Weierstraß-Institut
für Angewandte Analysis und Stochastik

im Forschungsverbund Berlin e.V.

Preprint ISSN 0946 – 8633

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

Denis Belomestny, [1] Christian Bender, [2] and John Schoenmakers [3]

submitted: 27th November 2006

[1] Weierstrass Institute
for Applied Analysis and Stochastics,
Mohrenstr. 39, 10117
Berlin, Germany
E-Mail: belomest@wias-berlin.de

[2] Institute for Mathematical Stochastics,
TU Braunschweig,
Pockelsstr. 14, 38106
Braunschweig, Germany
E-Mail: c.bender@tu-bs.de

[3] Weierstrass Institute
for Applied Analysis and Stochastics,
Mohrenstr. 39, 10117
Berlin, Germany
E-Mail: schoenma@wias-berlin.de

No. 1186
Berlin 2006

2000 *Mathematics Subject Classification.* 60H30, 65C05, 91B28.

Key words and phrases. Bermudan options, Monte Carlo method, primal-dual method, martingale representation theorem, regression.

Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Mohrenstraße 39
10117 Berlin
Germany

Fax: + 49 30 2044975
E-Mail: preprint@wias-berlin.de
World Wide Web: http://www.wias-berlin.de/

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 may be employed for computing 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) and nested consumption based methods (Belomestny and Milstein, 2006). Numerical experiments indicate the efficiency of the non-nested Monte Carlo algorithm and the variance reduced nested one.

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 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: by a relatively low computational costs an approximative exercise policy can be constructed, a straightforward Monte Carlo simulation giving thereafter 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 (2006). A comparative study of multiplicative and additive duals is provided in Chen and Glasserman (2005). 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 preserving martingale property of (1.1) is proposed. This algorithm, however, requires strong conditions on the basis functions, that 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 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 on basis functions without any restrictions on the basis (Section 3). The corresponding estimator $\widehat{M}$ for M is a martingale and consequently induces an upper bound. Moreover, if Y is constructed by linear regression, the same regression matrices can be used to estimate $\widehat{M}$. Hence, the construction of $\widehat{M}$ does require almost no computational costs in this situation (and of course, no time consuming nested simulations). Some results on the convergence of $\widehat{M}$ to M are presented in Theorem 2.1 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) and for another approach towards constructing upper bounds which was introduced in Belomestny and Milstein (2006). Moreover, the martingale $\widehat{M}$ can be used to derive estimates for the delta of the Bermudan option in a complete market, as is stressed in Section 5.

Finally we present numerical examples in Section 6. In our simulation study for a maximum call on several assets we find that the fast non-nested estimator introduced in this paper yields surprisingly good upper bounds. We also demonstrate a significant variance reduction effect of $\widehat{M}$, if used as 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

$$\begin{aligned} dX_t &= a(t, X_t)dt + b(t, X_t)dW_t && (2.1)\\ X_0 &= x. && (2.2)\end{aligned}$$

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

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

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 $Z_t = (Z_t^1, \ldots, Z_t^D)$ 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

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

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

Theorem 2.1. *(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

$$|Y^{up}(M^{(1)}) - Y^{up}(M^{(2)})|^2 \le E\left[\max_{0 \le j \le \mathcal{J}} |M^{(1)}_{T_j} - M^{(2)}_{T_j}|^2\right] \tag{2.9}$$

Hence, we obtain the following immediate corollary:

Corollary 2.2. *(i) It holds that*

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

(ii) Under the assumption of Theorem 2.1, (ii), we have

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

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

$$Y_{T_j} = u(T_j, X^\pi_{T_j}).$$

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 $X^\pi_{t_i}$. Precisely,

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

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 $\Delta^\pi W_i$, $i = 1, \ldots, \mathcal{I}$,

$$\Delta^\pi_\cdot \widetilde{W}_i := (\Delta^\pi_n \widetilde{W}_i)_{n=1,\ldots,\widetilde{N}} := ({}_n\widetilde{W}^d_{t_{i+1}} - {}_n\widetilde{W}^d_{t_i})_{n=1,\ldots,\widetilde{N}, d=1,\ldots,D}$$

(hence for a fixed time point t_{i+1}, interpreted as $\widetilde{N} \times D$ matrix). Given a row vector of (possibly time dependent) basis functions $\psi(t_i, \cdot) = (\psi_k(t_i, \cdot),\ k = 1, \ldots, K)$ and $\widetilde{N}$ independent samples $(t_i, {}_n\widetilde{X}^\pi_{t_i})$, $n = 1, \ldots, \widetilde{N}$ of the Euler scheme $X^\pi_{t_i}$ constructed from the above Brownian increments $\Delta^\pi_n \widetilde{W}_i$, $n = 1, \ldots, \widetilde{N}$, the corresponding regression matrix at time t_i is defined as the pseudo-inverse $A^\oplus_{t_i}$ of the matrix

$$A_{t_i} = \left(\psi_k(t_i, {}_n\widetilde{X}^\pi_{t_i})\right)_{n=1,\ldots,\widetilde{N}, k=1,\ldots,K}$$

(recall that the pseudo inverse $A^\oplus_{t_i}$ coincides with

$$(A^\top_{t_i} A_{t_i})^{-1} A^\top_{t_i},$$

if the matrix A_{t_i} has full rank). Then, the corresponding approximative regression mapping for $Z^{\pi}_{t_i}$ is defined by

$$\begin{aligned}\widehat{z}^{\pi}(t_i, x) &= \psi(t_i, x)\, A^{\oplus}_{t_i} \left(\frac{\Delta^{\pi}_{\cdot}\widetilde{W}_i}{\Delta^{\pi}_{i}} . \widetilde{Y}_{T_{j+1}} \right), \; T_j \le t_i < T_{j+1} \\ &=: \psi(t_i, x)\widehat{\beta}_{t_i}, \end{aligned} \tag{3.1}$$

using the suggestive notations

$$\left(\frac{\Delta^{\pi}_{\cdot}\widetilde{W}_i}{\Delta^{\pi}_{i}} . \widetilde{Y}_{T_{j+1}} \right) = \left(\frac{\Delta^{\pi}_{n}\widetilde{W}^d_i}{\Delta^{\pi}_{i}}\, {}_n\widetilde{Y}_{T_{j+1}} \right)_{n=1,\ldots,N,\, d=1,\ldots,D},$$

${}_n\widetilde{Y}_{T_{j+1}} := u(T_{j+1}, {}_n\widetilde{X}^{\bar\pi}_{T_{j+1}})$, and with $\widehat{\beta}_{t_i}$ being the $K \times D$ matrix of estimated regression coefficients.

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 system (2.1), which we suppose to be independent of the Brownian increments simulated above:

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

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

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

where $\mathcal{G}^{\widetilde{N}}_0 := \sigma(\Delta^{\pi}_n\widetilde{W}_i;\ i = 1, \ldots \mathcal{I}, n = 1, \ldots, \widetilde{N})$. 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 $(t_i, {}_nX^{\bar\pi}_{t_i})$, $n = 1, \ldots, N$, of $X^{\bar\pi}$ an unbiased estimator for $Y^{up}(\widehat{M}^{\pi})$ is obtained by

$$\widehat{Y}^{up}(\widehat{M}^{\pi}) = \frac{1}{N}\sum_{n=1}^{N} \max_{0 \le j \le \mathcal{J}} \left[h(T_j, {}_nX^{\bar\pi}_{T_j}) - \widehat{m}^{\pi}(T_j, {}_nX^{\bar\pi}_{T_j}, \Delta^{\pi}_nW) \right]. \tag{3.2}$$

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. While Corollary 2.2 suggests to choose a very fine partition π, such choice may cause an instable estimate of the approximate regression functions $\widehat{z}^{\pi}(t_i, x)$, unless the linear space spanned by the basis ψ and the number of simulated paths for the regression are "very large".

4 Variance reduced upper bound estimators

From Corollary 2.2 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 (which can still become computationally expensive, if a very fine partition is required). From (2.5) and the fact that $\widehat{M}^{\pi}$ is a martingale we see that

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

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) and Milstein and Schoenmakers (2002)) for the standard Monte Carlo estimator of $E^{T_{j-1}}Y_{T_j}$. Note that by (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 may write

$$\begin{aligned} Y^{up}(\widehat{M}^{\pi}) &= E\left[\max_{0\le i\le \mathcal{J}}\left(H_{T_i} - \sum_{j=1}^{i}\left(Y_{T_j} - \eta_j^{\pi}\right)\right)\right] \qquad (4.2)\\ &= E\left[\max_{0\le i\le \mathcal{J}}\left(H_{T_i} - \sum_{j=1}^{i}\left(Y_{T_j} - E^{T_{j-1}}Y_{T_j}\right) + \sum_{j=1}^{i}\varepsilon_j^{\pi}\right)\right] \\ &\le Y^{up}(M) + E\left[\sum_{j=1}^{\mathcal{J}}\left(\varepsilon_j^{\pi}\right)_+\right] =: Y^{up}(M) + E\left[\varepsilon_{sum}^{\pi}\right]. \end{aligned}$$

Obviously, also $E\left[\varepsilon_{sum}^{\pi}\right]$ will be closer to zero the finer the grid mesh $|\pi|$ and the larger the set of basis functions.

Instead of making partitions finer and finer while increasing the set of basis functions, one can alternatively take a comparably rough version of $\widehat{M}^{\pi}$, (i.e. with a rougher partition π and a small basis) and employ it as a control variate. This leads to variance reduced estimators as outlined below.

Variance reduced primal-dual algorithm

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

$$\eta_j := E^{T_{j-1}}\left[Y_{T_j}\right] + \varepsilon_j := Y_{T_j} - \left(\overline{M}_{T_j} - \overline{M}_{T_{j-1}}\right), \qquad j = 1, ..., \mathcal{J}. \tag{4.3}$$

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 (4.3) under the (regular) conditional measure $P^{X_{T_{j-1}}}$, and define the (pathwise) unbiased estimator

$$s_j^{(L)} := \frac{1}{L}\sum_{l=1}^{L} {}_l\eta_j \tag{4.4}$$

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

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

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

$$\widehat{Y}_{N,L}^{up}(\overline{M}):=\frac{1}{N}\sum_{n=1}^{N}{}_n\mathcal{U}^{(L)}. \tag{4.5}$$

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*

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

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

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

For the second inequality, let us write using (4.3) and (4.4),

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

It then follows that

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

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

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

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

$$(*)^2 \leq \mathcal{J}\sum_{j=1}^{\mathcal{J}}\left\{E\left[\left(\frac{1}{L}\sum_{l=1}^{L} {}_l\varepsilon_j\right)_+\right]\right\}^2 \leq \mathcal{J}\sum_{j=1}^{\mathcal{J}} E\left[\left(\frac{1}{L}\sum_{l=1}^{L} {}_l\varepsilon_j\right)^2\right]$$
$$= \frac{\mathcal{J}}{L}\sum_{j=1}^{\mathcal{J}} E\left[(\varepsilon_j)^2\right].$$

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

According to Theorem 4.1 the bias of the estimators (4.5) and (4.2) are commonly bounded by $E\left[\varepsilon_{sum}\right]$ when we take $\overline{M} = \widehat{M}^{\pi}$. Furthermore,

$$E\,\widehat{Y}^{up}_{N,L}(\overline{M}) \downarrow Y^{up}(M), \qquad \text{if} \qquad \left(L \to \infty \quad \text{or} \quad \overline{M}_{T_{\mathcal{J}}} \overset{\mathcal{L}_2}{\to} M_{T_{\mathcal{J}}}\right).$$

Variance reduced consumption based estimator

When Y is a lower approximation for the Snell envelope Belomestny and Milstein (2006) derived the following alternative upper bound estimator via the notion of consumption processes

$$Y^{up,BM} := Y_0 + \sum_{j=0}^{\mathcal{J}-1}\left(\max\{H_{T_j}, E^{T_j}Y_{T_{j+1}}\} - Y_{T_j}\right)$$
$$= E\left[H_{T_{\mathcal{J}}}\right] + \sum_{j=0}^{\mathcal{J}-1} E\left[\left(H_{T_j} - E^{T_j}Y_{T_{j+1}}\right)_+\right]$$
$$=: C_E + C^{up},$$

where C_E is the value of a European claim and C^{up} is called a consumption term. The estimation of conditional expectations can be done by standard Monte Carlo. In the same way as above for the primal-dual estimator, we obtain a variance reduced estimator for the consumption term,

$$\widehat{C}^{up}_{N,L} := \frac{1}{N}\sum_{n=1}^{N}\sum_{j=0}^{\mathcal{J}-1}\left({}_nH_{T_j} - {}_ns_j^{(L)}\right)_+$$
$$= \frac{1}{N}\sum_{n=1}^{N}\sum_{j=0}^{\mathcal{J}-1}\left({}_nH_{T_j} - \frac{1}{L}\sum_{l=1}^{L}{}_n\eta_j^{(l)}\right)_+$$

based on a sample of independent outer trajectories ${}_nX$, $n = 1, ..., N$, and L independent realizations ${}_n\eta_j^{(l)}$, $l = 1, ..., L$, of η_j given by (4.3) on each trajectory ${}_nX$. Obviously we have

$$E\left[\widehat{C}^{up,BM}_{N,L}\right] = E\left[\sum_{j=0}^{\mathcal{J}-1} E^{T_{\mathcal{J}}}\left[\left({}_1H_{T_j} - \frac{1}{L}\sum_{l=1}^{L}{}_1\eta_j^{(l)}\right)_+\right]\right] \tag{4.6}$$
$$= E\left[\sum_{j=0}^{\mathcal{J}-1} E^{T_j}\left[\left({}_1H_{T_j} - \frac{1}{L}\sum_{l=1}^{L}{}_1\eta_j^{(l)}\right)_+\right]\right]$$
$$\geq E\left[\sum_{j=0}^{\mathcal{J}-1}\left(H_{T_j} - E^{T_j}\left[Y_{T_{j+1}}\right]\right)_+\right]$$

by the convexity of the $()_+$ operator. Hence the (variance reduced) estimator (4.6) is biased up. In the spirit of Theorem 4.1 one can show for this kind of upper bound also that

$$C_E + E\left[\widehat{C}^{up}_{N,L}\right] \downarrow Y^{up,BM}, \qquad L \to \infty \qquad \textit{or} \qquad \left(\overline{M}_{T_j} \overset{\mathcal{L}_2}{\to} M_{T_j}, \quad j = 1, ..., \mathcal{J}\right).$$

Remark 4.2. The martingale estimator can also be applied to reduce the variance when estimating inner conditional expectations in the policy improvement procedure of Kolodko and Schoenmakers (2006). 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 (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, Schoenmakers (2005)).

In the 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 payoff 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 (see Schoenmakers (2005)), it then follows that

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

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

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

So, after estimating the function $z(\cdot, \cdot)$ by an independent regression procedure we may determine the hedge coefficients $\vartheta(\cdot, \cdot, \cdot)$ (usually called "deltas") 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

$$\vartheta(t, x, \mathfrak{n}) = \mathfrak{n}\frac{\partial}{\partial x}E^{t,x}Y_{T_j} = \mathfrak{n}z(t, x)(b(t, x) - x\sigma_{\mathcal{N}}(t, x))^{-1}.$$

Remark 5.1. The setup in this section covers the situation of a standard Libor (market) model, where X is a system of zero bonds defining the Libor rates, and the numeraire is taken to be the spot Libor measure or the terminal bond measure for instance. For details see Glasserman (2003) and Schoenmakers (2005).

6 Numerical example

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 method. 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 $\{\alpha_l\}$ are estimated, we can define

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

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. Since estimates $\widehat{C}_i$ for continuation values $C_i := E^{t_i}Y_{T_{j+1}}$ can be easily obtained by regression, we, while estimating $Z^{\pi}_{t_i}$, subtract $\widehat{C}_i$ from $Y_{T_{j+1}}$. This leads to the following equivalent definition of $Z^{\pi}_{t_i}$

$$Z^{\pi}_{t_i} := \frac{1}{\Delta^{\pi}_i} E^{t_i}\left[(\Delta^{\pi} W_i)^{\top}(Y_{T_{j+1}} - \widehat{C}_i)\right], \; T_j \le t_i < T_{j+1} \tag{6.1}$$

The subtraction of $\widehat{C}_i$ diminishes the variance and improves the quality of $\widehat{M}^{\pi}$. Another important issue is the choice of partition π. Theoretically, a finer partition implies better quality of $\widehat{M}^{\pi}$. However, in practice, the partition π should not be too fine in order to avoid a variance explosion. In our numerical study we have achieved quite good results by using two different partitions π and $\tilde{\pi}$ such that $\pi \subset \tilde{\pi}$. The first rougher partition is used to estimate regression coefficients β_{t_i}

$$\widehat{\beta}_{t_i} = A^{\oplus}_{t_i}\left(\frac{\Delta^{\pi}\widetilde{W}_i}{\Delta^{\pi}_i}.\widetilde{Y}_{T_{j+1}}\right), \quad t_i \in \pi, \quad T_j \le t_i \le T_{j+1}.$$

Thereafter $\widehat{\beta}_{t_i}$ are interpolated by a constant for points in $[t_i, t_{i+1}]$, that is $\widehat{\beta}_t = \widehat{\beta}_{t_i}$ for all $t \in [t_i, t_{i+1}]$. In such a way one can define $\widehat{z}^{\pi}(t,x) = \psi(t,x)\widehat{\beta}_t$ for all points $t \in \tilde{\pi}$ and construct, with a slight abuse of notation in the case $\pi \ne \tilde{\pi}$,

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

In all examples below we take as the finer partition, $\tilde{\pi} = \bar{\pi}$, i.e. the partition on which the Euler scheme is performed.

Bermudan max calls on D assets

This is a benchmark example studied in Glasserman (2003), Haugh and Kogan (2004) and Rogers (2001) among others. Specifically, the model with D identical assets is considered where each underlying has dividend yield δ. The risk-neutral dynamic of assets is given by

$$dX^d_t = (r - \delta)X^d_t dt + \sigma X^d_t dW^d_t, \quad d = 1, ..., D,$$

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

$$h(X_t) = (\max(X^1_t, ..., X^D_t) - \kappa)^+.$$

We consider an example when $T_j = jT/\mathcal{J}$, $j = 0, ..., \mathcal{J}$, with $T = 3$ and $\mathcal{J} = 9$. For estimating stopping rules $\{\tau_j\}$ we use 5×10^4 paths and take as a regression basis all polynomials of order less than or equal to 3 plus the payoff function h. The Euler scheme was performed on equidistant partition $\bar{\pi}$ with $|\bar{\pi}| = 0.01$. The same number of paths and the same basis functions have been used to estimate $\widehat{\beta}_{t_i}$, $t_i \in \pi$, where $\pi = \{T_0, \ldots, T_{\mathcal{J}}\}$. Now, local constant interpolation allows us to define $\widehat{\beta}_t$ and hence $\widehat{z}^{\pi}(t,x)$ for all $t \in \bar{\pi}$. Let us note that the complexity of the algorithm with interpolated $\widehat{\beta}_t$ corresponds in this case to the complexity of the usual Longstaff-Schwartz method because regression is only performed on the exercise grid. Moreover, matrices $A^{\oplus}_{t_i}$ computed during constructing the approximation Y can, in principle, be used here again provided that the same paths are used to estimate $\widehat{z}^{\pi}(t,x)$.

The results for $D = 2$ and $D = 5$ are presented in Table 1 in dependence on x_0 with $X_0 = (X_0^1, \ldots, X_0^D)^T$, $X_0^1 = \ldots = X_0^D = x_0$. Upper bounds $\widehat{Y}^{up}_{N,L}(0)$ are computed by primal-dual algorithm, hence by nested Monte Carlo, with N outer and L inner simulations without variance reduction (see for comparison Glasserman (2003)). As we see the standard primal-dual method requires in some cases more than 40 inner simulations to achieve the accuracy of the non-nested estimator. In fact, the latter one is regarding computation time comparable with the primal-dual using one inner simulation.

It is interesting to look at the dependence of the difference $\Delta := \widehat{Y}^{up}(\widehat{M}^\pi) - Y_0$ on the number of Monte Carlo paths N and the maximal order of regression polynomials p used for estimating coefficients β. In Fig.1 the corresponding curves for the two dimensional out of the money ($x_0 = 90$) Bermudan max call with the same parameters as before are presented. Note that the set of polynomial basis function is always extended by adding the pay-off function h. Fig. 1 indicates that the less N is the less improvement is observable with increasing p.

Let us turn now to the performance of our method in the setup of variance reduction. We compare upper bounds of the nested Monte Carlo estimator (primal-dual) with and without using control variates. In Fig. 2 the upper bound $\widehat{Y}^{up}_{N,L}(\overline{M})$ is shown as a function of L for the cases of the zero martingale $\overline{M} = 0$ (original primal-dual method) and $\overline{M} = \widehat{M}^\pi$ as estimated before. Again the example of 2-dimensional Bermudan max call with $x_0 = 90$ is considered and coefficients $\{\beta_{t_i}\}$ are estimated using 5×10^4 Monte Carlo simulations and all polynomials of order less than or equal to 3. Comparing Fig. 2 with Table 1 we conclude that the accuracy of $Y^{up}(M) \approx \widehat{Y}^{up}_{10^4,200}(0)$ is achieved by the variance reduced primal-dual estimator $\widehat{Y}^{up}_{10^4,L}(\widehat{M}^\pi)$ already with $L = 90$.

7 Conclusion

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 two alternatives to this algorithm. The first algorithm is fast, as it requires linear simulation cost only, and turns out to deliver good upper bounds. If nonetheless a higher accuracy is required, we suggest a variance reduced version of the primal-dual algorithm which allows to compute upper bounds with the same accuracy (as with the latter one) at lower costs.

A Proof of Theorem 2.1

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

$$
\begin{aligned}
Z^{\pi,d}_{t_i} &= \frac{1}{\Delta^\pi_i} E^{t_i}\left[\left(\Delta^\pi W^d_i\right)\left(Y_{T_{j+1}} - E^{T_j}[Y_{T_{j+1}}]\right)\right] \\
&= \frac{1}{\Delta^\pi_i} E^{t_i}\left[\left(\int_{t_i}^{t_{i+1}} dW^d_s\right)\left(\int_{T_j}^{T_{j+1}} Z_s dW_s\right)\right] \\
&= \frac{1}{\Delta^\pi_i} E^{t_i}\left[\int_{t_i}^{t_{i+1}} Z^d_s ds\right] \qquad (1.2)
\end{aligned}
$$

It follows from (1.2) that without any further assumptions,

$$
\lim_{|\pi|\to 0} E\left[\sum_{t_i\in\pi, T_j\leq t_i<T_j} \int_{t_i}^{t_{i+1}} |Z_s - Z^\pi_{t_i}|^2 ds\right] = 0 \qquad (1.3)
$$

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

$$
\begin{aligned}
E\left[\max_{0\leq j\leq \mathcal{J}} |M^\pi_{T_j} - M_{T_j}|^2\right] &\leq 4E\left[|M^\pi_T - M_T|^2\right] \\
= \quad 4E&\left[\sum_{j=0}^{\mathcal{J}-1} \sum_{t_i\in\pi, T_j\leq t_i<T_j} \int_{t_i}^{t_{i+1}} |Z_s - Z^\pi_{t_i}|^2 ds\right], \qquad (1.4)
\end{aligned}
$$

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 (1.2) these results imply that (1.3) can be strengthened to

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

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 (1.4).

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., under revision.

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 BSDEs. WIAS Preprint No. 1160, Berlin.

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

N. Chen and P. Glasserman (2005). Additive and Multiplicative Duals for American Option Pricing. Working paper, Finance and Stochastics, to appear.

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.

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. Opeations Research, **52**, No. 2, 258–270.

F. Jamshidian (2006). The duality of optimal exercise and domineering claims: A Doob-Meyer decomposition approach to the Snell envelope. Working paper, http://wwwhome.math.utwente.nl/ jamshidianf/.

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**, No. 1, 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 $\widehat{Y}^{up}(\widehat{M}^{\pi})$	Upper Bound $\widehat{Y}^{up}_{10^4,200}(0)$	Upper Bound $\widehat{Y}^{up}_{10^4,40}(0)$
	90	7.9751±0.139	8.6963±0.052	8.2311±0.091	8.621±0.092
2	100	13.883±0.177	14.515±0.073	14.182±0.011	15.23±0.013
	110	21.291±0.205	21.972±0.095	21.681±0.015	23.67±0.017
	90	16.523±0.194	18.134±0.069	17.163±0.012	17.53±0.014
5	100	26.042±0.232	27.976±0.085	27.216±0.016	27.87±0.016
	110	36.526±0.263	38.882±0.098	38.577±0.020	39.70±0.023

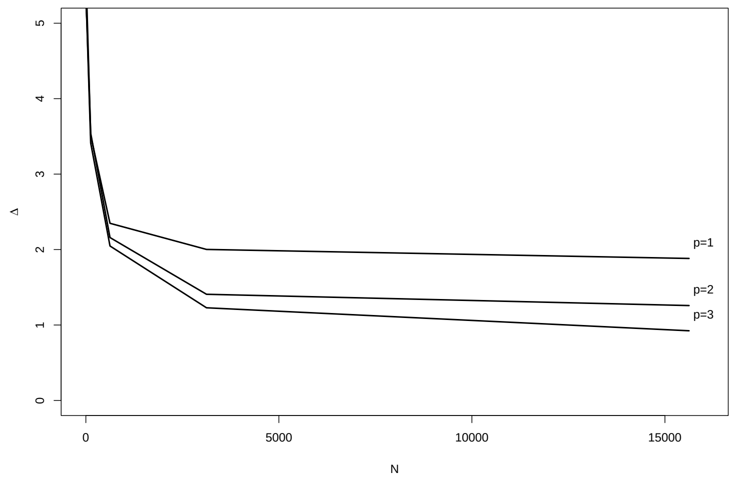

Figure 1: Difference $\Delta = \widehat{Y}^{up}(\widehat{M}^{\pi}) - Y_0$ in dependence on the number of Monte Carlo paths N and the maximal order p of polynomials used for regression.

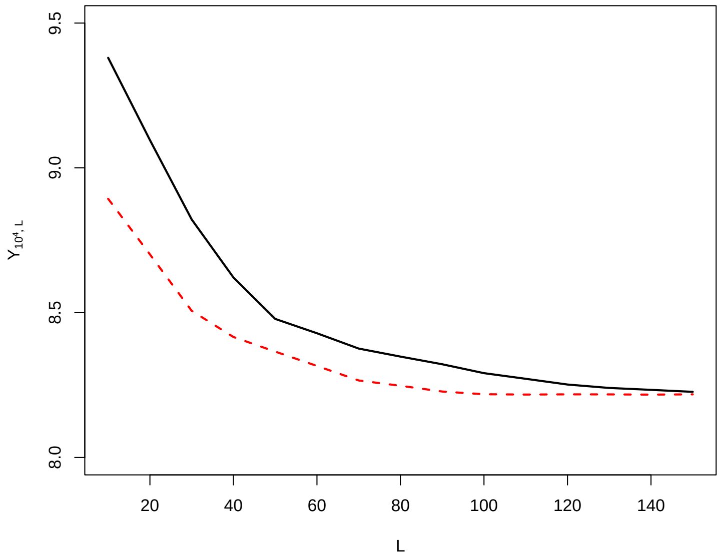

Figure 2: Upper bounds $\widehat{Y}^{up}_{N,L}(0)$ (solid line) and $\widehat{Y}^{up}_{N,L}(\widehat{M}^{\pi})$ (dash line) in dependence on the number of inner Monte Carlo paths L, the number of outer paths N being equal to 5×10^4.