffusion LIBOR market model, efficiency can be improved by using the log-Euler drift approximation instead of the more accurate predictor-corrector approximation from [21] for the sub-simulations. Although a very small bias is introduced if the Andersen–Broadie hedge control variate is used and the log-Euler drift approximation is only used for the sub-simulations, numerical tests indicate this has no noticeable effect on results.

5. Adaptive Basis Functions

The greatest frustration of the least-squares approach is choosing basis functions. For complicated products, this can be very tedious as it is largely based on trial and error, yet it has a significant impact on the results. While policy iteration reduces the importance of this choice, having a good set of basis functions is still important from a practical point of view, because the better the non-iterated strategy, the closer it is to the iterated strategy and the better the variance reduction achieved when using the Andersen–Broadie hedge control variate.

We propose taking a simple adaptive approach to choosing basis functions. The least-squares algorithm remains essentially unchanged. However, we introduce an adaptive component into the regression step. As part of the algorithm, the following need to be specified at each exercise time:

- a set of base explanatory variables,

- a set of additional explanatory variables,
- the maximum number of explanatory variables,
- the transformations from explanatory variables to basis functions,
- a criterion for comparing the fits obtained using different sets of explanatory variables.

Now, consider the algorithm:

(1) Generate a set of sample paths, storing the relevant information at each exercise time. Only the values of the explanatory variables at each exercise time should be stored, rather than the values of the basis functions as is often done in the standard least-squares method.

(2) Iterative step.
   - Regress the basis functions for each allowed set of explanatory variables against the observed continuation values at the relevant exercise time.
   - The set of explanatory variables that provides the best fit to the data according to the given criterion is then chosen.
   - The explanatory variables must include the base set plus any number (including zero) of the additional variables such that the maximum number is not exceeded.
   - Update path-wise observations of continuation value using the exercise strategy obtained by the best set of explanatory variables.
   - Repeat, moving backwards, until all exercise dates have been covered.

(3) Perform an independent simulation to obtain a bias-free estimate of a lower bound.

Based on the advice given in [31], we recommend using second order polynomials in the explanatory variables as basis functions. For comparing the fits obtained using different sets of basis functions, a simple adjusted r-squared value is sufficient. In particular, if SSE and SST give the error and total sum of squares respectively for a given regression, then the set of basis functions that maximises

$$R^2_{\text{adj}} = 1 - \left(1 - \frac{\text{SSE}}{\text{SST}}\right) \frac{k-1}{k-1-l},$$

should be chosen, where k gives the number of points included in the regression and l denotes the number of basis functions. See [15] for a discussion of adjusted r-squared values.

Although performing an exhaustive search could be time-consuming if more than two additional explanatory variables were included on top of the base variables, in the examples considered, excellent results were obtained including only a single explanatory variable over two simple base variables. The efficiency of this approach also benefits from the fact that a number of calculations can be re-used, performing the least-square regressions is very quick, and most importantly, when a two-pass approach is used to calculate lower bounds, the number of first pass paths needed is relatively small.

Consider the LIBOR market model. At an exercise time corresponding to tenor date T_j, using the base variables $f_j(T_j)$ and $SR_{j+1,n}(T_j)$ (the swap-rate starting at the next tenor date and running to the final tenor date, evaluated at T_j), together with the additional explanatory variables

$$P(T_j, T_{j+\beta}), P(T_j, T_{j+2\beta}), P(T_j, T_{j+3\beta}), \dots,$$

where the maximum number of variables is three, is particularly effective. In the case of a path-dependent product, an additional base variable should be used to encapsulate the path-dependence. We have found that using

$$\beta = \left\lfloor \frac{n - \eta(T_j)}{20} \right\rfloor + 1,$$

where $\lfloor x \rfloor$ denotes the integer part of x, provides a good trade-off between speed and accuracy.

6. Double Regression Enhancement

It is possible to generically improve least-squares type exercise strategies by using more than one regression at each exercise time. As least-squares regressions are relatively very quick to perform, the additional time required to perform an extra regression is negligible.

One method that takes advantage of this idea is based on the observation that least-squares exercise strategies are usually adequate to make the correct exercise decision when options are deeply in or out of the money, but have greater trouble when the decision is not clear. The same observation was used in [13] to improve policy iteration. As such, we suggest using an additional regression to obtain a better fit to continuation values when the option is not “deeply” in or out of the money, where the degree of moneyness is determined by the initial continuation value estimate. The algorithm is as follows:

(1) Generate a set of sample paths, storing the relevant information at each exercise time as per the least-squares method.
(2) Iterative step.

- Assume we are at exercise time t_j.
- Perform an initial regression as per the least-squares method. At each point k this gives an estimate of the discounted continuation value,

$$\hat{V}_j^k.$$

- Use this estimate to determine the moneyness of the Bermudan at that point according to the size of

$$\left|\hat{V}_j^k - H_j^k\right|.$$

- Perform a second regression, including only points that satisfy

$$\left|\hat{V}_j^k - H_j^k\right| < \gamma_j.$$

- This gives a new discounted continuation value estimate when the option is not deeply in or out of the money

$$\tilde{V}_j^k.$$

- Update path-wise observations of continuation value where exercise occurs according to

$$\begin{cases} \tilde{V}_j^k \leq H_j^k, & \text{if } \left|\hat{V}_j^k - H_j^k\right| < \gamma_j; \\ \hat{V}_j^k \leq H_j^k, & \text{otherwise.} \end{cases} \tag{6.1}$$

- Repeat, moving backwards, until all exercise dates have been covered.

(3) Perform an independent simulation to obtain a bias-free estimate of a lower bound. Extending the notation used in the first pass to the second pass, on path k, at exercise time t_j in the second simulation exercise according to (6.1).

When using the double regression approach with the adaptive basis functions of the previous section, the set of explanatory variables should be determined as usual in the first regression, and that set should then be used in the second regression.

7. Delta Hedge Control Variate

When working with a model that forms a complete market, such as the LIBOR market model, it is possible to trade continuously in the underlying assets of the model to perfectly replicate any contingent claim using Delta hedging; see [25]. If the Deltas can be calculated exactly and the control variate portfolio can be updated continuously, zero variance would be achievable. Although the Deltas must be estimated and continuous updating is not possible in a discrete simulation, useful variance reductions are still possible. Clearly, the key to this approach is obtaining Delta estimates quickly and easily, which will now be explained.

The innovation introduced here is using the least-squares continuation value estimates to approximate the Deltas needed to carry out Delta hedging. At each evolution time in a simulation, it is possible to use a simple least-squares regression to obtain an estimate of the continuation value at that point using the least-squares lower bound method. Note that the regressions are performed at each evolution time, regardless of whether or not it is an exercise time, so that the control variate can be updated as often as naturally possible. (Of course it is possible to introduce finer discretisations so that the control variate can be updated even more frequently. However, for models where each step involves a significant computational cost, such as the LIBOR market model, this will generally not be efficient.) These give estimates of the value of the option as a function of the basis functions. If the partial derivatives of the basis functions with respect to the model fundamental tradeable assets can be calculated, then it is possible to calculate the partial derivatives of the continuation value estimates to obtain Delta estimates. These estimates can then be used to form a Delta hedge portfolio across each step in the simulation, where the units of each asset held equals the corresponding Delta estimate. All cash flows are dealt with by investing in the numeraire asset. The Delta hedge portfolio can then be used as a control variate, with hedging taking place until the time of exercise, upon which the hedge portfolio is dissolved. If the hedging worked well, the control variate portfolio should have a similar value to that generated by the product and a sizable variance reduction should be realised overall.

Consider the LIBOR market model. Here the fundamental tradeable assets are the zero-coupon bonds associated with each tenor date. Consider the Delta hedge portfolio. At evolution time T_j, assume we are holding a portfolio worth a certain amount (this will be zero at T_0) and that we have an estimate of continuation value

$$\hat{V}_j\left(P(T_j, T_{j+1}), P(T_j, T_{j+2}), \ldots, P(T_j, T_{n+1})\right).$$

If the basis functions consist of polynomials of forward rates, swap rates and zero-coupon bonds as suggested in Section 5, then

$$\frac{\partial}{\partial P(T_j, T_k)} \hat{V}_j\left(P(T_j, T_{j+1}), P(T_j, T_{j+2}), \ldots, P(T_j, T_{n+1})\right)$$

can be calculated easily for all $k = j+1, \ldots, n+1$. We hold $\frac{\partial}{\partial P(T_j,T_k)}\hat{V}_j$ units of $P(., T_k)$ across the step, investing in the numeraire asset to take account of any additional cash. At the next evolution time, the portfolio is worth

$$\sum_{k=j+1}^{n+1} \left(\frac{\partial \hat{V}_j}{\partial P(T_j, T_k)} \right) P(T_{j+1}, T_k),$$

plus whatever was invested in the numeraire asset, where the arguments of $\hat{V}_j$ have been dropped to ease notation. If the exercise strategy says exercise or we are at the final evolution time, hedging is stopped and the portfolio is dissolved. Otherwise, we repeat. In particular, the total discounted value of the Delta hedge portfolio at time T_i, denoted B_i, is given by

$$B_i := \begin{cases} 0, & i = 0; \\ \sum_{j=1}^{i} \left(\frac{\sum_{k=j}^{n+1} \left(\frac{\partial \hat{V}_{j-1}}{\partial P(T_{j-1}, T_k)} \right) P(T_j, T_k)}{N(T_j)} - \right. \\ \left. \frac{\sum_{k=j}^{n+1} \left(\frac{\partial \hat{V}_{j-1}}{\partial P(T_{j-1}, T_k)} \right) P(T_{j-1}, T_k)}{N(T_{j-1})} \right), & i = 1, 2, \ldots n+1, \end{cases} \tag{7.1}$$

Rather than sample and average over observations of

$$N(0) \left(\sum_{j=1}^{\bar{\tau}} C_j + H_{\bar{\tau}} \right),$$

use

$$N(0) \left(\sum_{j=1}^{\bar{\tau}} C_j + H_{\bar{\tau}} - B_{\bar{\tau}} \right), \tag{7.2}$$

where to ease the notation we have assumed that exercise can occur at each tenor date. This is done so that the stopping time $\bar{\tau}$ is in terms of tenor dates rather than exercise dates. However, with a minor adjustment the following analysis goes over unchanged.

Although it has been argued that using (7.2) should lead to a variance reduction, it still remains to be shown that it does not introduce bias.

To show that there is no bias, we will extend the notation of Section 4.1.1. Let $\tilde{\mathcal{H}}_i$ contain the information generated by the driving Brownian motions and any sub-simulations up to and including time T_i. We have used this enlarged filtration to show that the control variate can be used with policy iteration, although we only use it for non-iterated lower bounds. Notice that B_i forms a martingale with respect to $\tilde{\mathcal{H}}_i$. This follows since the discounted value processes for the zero-coupon bonds are martingales with respect to $\tilde{\mathcal{H}}_i$, and as such,

$$\mathbb{E} \left[\frac{\sum_{k=j}^{n+1}(\frac{\partial}{\partial P(T_{j-1},T_k)}\hat{V}_{j-1})P(T_j,T_k)}{N(T_j)} - \frac{\sum_{k=j}^{n+1}(\frac{\partial}{\partial P(T_{j-1},T_k)}\hat{V}_{j-1})P(T_{j-1},T_k)}{N(T_{j-1})} | \; \tilde{\mathcal{H}}_{j-1} \right] = 0,$$

for each value of j, since the Deltas for the step T_{j-1} to T_j are $\tilde{\mathcal{H}}_{j-1}$-measurable.

Since the problem can be formulated to ensure that $\bar{\tau}$ is a finite stopping time with respect to $\tilde{\mathcal{H}}_i$, it then follows by the Optional Sampling theorem that

$$\begin{aligned} \mathbb{E}\left[B_{\bar{\tau}}\right] &= B_0, \\ &= 0, \end{aligned}$$

and therefore that the control variate is unbiased. Clearly, the above arguments easily translate to any other model where this method could be used. Note that in a discrete simulation, some models, such as the LIBOR market model, may exhibit discretisation bias which will mean we only have an approximate martingale. However, provided reasonable approximations are used, this discretisation bias can be assumed to be negligible; see [8] for the LIBOR market model.

In terms of practical implementation, there is an issue in that an estimate of the continuation value at T_0 as a function of the model assets is needed so that the Deltas for the initial step can be estimated. We suggest that when building the exercise strategy, start each path from a random initial forward rate curve so that one can perform a least-squares regression to obtain a Longstaff-Schwartz style estimate of the option value at T_0. A simple way to introduce randomness to the initial forward rate curve is to set

$$\hat{f}_j^i(T_0) = f_j(T_0)\exp\left(-\frac{1}{2}a^2 + aX_j^i\right),$$

where $X_j^i \overset{\mathrm{d}}{=} N(0,1)$, and $\hat{f}_j^i(T_0)$ denotes the j^{th} initial forward rate on the i^{th} path used to build the exercise strategy.

Another issue arises when sub-optimal points are excluded from the least-squares regressions; see [29], [9]. When this is done, the continuation value estimates for the exercise strategy are obtained by fitting to the regions where exercise may occur. However, with the Delta hedge control variate, it is important to have good estimates of continuation value when exercise is unlikely to occur because upon exercise, the

Delta hedging is also stopped. As such, if sub-optimal points are excluded from the exercise strategy regressions, we recommend performing a second regression including all points for the Delta hedge control variate at each exercise time. As mentioned in Section 6, performing an additional regression should not add much time.

However, the Delta hedge control variate also benefits from some of the other improvements suggested in this paper. In particular, the improved fits obtained using the adaptive basis functions of Section 5 lead to improved Delta estimates.

We note that there are certain similarities between our approach and that of Belomestny, Bender and Schoenmakers, [4], who develop a martingale that can be used to calculate upper bounds without sub-Monte Carlo simulations. In addition, they use their martingale as a control variate to improve the efficiency of the Andersen–Broadie method for upper bounds. Apart from the connection of both approaches to Delta hedging, our control variate portfolio also forms a martingale. Thus, the approximate Delta hedge portfolio introduced here can be used to find upper bounds via non-nested Monte Carlo simulation. However, we believe that the Andersen–Broadie approach to upper bounds is superior, and this approach is best suited to use as a control variate. Although in this paper we focus on the non-iterated lower bound and policy iteration, the Delta hedge control variate can be used to improve the efficiency of the sub-simulations for the Andersen–Broadie upper bound method just as easily, as in [24] and [4]. We also note that both the method in [4] and the one suggested here produce similar efficiency improvements in many cases.

8. Results

8.1. **Products Studied.** The improvements suggested in this paper are applied to three different products; cancellable CMS steepener, snowball and vanilla swaps (also referred to as Bermudan swaptions). In each case, coupons are paid with a natural time-lag and the underlying swap can be cancelled by the issuer any time after the second coupon payment. All swaps are assumed to have a \$1 notional. We focus on the position of the issuer whom it is assumed receives the floating LIBOR rate and pays the coupon specified by the particular product. Upon cancellation, no rebate is received. So,

$$S_j = \begin{cases} \left(f_{j-1}(T_{j-1}) - s_j\right)\delta_{j-1}, & j = 2, \dots n+1; \\ 0, & \text{otherwise}, \end{cases}$$

where s_j depends on the particular product.

We consider pricing products with 6, 10 and 20 year underlying swaps, where coupons are paid half-yearly with the first coupon due in one year. As such,

$$\delta_j = 0.5,$$

for all values of j and we consider $n = 12, 20, 40$ for both cancellable snowball and vanilla swaps, and correspondingly $n = 51, 59, 79$ for CMS steepeners. The larger values of n for CMS steepeners arise because a greater span is required in the tenor dates to cover the CMS swap rates required to determine the coupon payments.

Consider the coupon structures under the different products.

8.1.1. *CMS Steepener.* We assume the same structure for CMS steepener coupons as in [34]. In particular, s_j takes the form

$$s_j = \begin{cases} e, & j \leq 3; \\ \max\left[c\left(SR_{j-1,j-1+a}(T_{j-1}) - SR_{j-1,j-1+b}(T_{j-1})\right), d\right], & \text{otherwise,} \end{cases}$$

where $SR_{x,y}(t)$ denotes the swap rate from T_x to T_y evaluated at t. For parameters, it is assumed

$$c = 1.5, a = 40, b = 4, d = 0, e = 0.0955.$$

8.1.2. *Snowball.* When pricing cancellable snowball swaps, s_j takes the following form,

$$s_j = \begin{cases} c, & j \leq 3; \\ \max\left[s_{j-1} + a_{j-1} - f_{j-1}(T_{j-1}), b\right], & \text{otherwise.} \end{cases}$$

This is the same coupon structure as the one used in [5] and [9]. It is assumed the parameters take the following values,

$$c = 0.07, b = 0, a_j = 0.015 + 0.0025 \left\lfloor \frac{j-1}{2} \right\rfloor.$$

8.1.3. *Bermudan Swaption.* For Bermudan swaptions, the underlying swap is just a vanilla swap, and as such, s_j is a constant. We took

$$s_j = 0.04,$$

for all values of j.

8.2. **LIBOR Market Model Set-up.** We consider pricing cancellable exotic interest rate options in the displaced-diffusion LIBOR market model. An initially increasing forward rate curve is assumed, with

$$f_i(0) = 0.008 + 0.002i + x,$$

where x is varied depending on the length of the product considered. The values of x for the 6, 10 and 20 year products are 0.01, 0.005

and 0 respectively independent of the particular product. The common "abcd" time-dependent volatility structure is used, with

$$\sigma_i(t) = \begin{cases} 0, & t > T_i; \\ (0.05 + 0.09(T_i - t)) \exp\left(-0.44(T_i - t)\right) + 0.2, & \text{otherwise,} \end{cases}$$

and instantaneous correlation between the driving Brownian motions is assumed to be of the form

$$\rho_{i,j} = \exp\left(-\phi|i - j|\right),$$

with $\phi = 0.0669$. Displacements for all forward rates are assumed to be equal, with

$$\alpha_j = 1.5\%,$$

for all values of j.

In evolving the forward rates, the predictor-corrector drift approximation from [21] is used. However, the use of the log-Euler drift approximation is also used for the sub-simulations in the policy iteration method. Due to the accuracy of the predictor-corrector method in particular, we evolved the forward rates to each tenor date in a single step.

A five factor model is used in all examples, where the factor reduction is performed on the covariance matrices across each step in the simulation. Due to the significant factor reduction, the method for calculating drifts in [26] is used.

8.3. **Numerical Results.** We start by comparing two improvements to the least-squares method, the adaptive choice of basis functions described in Section 5 and the double regression approach of Section 6. The relevant results are contained in Tables 8.1, 8.2 and 8.3. In each of these tables, LS is used to indicate that the standard least-squares method described in [2] and [27] was used, SO to indicate that provably sub-optimal points were excluded as in [9], DR to indicate that the double regression approach was used, and LSA to indicate that the combined least-squares and Andersen method introduced in [5] was used. In each case, the lower bounds were obtained by using a two-pass approach. An initial pass using Mersenne Twister pseudo-random numbers was used to develop an approximate exercise strategy. For the 6 and 10 year products, 5000 first pass paths were used, and 10000 were used for the 20 year products. The number of paths was chosen to be close to what would be used in practice, as it represents a good trade-off between speed and accuracy. This is important, as in the context of the least-squares method it is not useful to have improvements that only work when an unpractically large number of first pass paths are required. A second independent simulation was then used to obtain an unbiased estimate of the lower bound. For the second pass, 2^{18} paths using Sobol quasi-random numbers and Brownian bridging (see [22]) were used to ensure a sufficient level of convergence. The variables used in the adaptive basis functions are those described in Section 5. However, when pricing snowballs, the current snowball coupon is also used as a base explanatory variable to capture the path-dependence. The upper bounds were calculated using the extension by Joshi, [27], to the Andersen–Broadie method for upper bounds. All duality gaps were calculated using 2500 paths for both the outer and inner simulations. The corresponding lower bounds were obtained using the adaptive basis functions, together with the double regression approach and the exclusion of provably sub-optimal points, and were developed using 50000 first pass paths. Such a large number of paths were used to obtain the best possible estimate of the upper bound to which the different lower bound methodologies could be compared. All computations were done on a laptop with a 2 GHz Intel Core2Duo processor.

For the double regression enhancement, we used

$$\gamma_j = 0.03,$$

for all values of j. Numerical tests not included here have found this to be particularly effective and robust.

It is easy to see that the adaptive basis functions perform very well when compared with using carefully chosen product-specific basis functions. The adaptive basis functions almost never do worse, and often provide significant improvements in accuracy, particularly for long-dated products and when no additional improvements are used with the least-squares method. In addition, it can be seen that the use of the adaptive basis functions adds little absolute computational time when building the approximate exercise strategy. These results are particularly pleasing since they demonstrate that the hand-crafting usually required when choosing basis functions can essentially be removed at almost no cost. However, for the 20 year products in particular, the duality gaps are still large and policy iteration is required to improve the lower bound further.

Tables 8.1, 8.2 and 8.3 also show the effectiveness of the double regression approach. As with the above results, using the double regression approach nearly always leads to an improvement in accuracy at almost zero additional computational cost (the only time it did not was when pricing the 20 year Bermudan swaption with the adaptive basis functions and the exclusion of provably sub-optimal points). In contrast, the combined least-squares and Andersen method often results

		LS	SO	DR	LSA	SO, DR	SO, LSA
6yr	Adapt	-363.7	-357.7	-356.8	-357.8	-357.1	-357.2
		(2.55)	(2.75)	(2.63)	(4.27)	(2.78)	(3.72)
	SJ	-370.5	-362.8	-359.8	-360.2	-359.2	-361.5
		(2.44)	(2.69)	(2.41)	(3.06)	(2.69)	(3.03)
10yr	Adapt	-233.4	-222.2	-221.4	-226.8	-220.3	-226.7
		(4.94)	(5.19)	(5.05)	(8.16)	(5.22)	(7.22)
	SJ	-247.7	-243.8	-231.1	-249.4	-231.7	-247.8
		(4.42)	(4.91)	(4.38)	(5.33)	(4.92)	(5.64)
20yr	Adapt	-64.2	-24.8	-30.9	-77.8	-16.7	-22.7
		(27.42)	(24.63)	(26.88)	(43.55)	(24.27)	(32.95)
	SJ	-143.4	-72.0	-89.1	-191.8	-60.0	-89.9
		(23.57)	(22.80)	(20.50)	(26.33)	(22.73)	(26.13)

TABLE 8.1. Lower bounds and corresponding first pass times for cancellable CMS steepener swaps comparing the adaptive method for choosing basis functions with the set used in [34], denoted SJ. Lower bounds are in basis points and times are in seconds. The corresponding upper bounds for the 6, 10 and 20 year products are -355.1, -212.4 and 9.5 basis points.

in a reduction in accuracy and can add significant computation time. Note that the optimizations for the Andersen part are carried out using the implementation of the Simplex method from *Quantlib*. As such, the double regression approach appears to be significantly more robust and effective (both in terms of accuracy and speed) compared to the combined least-squares and Andersen method. Combining the least-squares and Andersen method with the double regression approach did not lead to any significant additional improvements. In addition, for shorter dated products, the double regression approach generally removes the need to exclude sub-optimal points. This could be useful as the double regression approach is slightly more generic.

Table 8.4 shows the effectiveness of the Delta hedge control variate in the context of the standard least-squares method. In each example, 2^{18} second pass paths using Mersenne Twister pseudo-random numbers were used with and without a control variate, with the results then being used to calculate the data in Table 8.4. The efficiency improvement was calculated as follows. If χ_{CV} and χ denote the standard errors obtained with and without the control variate respectively, and t_{CV}

		LS	SO	DR	LSA	SO, DR	SO, LSA
6yr	Adapt	100.3	104.3	105.0	103.7	105.0	104.5
		(0.74)	(0.67)	(0.80)	(2.58)	(0.72)	(1.61)
	BJ	100.3	104.3	105.0	103.7	105.1	104.5
		(1.89)	(1.58)	(1.91)	(4.88)	(1.69)	(3.28)
10yr	Adapt	335.2	348.3	349.3	341.3	349.1	349.1
		(1.89)	(1.58)	(1.91)	(4.88)	(1.69)	(3.28)
	BJ	335.1	348.3	349.0	342.7	349.1	349.1
		(1.38)	(1.20)	(1.30)	(2.63)	(1.28)	(1.99)
20yr	Adapt	521.2	594.7	585.1	550.3	625.4	627.3
		(11.49)	(8.92)	(12.55)	(27.22)	(9.44)	(16.55)
	BJ	502.2	586.1	575.3	544.5	617.7	618.6
		(8.40)	(7.04)	(8.42)	(18.67)	(6.59)	(10.66)

TABLE 8.2. Lower bounds and corresponding first pass times for cancellable snowball swaps comparing the adaptive method for choosing basis functions with the set used in [9], denoted BJ. Lower bounds are in basis points and times are in seconds. The corresponding upper bounds for the 6, 10 and 20 year products are 106.5, 358.9 and 650.2 basis points.

and t denote the corresponding second pass simulation times, then

$$\text{EI} = \frac{\chi^2}{\chi^2_{\text{CV}}}\frac{t}{t_{\text{CV}}}.$$

So, the efficiency improvement represents the reduction in computation time that can be obtained using the control variate. The results were obtained using the adaptive basis functions, together with the double regression approach and the exclusion of provably sub-optimal points for the non-iterated exercise strategy. The least-squares approach with adaptive basis functions was used to develop the continuation value estimates used by the control variate, with no points being excluded from the regressions. To introduce randomness to the initial forward rate curve as outlined in Section 7, we used $a = 0.3$.

Table 8.4 indicates that the Delta hedge control variate can bring about useful efficiency improvements, even when only a small number of first pass paths is used to develop the continuation value estimates. In addition, the efficiency improvements can remain strong, even when the length of the underlying contract increases. This is the case for both CMS steepeners and Bermudan swaptions, but not so much for

		LS	SO	DR	LSA	SO, DR	SO, LSA
6yr	Adapt	28.4	28.2	28.9	27.3	28.8	27.4
		(0.64)	(0.63)	(0.70)	(2.41)	(0.72)	(1.61)
	Pit	22.63	28.2	28.9	27.4	28.9	27.3
		(2.44)	(2.69)	(2.41)	(3.06)	(2.69)	(3.03)
10yr	Adapt	181.9	183.0	182.9	179.8	183.2	180.0
		(1.59)	(1.42)	(1.61)	(4.92)	(1.50)	(3.20)
	Pit	170.5	179.7	182.8	178.7	183.0	183.2
		(1.00)	(1.03)	(1.03)	(1.75)	(1.06)	(1.55)
20yr	Adapt	1063.5	1085.6	1074.0	1057.1	1081.3	1085.2
		(13.92)	(7.75)	(12.89)	(28.49)	(8.16)	(13.98)
	Pit	934.6	1075.6	1014.0	927.6	1079.2	1081.3
		(8.76)	(5.80)	(6.19)	(8.67)	(5.88)	(7.86)

TABLE 8.3. Lower bounds and corresponding first pass times for cancellable vanilla swaps comparing the adaptive method for choosing basis functions with the set used in [31], denoted Pit. Lower bounds are in basis points and times are in seconds. The corresponding upper bounds for the 6, 10 and 20 year products are 29.1, 185.3 and 1095.6 basis points.

snowballs, although reasonable improvements are still possible. In particular, even in the most severe test, a 20 year cancellable snowball swap, efficiency improvements of over 3 are still possible, with efficiency improvements up to and over 10 possible elsewhere. It is worth noting that the control variate works very well for Bermudan swaptions, consistently producing efficiency improvements over 10 for contracts of all lengths when only 5000 first pass paths are used. While we have looked at the control variate in the context of the least-squares method, the improvements here translate to very similar improvements in the policy iteration method, which we will now discuss.

We look at the improvements that can be obtained using policy iteration in Table 8.5 and Figure 8.1. Figure 8.1 illustrates the bias that a naive implementation of the Andersen–Broadie hedge control variate introduces, together with the corresponding bias free estimates obtained using the different approaches of Section 4.1.2. The line labelled BKS is obtained using the method suggested in [5]. Similar graphs were obtained for all other examples. We see that the smaller the number of sub-simulation paths, the greater the Monte Carlo error of the sub-simulations and the greater the bias introduced. In particular, the bias can be significant. Even when using 100 paths for the sub-simulations,

	2000		5000		10000	
	EI	SER	EI	SER	EI	SER
CMS						
6yr	5.46	2.68	5.95	2.72	6.80	2.97
10yr	4.16	2.32	6.11	2.80	6.22	2.86
20yr	3.32	2.05	4.61	2.43	4.91	2.50
SB						
6yr	6.59	3.01	10.37	3.76	10.52	3.80
10yr	4.92	2.56	6.82	2.97	7.52	3.15
20yr	2.92	1.90	3.09	1.97	3.44	2.07
BS						
6yr	7.05	3.04	11.87	3.95	14.40	4.38
10yr	8.84	3.36	10.29	3.66	13.25	4.15
20yr	10.38	3.62	11.65	3.86	13.34	4.10

Table 8.4. Efficiency improvements obtained using the Delta hedge control variate. The top row indicates how many first pass paths were used to develop the continuation value estimates. EI and SER are used to denote efficiency improvement and standard error reduction respectively.

the difference between the biased and bias-free policy iteration improvements is approximately 15 basis points, which is much greater than the Monte Carlo error. It is also possible to see that the semi-analytical approach is effective in removing the bias. Even when only 20 sub-simulation paths are used in this long-dated example, the difference between the numerical or BKS methods and the semi-analytical approach is small compared to the Monte Carlo error.

Table 8.5 shows the accuracy and speed of the various methods and improvements. The first three columns use all improvements (i.e. the adaptive basis functions, double regression enhancement and exclusion of sub-optimal points are used for the least-squares exercise strategy, and the adaptive termination of sub-simulations and exclusion of sub-optimal points are used when assessing the iterated exercise strategy. In addition, the log-Euler drift approximation and Delta hedge control variate are used for all policy iteration sub-simulations), but different techniques for dealing with the bias in the Andersen–Broadie hedge control variate. The Num column uses the numerical bias removal, SA the semi-analytic, and BKS the technique used in [5]. The next three columns all use the numerical bias removal, but do not use a given

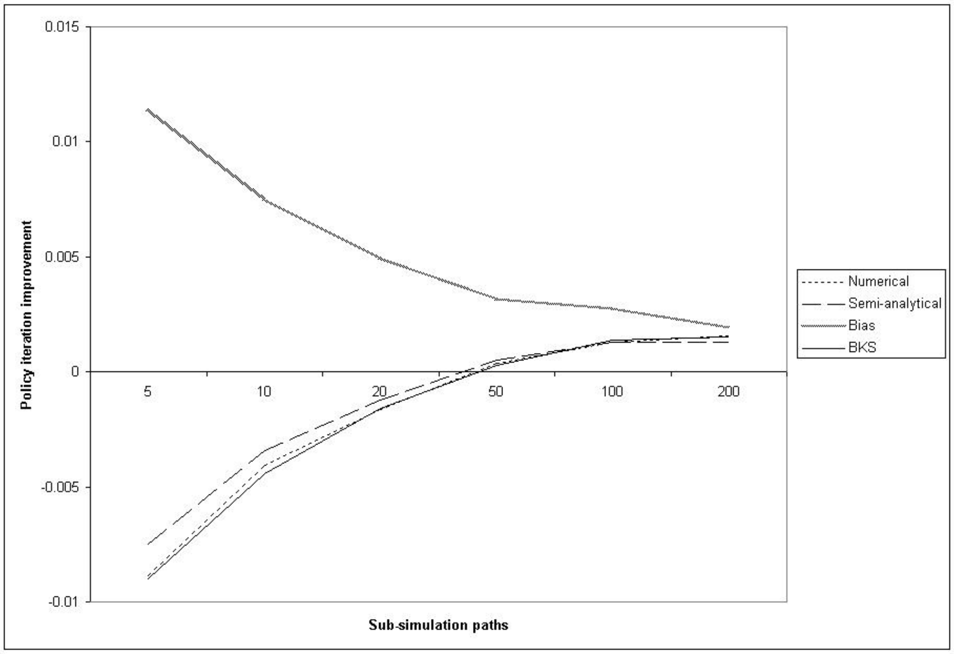

FIGURE 8.1. Estimated policy iteration improvements as a function of sub-simulation paths for a 20 year cancellable snowball swap. The Delta hedge control variate was used for the sub-simulations, together with the predictor-corrector drift approximation. Provably sub-optimal points were not excluded in the iterated exercise strategy. All estimated policy iteration improvements have a standard error of 2 basis points.

improvement. Non-Ad does not use the adaptive method for terminating sub-simulations, No SO does not exclude sub-optimal points in the iterated strategy and No LE uses the predictor-corrector drift approximation instead of the log-Euler approximation in the sub-simulations. For comparison, the final column gives the distance between the non-iterated lower bound and the upper bound. The upper bound is calculated using the same methodology as in Tables 8.1, 8.2 and 8.3. The distance between the iterated lower bound and the upper bound can be obtained by subtracting the policy iteration improvement from the Upper column. Note that the estimated policy iteration lower bounds can be obtained by adding the policy iteration improvement to the corresponding non-iterated lower bound in Tables 8.1, 8.2 and 8.3.

An additional implementation issue arises when using the numerical bias removal technique. In this case the sub-simulations used to value the control variate and assess the iterated exercise strategy are carried out separately. The adaptive termination technique of Section 4.2 is not relevant to the control variate sub-simulations, where instead a small standard error is required to reduce the variance of the lower bound estimator. As such, we used a fixed number of paths for the control variate sub-simulations. In particular, we used 50% more paths than the maximum number used for the iterative exercise strategy sub-simulations, as this was found to be efficient.

When using the adaptive method for terminating the sub-simulations to determine the iterated exercise strategy, we set the minimum and maximum number of paths at 15 and 200 respectively. For the non-adaptive method, 200 paths were used for all sub-simulations. For the numerical bias removal technique, we used 300 paths for the sub-simulations that determine the value of the Andersen–Broadie hedge control variate. In addition, the target value for (4.13) used to adaptively terminate the sub-simulations was taken to be 6 in all examples. The Delta hedge control variate was used for the sub-simulations in all examples.

The first thing to notice is that policy iteration can provide significant improvements in the accuracy of the lower bound. This is very useful, especially for the 20 year products. In particular, after one iteration, the greatest distance to the upper bound is approximately 10 basis points. We also emphasize that the upper bounds are undoubtedly imperfect and thus do not represent the true prices. As such, we expect the iterated lower bounds to be closer to the true prices than the upper bounds suggest.

As important is the time taken to calculate the policy iteration improvements. In every example, the simulation for the policy iteration improvement took well under half the time compared to the simulation for the non-iterated lower bound without convergence improvements, and in most cases took only a small fraction of the time. As such, policy iteration can be used to provide significant improvements in accuracy at modest increases in computation time. In particular, if control variates are used to increase the efficiency of the non-iterated lower bound, then the non-iterated lower bound and corresponding policy iteration improvement can be calculated in less time than the non-iterated lower bound without control variates and quasi-Monte Carlo. Even when the variance of the iterated lower bound (which, provided independent simulations are used, equals the variance of the non-iterated lower bound plus that of the policy iteration improvement) is taken into account, the iterated lower bound can still be estimated in similar computation times to the standard least-squares method. However, as demonstrated in [9], this sort of efficiency would not be possible without the improvements to the policy iteration algorithm introduced in this paper, that is, without practical policy iteration.

Table 8.5 also shows the improvements to the Andersen–Broadie hedge control variate to be effective, particularly for short-dated products where the non-iterated exercise strategy is generally very good. The increases in speed obtained by using the bias removal techniques suggested here over the method suggested in [5] range from factors of over 200 to 8.5. As expected, the semi-analytic bias removal technique is generally more efficient than the numerical removal, yet produces very similar lower bounds, mirroring the results in Figure 8.1. There were some cases where the semi-analytic approach took more time than the numerical bias removal. However, this only occurred for short-dated products, where the difference in absolute time was negligible.

On top of this, significant reductions in time are obtained by using the adaptive method for terminating sub-simulations discussed in Section 4.2. The fourth column of Table 8.5 shows that this consistently produced reductions in time of factors of approximately three.

The exclusion of sub-optimal points in the iterated exercise strategy and the use of the log-Euler drift approximation for sub-simulations also produced useful improvements, often providing reductions in time of 30 % and 20 % respectively.

By looking at the results obtained without the adaptive method for terminating sub-simulations and with using the predictor-corrector drift approximation for sub-simulations, it is easy to see that any bias introduced by applying either of the corresponding improvements is well within Monte Carlo error, and can assumed to be negligible.

We note one curiosity in the results of Table 8.5. For the six year Bermudan swaption, the estimated iterated lower bound is often above the upper bound. However, since the policy iteration improvements have a standard error of 0.5 basis points, all differences are well within one standard error of the policy iteration improvement. We therefore believe that these results are due to Monte Carlo error.

9. Conclusion

We have demonstrated that by using the practical policy iteration method introduced in this paper, very tight lower bounds can be obtained for long-dated Bermudan interest rate derivatives in the displaced-diffusion LIBOR market model. These lower bounds can be obtained in similar computation times to the standard least-squares method. The levels of accuracy and efficiency were only achievable due to a number

	Num	SA	BKS	Non-Ad	No SO	No LE	Upper
CMS							
6yr	1.77	1.96	1.22	1.45	1.87	1.58	2.0
	(1.37)	(2.40)	(210.45)	(3.14)	(1.90)	(1.61)	
10yr	6.26	6.87	5.37	6.26	6.51	5.57	7.9
	(7.10)	(6.32)	(234.63)	(20.04)	(9.50)	(8.19)	
20yr	14.12	12.94	15.27	13.91	11.57	13.34	26.2
	(42.89)	(31.63)	(363.82)	(113.21)	(47.50)	(55.59)	
SB							
6yr	0.70	0.35	0.66	0.70	0.33	0.70	1.5
	(0.09)	(0.26)	(196.09)	(0.27)	(0.11)	(0.10)	
10yr	7.04	7.73	8.08	7.57	6.91	7.04	9.8
	(7.48)	(6.78)	(227.50)	(24.35)	(10.98)	(9.37)	
20yr	13.90	13.63	13.33	15.09	15.34	13.12	24.8
	(21.27)	(15.89)	(335.49)	(52.06)	(31.35)	(26.28)	
BS							
6yr	0.70	0.43	-0.08	0.45	0.69	0.69	0.3
	(0.49)	(1.02)	(136.14)	(0.84)	(0.58)	(0.53)	
10yr	1.84	2.13	1.90	1.38	0.94	1.62	2.1
	(0.64)	(1.20)	(124.43)	(2.57)	(0.68)	(0.76)	
20yr	9.34	8.61	9.63	10.41	9.98	9.63	14.3
	(9.10)	(7.79)	(65.63)	(53.46)	(13.07)	(10.97)	

TABLE 8.5. Policy iteration improvements (in basis points) together with the percentage of time taken compared to that required to calculate the corresponding non-iterated lower bound without convergence improvements to the same level of standard error (in brackets). For the 6, 10 and 20 year products the policy iteration improvements have standard errors of 0.5, 1 and 2 basis points respectively.

of improvements introduced in this paper that resulted in significant acceleration to the policy iteration method, as well as improvements to the accuracy and efficiency of the least-squares method.

Our recommended algorithm for practical policy iteration is:

- use the least-squares approach with the adaptive basis functions of Section 5, the double regression enhancement of Section 6, and the exclusion of sub-optimal points as in [9], to develop the non-iterated exercise strategy,

- for the second pass of the non-iterated lower bound, use the Delta hedge control variate of Section 7 together with the use of quasi-Monte Carlo and Brownian bridging; see [22],
- use policy iteration, together with the Andersen–Broadie hedge control variate and the semi-analytic bias removal of Section 4.1.2, the adaptive termination for sub-simulations of Section 4.2, and the exclusion of sub-optimal points in the iterated exercise strategy, to calculate the policy iteration improvement
- for the sub-simulations, use the Delta hedge control variate and the log-Euler drift approximation (instead of predictor-corrector) to improve efficiency.

References

[1] F. M. Ametrano and M. S. Joshi. Smooth simultaneous calibration of the LIBOR market model to caplets and co-terminal swaptions, 2008. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1092665.

[2] A. Amin. Multi-factor cross currency LIBOR market model: implementation, calibration and examples, 2003. http://www.geocities.com/anan2999/.

[3] L. Andersen and M. Broadie. A primal-dual simulation algorithm for pricing multi-dimensional American options. *Management Science*, 50:1222–1234, 2004.

[4] D. Belomestny, C. Bender, and J. Schoenmakers. True upper bounds for Bermudan products via non-nested Monte Carlo. *Mathematical Finance*, 19:53–71, 2009.

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

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

[7] C. Bender and J. Schoenmakers. An iterative method for multiple stopping: convergence and stability. *Advances in Applied Probability*, 38:729–749, 2006.

[8] 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.

[9] C. J. Beveridge and M. S. Joshi. Juggling snowballs. *Risk*, December:100–104, 2008.

[10] A. Brace. *Engineering BGM*. Chapman and Hall, Sydney, 2007.

[11] A. Brace, D. Gatarek, and M. Musiela. The market model of interest rate dynamics. *Mathematical Finance*, 7:127–155, 1997.

[12] D. Brigo and F. Mercurio. *Interest Rate Models: Theory and Practice*. Springer Verlag, Heidelberg, 2001.

[13] M. Broadie and M. Cao. Improved lower and upper bound algorithms for pricing American options by simulation. *Quantitative Finance*, 8:845–861, 2008.

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

[15] N. R. Draper and H. Smith. *Applied Regression Analysis*. John Wiley and Sons Ltd, New York, 1998.
[16] C. P. Fries. Foresight bias and suboptimality correction in Monte-Carlo pricing of options with early exercise: classification, calculation and removal, 2005. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=839105.
[17] M. Giles and S. Xiaoke. Notes on using the nVidia 8800 GTX graphics card, 2008. http://people.maths.ox.ac.uk/ gilesm/hpc/NVIDIA/libor/report.pdf.
[18] P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, New York, 2004.
[19] J. M. Harrison and S. R. Pliska. A stochastic calculus model of continuous trading: complete markets. *Stochastic Processes and their Applications*, 15:313–316, 1983.
[20] M. Haugh and L. Kogan. Approximating pricing and exercising of high-dimensional American options: a duality approach. *Operations Research*, 52:258–270, 2004.
[21] C. Hunter, P. Jäckel, and M. S. Joshi. Getting the drift. *Risk*, July:81–84, 2001.
[22] P. Jäckel. *Monte Carlo Methods in Finance*. John Wiley and Sons Ltd., New York, 2001.
[23] F. Jamshidian. LIBOR and swap market models and measures. *Finance and Stochastics*, 1:293–330, 1997.
[24] M. S. Jensen and M. Svenstrup. Efficient control variates and strategies for Bermudan swaptions in a LIBOR market model. *Journal of Derivatives*, 12:20–33, 2005.
[25] M. S. Joshi. *The Concepts and Practice of Mathematical Finance*. Cambridge University Press, London, 2003.
[26] M. S. Joshi. Rapid drift computations in the LIBOR market model. *Wilmott Magazine*, May:84–85, 2003.
[27] 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.
[28] A. Kolodko and J. Schoenmakers. Iterative construction of the optimal Bermudan stopping time. *Finance and Stochastics*, 10:27–49, 2006.
[29] 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.
[30] M. Musiela and M. Rutkowski. Continuous-time term structure models: forward-measure approach. *Finance and Stochastics*, 1:261–292, 1997.
[31] 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.
[32] N. S. Rasmussen. Control variates for Monte Carlo valuation of American options. *Journal of Computational Finance*, 9:83–118, 2005.
[33] L. C. G. Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12:271–286, 2002.
[34] D. Skovmand and P. L. Jørgensen. The valuation of callable bonds with floored CMS-spread coupons. *Wilmott Magazine*, November:106–125, 2007.

Centre for Actuarial Studies, Dept of Economics, University of Melbourne, Victoria 3010, Australia
E-mail address: c.beveridge2@pgrad.unimelb.edu.au
E-mail address: mark.joshi@unimelb.edu.au