True upper bounds for Bermudan products via non-nested Monte Carlo*

Denis Belomestny[1], Christian Bender[2], John Schoenmakers[1]

July 26, 2007

Abstract

We present a generic non-nested Monte Carlo procedure for computing true upper bounds for Bermudan products, given an approximation of the Snell envelope. The pleonastic "true" stresses that, by construction, the estimator is biased above the Snell envelope. The key idea is a regression estimator for the Doob martingale part of the approximative Snell envelope, which preserves the martingale property. The so constructed martingale can be employed for computing tight dual upper bounds without nested simulation. In general, this martingale can also be used as a control variate for simulation of conditional expectations. In this context, we develop a variance reduced version of the nested primal-dual estimator (Andersen and Broadie, 2004). Numerical experiments indicate the efficiency of the proposed algorithms.

Keywords: Early exercise options, Monte Carlo simulation, linear regression.

1 Introduction

In recent years, much research on pricing of high-dimensional Bermudan derivatives was devoted to the approximation of the optimal exercise policy. Once a "good" but generally sub-optimal policy is known, a lower biased approximation of the Bermudan price can be found by straightforward Monte Carlo simulation of the underlying trajectories, stopped according to this policy. Most popular in this respect are the regression-based approaches of Carriere (1996), Longstaff and Schwartz (2001), Tsistsiklis and Van Roy (1999) and Clement et al. (2002). Another notable approach is backward

[1]Weierstrass Institute for Applied Analysis and Stochastics, Mohrenstr. 39, 10117 Berlin, Germany. {belomest,schoenma}@wias-berlin.de.

[2]Institute for Mathematical Stochastics, TU Braunschweig, Pockelsstr. 14, 38106 Braunschweig, Germany. c.bender@tu-bs.de (Corresponding author).

*Supported by the Deutsche Forschungsgemeinschaft through the SFB 649 Economic Risk, and the DFG Research Center MATHEON in Berlin.

construction of the exercise boundary using its suitable parametrization. This method is utilized by Andersen (2000) in the context of Bermudan swaptions. An important feature of these methods is their efficiency: At relatively low computational costs an approximative exercise policy can be constructed, and thereafter a straightforward Monte Carlo simulation gives a lower price.

The goal of this paper is an efficient method for computing an upper bound, given an approximation of the Snell envelope, for example, in the form of a pre-computed exercise boundary. Rogers (2001) and independently Haugh and Kogan (2004) developed a dual method which provides an upper bound for the Bermudan price, given an approximation of the Snell envelope. A multiplicative version of this method is studied by Jamshidian (2007). A comparative study of multiplicative and additive duals is provided in Chen and Glasserman (2007). Via the Doob martingale part of a "good" approximation of the Snell envelope, the dual approach gives a tight upper bound for the Bermudan price. The martingale part of the (generally unknown) Snell envelope would even result in the exact Bermudan price. Due to this fact the martingale part M of any "reasonable" approximation Y of the Snell envelope is a promising candidate for a "good" upper bound. Andersen and Broadie (2004) suggested to estimate this type of martingale upper bound by a simulation within a simulation approach. By the Doob decomposition we have

$$M_{T_{j+1}} - M_{T_j} = Y_{T_{j+1}} - E^{T_j}[Y_{T_{j+1}}]. \tag{1.1}$$

An inner Monte Carlo simulation is used to estimate the conditional expectation in (1.1), and an outer simulation is used to compute an outer expectation that determines the corresponding upper bound. Although the demand for nested simulation makes the Andersen and Broadie algorithm computationally extensive, it guarantees that the estimator for M, which fails to satisfy the martingale property in general, induces an upper bound estimate that is biased high. This important "biased high"-property is not shared in general, if faster estimation procedures such as regression methods are applied to estimate the conditional expectation in (1.1). The first attempt to overcome this difficulty was made in Glasserman and Yu (2005), where a special regression algorithm for preserving the martingale property of (1.1) is proposed. This algorithm, however, requires strong conditions on the basis functions, which may be hard to check in practice. As an alternative, Kolodko and Schoenmakers (2004) propose a different estimator which allows for a substantially reduced amount of inner simulations. While their procedure may be effective, it has a drawback: Their alternative estimator may fail to give an upper bound when the number of inner simulations used is too low.

In this paper we avoid estimating the conditional expectation in (1.1). Instead we construct, in a Brownian motion setting, an estimator for M that is based on the martingale representation theorem (Section 2). The main advantage is that the thus constructed estimators inherit the martingale property from M, if conditional expectations are estimated in a non-anticipative way. In particular the conditional expectations can be estimated by the popular linear regression method using basis functions without any restrictions on them (Section 3). The corresponding estimator $\widehat{M}$ for M is a martingale and consequently induces an upper bound. Therefore, if the lower bound Y is constructed by linear regression, a linear regression estimator $\widehat{M}$ for M can be implemented with comparable computational cost. Some results on the convergence of $\widehat{M}$ to M are presented in Theorem 2.2 and Remark 3.1.

In Section 4 we analyze how the estimator $\widehat{M}$ can alternatively be applied as control variate for the primal-dual algorithm of Andersen and Broadie (2004). Moreover, the martingale $\widehat{M}$ can be used to derive estimates for the hedge parameters of the Bermudan option in a complete market, as is stressed in Section 5.

Finally we present numerical examples in Section 6. In two case studies, the max-call on an asset basket and the Bermudan swaption in a multi-factor Libor market model, we find that the fast non-nested estimator introduced in this paper yields very good results. We also demonstrate a significant variance reduction effect when $\widehat{M}$ is used as a control variate for the primal-dual algorithm. Section 7 concludes.

2 Constructing dual upper bounds

We consider a Bermudan option that can be exercised at one date from the set $\mathcal{E} = \{T_0, \ldots, T_{\mathcal{J}}\}$. To simplify the notation we shall assume that $T_0 = 0$ and define $T := T_{\mathcal{J}}$. Let us further assume that we have a given pricing measure Q connected with a given discounting numeraire $\mathcal{N}$ on some filtered probability space. According to the Bermudan contract, when exercising at time $T_j \in \mathcal{E}$, the holder of the option receives a discounted payment of the form

$$H_{T_j} := h(T_j, X_{T_j}),$$

where $h(T_j, \cdot)$ is Lipschitz continuous and X_t is the solution of the SDE

$$dX_t = a(t, X_t)dt + b(t, X_t)dW_t \tag{2.1}$$

$$X_0 = x. \tag{2.2}$$

The coefficient functions $a : [0,T] \times \mathbb{R}^{\mathcal{D}} \to \mathbb{R}^{\mathcal{D}}$ and $b : [0,T] \times \mathbb{R}^{\mathcal{D}} \to \mathbb{R}^{\mathcal{D} \times D}$ are supposed to be Lipschitz in space and 1/2-Hölder continuous in time, with D denoting the dimension of the Brownian motion $W = (W^1, \ldots, W^D)^\top$ under the pricing measure Q. For now we do not assume additional regularity conditions on the diffusion coefficient $b(\cdot, \cdot)$. Throughout $(\mathcal{F}_t;\ 0 \le t \le T)$ is the augmented filtration generated by this Brownian motion. All expectations and conditional expectations are taken under the pricing measure Q. Conditional expectations under Q with respect to $\mathcal{F}_t$ will be denoted by $E^t[\cdot]$. The numeraire $\mathcal{N}$ is positive, adapted, and $\mathcal{N}_0 := 1$.

We think of X as a vector of financial quantities which is determined by some arbitrage free system of tradable quantities on the background. Of course all components of X may be tradable themselves, but for example X may be also a set of (Libor) interest rates which are determined by a system of (tradable) bonds.

Recall that for any martingale M_{T_j}, $0 \le j \le \mathcal{J}$ with respect to the filtration $(\mathcal{F}_{T_j};\ 0 \le j \le \mathcal{J})$ starting at $M_0 = 0$

$$Y^{up}(M) := E\left[\max_{0\le j\le \mathcal{J}} (H_{T_j} - M_{T_j})\right] \tag{2.3}$$

is an upper bound for the price of the Bermudan option with cash-flow H_{T_j}. Moreover, the Bermudan price is attained at the martingale part of the Doob decomposition of the discounted price process (Snell envelope). The latter process is denoted by $Y^*_{T_j}$.

Suppose some approximation Y_{T_j} of the Snell envelope is given. If Y is a good approximation and it is decomposed in its Doob decomposition

$$Y_{T_j} = Y_0 + M_{T_j} + U_{T_j}, \tag{2.4}$$

where the martingale M and the predictable process U start at zero, then we expect $Y^{up}(M)$ to be a close upper bound of Y^*_0. In principle, U and M can be found from Y via the relations

$$\begin{aligned} U_{T_{j+1}} - U_{T_j} &= E^{T_j}[Y_{T_{j+1}}] - Y_{T_j}, \\ M_{T_{j+1}} - M_{T_j} &= Y_{T_{j+1}} - E^{T_j}[Y_{T_{j+1}}]. \end{aligned} \tag{2.5}$$

If one estimates the conditional expectations in the above expressions – say, by standard regression methods –, the estimated version of M will loose the martingale property in general. In particular, it is not guaranteed that it induces an upper bound. We will now exploit the structure of the Brownian filtration to construct an approximation of M in a way that all conditional expectations can be estimated without loosing the martingale property.

Indeed, under the assumption that M_T is square integrable there is a square integrable (row vector valued) process $Z_t = (Z^1_t, \ldots, Z^D_t)$ satisfying

$$M_{T_j} = \int_0^{T_j} Z_t dW_t,\ j = 0, \ldots, \mathcal{J}. \tag{2.6}$$

Hence, our aim is to approximate Z instead of M and then make use of relation (2.6). Of course, we can estimate Z only at a finite number of time points. So we introduce a partition $\pi = \{t_0, \ldots, t_\mathcal{J}\}$ such that $t_0 = 0$, $t_\mathcal{J} = T$, and $\mathcal{E} \subset \pi$. We write formally, by (2.4) and (2.6),

$$Y_{T_{j+1}} - Y_{T_j} \approx \sum_{t_l \in \pi; T_j \le t_l < T_{j+1}} Z_{t_l}(W_{t_{l+1}} - W_{t_l}) + U_{T_{j+1}} - U_{T_j}.$$

Multiplying by the increment of the dth Brownian motion $(W^d_{t_{i+1}} - W^d_{t_i})$ and taking conditional expectations we obtain, by the $(\mathcal{F}_{T_j})_{j=1\ldots,\mathcal{J}}$-predictability of U

$$Z^d_{t_i} \approx \frac{1}{t_{i+1} - t_i} E^{t_i}\left[(W^d_{t_{i+1}} - W^d_{t_i})Y_{T_{j+1}}\right],\ T_j \le t_i < T_{j+1}.$$

This formal argumentation motivates the definition

$$Z^\pi_{t_i} := \frac{1}{\Delta^\pi_i} E^{t_i}\left[(\Delta^\pi W_i)^\top Y_{T_{j+1}}\right],\ T_j \le t_i < T_{j+1} \tag{2.7}$$

with an obvious definition of the increments, e.g. $\Delta^\pi W^d_i := W^d_{t_{i+1}} - W^d_{t_i}$. The corresponding approximation of the martingale M is

$$M^\pi_{T_j} := \sum_{t_i \in \pi; 0 \le t_i < T_j} Z^\pi_{t_i}(\Delta^\pi W_i). \tag{2.8}$$

Remark 2.1. In practice one often has an approximation $(\tau_1, \ldots, \tau_\mathcal{J})$ of the optimal stopping policy at hand. Such family of stopping times induces an approximation of the Snell envelope via $Y_{T_j} := E^{T_j}[H_{\tau_j}]$. Thanks to the tower property of conditional expectations, (2.7) can be rewritten as

$$Z^\pi_{t_i} := \frac{1}{\Delta^\pi_i} E^{t_i}\left[(\Delta^\pi W_i)^\top H_{\tau_{j+1}}\right],\ T_j \le t_i < T_{j+1}. \tag{2.9}$$

Hence, there is no need to compute the conditional expectations in the definition of Y.

The following theorem shows that the martingale M^π based on the discretized Itô integral converges to the original one, M.

Theorem 2.2. *(i) We have,*

$$\lim_{|\pi|\to 0} E\left[\max_{0\le j\le \mathcal{J}} |M^\pi_{T_j} - M_{T_j}|^2\right] = 0$$

where $|\pi|$ denotes the mesh of π.
(ii) Suppose that either $Y_{T_j} = u(T_j, X_{T_j})$ or $Y_{T_j} = u(T_j, X^{\bar\pi}_{T_j})$, $j = 1, \ldots, \mathcal{J}$, where the functions $u(T_j, \cdot)$ are Lipschitz continuous and $X^{\bar\pi}_{t_i}$ is the Euler approximation of X_t corresponding to a partition $\bar\pi \supset \pi$. Then there exists a constant $C > 0$ such that

$$E\left[\max_{0\le j\le \mathcal{J}} |M^\pi_{T_j} - M_{T_j}|^2\right] \le C|\pi|$$

The proof is postponed to the Appendix.

Note that, for two martingales $M^{(1)}$ and $M^{(2)}$ starting in 0, one can obtain by straightforward manipulations

$$|Y^{up}(M^{(1)}) - Y^{up}(M^{(2)})|^2 \le E\left[\max_{0\le j\le \mathcal{J}} |M^{(1)}_{T_j} - M^{(2)}_{T_j}|^2\right] \tag{2.10}$$

Hence, we obtain the following immediate corollary:

Corollary 2.3. *(i) It holds that*

$$\lim_{|\pi|\to 0} Y^{up}(M^\pi) = Y^{up}(M)$$

(ii) Under the assumption of Theorem 2.2, (ii), we have

$$|Y^{up}(M^\pi) - Y^{up}(M)|^2 \le C|\pi|$$

The above corollary states that the upper bounds due to M and M^π do not differ much, when the mesh of the partition π is sufficiently small. The main advantage of M^π is that (2.8) remains a martingale, even if the conditional expectations in (2.7) are estimated (of course in a non-anticipative manner). Denoting such martingale (with the conditional expectations in (2.7) estimated) by $\widehat{M}^\pi$, $Y^{up}(\widehat{M}^\pi)$ therefore always defines an upper bound of the Bermudan price Y_0^*. This is in contrast to the representation of M in (2.5). Estimating the conditional expectations in (2.5) can in general destroy the martingale property and so the estimated version may not induce an upper bound.

3 Upper bounds without nested Monte Carlo

We now describe an algorithm based on the construction of the martingales M^π that allows to calculate dual upper bounds without nested Monte Carlo. To this end we suppose that the approximative Snell envelope Y_{T_j} is of the form

$$Y_{T_j} = u(T_j, X^{\bar\pi}_{T_j}).$$

We emphasize that numerical methods to approximate the Snell envelope typically yield approximations of this form. It is then straightforward that the conditional expectations in the definition of Z are, in fact, regressions on $X^{\bar\pi}_{t_i}$. Precisely,

$$Z^\pi_{t_i} = \frac{1}{\Delta^\pi_i} E^{X^{\bar\pi}_{t_i}}\left[(\Delta^\pi W_i)^\top u(T_{j+1}, X^{\bar\pi}_{T_{j+1}})\right],\ T_j \le t_i < T_{j+1}.$$

Next we approximate $Z^\pi_{t_i}$ by simulation based least squares regression on basis functions as was suggested by Longstaff and Schwartz (2001) for lower bounds. To this end we simulate $\widetilde{N}$ independent samples of the Brownian increments $\Delta^\pi W_i^d$, $i = 1, \dots, \mathcal{I}$, $d = 1, \dots, D$,

$$\Delta^\pi_\cdot \widetilde{W}_i^d := (\Delta^\pi_n \widetilde{W}_i^d)_{n=1,\dots,N} .$$

Given row vectors of (possibly time dependent) basis functions $\psi^d(t_i, \cdot) = (\psi_k^d(t_i, \cdot),\ k = 1, \dots, K)$, $d = 1, \dots, D$, and $\widetilde{N}$ independent samples $(t_i,\ {}_n\widetilde{X}_{t_i}^{\bar\pi})$, $n = 1, \dots, \widetilde{N}$ of the Euler scheme $X_{t_i}^{\bar\pi}$ constructed from the above Brownian increments $\Delta_n^\pi \widetilde{W}_i$, we define, for $T_j < t_i \le T_{i+1}$,

$$\widehat{\beta}_{t_i}^d := \underset{\beta \in \mathbb{R}^K}{\arg\min} \left\{ \sum_{n=1}^N \left| \frac{\Delta_n^\pi \widetilde{W}_i^d}{\Delta_i^\pi} u(T_{j+1},\ {}_n\widetilde{X}_{T_{j+1}}^{\bar\pi}) - \psi^d(t_i,\ {}_n\widetilde{X}_{t_i}^{\bar\pi})\beta \right|^2 \right\}.$$

Then, the corresponding approximative regression mapping $\widehat{z}^\pi(t_i, x) = (\widehat{z}_1^\pi(t_i, x), \dots, \widehat{z}_D^\pi(t_i, x))$ for $Z_{t_i}^\pi$ is defined by

$$\widehat{z}_d^\pi(t_i, x) = \psi(t_i, x)\widehat{\beta}_{t_i}^d,$$

After having obtained the functions $\widehat{z}^\pi(t_i, x)$ in (3.1) by the above described regression procedure, we next construct an approximation of M^π by plugging in the Euler approximation of the system (2.1), which we suppose to be independent of the Brownian increments simulated above:

$$\widehat{M}_{T_j}^\pi := \widehat{m}^\pi(T_j, X^{\bar\pi}, \Delta^\pi W) := \sum_{t_i \in \pi; 0 \le t_i < T_j} \widehat{z}^\pi(t_i, X_{t_i}^{\bar\pi})(\Delta^\pi W_i).$$

Clearly $\widehat{M}_{T_j}^\pi$ is a martingale with respect to the enlarged filtration

$$\mathcal{F}_{T_j}^N := \mathcal{F}_{T_j} \vee \mathcal{G}_0^N, \qquad j = 0, ..., \mathcal{J},$$

where $\mathcal{G}_0^N := \sigma(\Delta_n^\pi \widetilde{W}_i;\ i = 1, \dots \mathcal{I}, n = 1, \dots, \widetilde{N})$. Obviously, the underlying stopping problem does not change by this enlargement of filtration and, consequently, $Y^{up}(\widehat{M}^\pi)$ is an upper bound for the discounted Bermudan option price. By sampling a new set of N independent trajectories $(t_i,\ {}_nX_{t_i}^{\bar\pi})$, $n = 1, \dots, N$, of $X^{\bar\pi}$ an unbiased estimator for $Y^{up}(\widehat{M}^\pi)$ is obtained by

$$\widehat{Y}^{up}(\widehat{M}^\pi) = \frac{1}{N} \sum_{n=1}^N \max_{0 \le j \le \mathcal{J}} \left[h(T_j,\ {}_nX_{T_j}^{\bar\pi}) - \widehat{m}^\pi(T_j,\ {}_nX_{T_j}^{\bar\pi}, \Delta_n^\pi W) \right]. \tag{3.1}$$

Remark 3.1. If the functions $u(T_j, \cdot)$ are Lipschitz continuous, it can be deduced from the results on simulation of forward backward SDE by Lemor et al. (2006) and Bender and Denk (2006) that the error

$$|Y^{up}(\widehat{M}^\pi) - Y^{up}(M^\pi)|$$

becomes arbitrarily small, provided the basis is appropriately chosen and the number N of simulated trajectories is sufficiently large. It is, however, well understood that the quality of this approximation heavily depends on the choice of π. The finer one chooses π, the better the quality of the conditional expectation estimator (i.e. better basis, more simulated paths) is required, in order to obtain a stable estimate. In Section 6 we suggest a way to get around this problem in practice.

4 Variance reduced primal-dual algorithm

From Corollary 2.3 and Remark 3.1 we may deduce that $Y^{up}(\widehat{M}^\pi)$ is a close approximation of $Y^{up}(M)$, provided the partition π is sufficiently fine and the numerical regression is appropriately tailored. From (2.5) and the fact that $\widehat{M}^\pi$ is a martingale we see that η_j^π defined by

$$\begin{aligned} \eta_j^\pi &:= E^{T_{j-1}}\left[Y_{T_j}\right] + \varepsilon_j^\pi \\ \varepsilon_j^\pi &= Y_{T_j} - E^{T_{j-1}}\left[Y_{T_j}\right] - (\widehat{M}_{T_j}^\pi - \widehat{M}_{T_{j-1}}^\pi) \end{aligned} \tag{4.1}$$

is an unbiased estimator of $E^{T_{j-1}}Y_{T_j}$. Thus, $\widehat{M}_{T_j}^\pi - \widehat{M}_{T_{j-1}}^\pi$ may be seen as a control variate (see for example Glasserman, 2003; Milstein and Schoenmakers, 2002) for the standard Monte Carlo estimator of $E^{T_{j-1}}Y_{T_j}$. Note that by (2.5) and (4.1),

$$\varepsilon_j^\pi = (M_{T_j} - \widehat{M}_{T_j}^\pi) - (M_{T_{j-1}} - \widehat{M}_{T_{j-1}}^\pi).$$

Clearly, for any partition π we have $E^{T_{j-1}}\varepsilon_j^\pi = 0$ and, loosely speaking, the variance of ε_j^π is closer to zero the more effort one puts into the construction of $\widehat{M}^\pi$.

We will now discuss, how the martingale $\widehat{M}^\pi$ can be used to reduce the variance in the primal-dual algorithm of Andersen and Broadie (2004). To this end let $\overline{M}$ be a martingale such that

$$E^{T_{j-1}}\left[\overline{M}_{T_j}\right] = E^{X_{T_{j-1}}}\left[\overline{M}_{T_j}\right] = \overline{M}_{T_{j-1}},$$

and let

$$\eta_j := E^{T_{j-1}}\left[Y_{T_j}\right] + \varepsilon_j := Y_{T_j} - \left(\overline{M}_{T_j} - \overline{M}_{T_{j-1}}\right), \qquad j = 1, ..., \mathcal{J}. \tag{4.2}$$

(We will specialize from $\overline{M}$ to $\widehat{M}^\pi$ later). On a given trajectory X we consider for each j, $j = 1, ..., \mathcal{J}$, independent copies ${}_l\eta_j = E^{T_{j-1}}\left[Y_{T_j}\right] + {}_l\varepsilon_j$, $l = 1, ..., L$, of (4.2) under the (regular) conditional measure $P^{X_{T_{j-1}}}$, and define the (pathwise) unbiased estimator

$$s_j^{(L)} := \frac{1}{L}\sum_{l=1}^{L} {}_l\eta_j \tag{4.3}$$

for $E^{T_{j-1}}\left[Y_{T_j}\right]$. It thus holds,

$$E^{T_{\mathcal{J}}}\left[s_j^{(L)}\right] = E^{T_{j-1}}\left[Y_{T_j}\right], \quad Var^{T_{\mathcal{J}}}\left[s_j^{(L)}\right] = \frac{1}{L}Var^{T_{j-1}}\left[\varepsilon_j\right], \quad E^{T_{j-1}}\left[\varepsilon_j\right] = 0.$$

Naturally we next consider the (pathwise) estimator

$$\mathcal{U}^{(L)} := \max_{0\le i\le \mathcal{J}}\left(H_{T_i} - \sum_{j=1}^{i}\left(Y_{T_j} - s_j^{(L)}\right)\right),$$

and, based on N independent copies ${}_n\mathcal{U}^{(L)}$, $1 \le n \le N$, the estimator

$$\widehat{Y}_{N,L}^{up}(\overline{M}) := \frac{1}{N}\sum_{n=1}^{N} {}_n\mathcal{U}^{(L)}. \tag{4.4}$$

Note that $\widehat{Y}_{N,L}^{up}(0)$ is the estimator introduced in the primal-dual algorithm of Andersen and Broadie (2004). So $\widehat{Y}_{N,L}^{up}(\overline{M})$ may be considered a variance reduced version of this algorithm with control variate $\overline{M}$.

Theorem 4.1. *It holds that*

$$\begin{aligned} Y^{up}(M) \quad \le \quad & E\left[\widehat{Y}_{N,L}^{up}(\overline{M})\right] \le Y^{up}(M) + \min\left(E\left[\varepsilon_{sum}\right], \sqrt{\frac{\mathcal{J}}{L}\sum_{j=1}^{\mathcal{J}} E\left[\varepsilon_j^2\right]}\right) \\ & = Y^{up}(M) + \min\left(E\left[\varepsilon_{sum}\right], \sqrt{\frac{\mathcal{J}}{L}E\left[\left(M_{T_{\mathcal{J}}} - \overline{M}_{T_{\mathcal{J}}}\right)^2\right]}\right), \end{aligned}$$

where $\varepsilon_{sum} := \sum_{j=1}^{\mathcal{J}}\left(\varepsilon_j\right)_+$.
In particular, the estimator $\widehat{Y}_{N,L}^{up}(\overline{M})$ *is biased up.*

Proof. To prove the first inequality we note that

$$\begin{aligned} E\left[\widehat{Y}_{N,L}^{up}(\overline{M})\right] &= E\left[E^{T_{\mathcal{J}}}\left[\mathcal{U}^{(L)}\right]\right] \ge E\left[\max_{0\le i\le \mathcal{J}} E^{T_{\mathcal{J}}}\left[H_{T_i} - \sum_{j=1}^{i}\left(Y_{T_j} - s_j^{(L)}\right)\right]\right] \\ &= E\left[\max_{0\le i\le \mathcal{J}}\left(H_{T_i} - \sum_{j=1}^{i}\left(Y_{T_j} - E^{T_{j-1}}Y_j\right)\right)\right] = Y^{up}(M). \end{aligned}$$

For the second inequality, let us write using (4.2) and (4.3),

$$
\begin{aligned}
E^{T_{\mathcal{J}}}\left[\mathcal{U}^{(L)}\right] =& E^{T_{\mathcal{J}}}\left[\max_{0\le i\le \mathcal{J}}\left(H_{T_i}-\sum_{j=1}^{i}\left(Y_{T_j}-s_j^{(L)}\right)\right)\right]\\
&\le \max_{0\le i\le \mathcal{J}}\left(H_{T_i}-\sum_{j=1}^{i}\left(Y_{T_j}-E^{T_{j-1}}\left[Y_j\right]\right)\right)\\
&\quad+\sum_{j=1}^{\mathcal{J}}E^{T_{j-1}}\left[\left(s_j^{(L)}-E^{T_{j-1}}\left[Y_j\right]\right)_+\right]\\
&=\max_{0\le i\le \mathcal{J}}\left(H_{T_i}-\sum_{j=1}^{i}\left(Y_{T_j}-E^{T_{j-1}}\left[Y_j\right]\right)\right)+\sum_{j=1}^{\mathcal{J}}E^{T_{j-1}}\left[\left(\frac{1}{L}\sum_{l=1}^{L}{}_{l}\varepsilon_j\right)_+\right].
\end{aligned}
$$

It then follows that

$$
\begin{aligned}
E\left[\widehat{Y}^{up}_{N,L}(\overline{M})\right]&\le Y^{up}(M)+\sum_{j=1}^{\mathcal{J}}E\left[\left(\frac{1}{L}\sum_{l=1}^{L}{}_{l}\varepsilon_j\right)_+\right]\\
&=: Y^{up}(M)+(*).
\end{aligned}
$$

So we have on the one hand, by convexity of the $()_+$ operator,

$$
(*)\le\sum_{j=1}^{\mathcal{J}}E\left[(\varepsilon_j)_+\right]=E\left[\varepsilon_{sum}\right].
$$

On the other hand, by respectively Cauchy-Schwartz and Jensen's inequality, we have

$$
\begin{aligned}
(*)^2&\le \mathcal{J}\sum_{j=1}^{\mathcal{J}}\left\{E\left[\left(\frac{1}{L}\sum_{l=1}^{L}{}_{l}\varepsilon_j\right)_+\right]\right\}^2\le \mathcal{J}\sum_{j=1}^{\mathcal{J}}E\left[\left(\frac{1}{L}\sum_{l=1}^{L}{}_{l}\varepsilon_j\right)^2\right]\\
&=\frac{\mathcal{J}}{L}\sum_{j=1}^{\mathcal{J}}E\left[(\varepsilon_j)^2\right].
\end{aligned}
$$

The last equality follows by a telescoping sum using $E\left[\varepsilon_j^2\right]=$ $E\left[(M_{T_j}-\overline{M}_{T_j})^2-(M_{T_{j-1}}-\overline{M}_{T_{j-1}})^2\right].$ □

Choosing $\overline{M}=\widehat{M}^{\pi}$, we deduce from Theorem 4.1 that $E\left[\widehat{Y}^{up}_{N,L}(\widehat{M}^{\pi})\right]$ converges to the true dual upper bound $Y^{up}(M)$, when the number of paths L tends to infinity or the numerical effort for computing the regression martingale $\widehat{M}^{\pi}$ is increased.

Remark 4.2. The martingale estimator can also be applied to reduce the variance when estimating inner conditional expectations for the consumption based upper bound by Belomestny and Milstein (2006) or in the policy improvement procedure of Kolodko and Schoenmakers (2006). For the policy improvement, this looks promising in particular in combination with the variance reduction for the outer simulation suggested in Bender et al. (2006).

5 Connection with hedge controls and 'deltas'

Let us now suppose that X in (2.1) is a system of tradable securities with $D \leq \mathcal{D}$ (not more Brownian motions than securities) and that the numeraire $\mathcal{N}$ is tradable also. As $\mathcal{N}$ should be positive, we additionaly assume that its dynamics are given by

$$\frac{d\mathcal{N}_t}{\mathcal{N}_t} = \mu_{\mathcal{N}}(t, X_t)dt + \sigma_{\mathcal{N}}(t, X_t)dW_t, \qquad \mathcal{N}_0 = 1,$$

for some smooth and bounded scalar function $\mu_{\mathcal{N}}(\cdot, \cdot)$ and row vector function $\sigma_{\mathcal{N}}(\cdot, \cdot)$. Thus, by assumption, $X/\mathcal{N}$ is a martingale under Q. We moreover assume some extra structural assumptions on the coefficient functions $a, b, \mu_{\mathcal{N}}$, and $\sigma_{\mathcal{N}}$, such that the system $(X, \mathcal{N})$ constitutes a complete market (see e.g. Schoenmakers, 2005, Ch. 1 for details).

In this case of a complete market there is a direct connection between the process Z in (2.6) and the hedge coefficients for replication of the European claim with discounted pay-off Y_{T_j} in the interval $[T_{j-1}, T_j]$. Let us assume that Y_{T_j} is a function of X_{T_j}. Then, by completeness, the claim with pay-off $\mathcal{N}_{T_j} Y_{T_j}$ can be perfectly hedged by a self-financing portfolio $(\vartheta, \theta; X, \mathcal{N})$ with coefficients ϑ, θ being functions $(t, X, \mathcal{N})$. The i-th component of the $\mathcal{D}$-dimensional row vector function $\vartheta(t, X, \mathcal{N})$ denotes the number of shares to hold in X^i and $\theta(t, X, \mathcal{N})$ the amount of units to carry in $\mathcal{N}$, for realizing a perfect dynamic hedge in a self-financing way. We thus have

$$\mathcal{N}_{T_j} Y_{T_j} = \mathcal{N}_{T_{j-1}} E^{T_{j-1}} Y_{T_j} + \int_{T_{j-1}}^{T_j} \vartheta(t, X_t, \mathcal{N}) dX_t + \int_{T_{j-1}}^{T_j} \theta(t, X_t, \mathcal{N}) d\mathcal{N}_t.$$

By a standard lemma connected with Itô's formula (e.g. Lemma 1.1 in Schoenmakers, 2005), it then follows that

$$\begin{aligned} (5.1) \qquad Y_{T_j} &= E^{T_{j-1}} Y_{T_j} + \textstyle\int_{T_{j-1}}^{T_j} \vartheta(t, X, \mathcal{N}) d(\mathcal{N}_t^{-1} X_t) \\ &= E^{T_{j-1}} Y_{T_j} + \textstyle\int_{T_{j-1}}^{T_j} \mathcal{N}_t^{-1} \vartheta(t, X_t, \mathcal{N})(b(t, X_t) - X_t \sigma_{\mathcal{N}}(t, X_t)) dW_t. \end{aligned}$$

We note that the latter equation follows easily from Itô's lemma using the fact that $\mathcal{N}^{-1} X$ is a martingale. From (2.6) and (5.1) we conclude that

$$(5.2) \qquad \mathcal{N}_t^{-1} \vartheta(t, X_t, \mathcal{N})(b(t, X_t) - X_t \sigma_{\mathcal{N}}(t, X_t)) = Z_t =: z(t, X_t).$$

So, after estimating the function $z(\cdot,\cdot)$ by an independent regression procedure we may determine the hedge coefficients $\vartheta(\cdot,\cdot,\cdot)$ from (5.2). For example, if $\mathcal{D} = D$ and the matrix b is invertible, completeness implies that also $b - x\sigma_{\mathcal{N}}$ is invertible, so then the hedge coefficients are unique and follow from

$$\vartheta(t,x,\mathfrak{n}) \quad = \quad \mathfrak{n}z(t,x)(b(t,x) - x\sigma_{\mathcal{N}}(t,x))^{-1}.$$

Given that Y is a reasonably good approximation of the Snell envelope, the linear regression estimator for z hence provides us with an estimator for the Bermudan hedge at all times and states.

If we drop the assumption of $(\mathcal{N}, X)$ being a complete system of tradable assets, we can still find an analogous connection between $z(t,x)$ and the delta. By the Markovianity of X and by the fact that $E^{(t,X_t)}Y_{T_j} := E^t Y_{T_j}$ is a martingale from $T_{j-1} \le t \le T_j$, we obtain via application of Itô's formula to $E^{(t,X_t)}Y_{T_j}$,

$$z(t,x) = \nabla_x E^{t,x} Y_{T_j}\, b(t,x),. \tag{5.3}$$

As Y is an approximation of the Snell envelope, we get an estimate for the Bermudan delta at all times and states via the linear regression estimate for z. This delta estimate based on (2.7) can be viewed as a linear regression version of the so called covariation delta method. This method is discussed and compared to other methods in the framework of a plain Monte Carlo implementation in Detemple et al. (2005).

6 Numerical examples

In our implementation study we first construct a family of stopping rules $\tau_j : \Omega \to \{T_j, \ldots, T_{\mathcal{J}}\}$ by the Longstaff-Schwartz algorithm. This basically boils down to choosing a basis $(\phi_k(t,x),\, k = 1, \ldots, K)$ and estimating vectors of regression coefficients $(\alpha_l \in \mathbb{R}^K,\, l = 0, \ldots, \mathcal{J})$. Once the coefficients $\{\alpha_l\}$ are estimated, we can define

$$\tau_j := \min\{j \le l \le \mathcal{J} :\, \alpha_l^\top \phi(T_l, X_{T_l}) \le H_{T_l}\}$$

and

$$Y_{T_j} := E^{T_j} H_{\tau_j}, \quad j = 1, \ldots, \mathcal{J}.$$

We stress that stopping rules $\{\tau_j\}$ are estimated only once and remain fixed thereafter. Having $\{Y_{T_j}\}$ at hand we proceed generally as described in Section 3. Recall that we do not need to calculate the conditional expectations in the definition of Y as explained in Remark 2.1, but may consider (2.9) instead. To ensure a stable estimate based on regression in practice, we suggest the following generic procedure: Firstly one can easily compute regression estimates $\widehat{C}_i$ of $C_i := E^{t_i} H_{\tau_{j+1}}$, for (i,j) such that $T_j < i \le T_{j+1}$, while running the Longstaff-Schwartz algorithm. Then, replacing $H_{\tau_{j+1}}$ by $H_{\tau_{j+1}} - \widehat{C}_i$ in (2.9) is expected to significantly reduce the conditional variance in (2.9), but does not change the conditional expectation. Hence, we actually solve the least squares problems, for $T_j < t_i \leq T_{j+1}$,

$$\widehat{\beta}^d_{t_i} := \arg\min \left\{ \sum_{n=1}^{N} \left| \frac{\Delta^\pi_n \widetilde{W}^d_i}{\Delta^\pi_i} \left({}_nH_{\tau_{j+1}} - {}_n\widehat{C}_i \right) - \psi^d(t_i,\ {}_n\widetilde{X}^{\bar\pi}_{t_i})\beta \right|^2 ;\ \beta \in \mathbb{R}^K \right\}$$

along the simulated paths ${}_n\widetilde{X}^{\bar\pi}$, $n = 1, \dots, \tilde{N}$. Secondly, the choice of partition π is important to guarantee a stable estimate. In this respect, we recommend to calculate the regression coefficients on a comparably rough partition only, and then interpolate them locally constant to a very fine partition. To be precise, we run the regression only on the set of exercise dates (hence $\pi = \mathcal{E}$) to obtain $\hat\beta_{T_i}$, $T_i \in \mathcal{E}$, and set $\hat\beta_t = \hat\beta_{T_i}$ for $t \in [T_i, T_{i+1})$. Then we define $z^{\mathcal{E}}_d(t,x) := \psi^d(t,x)\widehat{\beta}^d_t$ for all $t \in [0,T]$, and we can consider an estimated martingale by summing over the finer partition $\bar\pi$ (on which the Euler scheme is performed). That is, with a slight abuse of notation, we define

$$\widehat{M}^{\mathcal{E}}_{T_j} := \widehat{M}^{\mathcal{E},\bar\pi}_{T_j} := \sum_{t\in\bar\pi; 0\leq t<T_j} \hat{z}^{\mathcal{E}}(t, X^{\bar\pi}_t)(\Delta^{\bar\pi} W_t), \quad T_j \in \mathcal{E}.$$

Moreover, this specific implementation guarantees that the complexity of the algorithm with interpolated $\widehat{\beta}_t$ corresponds to the one of the usual Longstaff-Schwartz method, because regression is only performed on the exercise grid.

Finally, the choice of the basis functions is essential to achieve tight upper bounds. As one expects the integrand of the martingale representation to be less regular than Y itself, more care should be taken when choosing the basis for the upper bound procedure. We next suggest a basis, when some soft problem information is available. To be precise, we suppose that an analytical (approximation) formula for the discounted price of the corresponding European options is available. This is the case for many practically relevant derivatives such as the max-call-option and swaptions. Then, by Itô's formula, for $0 \leq t \leq T_j$,

$$E^t[H_{T_j}] =: u^E_{T_j}(t, X_t) = E[H_{T_j}] + \int_0^t \nabla_x u^E_{T_j}(s, X_s) b(s, X_s) dW_s. \tag{6.1}$$

If the analytical formula for u^E is not exact this identity may still hold approximately. Equation (6.1) suggests to include some of the integrands on the right hand side into the basis. As a minimal basis for calculating the regression coefficient $\hat\beta^d_{T_j}$ of the dth component over the period $[T_j, T_{j+1}]$ we propose the three functions, 1, the dth component of $\nabla_x u^E_{T_{j+1}}(s,x)b(s,x)$, and the dth component of $\nabla_x u^E_{T_\jmath}(s,x)b(s,x)$, which contain information about the still alive European options with the shortest and with the longest maturity. This simple basis, consisting of three functions only, works very well in our numerical examples. Of course, this basis can be easily extended by adding information about more European options and/or interacting terms.

Bermudan max-calls on D assets

This is a benchmark example studied in Andersen and Broadie (2004), Glasserman (2003), Haugh and Kogan (2004) Rogers (2001), and others. Specifically, a model with D identically distributed assets is considered where each asset yields dividends with rate δ. The risk-neutral dynamic of this asset system is given by

$$dX_t^d = (r-\delta)X_t^d dt + \sigma X_t^d dW_t^d, \quad d = 1, ..., D,$$

where W_t^d, $k = 1, ..., D$, are independent one-dimensional Brownian motions and r, δ, σ are constants. At one time $t \in \{T_0, ..., T_{\mathcal{J}}\}$ the holder of the option may exercise to receive the payoff

$$h(X_t) = (\max(X_t^1, ..., X_t^D) - \kappa)^+.$$

We consider an example when $T_j = jT/\mathcal{J}$, $j = 0, ..., \mathcal{J}$, with $T = 3$ and $\mathcal{J} = 9$. The stopping rules $\{\tau_j\}$ are estimated by regression using 5×10^4 Monte Carlo paths and (following Andersen and Broadie (2004)) a set of basis functions consisting of all monomials in the underlying of degree less than or equal to three, the European max-call option with maturity T, its square, and its cube. The Euler scheme is performed on an equidistant partition $\bar{\pi}$ with $|\bar{\pi}| = 0.01$. The coefficients $\widehat{\beta}_{t_i}$, $t_i \in \pi$ are estimated using the same number of paths, and the three-function basis motivated by (6.1) as described (generically) above. Here, we utilize the analytical expression

$$\sum_{l=1}^{D} X_0^l \frac{e^{-\delta T}}{\sqrt{2\pi}} \int_{(-\infty, d_+^l]} \exp[-\frac{1}{2}z^2] \prod_{\substack{l'=1 \\ l' \neq l}}^{D} \mathcal{N}\left(\frac{\ln \frac{X_0^l}{X_0^{l'}}}{\sigma\sqrt{T}} - z + \sigma\sqrt{T} \right) dz$$
$$- \kappa e^{-rT} + \kappa e^{-rT} \prod_{l=1}^{D} \left(1 - \mathcal{N}\left(d_-^l\right)\right),$$

where

$$d_-^l := \frac{\ln \frac{X_0^l}{\kappa} + (r - \delta - \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}, \quad d_+^l = d_-^l + \sigma\sqrt{T},$$

and $\mathcal{N}$ denotes the cumulative standard Gaussian distribution, for the price of the European max-call option.

The results for $D = 2$ and $D = 5$ are presented in Table 1 in dependence of x_0 with $X_0 = (X_0^1, \dots, X_0^D)^T$, $X_0^1 = ... = X_0^D = x_0$. The price intervals in the last column are quoted from Andersen and Broadie (2004), where the upper bounds are computed by the primal-dual algorithm, hence by nested Monte Carlo. These values are almost exact, however, computed with considerable computational effort (10 000 inner simulations used for the upper bounds). As we see, the results for the fast non-nested estimator are quite close to these almost exact values.

It is instructive to look at the quality of dual upper bounds in dependence of the choice of a regression basis. Figure 1 shows upper bounds (with 95% confidence intervals) corresponding to 5 different choices of regression bases. The description of the sets is given in the Table 2, where $\mathrm{EP}(t, X; T)$ is the price of a European max-call option with maturity T at time t, and $\mathrm{Pol}_n(y)$ is the set of monomials of degree less than or equal to n in the components of a vector y. The set SI is used for the computation of Table 1. The set SV corresponds to estimating constant volatility parameters between the exercise times. This simplistic choice of a martingale results in a rather poor upper bound. Standard bases, built upon polynomials (SIII, SIV), significantly improve the quality of the upper bound estimator. They are reasonably tight, but nonetheless they are still about 5% off relatively. Finally the bases consisting of European deltas (SI, SII) yield almost the exact values. Hence, the main conclusion, drawn from this figure, is that the problem specific information captured in the European deltas, is extremely useful when estimating the shape of the martingale integrand of Y by regression.

Finally, Figure 2 compares the Andersen-Broadie algorithm with its variance reduced version suggested in Section 4. Again we consider the out-of-the-money Bermudan max-call on two assets. The martingale for the variance reduction is calculated using basis SI. The figure illustrates a reasonable performance of the proposed control variate.

Bermudan swaptions in the LIBOR market model

We here recall the example of Bermudan swaptions in the Libor market model considered in Kolodko and Schoenmakers (2006). For a fixed tenor structure $0 = \mathcal{T}_0 < \mathcal{T}_1 < \ldots < \mathcal{T}_{\mathcal{D}+1}$, the forward Libor $L_i(t)$, $1 \leq i \leq \mathcal{D}$, (that is the annualized effective rate on a loan over period $[\mathcal{T}_i, \mathcal{T}_{i+1}]$ contracted at date t which is to be payed at T_{i+1}) is governed by the SDE (e.g., Jamshidian (1997))

$$dL_i = \sum_{j=\kappa(t)}^{i} \frac{\delta_j L_i L_j \, \gamma_i \cdot \gamma_j}{1 + \delta_j L_j} \, dt + L_i \, \gamma_i \cdot dW^*, \quad 0 \leq t \leq T_i, \tag{6.2}$$

where $\delta_i = \mathcal{T}_{i+1} - \mathcal{T}_i$ are day count fractions, $t \to \gamma_i(t) = (\gamma_{i,1}(t), \ldots, \gamma_{i,D}(t))$ are deterministic volatilities defined in $[0, \mathcal{T}_i]$ (factor loadings), and $\kappa(t) := \min\{m : T_m \geq t\}$ is the next reset date after time t. In (6.2), $(W^*(t) \mid 0 \leq$

$t \le \mathcal{T}_{n-1}$) is a standard Wiener process in $\mathbb{R}^D$ under the spot Libor measure P^* where $1 \le D \le \mathcal{D}$. This measure is induced by the numeraire

$$B_*(0) := 1, \quad B_*(t) := \frac{B_{\kappa(t)}(t)}{B_1(0)} \prod_{i=1}^{\kappa(t)-1} \left(1 + \delta_i L_i(\mathcal{T}_i)\right), \ t > 0,$$

with $\prod_{i=1}^{0} := 1$, and $B_i(t)$ being the value of a zero coupon bond with face value \$1 at time $t \le \mathcal{T}_i$.

A Bermudan (payer) swaption with strike θ on a \$1 principal issued at $t = 0$, gives the right to exercise (once) the cash-flow of a swap,

$$S_{\mathcal{T}_i} := \sum_{j=i}^{\mathcal{D}} B_{j+1}(\mathcal{T}_i)\delta_j \left(L_j(\mathcal{T}_i) - \theta\right) \tag{6.3}$$

at a date $\mathcal{T}_i \in \mathcal{E} := \{T_1, ..., T_{\mathcal{J}}\} :\subset \{\mathcal{T}_1, ..., \mathcal{T}_{\mathcal{D}+1}\}$, where $\mathcal{E}$ is a given set of exercise dates. (Of course when $S_{\mathcal{T}_i} < 0$ there will be no exercise.)

As in Kolodko and Schoenmakers (2006) we take a Libor volatility structure,

$$\gamma_i(t) = cg(T_i - t)e_i, \quad \text{where} \quad g(s) = g_\infty + (1 - g_\infty + as)e^{-bs}$$

and e_i are D-dimensional unit vectors, decomposing some input correlation matrix of rank D. From a full-rank basis correlation structure

$$\rho_{ij} = \exp(-\varphi|i-j|), \qquad i, j = 1, \ldots, \mathcal{D}, \qquad \varphi > 0, \tag{6.4}$$

we deduce for different D, $D \le \mathcal{D}$, a rank-D correlation matrices ρ^D with decomposition $\rho^D_{ij} = e_i \cdot e_j$, $1 \le i, j \le \mathcal{D}$, by principal component analysis in order to generate a Libor models with D driving Brownian motions (factors). For the numerical experiments we take a flat 10% initial curve over a 40 period quarterly tenor structure, the parameters

$$\mathcal{D} = 40, \ \delta_i = 0.25, \ c = 0.2, \ a = 1.5, \ b = 3.5, \ g_\infty = 0.5, \ \varphi = 0.0413, \tag{6.5}$$

and consider Bermudan swaptions with yearly exercise opportunities, i.e. $T_i = \mathcal{T}_{4i}$, $i = 1, \ldots, 10$. For the numerical integration of the SDE (6.2) the log-Euler scheme is used with $\Delta t = \delta/5$.

A lower approximation Y_0 is constructed by the Longstaff-Schwartz algorithm using a basis which consists (at exercise date T_i) of the European swaption

$$\mathcal{S}_{T_i}(L_i(T_i), \ldots, L_{\mathcal{D}}(T_i)) := B_*(T_i)E^{T_i}\left(\frac{S_{T_{i+1}}(L_{i+1}(T_{i+1}), \ldots, L_{\mathcal{D}}(T_{i+1}))}{B_*(T_{i+1})}\right),$$

which can be exercised at T_{i+1}, and all powers up to second order of the payoff S_{T_i}. Although closed form expressions for European swaptions do not exist in a Libor market model, there do exist very accurate (typically better than 0.3% relative error) formulas which we use for the computation of $\mathcal{S}_i$ (e.g. see Schoenmakers (2005)). Following our previous recommendations we use the derivatives of these European swaptions with respect to Libors in the set of basis functions for estimating the martingale in the representation of the dual upper bound.

Based on 5×10^4 simulated Libor trajectories, we estimate the coefficients β_t on the exercise partition $\mathcal{E}$ using regression and then interpolate them constantly on the tenor grid $\{\mathcal{T}_1, ..., \mathcal{T}_{\mathcal{D}+1}\}$. A new simulation of 10^5 trajectories is used to get values for $Y_0^{up}(\widehat{M}^\pi)$, which are presented in Table 3 for different strikes θ and numbers D of Brownian motions. The price intervals in the last column are quoted from Kolodko and Schoenmakers (2006), where the lower bounds are obtained via a one step policy iteration procedure and the (dual) upper bounds correspond to Andersen (2000) lower bounds (strategy I) using 500 inner simulations. As we see all estimates $Y_0^{up}(\widehat{M}^\pi)$ are about 1% (relative to price, one exception of 2%) close to the 'true' values.

7 Concluding remarks

Nowadays the primal-dual algorithm is likely to be the most popular algorithm to compute Bermudan upper bounds, although its requirement for nested simulations does make it computationally extensive. In this paper we presented, in a Brownian motion setting, an alternative to this algorithm which is fast, as it requires linear simulation cost only. In our simulation study for two high-dimensional case studies we found that the new algorithm yields surprisingly tight upper bounds particularly if closed form approximations of the corresponding European option prices are at hand. If nonetheless for a certain product a higher accuracy is required (e.g. if such information about Europeans is not available), we propose a variance reduced version of the primal-dual algorithm which allows to compute upper bounds with the same accuracy as the latter one at lower costs.

A Proof of Theorem 2.2

Fix some $T_j < T$ and consider $t_i,\ T_j \le t_i < T_{j+1}$. Then, by (2.6) and Itô's isometry, we get for the dth component of $Z^\pi_{t_i}$

$$
\begin{aligned}
\text{(A.1)} \qquad Z^{\pi,d}_{t_i} &= \frac{1}{\Delta^\pi_i} E^{t_i}\left[(\Delta^\pi W^d_i)\left(Y_{T_{j+1}} - E^{T_j}[Y_{T_{j+1}}]\right)\right] \\
&= \frac{1}{\Delta^\pi_i} E^{t_i}\left[\left(\int_{t_i}^{t_{i+1}} dW^d_s\right)\left(\int_{T_j}^{T_{j+1}} Z_s dW_s\right)\right] \\
&= \frac{1}{\Delta^\pi_i} E^{t_i}\left[\int_{t_i}^{t_{i+1}} Z^d_s ds\right].
\end{aligned}
$$

It follows from (A.1) that without any further assumptions,

$$
\text{(A.2)} \qquad \lim_{|\pi|\to 0} E\left[\sum_{t_i\in\pi, T_j\le t_i<T_j} \int_{t_i}^{t_{i+1}} |Z_s - Z^\pi_{t_i}|^2 ds\right] = 0
$$

as noted e.g. in Lemor et al. (2006). Since, by Doob's inequality and Itô's isometry,

$$
\begin{aligned}
& E\left[\max_{0\le j\le \mathcal{J}} |M^\pi_{T_j} - M_{T_j}|^2\right] \le 4E\left[|M^\pi_T - M_T|^2\right] \\
\text{(A.3)} \qquad &= 4E\left[\sum_{j=0}^{\mathcal{J}-1} \sum_{t_i\in\pi, T_j\le t_i<T_j} \int_{t_i}^{t_{i+1}} |Z_s - Z^\pi_{t_i}|^2 ds\right],
\end{aligned}
$$

assertion (i) immediately follows.

We now prove (ii) and first consider the case $Y_{T_j} = u(T_j, X_{T_j})$. Note that, on $[T_j, T_{j+1}]$, Z is the control part of the simple forward-backward SDE (FBSDE)

$$
\begin{aligned}
X_t &= X_{T_j} + \int_{T_j}^t b(s, X_s) ds + \int_{T_j}^t b(s, X_s) dW_s \\
\bar{Y}_t &= u(T_{j+1}, X_{T_{j+1}}) - \int_t^{T_{j+1}} Z_t dW_t.
\end{aligned}
$$

Due to the Lipschitz continuity of $u(T_{j+1}, \cdot)$ results on L^2-regularity obtained for the control part of FBSDEs in more general situations by Zhang (2004) and Bender and Zhang (2006) can be applied. In combination with (A.1) these results imply that (A.2) can be strengthened to

$$
E\left[\sum_{t_i\in\pi, T_j\le t_i<T_j} \int_{t_i}^{t_{i+1}} |Z_s - Z^\pi_{t_i}|^2 ds\right] \le C_j |\pi|
$$

for some constant C_j. Hence, (ii) follows in the case $Y_{T_j} = u(T_j, X_{T_j})$ with constant $C = \sum_j C_j$ thanks to (A.3).

To prove (ii) in the case $Y_{T_j} = u(T_j, X^{\bar{\pi}}_{T_j})$, denote the martingale part in the Doob decomposition of $\bar{Y}_{T_j} = u(T_j, X_{T_j})$ by $\bar{M}$. Moreover, define

$$\begin{aligned}\bar{Z}^{\pi}_{t_i} &:= \frac{1}{\Delta^{\pi}_i} E^{t_i}\left[(\Delta^{\pi} W_i)^{\top} u(T_j, X_{T_j})\right],\ T_j \le t_i < T_{j+1} \\ \bar{M}^{\pi}_{T_j} &:= \sum_{t_i \in \pi; 0 \le t_i < T_j} \bar{Z}^{\pi}_{t_i}(\Delta^{\pi} W_i).\end{aligned}$$

Then,

$$\begin{aligned}E\left[\max_{0 \le j \le \mathcal{J}} |M^{\pi}_{T_j} - M_{T_j}|^2\right] &\le 12E\left[|M^{\pi}_T - \bar{M}^{\pi}_T|^2 + |\bar{M}_T - \bar{M}^{\pi}_T|^2 + |M_T - \bar{M}_T|^2\right] \\ &= 12[(I) + (II) + (III)]\end{aligned}$$

From the previous case, the second term is of order $|\pi|$. From the Lipschitz continuity of $u(T_j, \cdot)$ we get

$$\begin{aligned}(III) &= E\left[|\sum_{j=1}^{\mathcal{J}} u(T_j, X^{\bar{\pi}}_{T_j}) - u(T_j, X_{T_j}) - E^{T_{j-1}}[u(T_j, X^{\bar{\pi}}_{T_j}) - u(T_j, X_{T_j})]|^2\right] \\ &\le K \sum_{j=1}^{\mathcal{J}} E\left[|X^{\bar{\pi}}_{T_j} - X_{T_j}|^2\right] \le K|\bar{\pi}| \le K|\pi|\end{aligned}$$

where the generic constant K may differ from application to application. To estimate (I), note that, for $T_j \le t_i < T_{j+1}$,

$$\begin{aligned}&E^{t_i}\left[(\Delta^{\pi} W_i)^{\top}(u(T_{j+1}, X^{\bar{\pi}}_{T_{j+1}}) - u(T_{j+1}, X_{T_{j+1}}))\right]^2 (\Delta^{\pi}_i)^{-1} \\ =\ &E^{t_i}\Big[(\Delta^{\pi} W_i)^{\top}(E^{t_{i+1}}[u(T_{j+1}, X^{\bar{\pi}}_{T_{j+1}}) - u(T_{j+1}, X_{T_{j+1}})] \\ &-E^{t_i}[u(T_{j+1}, X^{\bar{\pi}}_{T_{j+1}}) - u(T_{j+1}, X_{T_{j+1}})])\Big]^2 (\Delta^{\pi}_i)^{-1} \\ \le\ &E^{t_i}\Big[E^{t_{i+1}}[u(T_{j+1}, X^{\bar{\pi}}_{T_{j+1}}) - u(T_{j+1}, X_{T_{j+1}})]^2 \\ &-E^{t_i}[u(T_{j+1}, X^{\bar{\pi}}_{T_{j+1}}) - u(T_{j+1}, X_{T_{j+1}})]^2\Big].\end{aligned}$$

Thus,

$$\begin{aligned}(I) &= \sum_{j=0}^{\mathcal{J}-1} \sum_{T_j \le t_i < T_{j+1}} E\left[(\Delta^{\pi} W_i)^{\top}(u(T_{j+1}, X^{\bar{\pi}}_{T_{j+1}}) - u(T_{j+1}, X_{T_{j+1}}))\right]^2 (\Delta^{\pi}_i)^{-1} \\ &\le \sum_{j=0}^{\mathcal{J}-1} E\left[|u(T_{j+1}, X^{\bar{\pi}}_{T_{j+1}}) - u(T_{j+1}, X_{T_{j+1}})|^2\right] \\ &\le KE[|X^{\bar{\pi}}_{T_{j+1}} - X_{T_{j+1}}|^2] \le K|\bar{\pi}| \le K|\pi|.\end{aligned}$$

References

L. Andersen (2000). A simple approach to the pricing of Bermudan swaptions in the multi-factor Libor Market Model. Journal of Computational Finance, **3**, 5-32.

L. Andersen, M. Broadie (2004). A primal-dual simulation algorithm for pricing multidimensional American options. Management Sciences, **50**, No. 9, 1222-1234.

D. Belomestny, G.N. Milstein (2006). Monte Carlo evaluation of American options using consumption processes. International Journal of Theoretical and Applied Finance, **9**, No. 4, 1-27.

C. Bender, R. Denk (2006). A Forward Scheme for Backward SDEs. Stoch. Process. Appl., forthcoming.

C. Bender, A. Kolodko, J. Schoenmakers (2006). Iterating cancellable snowballs and related exotics. RISK, September 2006 pp. 126–130.

C. Bender, J. Zhang (2006). Time discretization and Markovian iteration for coupled FBSDEs. Ann. Appl. Probab., forthcoming.

J. Carriere (1996). Valuation of early-exercise price of options using simulations and nonparametric regression. Insurance: Mathematics and Economics, **19**, 19-30.

N. Chen and P. Glasserman (2007). Additive and Multiplicative Duals for American Option Pricing. Finance and Stochastics, **11**, 153Ű-179.

E. Clément, D. Lamberton, P. Protter (2002). An analysis of a least squares regression algorithm for American option pricing. Finance and Stochastics, **6**, 449-471.

J. Detemple, R. Garcia, M. Rindisbacher (2005). Asymptotic Properties of Monte Carlo Estimators of Derivatives. Manag. Sci., **51**, 1657–1675.

P. Glasserman (2003). Monte Carlo Methods in Financial Engineering. Springer.

P. Glasserman and B. Yu (2005). Pricing American Options by Simulation: Regression Now or Regression Later?, Monte Carlo and Quasi-Monte Carlo Methods, (H. Niederreiter, ed.), Springer, Berlin.

M. Haugh, L. Kogan (2004). Pricing American options: a duality approach. Operations Research, **52**, No. 2, 258–270.

F. Jamshidian, (1997) LIBOR and swap market models and measures. *Finance and Stochastics*, **1**, 293–330.

F. Jamshidian (2007). The duality of optimal exercise and domineering claims: A Doob-Meyer decomposition approach to the Snell envelope. Stochastics, **79**, No. 1–2, 27–60.

A. Kolodko, J. Schoenmakers (2004). Upper bounds for Bermudan style derivatives. Monte Carlo Methods and Appl., **10**, No. 3–4, 331-343.

A. Kolodko, J. Schoenmakers (2006). Iterative construction of the optimal Bermudan stopping time. Finance and Stochastics, **10**, 27–49.

J. Lemor, E. Gobet, X. Warin (2006). Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations. Bernoulli, **12**, 889–916.

F.A. Longstaff, E.S. Schwartz (2001). Valuing American options by simulation: a simple least-squares approach. Review of Financial Studies, **14**, 113–147.

G.N. Milstein, J. Schoenmakers (2002). Monte Carlo construction of hedging strategies against multi-asset European claims. Stoch. Stoch. Rep., **73**, 125–157.

L.C.G. Rogers (2001). Monte Carlo valuation of American options. Mathematical Finance, **12**, 271–286.

J. Schoenmakers (2005). Robust Libor Modelling and Pricing of Derivative Products. Chapman & Hall/CRC.

J. Tsitsiklis, B. Van Roy (1999). Regression methods for pricing complex American style options. IEEE Trans. Neural. Net., **12**, 694–703.

J. Zhang (2004). A numerical scheme for BSDEs. Ann. Appl. Probab., **14**, 459–488.

Table 1: Bounds (with 95% confidence intervals) for Bermudan max call with parameters $\kappa = 100$, $r = 0.05$, $\sigma = 0.2$, $\delta = 0.1$ and different D and x_0

D	x_0	Lower Bound Y_0	Upper Bound $Y_0^{up}(\widehat{M}^\pi)$	A&B Price Interval
	90	8.0242±0.075	8.0891±0.068	[8.053, 8.082]
2	100	13.859±0.094	13.958±0.085	[13.892, 13.934]
	110	21.330±0.109	21.459±0.097	[21.316, 21.359]
	90	16.575±0.072	16.681±0.070	[16.602, 16.655]
5	100	26.104±0.085	26.273±0.079	[26.109, 26.292]
	110	36.701±0.098	36.902±0.091	[36.704, 36.832]

Table 2: Sets of basis functions in Fig. 1

	Basis functions at (t, x) with $T_{j-1} \leq t < T_j$
SI	$\left\{1,\ X^k \frac{\partial\, \mathrm{EP}(t,X;T_j)}{\partial x^k},\ X^k \frac{\partial\, \mathrm{EP}(t,X;T_J)}{\partial x^k}\right\}$, $k = 1, \ldots, D$
SII	$\left\{1,\ X^k \frac{\partial\, \mathrm{EP}(t,X;T_j)}{\partial x^k}\right\}$, $k = 1, \ldots, D$
SIII	$\mathrm{Pol}_3(X, \mathrm{EP}(t, x;\, T_j))$
SIV	$\mathrm{Pol}_3(X)$
SV	1

Table 3: Prices of bermudan swaptions $\times 10^4$

θ	D	Lower Bound Y_0	Upper Bound $Y_0^{up}(\widehat{M}^\pi)$	K&S Price Interval
	1	1108.8±1.41	1109.6±0.86	[1108.9±2.4, 1109.4±0.7]
0.08	2	1101.6±1.53	1104.7±0.91	[1100.5±2.4, 1103.7±0.7]
	10	1096.4±1.61	1103.2±0.98	[1096.9±2.1, 1098.1±0.6]
	1	121.0±0.71	122.4±0.87	[121.0±0.6, 121.3±0.4]
0.12	2	113.3±0.75	115.2±0.89	[113.8±0.5, 114.9±0.4]
	10	100.1±0.83	103.4±0.96	[100.7±0.4, 101.5±0.3]

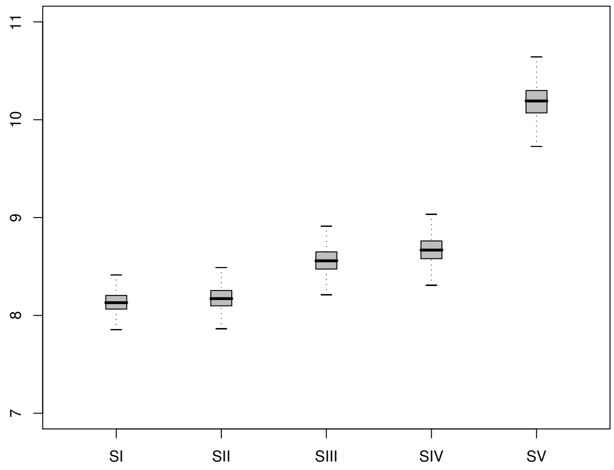

Figure 1: Upper bounds $Y_0^{up}(\widehat{M}^\pi)$, $x_0 = 90$, $D = 2$, with 95% confidence intervals for five different regression bases used for estimating $\widehat{M}^\pi$.

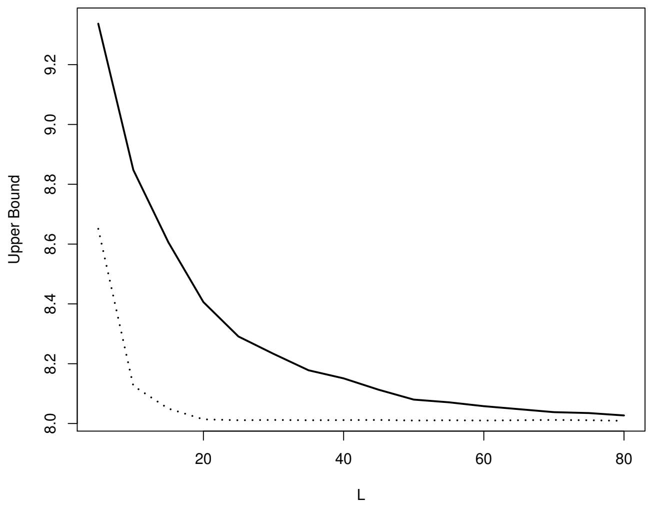

Figure 2: Upper bounds $Y_0^{up}(0)$ (solid line) and $Y_0^{up}(\hat{M}^\pi)$ (dotted line) in dependence on the number of inner Monte Carlo paths L, the number of outer paths N being equal to 10000.