

A primal-dual algorithm for BSDEs

Christian Bender¹, Nikolaus Schweizer¹, and Jia Zhuo²

September 26, 2013

Abstract

We generalize the primal-dual methodology, which is popular in the pricing of early-exercise options, to a backward dynamic programming equation associated with time discretization schemes of (reflected) backward stochastic differential equations (BSDEs). Taking as an input some approximate solution of the backward dynamic program, which was pre-computed, e.g., by least-squares Monte Carlo, our methodology allows to construct a confidence interval for the unknown true solution of the time discretized (reflected) BSDE at time 0. We numerically demonstrate the practical applicability of our method in two five-dimensional nonlinear pricing problems where tight price bounds were previously unavailable.

Keywords: Backward SDE, numerical approximation, Monte Carlo, option pricing.

AMS classification: 65C30, 65C05, 91G20, 91G60.

1 Introduction

In this paper we aim at constructing tight confidence intervals for the solution $(Y_i)_{i=0,\dots,n}$ of a dynamic programming equation of the form

$$Y_i = \max\{S_i, E_i[Y_{i+1}] + f(i, Y_i, E_i[\beta_{i+1}Y_{i+1}])\Delta_i\}, \quad Y_n = S_n \quad (1)$$

at time $i = 0$. Dynamic programming equations of the form (1) naturally arise in time discretization schemes for (reflected) BSDEs. We assume that $(\Omega, \mathcal{F}, (\mathcal{F}_i)_{i=0,\dots,n}, P)$ is a filtered probability space and $E_i[\cdot]$ denotes conditional expectation given \mathcal{F}_i . The reflecting barrier S is an adapted $\mathbb{R} \cup \{-\infty\}$ -valued process, β is an adapted \mathbb{R}^D -valued process related to the driver of the BSDE (e.g. suitably truncated and normalized increments of a D -dimensional Brownian motion), the generator $f : \Omega \times \{0, \dots, n-1\} \times \mathbb{R} \times \mathbb{R}^D \rightarrow \mathbb{R}$ is an adapted random field, and Δ_i are constants which can be thought of as the stepsizes of a time discretization scheme. Appropriate integrability and continuity assumptions will be specified later on.

The special case $f \equiv 0$ of (1) is the well-known recursion for the valuation of Bermudan options. In the wake of the financial crisis, there is an increased interest in ‘small’ nonlinearities in pricing. These are due, e.g., to counterparty risk or funding risk – and had largely been neglected in practice. Building on the BSDE literature and its early pricing applications such as Bergman

¹Saarland University, Department of Mathematics, Postfach 151150, D-66041 Saarbrücken, Germany, bender@math.uni-sb.de; schweizer@math.uni-sb.de.

Financial support by the Deutsche Forschungsgemeinschaft under grant BE3933/5-1 is gratefully acknowledged.

²Department of Mathematics, University of Southern California, 3620 S. Vermont Ave., KAP 104 Los Angeles, CA 90089-2532, jiazhuo@usc.edu.

(1995), Duffie et al. (1996) or the examples in El Karoui et al. (1997), the number of pricing problems which have been formulated as BSDEs – and thus have a discretization of the form (1) – is steadily growing. Recent examples include funding risk (Laurent et al., 2012; Crépey et al., 2013), counterparty risk (Crépey et al., 2013; Henry-Labordère, 2012), model uncertainty (Guyon and Henry-Labordère, 2010; Alanko and Avellaneda, 2013), and hedging under transaction costs (Guyon and Henry-Labordère, 2010). In some of these examples the nonlinearity depends on the delta or the gamma of the option, which can be incorporated in our discrete time setting by choosing the weights β appropriately. The aim of the present paper is to provide a unified and numerically efficient framework for calculating upper and lower price bounds for these problems – parallel to the well-known primal-dual bounds in Bermudan option pricing.

The error due to the time discretization (1) for BSDEs driven by a Brownian motion has been thoroughly analyzed in the literature under various regularity conditions. We refer to Zhang (2004); Bouchard and Touzi (2004); Gobet and Labart (2007); Gobet and Makhlof (2010) for the non-reflected case (corresponding to $S_i \equiv -\infty$ for $i < n$) and to Bally and Pagès (2003); Ma and Zhang (2005); Bouchard and Chassagneux (2008) for the reflected case. We emphasize that the results in the present paper can also be applied to the time discretization schemes for BSDEs driven by a Brownian motion with generators with quadratic growth as in Chassagneux and Richou (2013), time discretization schemes for BSDEs with jumps considered in Bouchard and Elie (2008), and the time discretization scheme for fully nonlinear parabolic PDEs by Fahim et al. (2011).

A standard procedure for solving an equation of type (1) numerically is the so-called approximate dynamic programming approach. Here, the conditional expectations in (1) are replaced by some approximate conditional expectations operator. The main difficulty of this approach is, that in each step backwards in time a conditional expectation must be computed numerically, building on the approximate solution one step ahead. This leads to a high order nesting of conditional expectations. Hence, it is crucial that the approximate conditional expectations operator can be nested several times without exploding computational cost. Among the techniques which have been applied and analyzed in the context of BSDEs driven by a (high-dimensional) Brownian motion are least-squares Monte Carlo (Lemor et al., 2006; Bender and Denk, 2007), quantization (Bally and Pagès, 2003), Malliavin Monte Carlo (Bouchard and Touzi, 2004), cubature on Wiener space (Crisan and Manolarakis, 2012), and sparse grid methods (Zhang et al., 2013).

Although convergence rates are available in the literature for these different methods, the quality of the numerical approximation in the practically relevant pre-limit situation is typically difficult to assess. Generalizing the primal-dual methodology, which was introduced by Andersen and Broadie (2004) in the context Bermudan option pricing, we suggest to take the numerical solution of the approximate dynamic program as an input, in order to construct a confidence interval for Y_0 via a Monte Carlo approach. In a nutshell, the rationale is to find a maximization problem and a minimization problem with value Y_0 , for which optimal controls are available in terms of the true solution $(Y_i)_{i=0,\dots,n}$ of the dynamic program (1). Using the approximate solution instead of the true one, then yields suboptimal controls for these two optimization problems. If the numerical procedure in the approximate dynamic program was successful, these controls are close to optimal and lead to tight lower and upper bounds for Y_0 . Unbiased estimators for the lower and the upper bound can finally be computed by plain Monte Carlo, which results in a confidence interval for Y_0 . Bender and Steiner (2013) provides a different a posteriori criterion for BSDEs which is better suited for qualitative convergence analysis than for deriving quantitatively meaningful bounds on Y_0 . The two approaches are thus complimentary.

The paper is organized as follows: In Section 2 we briefly discuss some basic properties of the dynamic programming equation (1) and show how it arises in our two numerical examples, funding risk and counterparty risk. The case of a convex or concave generator f is treated in Section 3. In Section 3.1 we first suggest a pathwise approach to the dynamic programming equation (1) which avoids the evaluation of conditional expectations in the backward recursion in time. This pathwise approach depends on the choice of a $(D + 1)$ -dimensional martingale. For convex generators it leads to the construction of supersolutions for (1) and to a minimization problem over martingales with value process Y_i . We then note that, due to convexity, Y_i can also be represented as the supremum over a class of classical optimal stopping problems. This representation can be thought of as a discrete time, reflected analogue of a result in El Karoui et al. (1997) for continuous time, non-reflected BSDEs driven by a Brownian motion. If we think of this maximization problem as the primal problem, then the pathwise approach can be interpreted as a dual minimization problem in the sense of information relaxation. This type of duality was first introduced independently by Rogers (2002) and Haugh and Kogan (2004) in the context of Bermudan option pricing, and was later extended by Brown et al. (2010) to general discrete time stochastic control problems. The situation for concave generators is slightly different, because the reflecting barrier is genuinely convex. Nonetheless, similar ideas can be applied in order to derive a minimization problem and a maximization problem with value process Y_i .

In Section 3.2 we explain, how the representations for Y_0 as the value of a maximization and a minimization problem can be exploited in order to construct confidence intervals for Y_0 via Monte Carlo simulation. This algorithm generalizes the primal-dual algorithm of Andersen and Broadie (2004) from optimal stopping problems (i.e., the case $f \equiv 0$) to the case of convex or concave generators. We also suggest some generic control variates which turn out to be powerful in our numerical examples. Numerical examples for the pricing of a European and a Bermudan option on the maximum of five assets under different interest rates for borrowing and lending (funding risk) are presented in Section 3.3. For constructing the input approximations, we apply the least-squares Monte Carlo algorithm of Lemor et al. (2006) and its martingale basis variant by Bender and Steiner (2012) with just a few (up to seven) basis functions. This turns out to be sufficient for achieving very tight 95% confidence intervals with relative error of typically less than 1% between lower and upper confidence bound in our five-dimensional test examples.

For general generators f which are neither convex nor concave, we suggest in Section 4.1 to apply the input approximation of the approximate dynamic program in order to construct an auxiliary generator f^{up} , which is convex and dominates f , and another one f^{low} , which is concave and dominated by f . This construction can be done in a way that f^{up} and f^{low} converge to f , when the input approximation of the approximate dynamic program approaches the true solution. The methods of Section 3 can then be applied to the convex generator f^{up} and the concave generator f^{low} in order to build a confidence interval for Y_0 in the general case. In Section 4.2, we test the performance of this algorithm in two applications, the previous example of funding risk and a model of counterparty credit risk where the driver is neither concave nor convex. Again, tight price bounds can be achieved. Appendix A sets our discrete time results into the context of their continuous time analogues.

2 Discrete time reflected BSDEs

Suppose $(\Omega, \mathcal{F}, (\mathcal{F}_i)_{i=0, \dots, n}, P)$ is a filtered probability space in discrete time. We consider the discretized version of a reflected BSDE of the form (1). Throughout the paper we make the

following assumptions: The time increments Δ_i , $i = 0, \dots, n-1$, are positive real numbers. S is an adapted process with values in $\mathbb{R} \cup \{-\infty\}$ such that

$$\sum_{i=0}^{n-1} E[|S_i \mathbf{1}_{\{S_i > -\infty\}}|] + E[|S_n|] < \infty;$$

The random field $f : \Omega \times \{0, \dots, n-1\} \times \mathbb{R} \times \mathbb{R}^D \rightarrow \mathbb{R}$ is measurable, $f(\cdot, y, z)$ is adapted for every $(y, z) \in \mathbb{R} \times \mathbb{R}^D$, and $\sum_{i=1}^{n-1} E[|f(i, 0, 0)|] < \infty$. Moreover, there are adapted, nonnegative processes $\alpha^{(d)}$, $d = 0, \dots, D$ such that the stochastic Lipschitz conditions

$$|f(i, y, z) - f(i, y', z')| \leq \alpha_i^{(0)} |y - y'| + \sum_{d=1}^D \alpha_i^{(d)} |z_d - z'_d|$$

hold for every $(y, z), (y', z') \in \mathbb{R} \times \mathbb{R}^D$. Finally, β is a bounded, adapted \mathbb{R}^D -valued process and the following relations hold:

$$\alpha_i^{(0)} < \frac{1}{\Delta_i}, \quad \sum_{d=1}^D \alpha_i^{(d)} |\beta_{d,i+1}| \leq \frac{1}{\Delta_i}. \quad (2)$$

A straightforward contraction mapping argument shows that under these assumptions there exists a unique adapted and integrable process Y such that (1) is satisfied.

Example 2.1. To illustrate the setting let us introduce the two nonlinear pricing problems which also appear in our numerical experiments: Pricing with different interest rates for borrowing and lending, and pricing in a reduced-form model of counterparty credit risk. Going back to Bergman (1995), the first one is a standard example in the BSDE literature. Laurent et al. (2012) have recently emphasized its practical relevance in the context of funding risk. Following the financial crisis there has also been increased interest in credit risk models similar to the second example, see Pallavicini et al. (2012); Crépey et al. (2013); Henry-Labordère (2012) and the references therein.

(i) Let there be a financial market with two riskless and D risky assets. The prices of the risky assets X_t^1, \dots, X_t^D evolve according to

$$dX_t^d = X_t^d \left(\mu_t^d dt + \sum_{k=1}^D \sigma_t^{d,k} dW_t^k \right),$$

where W is a standard D -dimensional Brownian motion, and where μ and σ are predictable and bounded processes. Moreover, σ is assumed to be a.s. invertible with bounded inverse. The filtration is given by the usual augmented Brownian filtration. The two riskless assets have bounded and predictable short rates R_t^l and R_t^b with $R_t^l \leq R_t^b$ a.s. These are the interest rates for lending and borrowing, i.e., an investor can only hold positive positions in the first one, and only negative ones in the second. Consider a square-integrable European claim $h(X_T)$ with maturity T . It is well-known that a replicating portfolio for $h(X_T)$ is characterized by two processes $Y_t \in \mathbb{R}$ and $Z_t \in \mathbb{R}^D$ which solve the BSDE

$$dY_t = -f(t, Y_t, Z_t) dt + Z_t^\top dW_t \quad (3)$$

with terminal condition $Y_T = h(X_T)$ where

$$f(t, y, z) = -R_t^l y - z^\top \sigma_t^{-1} \left(\mu_t - R_t^l \bar{\mathbf{1}} \right) + (R_t^b - R_t^l) \left(y - z^\top \sigma_t^{-1} \bar{\mathbf{1}} \right)_-,$$

see Bergman (1995) or the survey paper of El Karoui et al. (1997). Here, $\bar{1}$ denotes the vector containing only ones in \mathbb{R}^D and $^\top$ denotes matrix transposition. The function f is both convex and Lipschitz continuous. Our key quantity of interest is the claim's fair price at time 0 given by Y_0 . $Z_t^\top \sigma_t^{-1}$ corresponds to the vector of amounts of money invested in the risky assets at time t . Discretizing time and taking conditional expectations gives a recursion of the form (1) for Y , see Zhang (2004); Bouchard and Touzi (2004). In this European case $S_i \equiv -\infty$ for $i < n$ and $S_n \equiv h(X_T)$. The discretization of the Z -part is given by $Z_i = E_i \left[\frac{W_{t_{i+1}} - W_{t_i}}{t_{i+1} - t_i} Y_{i+1} \right]$. This corresponds to a vector of Malliavin derivatives and is thus naturally related to a delta hedge. In view of (1) we would thus like to choose $\beta_{i+1} = \frac{W_{t_{i+1}} - W_{t_i}}{t_{i+1} - t_i}$. However, in order to fulfill condition (2) we have to truncate the Brownian increments at some value. Since $\frac{W_{t_{i+1}} - W_{t_i}}{t_{i+1} - t_i}$ is a vector of normal random variables with standard deviation of order $(t_{i+1} - t_i)^{\frac{1}{2}}$, we can make this truncation error small as the time discretization gets finer, see e.g. Lemor et al. (2006). For a Bermudan or American claim, (3) is replaced by a suitable reflected BSDE. In the discretization, S_i is then the payoff from exercising at time t_i .

(ii) The second example is a special case of the model of counterparty credit risk due to Duffie et al. (1996). We change the setting of (i) by assuming there is only one riskless asset with rate R_t which can be both borrowed and lent. Moreover, we consider a risk-neutral valuation framework, i.e., $\mu_t = R_t \bar{1}$. Given a square-integrable European claim $h(X_T)$ with maturity T , we denote by Y_t the claim's fair price at time t conditional on no default having occurred yet. The claim's possible default is modelled through a stopping time which is the first jump time of a Poisson process with intensity Q_t . Here, $Q_t = Q(Y_t)$ is a decreasing, continuous and bounded function of Y_t , i.e., if the claim's value is low, default becomes more likely. If default occurs at time t , the claim's holder receives a fraction $\delta \in (0, 1)$ of the current value Y_t . Following Proposition 3 in Duffie et al. (1996) the value process is then characterized by the nonlinear relation

$$Y_t = E_t \left[\int_t^T f(s, Y_s) ds + h(X_T) \right],$$

where $f(t, y) = -(1 - \delta)Q(y)y - R_t y$. Discretizing naturally leads to an equation of type (1) with $\beta \equiv 0$. Condition (2) then reduces to the requirement that the time discretization is sufficiently fine.

The dynamic programming equation (1) implies that the solution Y also solves the optimal stopping problem

$$Y_i = \operatorname{esssup}_{\tau \in \mathcal{S}_i} E_i \left[S_\tau + \sum_{j=i}^{\tau-1} f(j, Y_j, E_j[\beta_{j+1} Y_{j+1}]) \Delta_j \right], \quad i = 0, \dots, n,$$

where \mathcal{S}_i is the set of stopping times with values bigger or equal to i . This optimal stopping problem is unusual in the sense that the reward upon stopping depends on the Snell envelope Y_i . Note that one can pose restrictions on the set of admissible stopping times by choosing the set $\{(i, \omega); S_i(\omega) = -\infty\}$, at which exercise is never optimal. We can hence restrict the supremum in this optimal stopping problem to the subset $\bar{\mathcal{S}}_i \subset \mathcal{S}_i$ of stopping times τ which take values in $\mathcal{E}(\omega) = \{i; S_i(\omega) > -\infty\}$. An optimal stopping time is given by

$$\tau_i^* = \inf \{j \geq i; S_j \geq E_j[Y_{j+1}] + f(j, Y_j, E_j[\beta_{j+1} Y_{j+1}]) \Delta_j\} \quad (4)$$

We also note the following alternative representation of Y_i via optimal stopping of a nonlinear functional.

Proposition 2.2. *For every $i = 0, \dots, n$,*

$$Y_i = \operatorname{esssup}_{\tau \in \bar{S}_i} Y_i^{(\tau)}$$

where $(Y_j^{(\tau)})_{j \geq i}$ solves the dynamic programming equation

$$Y_j^{(\tau)} = E_j[Y_{j+1}^{(\tau)}] + f(j, Y_j^{(\tau)}, E_j[\beta_{j+1} Y_{j+1}^{(\tau)}]) \Delta_j, \quad i \leq j < \tau, \quad Y_\tau^{(\tau)} = S_\tau$$

Moreover, the stopping time τ_i^* , defined in (4) is optimal.

This representation is a direct consequence of the following simple, but useful, comparison theorem.

Proposition 2.3. *Suppose there are stopping times $\sigma \leq \tau$ such that for every $\sigma \leq i < \tau$*

$$\begin{aligned} Y_i^{up} &\geq \max\{S_i, E_i[Y_{i+1}^{up}] + f(i, Y_i^{up}, E_i[\beta_{i+1} Y_{i+1}^{up}]) \Delta_i\} \\ Y_i^{low} &\leq \max\{S_i, E_i[Y_{i+1}^{low}] + f(i, Y_i^{low}, E_i[\beta_{i+1} Y_{i+1}^{low}]) \Delta_i\} \end{aligned}$$

and $Y_\tau^{up} \geq Y_\tau^{low}$. Then, under the standing assumptions, $Y_i^{low} \leq Y_i^{up}$ holds for every $\sigma \leq i \leq \tau$.

Proof. We define, for $i = 1, \dots, n$,

$$\Delta Y_i = (Y_i^{up} - Y_i^{low}) \mathbf{1}_{\{\sigma \leq i \leq \tau\}}.$$

It is sufficient to show that $\Delta Y_i \geq 0$ for every $i = 1, \dots, n$. We prove this assertion by backward induction and note that it holds in the case $i = n$ by assumption. Now, suppose that $\Delta Y_{i+1} \geq 0$ is already shown. Then, on the set $\{Y_i^{low} > S_i\} \cap \{\sigma \leq i < \tau\}$ we obtain by the Lipschitz assumption on f ,

$$\begin{aligned} \Delta Y_i &\geq E_i[\Delta Y_{i+1}] + (f(i, Y_i^{up}, E_i[\beta_{i+1} Y_{i+1}^{up}]) - f(i, Y_i^{low}, E_i[\beta_{i+1} Y_{i+1}^{low}])) \Delta_i \\ &\geq E_i \left[\Delta Y_{i+1} \left(1 - \sum_{d=1}^D \alpha_i^{(d)} |\beta_{d,i+1}| \Delta_i \right) \right] - \alpha_i^{(0)} |\Delta Y_i| \Delta_i \\ &\geq -\alpha_i^{(0)} |\Delta Y_i| \Delta_i, \end{aligned}$$

which yields $\Delta Y_i \geq 0$. On the set $\{Y_i^{low} \leq S_i\} \cup \{i \geq \tau\} \cup \{i < \sigma\}$, the inequality $\Delta Y_i \geq 0$ is obvious. \square

3 The case of a convex generator

3.1 Optimization problems related to the dynamic programming equation

In this section we discuss how to construct ‘tight’ supersolutions and subsolutions to the dynamic programming equation (1) when the generator f is convex in $(y, z) \in \mathbb{R}^{1+D}$. We first explain a pathwise approach, which leads to supersolutions due to the convexity of f . Roughly speaking, the idea is to remove all conditional expectations from equation (1) and subtract martingale

increments, wherever conditional expectations were removed. To this end, let us fix a one-dimensional martingale M^0 and an \mathbb{R}^D -valued martingale M . We define the non-adapted process $\theta_i^{up} = \theta_i^{up}(M^0, M)$ via

$$\theta_i^{up} = \max\{S_i, \theta_{i+1}^{up} - (M_{i+1}^0 - M_i^0) + f_i(\theta_i^{up}, \beta_{i+1}\theta_{i+1}^{up} - (M_{i+1} - M_i))\Delta_i\}, \quad \theta_n^{up} = S_n. \quad (5)$$

Once the martingales are chosen, this recursion can be solved path by path. After solving the recursion, we take conditional expectations. Exploiting the convexity of f and of the max-operator as well as the martingale property, we obtain,

$$E_i[\theta_i^{up}] \geq \max\{S_i, E_i[E_{i+1}[\theta_{i+1}^{up}]] + f_i(E_i[\theta_i^{up}], E_i[\beta_{i+1}E_{i+1}[\theta_{i+1}^{up}]])\Delta_i\}$$

Consequently, $E_i[\theta_i^{up}(M^0, M)]$ is a supersolution of (1) and by the comparison result of Proposition 2.3

$$E_i[\theta_i^{up}(M^0, M)] \geq Y_i.$$

We now choose $M^{0,*}$ and M^* as the Doob martingales of Y and βY , respectively, and claim that

$$\theta_i^{up,*} := \theta_i^{up}(M^{0,*}, M^*) = Y_i$$

almost surely. This can easily be shown by backward induction on i , with the case $i = n$ being trivial. Suppose that the claim is true for $i + 1$. Then, making use of the definition of the Doob decomposition,

$$\begin{aligned} \theta_i^{up,*} &= \max\{S_i, Y_{i+1} - (Y_{i+1} - E_i[Y_{i+1}]) + f_i(\theta_i^{up,*}, \beta_{i+1}Y_{i+1} - (\beta_{i+1}Y_{i+1} - E_i[\beta_{i+1}Y_{i+1}]))\Delta_i\} \\ &= \max\{S_i, E_i[Y_{i+1}] + f_i(\theta_i^{up,*}, E_i[\beta_{i+1}Y_{i+1}])\Delta_i\}. \end{aligned}$$

By the Lipschitz property of f in the y -variable, a straightforward contraction mapping argument shows that $\theta_i^{up,*} = Y_i$. We summarize these findings in the following theorem.

Theorem 3.1. *Suppose f is convex in (y, z) . Then, for every $i = 0, \dots, n$,*

$$Y_i = \operatorname{ess\,inf}_{(M^0, M) \in \mathcal{M}_{1+D}} E_i[\theta_i^{up}(M^0, M)]$$

where \mathcal{M}_{1+D} denotes the set of \mathbb{R}^{1+D} -valued martingales and $\theta^{up}(M^0, M)$ is defined by the pathwise dynamic programming equation (5). Moreover, the martingale $(M^{0,*}, M^*)$, where $M^{0,*}$ and M^* are the Doob martingales of Y and βY , is optimal even in the sense of pathwise control, i.e.

$$\theta_i^{up}(M^{0,*}, M^*) = Y_i, \quad P\text{-a.s.}$$

The previous theorem can be applied to compute upper confidence bounds on Y_0 . To this end one first chooses a $(1 + D)$ -dimensional martingale, which one thinks is close to the Doob martingale of $(Y, \beta Y)$. This can e.g. be (related to) the Doob martingale of an approximation \tilde{Y} of Y which was pre-computed by an algorithm of one's choice. Then one solves the pathwise dynamic program in (5) and finally approximates the expectation by averaging over sample paths. The details of such an implementation are discussed in Section 3.2 below. One issue, which arises in this approach, is that the pathwise dynamic program is not explicit in time, as θ_i^{up} appears on both sides of the equation. It can be solved by a Picard iteration to a given precision. In some situations, the following explicit expression in terms of a pathwise maximization problem is advantageous.

Proposition 3.2. Suppose f is convex in (y, z) , and define the convex conjugate in the y -variable by

$$f^{\#y}(\omega, i, r, z) = \sup\{ry - f(\omega, i, y, z); \quad y \in \mathbb{R}\}$$

which is defined on

$$D_{f^{\#y}}^{(i, \omega, z)} := \{r \in \mathbb{R}; f^{\#y}(\omega, i, r, z) < \infty\} \subset [-\alpha_i^{(0)}(\omega), \alpha_i^{(0)}(\omega)].$$

Then, for $(M^0, M) \in \mathcal{M}_{1+D}$ and $i = 0, \dots, n-1$, θ_i^{up} as defined in (5) can be rewritten as

$$\theta_i^{up} = \max \left\{ S_i, \sup_{r \in D_{f^{\#y}}^{(i, \omega, z)}} \frac{1}{1 - r\Delta_i} \left(\theta_{i+1}^{up} - (M_{i+1}^0 - M_i^0) - f^{\#y}(i, r, \beta_{i+1}\theta_{i+1}^{up} - (M_{i+1} - M_i)\Delta_i) \right) \right\}. \quad (6)$$

Proof. By convexity, we have $f(i, \cdot) = (f(i, \cdot)^{\#y})^{\#r}$, where $\#r$ denotes the convex conjugate in the r -variable of $f^{\#y}(i, \cdot)$. Hence,

$$\theta_i^{up} = \max \left\{ S_i, \sup_{r \in D_{f^{\#y}}^{(i, \omega, z)}} \left(\theta_{i+1}^{up} - (M_{i+1}^0 - M_i^0) + r\theta_i^{up}\Delta_i - f^{\#y}(i, r, \beta_{i+1}\theta_{i+1}^{up} - (M_{i+1} - M_i)\Delta_i) \right) \right\}$$

By the same argument as on p. 36 in El Karoui et al. (1997) the supremum is achieved at some r^* . Hence,

$$\theta_i^{up} = \max\{S_i, \theta_{i+1}^{up} - (M_{i+1}^0 - M_i^0) + r^*\theta_i^{up}\Delta_i - f^{\#y}(i, r^*, \beta_{i+1}\theta_{i+1}^{up} - (M_{i+1} - M_i)\Delta_i)\}$$

For $\theta_i^{up} > S_i$ we, thus obtain

$$\theta_i^{up} = \frac{1}{1 - r^*\Delta_i} \left(\theta_{i+1}^{up} - (M_{i+1}^0 - M_i^0) - f^{\#y}(i, r^*, \beta_{i+1}\theta_{i+1}^{up} - (M_{i+1} - M_i)\Delta_i) \right)$$

Consequently, θ_i^{up} is dominated by the right hand side of the assertion. The reverse inequality can be shown in the same way. \square

Example 3.3. In Example 2.1 (i),

$$f^{\#y}(i, r, z) = z^\top \sigma_i^{-1}(\mu_i + r\bar{1})$$

and the maximizer must belong to the set $\{-R_i^b, -R_i^l\}$, because $f(i, \cdot) = (f(i, \cdot)^{\#y})^{\#r}$. Hence, for the European option case, a recursion for θ^{up} , which is explicit in time, reads

$$\theta_i^{up} = \sup_{r \in \{-R_i^b, -R_i^l\}} \frac{1}{1 - r\Delta_i} \left(\theta_{i+1}^{up} - (M_{i+1}^0 - M_i^0) - [\beta_{i+1}\theta_{i+1}^{up} - (M_{i+1} - M_i)]^\top \sigma_i^{-1}(\mu_i + r\bar{1})\Delta_i \right).$$

In order to derive a maximization problem with value process given by Y_i , we denote by $f^\#$ the convex conjugate of f in (y, z) , i.e.

$$f^\#(\omega, i, r, \rho) = \sup\{ry + \rho^\top z - f(\omega, i, y, z); \quad (y, z) \in \mathbb{R}^{1+D}\}$$

which is defined on

$$D_{f^\#}^{(i,\omega)} := \{(r, \rho) \in \mathbb{R}^{1+D}; f^\#(\omega, i, r, \rho) < \infty\} \subset \prod_{d=0}^D [-\alpha_i^{(d)}(\omega), \alpha_i^{(d)}(\omega)].$$

We also define

$$\mathcal{U}_i(f^\#) := \{(r_j, \rho_j)_{j \geq i} \text{ adapted}; \sum_{j=i}^{n-1} E[|f^\#(j, r_j, \rho_j)|] < \infty\}$$

The following result is a discrete time reflected analogue of Proposition 3.4 in El Karoui et al. (1997). For discrete time (non-reflected) BSDEs a similar result (for convex generators in z only) can be found in Cheridito and Stadje (2013) under a different set of assumptions.

Theorem 3.4. *Suppose f is convex in (y, z) . Let*

$$\theta_i^{\text{low}}(\tau, r, \rho) := \Gamma_{i,\tau}(r, \rho) S_\tau - \sum_{j=i}^{\tau-1} \Gamma_{i,j}(r, \rho) \frac{f^\#(j, r_j, \rho_j) \Delta_j}{1 - r_j \Delta_j} \quad \text{where} \quad \Gamma_{i,j}(r, \rho) := \prod_{k=i}^{j-1} \frac{1 + \rho_k^\top \beta_{k+1} \Delta_k}{1 - r_k \Delta_k}$$

Then,

$$Y_i = \text{esssup}_{\tau \in \bar{\mathcal{S}}_i} \text{esssup}_{(r,\rho) \in \mathcal{U}_i(f^\#)} E_i[\theta_i^{\text{low}}(\tau, r, \rho)]$$

Maximizers are given by any $(r_j^*, \rho_j^*)_{j \geq i}$ such that for $j = i, \dots, n-1$,

$$r_j^* Y_j + \rho_j^{*\top} E_j[\beta_{j+1} Y_{j+1}] - f^\#(j, r_j^*, \rho_j^*) = f(j, Y_j, E_j[\beta_{j+1} Y_{j+1}]) \quad (7)$$

and τ_i^* as defined in (4).

Proof. Fix $i \in \{0, \dots, n\}$. Given a stopping time $\tau \in \bar{\mathcal{S}}_i$ and a pair $(r, \rho) \in \mathcal{U}_i(f^\#)$ define

$$Y_j(\tau, r, \rho) := E_j \left[\theta_j^{\text{low}}(\tau, r, \rho) \right], \quad i \leq j \leq \tau.$$

Then, $Y_\tau(\tau, r, \rho) = S_\tau$ and, for $i \leq j < \tau$,

$$\begin{aligned} Y_j(\tau, r, \rho) &= E_j[Y_{j+1}(\tau, r, \rho)] + (\rho_j^\top E_j[\beta_{j+1} Y_{j+1}(\tau, r, \rho)] + r_j Y_j(\tau, r, \rho) - f^\#(j, r_j, \rho_j)) \Delta_j \\ &\leq E_j[Y_{j+1}(\tau, r, \rho)] + f(j, Y_j(\tau, r, \rho), E_j[\beta_{j+1} Y_{j+1}(\tau, r, \rho)]) \Delta_j, \end{aligned}$$

where the last estimate is due to the fact that $f^{\#\#} = f$ by convexity. Now the comparison result in Proposition 2.3 and Proposition 2.2 imply

$$Y_i(\tau, r, \rho) \leq Y_i^{(\tau)} \leq Y_i.$$

For the converse inequality, we first notice that the argument in the Lemma on p. 36 in El Karoui et al. (1997) implies that there is a pair of processes $(r_j^*, \rho_j^*)_{j \geq i} \in \mathcal{U}_i(f^\#)$ such that (7) holds. Then, by the definition of τ_i^* , we obtain for $i \leq j < \tau_i^*$

$$\begin{aligned} Y_j &= E_j[Y_{j+1}] + f(j, Y_j, E_j[\beta_{j+1} Y_{j+1}]) \Delta_j \\ &= E_j[Y_{j+1}] + r_j^* Y_j + \rho_j^{*\top} E_j[\beta_{j+1} Y_{j+1}] - f^\#(j, r_j^*, \rho_j^*) \end{aligned}$$

As $Y_{\tau_i^*} = S_{\tau_i^*}$, we conclude that by uniqueness for this dynamic programming equation,

$$Y_i = Y_i(\tau_i^*, r^*, \rho^*).$$

□

Remark 3.5. If we think of the representation in Theorem 3.4 as a ‘primal’ maximization problem, then the representation in Theorem 3.1 can be interpreted as a dual minimization problem in the sense of information relaxation. This dual approach was introduced for Bermudan option pricing by Rogers (2002) and Haugh and Kogan (2004), and was further developed for discrete time stochastic control problems by Brown et al. (2010). Indeed, given a martingale $(M^0, M) \in \mathcal{M}_{1+D}$, we define

$$\begin{aligned} \mathfrak{p}_{M^0, M} : \{i, \dots, n\} \times \prod_{j=i}^{n-1} D_{f^\#}^{(j, \omega)} &\rightarrow L^1(\Omega, P), \\ (k, (r, \rho)) &\mapsto \sum_{j=i}^{k-1} \Gamma_{i,j}(r, \rho) \frac{(M_{j+1}^0 - M_j^0) + \rho_j^\top (M_{j+1} - M_j) \Delta_j}{1 - r_j \Delta_j}. \end{aligned}$$

Then, for every $(\tau, (r, \rho)) \in \bar{\mathcal{S}}_i \times \mathcal{U}_i(f^\#)$,

$$E_i[\mathfrak{p}_{M^0, M}(\tau, r, \rho)] = 0. \quad (8)$$

We next relax the adaptedness property of the controls $(\tau, (r, \rho))$ and observe that, by Theorem 3.4 and (8),

$$\begin{aligned} Y_i &= \operatorname{esssup}_{\tau \in \bar{\mathcal{S}}_i} \operatorname{esssup}_{(r, \rho) \in \mathcal{U}_i(f^\#)} E_i[\theta_i^{\text{low}}(\tau, r, \rho) - \mathfrak{p}_{M^0, M}(\tau, r, \rho)] \\ &\leq E_i \left[\max_{k=i, \dots, n} \max_{\substack{(r_j, \rho_j) \in D_{f^\#}^{(j, \omega)} \\ j=i, \dots, k-1}} \left(\theta_i^{\text{low}}(k, r, \rho) - \mathfrak{p}_{M^0, M}(k, r, \rho) \right) \right] \\ &=: E_i[\tilde{\theta}_i(M^0, M)]. \end{aligned}$$

Notice that the maximum on the right hand side of the inequality is taken pathwise, which means that we may now choose anticipating controls. The rationale of the information relaxation approach is that one allows for anticipating controls, but subtracts a penalty, here $\mathfrak{p}_{M^0, M}$. The penalty does not penalize non-anticipating controls by (8). We say that a penalty \mathfrak{p}^* is optimal, if it penalizes anticipating controls in a way that the pathwise maximum is achieved at a non-anticipating control. This implies

$$Y_i = E_i \left[\max_{k=i, \dots, n} \max_{\substack{(r_j, \rho_j) \in D_{f^\#}^{(j, \omega)} \\ j=i, \dots, k-1}} \left(\theta_i^{\text{low}}(k, r, \rho) - \mathfrak{p}^*(k, r, \rho) \right) \right].$$

In the present setting, one can show that

$$\tilde{\theta}_i(M^0, M) = \theta_i^{\text{up}}(M^0, M).$$

To see this, one first derives a recursion formula for $\tilde{\theta}_i(M^0, M)$ and then follows the arguments behind Proposition 3.2. In particular, Theorem 3.1 shows that an optimal penalty is given by $\mathfrak{p}_{M^0, *, M^*}$.

When f is concave in (y, z) one can prove the following theorem by essentially the same arguments as in the convex case. It is not completely symmetric to the convex case, because the reflecting barrier is genuinely convex.

Theorem 3.6. *Suppose f is concave in (y, z) .*

(i) *Then, for every $i = 0, \dots, n$,*

$$Y_i = \operatorname{ess\,inf}_{M^0 \in \mathcal{M}_1} \operatorname{ess\,inf}_{(r, \rho) \in \mathcal{U}_i((-f)^\#)} E_i[\vartheta^{up}(r, \rho, M^0)], \quad \text{where}$$

$$\vartheta^{up}(r, \rho, M^0) = \max_{k=i, \dots, n} \Gamma_{i,k}(-r, -\rho) S_k + \sum_{j=i}^{k-1} \Gamma_{i,j}(-r, -\rho) \frac{(-f)^\#(j, r_j, \rho_j) \Delta_j}{1 + r_j \Delta_j} - (M_k^0 - M_i^0).$$

Minimizers are given by $(r_j^, \rho_j^*)_{j \geq i}$ satisfying*

$$-r_j^* Y_j - \rho_j^{*\top} E_j[\beta_{j+1} Y_{j+1}] + (-f)^\#(j, r_j^*, \rho_j^*) = f(j, Y_j, E_j[\beta_{j+1} Y_{j+1}])$$

and $M^{0,}$ being the martingale part of the Doob decomposition of $(Y_j \Gamma_{i,j}(-r^*, -\rho^*))_{j \geq i}$.*

(ii) *Given a stopping time $\tau \in \bar{S}_i$ and a martingale $(M^0, M) \in \mathcal{M}_{1+D}$, define $\vartheta_j^{low} = \vartheta_j^{low}(\tau, M^0, M)$ for $i \leq j < \tau$ via*

$$\vartheta_j^{low} = \vartheta_{j+1}^{low} - (M_{i+1}^0 - M_i^0) + f(j, \vartheta_j^{low}, \beta_{j+1} \vartheta_{j+1}^{low} - (M_{j+1} - M_j) \Delta_j), \quad \vartheta_\tau^{low} = S_\tau.$$

Then,

$$Y_i = \operatorname{ess\,sup}_{\tau \in \bar{S}_i} \operatorname{ess\,sup}_{(M^0, M) \in \mathcal{M}_{1+D}} E_i[\vartheta_i^{low}(\tau, M^0, M)]$$

A maximizer is given by the triplet $(\tau_i^, M^{0,*}, M^*)$, where τ_i^* was defined in (4) and $M^{0,*}$, M^* are the Doob martingales of Y and βY , respectively.*

3.2 A primal-dual algorithm

In this section we explain, how the results of the previous subsection can be applied in order to construct confidence intervals for Y_0 in the spirit of the Andersen and Broadie (2004) algorithm for Bermudan option pricing, when f is convex in (y, z) . To this end we suppose that we are in a Markovian setting, i.e. $f(i, \cdot) = F(i, X_i, \cdot)$ and $S_i = G_i(X_i)$ depend on ω only through an \mathbb{R}^N -valued Markovian process X_i , and β_{i+1} is independent of \mathcal{F}_i . Then, there are deterministic functions $y_i(x), q_i(x), z_{d,i}(x), d = 1, \dots, D$, such that

$$Y_i = y_i(X_i), \quad E_i[Y_{i+1}] = q_i(X_i), \quad E_i[\beta_{d,i+1} Y_{i+1}] = z_{d,i}(X_i)$$

We assume that approximations $\tilde{y}_i(x), \tilde{q}_i(x)$ and $\tilde{z}_{d,i}(x)$ for these functions are pre-computed by some numerical algorithm. In our numerical experiments below, a least-squares Monte Carlo estimator for the conditional expectations in (1) is applied in order to construct these approximations, but other choices are possible.

Given these approximations, we sample Λ^{out} independent copies

$$(X_i(\lambda), \beta_i(\lambda); i = 0, \dots, n)_{\lambda=1, \dots, \Lambda^{out}}$$

of $(X_i, \beta_i; i = 0, \dots, n)$, to which we refer as ‘outer’ paths. For the upper confidence bound we apply Theorem 3.1. We thus wish to calculate $\theta_i^{up}(M^0, M)$ for some martingales M^0, M , which are ‘close’ to the unknown Doob martingales of Y and βY . We apply instead the Doob

martingales of the approximations $\tilde{y}(X)$ and $\beta\tilde{y}(X)$ to Y and βY . Along the λ th outer path this leads in view of (5) to

$$\begin{aligned} & \theta_i^{up}(\lambda) \\ = & \max\{G_i(X_i(\lambda)), \theta_{i+1}^{up}(\lambda) - (\tilde{y}_{i+1}(X_{i+1}(\lambda)) - E[\tilde{y}_{i+1}(X_{i+1})|X_i = X_i(\lambda)]) \\ & + f(i, \theta_i^{up}(\lambda), \beta_{i+1}(\lambda)\theta_{i+1}^{up}(\lambda) - (\beta_{i+1}(\lambda)\tilde{y}_{i+1}(X_{i+1}(\lambda)) - E[\beta_{i+1}\tilde{y}_{i+1}(X_{i+1})|X_i = X_i(\lambda)])\Delta_i\} \end{aligned} \quad (9)$$

Then, by Theorem 3.1, the estimator

$$\hat{Y}^{up} := \frac{1}{\Lambda^{out}} \sum_{\lambda=1}^{\Lambda^{out}} \theta_0^{up}(\lambda)$$

for Y_0 , which is obtained by averaging over the outer paths, has a positive bias. In general, we cannot expect that the conditional expectations in (9) can be calculated in closed form. Instead we apply a conditionally unbiased estimator for these conditional expectations by averaging over a set of ‘inner’ samples. For each i and each outer path $X(\lambda)$ generate Λ^{in} independent copies of (X_{i+1}, β_{i+1}) under the conditional law given that $X_i = X_i(\lambda)$. These samples are denoted by $(X_{i+1}(\lambda, l), \beta_{i+1}(\lambda, l))$, $l = 1, \dots, \Lambda^{in}$. We then define the plain Monte Carlo estimators for the conditional expectations in (9) along the λ th outer paths by

$$\begin{aligned} \hat{E}[\tilde{y}_{i+1}(X_{i+1})|X_i = X_i(\lambda)] &= \frac{1}{\Lambda^{in}} \sum_{l=1}^{\Lambda^{in}} \tilde{y}_{i+1}(X_{i+1}(\lambda, l)) \\ \hat{E}[\beta_{i+1}\tilde{y}_{i+1}(X_{i+1})|X_i = X_i(\lambda)] &= \frac{1}{\Lambda^{in}} \sum_{l=1}^{\Lambda^{in}} \beta_{i+1}(\lambda, l)\tilde{y}_{i+1}(X_{i+1}(\lambda, l)) \end{aligned} \quad (10)$$

Then, in the recursive construction for $\theta_i^{up}(\lambda)$ we replace the conditional expectations in (9) by the plain Monte Carlo estimators (10) in all instances and apply the notation $\theta_i^{up, AB}(\lambda)$. The corresponding upper bound estimator for Y_0 is obtained by averaging over the outer paths

$$\hat{Y}^{up, AB} := \frac{1}{\Lambda^{out}} \sum_{\lambda=1}^{\Lambda^{out}} \theta_0^{up, AB}(\lambda).$$

Here, the superscript ‘ AB ’ stands for Andersen and Broadie, who suggested this method for Bermudan option in 2004. By a straightforward application of Jensen’s inequality we observe that, by convexity of the max-operator and of f , $\hat{Y}^{up, AB}$ has an additional positive bias compared to \hat{Y}^{up} , which is due to the inner simulations. In particular, $\hat{Y}^{up, AB}$ has a positive bias as an estimator for Y_0 .

The numerical experiments below illustrate that the additional bias due to the inner simulations may be substantial with a moderate number of inner paths (say 1,000). It therefore appears to be essential to apply variance reduction techniques for the estimation of the conditional expectations in (9) by Monte Carlo. We suggest some control variates, for which we merely require that

$$E[\beta_{d, i+1}], \quad E[\beta_{d, i+1} \beta_{d', i+1}], \quad d, d' = 1, \dots, D$$

are available in closed form. This is e.g. the case when $\beta_{d,i+1}$ is (up to a constant) given by truncated increments of independent Brownian motions. In this case we perform an orthogonal projection of $\tilde{y}_{i+1}(X_{i+1})$ on the span of the random variables $(\beta_{1,i+1}, \dots, \beta_{D,i+1})$ under the conditional probability given X_i . This orthogonal projection is given by

$$\beta_{i+1}^\top B_{i+1}^+ E[\beta_{i+1} \tilde{y}_{i+1}(X_{i+1}) | X_i = x],$$

where B_{i+1}^+ is the Moore-Penrose pseudoinverse of the matrix

$$B_{i+1} = (E[\beta_{d,i+1} \beta_{d',i+1}])_{d,d'=1,\dots,D}.$$

Here, we made use of the assumption that β_{i+1} is independent of \mathcal{F}_i . If \tilde{y} and \tilde{z} are good approximations of y and z , then $\tilde{z}_i(X_i)$ is also expected to be a good approximation of $E[\beta_{i+1} \tilde{y}_{i+1}(X_{i+1}) | X_i]$. These considerations motivate us to replace the estimators (9) for the conditional expectations in (10) by

$$\begin{aligned} & \hat{E}^C[\tilde{y}_{i+1}(X_{i+1}) | X_i = X_i(\lambda)] \\ &= E[\beta_{i+1}]^\top B_{i+1}^+ \tilde{z}_i(X_i(\lambda)) + \frac{1}{\Lambda^{in}} \sum_{l=1}^{\Lambda^{in}} \left(\tilde{y}_{i+1}(X_{i+1}(\lambda, l)) - \beta_{i+1}(\lambda, l)^\top B_{i+1}^+ \tilde{z}_i(X_i(\lambda)) \right) \end{aligned}$$

and

$$\begin{aligned} \hat{E}^C[\beta_{i+1} \tilde{y}_{i+1}(X_{i+1}) | X_i = X_i(\lambda)] &= E[\beta_{i+1}] \tilde{q}_i(X_i(\lambda)) + B_{i+1} B_{i+1}^+ \tilde{z}_i(X_i(\lambda)) \\ &+ \frac{1}{\Lambda^{in}} \sum_{l=1}^{\Lambda^{in}} \beta_{i+1}(\lambda, l) (\tilde{y}_{i+1}(X_{i+1}(\lambda, l)) - \tilde{q}_i(X_i(\lambda)) - \beta_{i+1}(\lambda, l)^\top B_{i+1}^+ \tilde{z}_i(X_i(\lambda))), \end{aligned} \quad (11)$$

which are still conditionally unbiased. The estimator $\hat{Y}^{up,ABC}$ is then calculated analogously to $\hat{Y}^{up,AB}$, but applying (11) instead of (10). Again, by Jensen's inequality, the 'up'-estimator has a positive bias. For a classical optimal stopping problem, a similar control variate for inner simulations was suggested by Belomestny et al. (2009) in the special case when β_{i+1} are increments of independent Brownian motions. For Bermudan option pricing problems various other constructions for the input martingales have been introduced in the literature, see e.g. Belomestny et al. (2009), Desai et al. (2012) and Schoenmakers et al. (2013). These constructions can also be adapted to the present BSDE setting.

In order to construct an estimator for Y_0 with a negative bias, we define $\tilde{\tau}(\lambda)$ by

$$\tilde{\tau}(\lambda) = \inf\{j \geq 0; G_j(X_j(\lambda)) \geq \tilde{q}_j(X_j(\lambda)) + F(j, X_j(\lambda), \tilde{y}_j(X_j(\lambda)), \tilde{z}_j(X_j(\lambda))) \Delta_j\}$$

and $(\tilde{r}_i(\lambda), \tilde{\rho}_i(\lambda))_{i=0,\dots,n-1}$ as (approximate) solutions of

$$\begin{aligned} & \tilde{r}_j(\lambda) \tilde{y}_j(X_j(\lambda)) + \tilde{\rho}_j(\lambda) \tilde{z}_j(X_j(\lambda)) - F^\#(j, X_j(\lambda), \tilde{r}_j(\lambda), \tilde{\rho}_j(\lambda)) \\ &= F(j, X_j(\lambda), \tilde{y}_j(X_j(\lambda)), \tilde{z}_j(X_j(\lambda))). \end{aligned} \quad (12)$$

Then, by Theorem 3.4, the plain Monte Carlo estimator

$$\begin{aligned} & \frac{1}{\Lambda^{out}} \left(\sum_{\lambda=1}^{\Lambda^{out}} \Gamma_{0,\tilde{\tau}(\lambda)}(\tilde{r}(\lambda), \tilde{\rho}(\lambda)) G(\tilde{\tau}(\lambda), X_{\tilde{\tau}(\lambda)}(\lambda)) \right. \\ & \left. + \sum_{j=0}^{\tilde{\tau}(\lambda)-1} \Gamma_{0,j}(\tilde{r}(\lambda), \tilde{\rho}(\lambda)) \frac{F^\#(j, X_j(\lambda), \tilde{r}_j(\lambda), \tilde{\rho}_j(\lambda)) \Delta_j}{1 - \tilde{r}_j(\lambda) \Delta_j} \right) \end{aligned}$$

for Y_0 has a negative bias. We recommend to run this estimator with a control variate in order to reduce the number of samples Λ^{out} . In this regard, we suggest the use of

$$\begin{aligned} & \sum_{j=0}^{\tilde{\tau}-1} \Gamma_{0,j}(\tilde{r}, \tilde{\rho}) \frac{\tilde{y}_{j+1}(X_{j+1}) - E[\tilde{y}_{j+1}(X_{j+1})|X_j]}{1 - \tilde{r}_j \Delta_j} \\ & + \Gamma_{0,j}(\tilde{r}, \tilde{\rho}) \frac{\tilde{\rho}_j^\top \Delta_j (\beta_{j+1} \tilde{y}_{j+1}(X_{j+1}) - E[\beta_{j+1} \tilde{y}_{j+1}(X_{j+1})|X_j])}{1 - \tilde{r}_j \Delta_j}, \end{aligned} \quad (13)$$

if the conditional expectations are available in closed form. If not, a set of ‘inner’ simulations will be required for the construction of the upper bound estimator anyway, and this inner sample can be used to estimate the conditional expectations in the control variate (13) via (11). The resulting estimator with a negative bias is denoted $\hat{Y}^{low,ABC}$. Finally, an (asymptotic) 95% confidence interval for Y_0 can be constructed by adding (resp. subtracting) 1.96 empirical standard deviations to the upper estimator (from the lower estimator).

3.3 Numerical examples

We apply the above algorithm in the context of adjusting the option price value due to funding constraints in the context of Example 2.1 (i). We consider the pricing problem of a European and a Bermudan call spread option with maturity T on the maximum of D assets, which are modeled by independent, identically distributed geometric Brownian motions with drift μ and volatility σ whose values at time $t_i = Ti/n$, $i = 0, \dots, n$ are denoted by $X_{d,i}$. The interest rates R^b and R^l are constant over time. The generator f is then given by

$$F(i, x, y, z) = -R^l y - \frac{\mu - R^l}{\sigma} \sum_{d=1}^D z_d + (R^b - R^l) \left(y - \frac{1}{\sigma} \sum_{d=1}^D z_d \right)_-$$

We define $\beta_{d,i+1}(t_{i+1} - t_i)$ as the truncated Brownian increment driving the d th stock over the period $[t_i, t_{i+1}]$. The payoff of the option is given by

$$G_i(x) = \begin{cases} (\max_{d=1, \dots, D} x_d - K_1)_+ - 2(\max_{d=1, \dots, D} x_d - K_2)_+, & i \in \mathcal{E} \\ -\infty, & i \notin \mathcal{E}. \end{cases}$$

for strikes K_1, K_2 and a set of time points \mathcal{E} at which the option can be exercised. Hence, $\mathcal{E} = \{n\}$ gives a European option. For the Bermudan option case we consider the situation of four exercise dates which are equidistant over the time horizon, i.e. $\mathcal{E} = \{n/4, n/2, 3n/4, n\}$. Unless otherwise noted, we use the following parameter values:

$$D = 5, T = 0.25, R^l = 0.01, R^b = 0.06, X_{d,0} = 100, \mu = 0.05, \sigma = 0.2, K_1 = 95, K_2 = 115.$$

We first generate approximations $\tilde{y}^{LGW}, \tilde{q}^{LGW}, \tilde{z}^{LGW}$ by the least-squares Monte Carlo algorithm of Lemor et al. (2006). This algorithm requires the choice of a set of basis functions. Then an empirical regression on the span of these basis functions is performed with a set of Λ^{reg} sample paths, which are independent of the outer and inner samples required for the primal-dual algorithm later on. In the European option case we apply the following sets of basis functions: For the implementation with $b^y = 2$ basis functions we choose 1 and $E[G_n(X_n)|X_i = x]$ for the computation of $\tilde{y}_i^{LGW}(x)$, $\tilde{q}_i^{LGW}(x)$ and $x_d \frac{d}{dx_d} E[G_n(X_n)|X_i = x]$ for the computation of $\tilde{z}_{d,i}^{LGW}(x)$,

$d = 1, \dots, D$. For call options on the maximum of D Black-Scholes stocks, closed form expressions for the option price and its delta in terms of a multivariate normal distribution are derived in Johnson (1987). In the present setting, this formula can be simplified to an expectation of a function of a one-dimensional standard normal random variable, see e.g. Belomestny et al. (2009). As a trade-off between computational time and accuracy, we approximate this expectation via quantization of the one-dimensional standard normal distribution with 21 grid points. In the implementation with $b^y = 7$ basis functions we additionally apply x_1, \dots, x_5 as basis functions for $\tilde{y}_i^{LGW}(x)$, $\tilde{q}_i^{LGW}(x)$, and x_d as a basis function for $\tilde{z}_{d,i}^{LGW}(x)$. For the Bermudan option case we use six basis functions for $\tilde{y}_i^{LGW}(x)$, $\tilde{q}_i^{LGW}(x)$, namely 1, $E[G_j(X_j)|X_i = x]$, $j \in \mathcal{E}$, and $\max_{j \in \mathcal{E}, j \geq i} E[G_j(X_j)|X_i = x]$. The corresponding deltas $x_d \frac{d}{dx_d} E[G_j(X_j)|X_i = x]$, $j \in \mathcal{E}, j \geq i$, are chosen as basis functions for $\tilde{z}_{d,i}^{LGW}(x)$.

In the European option case, this choice of basis functions also allows to apply the martingale basis algorithm of Bender and Steiner (2012), although a slight bias is introduced due to the approximation of the basis functions by the quantization approach. Compared to the generic least-squares Monte Carlo algorithm the use of martingale basis functions allows to compute some conditional expectations in the approximate backward dynamic program explicitly. These closed form computations can be thought of as a perfect control variate within the regression algorithm.

For the computation of the upper confidence bounds we use the explicit recursion for θ^{up} derived in Example 3.3. For the computation of the lower confidence bound we note that the defining equation (12) for the approximate controls $(\tilde{r}, \tilde{\rho})$ for the lower bound can be solved explicitly as

$$\begin{aligned}\tilde{r}_i &= -R^b \mathbf{1}_{\{\tilde{y}(i, X_i) \leq \sigma^{-1} \sum_{d=1}^D \tilde{z}_s(i, X_i)\}} - R^l \mathbf{1}_{\{\tilde{y}(i, X_i) > \sigma^{-1} \sum_{d=1}^D \tilde{z}_s(i, X_i)\}} \\ \tilde{\rho}_{d,i} &= -\sigma^{-1}(\tilde{r}_i + \mu).\end{aligned}$$

Figure 1 illustrates the effectiveness of the control variate for the inner samples in the computation of the upper bounds for the European option case with $n = 40$ time steps. The input approximation is generated by the martingale basis algorithm with seven basis functions and $\Lambda^{reg} = 1,000$ sample paths for the empirical regression. The figure depicts the corresponding upper bound estimator for the option price Y_0 with $\Lambda^{out} = 10,000$ sample paths as a function of the number of inner samples Λ^{in} . From top to bottom, it shows the upper estimators $\hat{Y}^{up,AB}$ (i.e. without inner control variate), $\hat{Y}^{up,ABC}$ (i.e. with inner control variate), and for comparison the lower bound estimator $\hat{Y}^{low,AB}$. We immediately observe that the predominant part of the upper bias in $\hat{Y}^{up,AB}$ stems from the subsampling in the approximate construction of the Doob martingales. Without the use of inner control variates, the relative error between upper and lower estimator is about 6% for $\Lambda^{in} = 100$ inner samples and decreases to about 1.5% for $\Lambda^{in} = 1,000$ inner samples. Application of the inner control variates reduces this relative error to less than 0.25% even in the case of only $\Lambda^{in} = 100$ inner samples.

Table 1 illustrates the influence of different input approximations. It shows realizations of the lower estimator $\hat{Y}^{low,ABC}$ and the upper estimator $\hat{Y}^{up,ABC}$ for the option price Y_0 as well as the empirical standard deviations, as the number of time steps increases from $n = 40$ to $n = 160$. The column on the left explains which algorithm is run for the input approximation. Here, LGW stands for the Lemor-Gobet-Warin algorithm and MB for the martingale basis algorithm. It also states the number of regression samples and the number of basis functions b^y , which are applied in the least-squares Monte Carlo. The lower and upper price estimates for the Bermudan option case are presented in the last two lines. In this case, the martingale basis algorithm is not available,




Figure 1: Influence of the number of inner simulations and the control variate: upper bound without inner control variate, upper bound with inner control variate, and lower bound (from the top to the bottom).

and the Lemor-Gobet-Warin algorithm is run with the six basis functions stated above. We apply $\Lambda^{out} = 10,000$ and $\Lambda^{in} = 100$ samples in all cases.

By and large, the table shows that in this 5-dimensional example extremely tight 95% confidence intervals can be computed by the primal-dual algorithm, although the input approximations are based on very few, but well chosen, basis functions. For the martingale basis algorithm as input approximation with just two basis functions and 100 regression paths the relative error between lower and upper 95%-confidence bound is about 0.7% even for $n = 160$ steps in the time discretization. It can be further decreased to less than 0.5%, when seven basis functions and 1,000 regression paths are applied. If one takes the input approximation of the Lemor-Gobet-Warin algorithm with the same set of basis functions, then the primal-dual algorithm can in principle produce confidence intervals of about the same length as in the case of the martingale basis algorithm. However, in our simulation study the number of regression paths must be increased by a factor of 1,000 in order to obtain input approximations which have the same quality as those computed by the martingale basis algorithm. Hence our numerical results demonstrate the huge variance reduction effect of the martingale basis algorithm. In the Bermudan option case, the primal-dual algorithm still yields 95%-confidence intervals with a relative width of less than 1% for up to $n = 160$ time steps, when the input approximation is computed by the Lemor-Gobet-Warin algorithm with 6 basis functions and 1 million regression paths.

4 The case of a non-convex generator

4.1 Optimization problems related to the dynamic programming equation

In this section we skip the assumption on the convexity (or concavity) of the generator f and merely assume that the standing assumptions are in force. In this situation the construction of confidence bounds for Y_0 can be based on approximations of f by convex and concave genera-

Algorithm \ n	40		80		120		160	
LGW $\Lambda^{reg}=10^4, b^y=2$	13.7786 (0.0028)	13.8339 (0.0031)	13.7597 (0.0033)	13.8858 (0.0041)	13.7583 (0.0037)	13.9482 (0.0051)	13.7478 (0.0043)	14.0149 (0.0062)
LGW $\Lambda^{reg}=10^5, b^y=2$	13.7783 (0.0022)	13.8172 (0.0024)	13.7817 (0.0022)	13.8443 (0.0027)	13.7848 (0.0024)	13.8682 (0.0029)	13.7855 (0.0025)	13.8967 (0.0033)
MB $\Lambda^{reg}=10^2, b^y=2$	13.7850 (0.0022)	13.8185 (0.0023)	13.7898 (0.0021)	13.8435 (0.0025)	13.7863 (0.0022)	13.8578 (0.0025)	13.7904 (0.0022)	13.8779 (0.0026)
LGW $\Lambda^{reg}=10^5, b^y=7$	13.7818 (0.0020)	13.8140 (0.0021)	13.7767 (0.0020)	13.8321 (0.0022)	13.7789 (0.0022)	13.8560 (0.0025)	13.7764 (0.0025)	13.8902 (0.0031)
LGW $\Lambda^{reg}=10^6, b^y=7$	13.7829 (0.0017)	13.8079 (0.0018)	13.7867 (0.0016)	13.8233 (0.0018)	13.7884 (0.0017)	13.8393 (0.0020)	13.7867 (0.0017)	13.8515 (0.0022)
MB $\Lambda^{reg}=10^3, b^y=7$	13.7844 (0.0017)	13.8077 (0.0017)	13.7897 (0.0016)	13.8245 (0.0017)	13.7887 (0.0016)	13.8353 (0.0019)	13.7880 (0.0017)	13.8485 (0.0021)
$LGW^{Bermudan}$ $\Lambda^{reg}=10^5$	15.5362 (0.0028)	15.5664 (0.0028)	15.5441 (0.0037)	15.6160 (0.0035)	15.5246 (0.0041)	15.6396 (0.0042)	15.5342 (0.0041)	15.6886 (0.0048)
$LGW^{Bermudan}$ $\Lambda^{reg}=10^6$	15.5422 (0.0028)	15.5684 (0.0026)	15.5482 (0.0032)	15.6050 (0.0033)	15.5441 (0.0035)	15.6364 (0.0039)	15.5443 (0.0039)	15.6694 (0.0042)

Table 1: Upper and lower price bounds for different time discretizations and input approximations in the European and Bermudan case. Standard deviations are in brackets.

tors. We assume that some integrable approximation $(\tilde{Y}_i, \tilde{Z}_i)$ of $(Y_i, E_i[\beta_{i+1}Y_{i+1}])$ is given. This approximation can be pre-computed by any algorithm. We now choose a measurable function

$$h^{up} : \Omega \times \{0, \dots, n\} \times \mathbb{R} \times \mathbb{R}^D \times \mathbb{R} \times \mathbb{R}^D \rightarrow \mathbb{R}$$

with the following properties:

- a) $h^{up}(\cdot, \tilde{y}, \tilde{z}; y, z)$ is adapted for every $(\tilde{y}, \tilde{z}), (y, z) \in \mathbb{R} \times \mathbb{R}^D$. Moreover h^{up} satisfies the stochastic Lipschitz condition

$$|h^{up}(i, \tilde{y}, \tilde{z}; y, z) - h^{up}(i, \tilde{y}, \tilde{z}; y', z')| \leq \alpha_i^{(0)} |y - y'| + \sum_{d=1}^D \alpha_i^{(d)} |z_d - z'_d|$$

for every $(\tilde{y}, \tilde{z}), (y, z), (y', z') \in \mathbb{R} \times \mathbb{R}^D$.

- b) $h^{up}(i, \tilde{y}, \tilde{z}; y, z)$ is convex in (y, z) , $h^{up}(i, \tilde{y}, \tilde{z}; 0, 0) = 0$ for every $(\tilde{y}, \tilde{z}) \in \mathbb{R} \times \mathbb{R}^D$, and

$$h^{up}(i, \tilde{y}, \tilde{z}; \tilde{y} - y, \tilde{z} - z) \geq f(i, y, z) - f(i, \tilde{y}, \tilde{z})$$

for every $(\tilde{y}, \tilde{z}), (y, z) \in \mathbb{R} \times \mathbb{R}^D$.

A generic choice is the function

$$h^{|up|}(i, \tilde{y}, \tilde{z}; y, z) = \alpha_i^{(0)} |y| + \sum_{d=1}^D \alpha_i^{(d)} |z_d|,$$

which obviously satisfies these properties. We will illustrate in the numerical examples below, that it might be beneficial to tailor the function h^{up} to the specific problem instead of applying the generic choice $h^{|up|}$.

Given h^{up} , (\tilde{Y}, \tilde{Z}) we define $\Theta_i^{h^{up}} = \Theta_i^{h^{up}}(\tilde{Y}, \tilde{Z})$ via

$$\begin{aligned} \Theta_i^{h^{up}} = & \max\{S_i, \Theta_{i+1}^{h^{up}} - (\tilde{Y}_{i+1} - E_i[\tilde{Y}_{i+1}]) + f_i(\tilde{Y}_i, \tilde{Z}_i)\Delta_i \\ & + h^{up}(i, \tilde{Y}_i, \tilde{Z}_i; \tilde{Y}_i - \Theta_i^{h^{up}}, \tilde{Z}_i - \beta_{i+1}\Theta_{i+1}^{h^{up}} + \beta_{i+1}\tilde{Y}_{i+1} - E_i[\beta_{i+1}\tilde{Y}_{i+1}])\Delta_i\}, \end{aligned} \quad (14)$$

initiated at $\Theta_n^{h^{up}} = S_n$. We then obtain the following minimization problem with value process Y_i in terms of $\Theta_i^{h^{up}}(\tilde{Y}, \tilde{Z})$.

Theorem 4.1. *For every $i = 0, \dots, n$,*

$$Y_i = \underset{(\tilde{Y}_j, \tilde{Z}_j)_{j \geq i} \text{ adapted and integrable}}{\text{essinf}} E_i[\Theta_i^{h^{up}}(\tilde{Y}, \tilde{Z})].$$

Moreover, a minimizing pair is given by $(Y_j^*, Z_j^*) = (Y_j, E_j[\beta_{j+1}Y_{j+1}])$ which even satisfies the principle of pathwise optimality.

Proof. We fix a pair of adapted and integrable processes (\tilde{Y}, \tilde{Z}) and define Y_j^{up} , $j \geq i$, as

$$Y_j^{up} = \max\{S_j, E_j[Y_{j+1}^{up}] + [f(j, \tilde{Y}_j, \tilde{Z}_j) + h^{up}(j, \tilde{Y}_j, \tilde{Z}_j, \tilde{Y}_j - Y_j^{up}, \tilde{Z}_j - E_j[\beta_{j+1}Y_{j+1}^{up}])]\Delta_j\}, \quad Y_n^{up} = S_n,$$

which satisfies $Y_i^{up} \geq Y_i$ by the comparison result in Proposition 2.3. Then, an application of Theorem 3.1, with Y_i replaced by Y_i^{up} yields $E_i[\Theta_i^{h^{up}}(\tilde{Y}, \tilde{Z})] \geq Y_i^{up}$. Hence,

$$Y_i \leq \underset{(\tilde{Y}_j, \tilde{Z}_j)_{j \geq i} \text{ adapted and integrable}}{\text{essinf}} E_i[\Theta_i^{h^{up}}(\tilde{Y}, \tilde{Z})].$$

It now suffices to show that

$$Y_j = \Theta_j^{h^{up}}(Y, E_j[\beta_{j+1}Y_{j+1}]) =: \Theta_j^{h^{up},*},$$

P -almost surely for every $j = i, \dots, n$. This is certainly true for $j = n$. Going backwards in time we obtain by induction

$$\begin{aligned} & \Theta_j^{h^{up},*} \\ = & \max\{S_j, Y_{j+1} - (Y_{j+1} - E_j[Y_{j+1}]) + f(j, Y_j, E_j[\beta_{j+1}Y_{j+1}])\Delta_j \\ & + h^{up}(j, Y_j, E_j[\beta_{j+1}Y_{j+1}]; Y_j - \Theta_j^{h^{up},*}, E_j[\beta_{j+1}Y_{j+1}] - \beta_{j+1}Y_{j+1} + \beta_{j+1}Y_{j+1} - E_j[\beta_{j+1}Y_{j+1}])\Delta_j\} \\ = & \max\{S_j, E_j[Y_{j+1}] + (f(j, Y_j, E_j[\beta_{j+1}Y_{j+1}]) + h^{up}(j, Y_j, E_j[\beta_{j+1}Y_{j+1}]; Y_j - \Theta_j^{h^{up},*}, 0))\Delta_j\} \end{aligned}$$

As $h^{up}(j, Y_j, E_j[\beta_{j+1}Y_{j+1}]; 0, 0) = 0$, we observe that Y_j also solves the above equation. Hence, by uniqueness (due to the Lipschitz assumption on h^{up}), we obtain $Y_j = \Theta_j^{h^{up},*}$. \square

A maximization problem with value process Y_i can be constructed analogously. We denote by h^{low} any mapping which satisfies the same properties as h^{up} but with condition b) replaced by

b') $h^{low}(i, \tilde{y}, \tilde{z}; y, z)$ is concave in (y, z) , $h^{low}(i, \tilde{y}, \tilde{z}; 0, 0) = 0$ for every $(\tilde{y}, \tilde{z}) \in \mathbb{R} \times \mathbb{R}^D$, and

$$h^{low}(i, \tilde{y}, \tilde{z}; \tilde{y} - y, \tilde{z} - z) \leq f(i, y, z) - f(i, \tilde{y}, \tilde{z})$$

for every $(\tilde{y}, \tilde{z}), (y, z) \in \mathbb{R} \times \mathbb{R}^D$.

The generic choice is now

$$h^{|low|}(i, \tilde{y}, \tilde{z}; y, z) = -\alpha_i^{(0)}|y| - \sum_{d=1}^D \alpha_i^{(d)}|z_d|.$$

Given h^{low} , a pair of adapted processes (\tilde{Y}, \tilde{Z}) and a stopping time $\tau \in \bar{\mathcal{S}}_0$ we define $\Theta_i^{h^{low}} = \Theta_i^{h^{low}}(\tilde{Y}, \tilde{Z}, \tau)$ via

$$\begin{aligned} \Theta_i^{h^{low}} &= \Theta_{i+1}^{h^{low}} - (\tilde{Y}_{i+1} - E_i[\tilde{Y}_{i+1}]) + f_i(\tilde{Y}_i, \tilde{Z}_i)\Delta_i \\ &\quad + h^{low}(i, \tilde{Y}_i, \tilde{Z}_i; \tilde{Y}_i - \Theta_i^{h^{low}}, \tilde{Z}_i - \beta_{i+1}\Theta_{i+1}^{h^{low}} + \beta_{i+1}\tilde{Y}_{i+1} - E_i[\beta_{i+1}\tilde{Y}_{i+1}])\Delta_i, \end{aligned} \quad (15)$$

for $i < \tau$ initiated at $\Theta_\tau^{h^{low}} = S_\tau$. Making use of Theorem 3.6 and the same arguments as in the previous theorem we obtain:

Theorem 4.2. *For every $i = 0, \dots, n$,*

$$Y_i = \operatorname{esssup}_{\tau \in \bar{\mathcal{S}}_i} \operatorname{esssup}_{(\tilde{Y}_j, \tilde{Z}_j)_{j \geq i} \text{ adapted and integrable}} E_i[\Theta_i^{h^{low}}(\tilde{Y}, \tilde{Z}, \tau)].$$

Moreover, a minimizing triplet is given by $(Y_j^*, Z_j^*, \tau^*) = (Y_j, E_j[\beta_{j+1}Y_{j+1}], \tau_j^*)$ which even satisfies the principle of pathwise optimality. (We recall that τ_j^* was defined in (4)).

Example 4.3. For the generic choices $h^{|up|}$ and $h^{|low|}$, we can apply Proposition 3.2 in order to make the recursion formulas in (14) and (15) explicit. They read

$$\begin{aligned} \Theta_i^{h^{|up|}} &= \max\{S_i, \sup_{r \in \{-\alpha_i^{(0)}, \alpha_i^{(0)}\}} \frac{1}{1+r\Delta_i} \left(\Theta_{i+1}^{h^{|up|}} - (\tilde{Y}_{i+1} - E_i[\tilde{Y}_{i+1}]) + f(i, \tilde{Y}_i, \tilde{Z}_i)\Delta_i \right. \\ &\quad \left. + \sum_{d=1}^D |\tilde{Z}_{d,i} - \beta_{d,i+1}\Theta_{i+1}^{h^{|up|}} + \beta_{d,i+1}\tilde{Y}_{i+1} - E_i[\beta_{d,i+1}\tilde{Y}_{i+1}]| \right) \Delta_i\}, \end{aligned}$$

and

$$\begin{aligned} \Theta_i^{h^{|low|}} &= \inf_{r \in \{-\alpha_i^{(0)}, \alpha_i^{(0)}\}} \frac{1}{1+r\Delta_i} \left(\Theta_{i+1}^{h^{|low|}} - (\tilde{Y}_{i+1} - E_i[\tilde{Y}_{i+1}]) + f(i, \tilde{Y}_i, \tilde{Z}_i)\Delta_i \right. \\ &\quad \left. - \sum_{d=1}^D \alpha_i^{(d)} |\tilde{Z}_{d,i} - \beta_{d,i+1}\Theta_{i+1}^{h^{|low|}} + \beta_{d,i+1}\tilde{Y}_{i+1} - E_i[\beta_{d,i+1}\tilde{Y}_{i+1}]| \right) \Delta_i. \end{aligned}$$

The main advantage of the corresponding upper and lower bounds is that they can be calculated generically without any extra information on f (such as the convex conjugates which were required in the section on convex generators). There is, however, a price to pay for this generic approach. Indeed, given the Lipschitz process $\alpha_i^{(d)}$, the choice $h^{|up|}$, $h^{|low|}$ can be shown to lead to the crudest upper and lower bounds among all admissible functions h^{up} , h^{low} , i.e.

$$E_i[\Theta_i^{h^{|up|}}(\tilde{Y}, \tilde{Z})] \geq E_i[\Theta_i^{h^{up}}(\tilde{Y}, \tilde{Z})]$$

for every pair (\tilde{Y}, \tilde{Z}) , and analogously for the lower bounds. In practice, the generic bounds may be too crude, when D is large and the approximation \tilde{Z}_i of $E_i[\beta_{i+1}Y_{i+1}]$ is not yet very good. In general we therefore recommend to choose the functions h^{up} and h^{low} in a way that $h^{up}(i, \tilde{Y}_i, \tilde{Z}_i; y, z)$ and $h^{low}(i, \tilde{Y}_i, \tilde{Z}_i; y, z)$ are close to zero in a neighborhood of zero in the (y, z) -coordinates, in which one expects the residuals $(\tilde{Y}_i - Y_i, \tilde{Z}_i - E_i[\beta_{i+1}Y_{i+1}])$ to be typically located.

$\delta \setminus n$	40		80		120		160	
0	71.6551 (0.0071)	71.8589 (0.0068)	71.6774 (0.0072)	71.8828 (0.0068)	71.6664 (0.0070)	71.8656 (0.0068)	71.6621 (0.0069)	71.8659 (0.0072)
$\frac{1}{3}$	74.1023 (0.0062)	74.2241 (0.0060)	74.1010 (0.0065)	74.2225 (0.0062)	74.1032 (0.0062)	74.2229 (0.0061)	74.1187 (0.0065)	74.2391 (0.0063)
$\frac{2}{3}$	76.3335 (0.0057)	76.3865 (0.0057)	76.3364 (0.0057)	76.3886 (0.0057)	76.3416 (0.0059)	76.3943 (0.0058)	76.3290 (0.0061)	76.3814 (0.0059)

Table 2: Upper and lower price bounds for different recovery rates and time discretizations. Standard deviations are in brackets.

4.2 Numerical examples

Once the functions h^{low} and h^{up} are chosen, an algorithm for computing confidence intervals for Y_0 based on Theorems 4.1 and 4.2 can be designed analogously to the primal-dual algorithm in Section 3.2 for the convex case.

We first illustrate the algorithm in the context of Example 2.1 (ii). For the underlying, we choose the same five-dimensional geometric Brownian motion as in Section 3.3 except that $T = 1$ and the drift and risk-free rate equal $R = 0.02$. The payoff of the (European) claim is given by $G_n(x) = \min_{d=1,\dots,D} x_d$. For the default risk function Q , we assume that there are three regimes, high risk, intermediate risk and low risk: There are thresholds $v^h < v^l$ and rates $\gamma^h > \gamma^l$ such that $Q(y) = \gamma^h$ for $y < v^h$ and $Q(y) = \gamma^l$ for $y > v^l$. Over $[v^h, v^l]$, Q interpolates linearly. The resulting function f is Lipschitz continuous but generally neither convex nor concave. The candidates for the Lipschitz constant $\alpha^{(0)}$ are the absolute values of the left and right derivatives of f in v^h and v^l . In the implementation, we stick to the generic choice

$$-h^{low}(i, \tilde{y}; y) = h^{up}(i, \tilde{y}; y) = \alpha^{(0)}|y|,$$

using that the nonlinearity is independent of the Z -part in this example. We choose

$$v^h = 54, \quad v^l = 90, \quad \gamma^h = 0.2, \quad \gamma^l = 0.02.$$

For the calculation of \tilde{y} , we use the Lemor-Gobet-Warin algorithm with two basis functions, 1 and $E[G_n(X_n)|X_i = x]$, and $\Lambda^{reg} = 100,000$. Moreover, $\Lambda^{out} = 4,000$, $\Lambda^{in} = 1,000$.

In the absence of default risk, the derivative value is given by 78.37. Table 2 displays upper and lower price bounds for different time discretizations and recovery rates δ . As expected, a smaller recovery rate leads to a smaller option value. The relative width of the confidence intervals is well below 0.5% in all cases. For the larger values of δ , the bounds are even tighter: Larger values of δ lead to less nonlinearity in the pricing problem and to smaller Lipschitz constants ($\alpha^{(0)} = 0.41, 0.27, 0.12$ for $\delta = 0, 1/3, 2/3$). Compared to the example of Section 3.3, the bounds are much less dependent on the time discretization. This is due to the fact, that no Z -part has to be approximated, as is the case for many BSDEs in the credit risk literature, see Crépey et al. (2013); Henry-Labordère (2012). To sum up, the generic approach is perfectly sufficient in this example.

We finally revisit the example of Section 3.3. For the input approximation we run the martingale basis algorithm with seven basis functions for Y and 1,000 regression paths as specified there. The confidence bounds for the European call spread option on the maximum of five Black-Scholes stocks are calculated with $\Lambda^{in} = \Lambda^{out} = 1,000$ paths based on the following choices

Algorithm \ n	40		80		120		160	
fully generic	13.3604 (0.0132)	14.1774 (0.0169)	12.7905 (0.0332)	14.7496 (0.0407)	12.0148 (0.0612)	15.8512 (0.0834)	10.7872 (0.1005)	17.5326 (0.1504)
semi-generic	13.7259 (0.0041)	13.8505 (0.0046)	13.6984 (0.0053)	13.8801 (0.0059)	13.6811 (0.0059)	13.9136 (0.0071)	13.6686 (0.0065)	13.9459 (0.0078)

Table 3: Upper and lower price bounds for different time discretizations under the fully generic and semi-generic algorithms. Standard deviations are in brackets.

of h^{low} and h^{up} . For the *fully generic* implementation we apply

$$-h^{low}(i, \tilde{y}, \tilde{z}; y, z) = h^{up}(i, \tilde{y}, \tilde{z}; y, z) = R^b|y| + \frac{\max\{R^b - \mu, |R^l - \mu|\}}{\sigma} \sum_{d=1}^5 |z_d|.$$

For the *semi-generic* implementation we choose

$$\begin{aligned} h^{low}(i, \tilde{y}, \tilde{z}; y, z) &= R^l y + \frac{\mu - R^l}{\sigma} \sum_{d=1}^5 z_d - (R^b - R^l) \left(y - \frac{1}{\sigma} \sum_{d=1}^5 z_d \right)_-, \\ h^{up}(i, \tilde{y}, \tilde{z}; y, z) &= R^l y + \frac{\mu - R^l}{\sigma} \sum_{d=1}^5 z_d + (R^b - R^l) \left(y - \frac{1}{\sigma} \sum_{d=1}^5 z_d \right)_+. \end{aligned}$$

This choice only partially exploits the structure of the generator. It can be applied to any generator which is a linear function of (y, z) plus a nondecreasing $(R^b - R^l)$ -Lipschitz continuous function of a linear combination of (y, z) . The specific form of the Lipschitz function is not used in this construction of h^{low} and h^{up} , but, of course, the coefficients for the linear combinations must be adjusted to the generator in the obvious way. For this semi-generic case the pathwise recursion formulas for $\Theta^{h^{up}}$ and $\Theta^{h^{low}}$ can be made explicit in time analogously to the generic case, which was discussed in Example 4.3.

Table 3 shows the resulting low-biased and high-biased estimates for the option price Y_0 as well as their empirical standard deviations. We observe that the generic bounds are not satisfactory in this example. The relative width of the 95% confidence intervals ranges from about 6.5% for $n = 40$ to more than 65% for $n = 160$ time steps. This can be explained by the fact that the approximation of $E_i[\beta_{i+1} Y_{i+1}]$ by \tilde{Z}_i (which is expressed in terms of just two basis functions) is not yet good enough. The quality of \tilde{Z} plays an all important role for the generic bounds due to the appearance of the terms $\sum_{d=1}^5 |z_d|$ in the definitions of h^{low} and h^{up} . In the semi-generic setting the expressions of the form $(y - \frac{1}{\sigma} \sum_{d=1}^5 z_d)_\pm$ in h^{up} and h^{low} are much more favorable concerning the approximation error of $E_i[\beta_{i+1} Y_{i+1}]$ by \tilde{Z}_i . Therefore, the semi-generic implementation yields much better 95% confidence intervals with a relative width of about 1% for $n = 40$ and still less than 2.5% for $n = 160$ time steps.

By and large, this example shows that the generic bounds may be too crude, if applied to good but not excellent approximations (\tilde{Y}, \tilde{Z}) , in particular when the z -variable of the generator is high-dimensional. Nonetheless very acceptable confidence intervals can still be obtained based on the same approximation (\tilde{Y}, \tilde{Z}) , if some information about the generator is incorporated in the choice of h^{up} and h^{low} .

A Continuous time analogues

In this appendix we consider BSDEs driven by a Brownian motion W of the form

$$Y_t = \xi + \int_t^T f(s, Y_s, Z_s) ds - \int_t^T Z_s^\top dW_s. \quad (16)$$

We assume that the pair (f, ξ) are standard parameters in the sense of El Karoui et al. (1997), p. 18, i.e. square-integrability conditions and a uniform Lipschitz condition on f are in force. Moreover, f is supposed to be convex in (y, z) .

Then, by Proposition 3.4 in El Karoui et al. (1997)

$$Y_t = \operatorname{esssup}_{(r, \rho) \in \mathcal{U}_t^2(f^\#)} E \left[\gamma_{t,T}(r, \rho) \xi - \int_t^T \gamma_{t,s}(r, \rho) f^\#(s, r_s, \rho_s) ds \middle| \mathcal{F}_t^W \right],$$

where $(\mathcal{F}_t^W)_{t \in [0, T]}$ is the augmented filtration generated by the driving Brownian motion,

$$\gamma_{t,s} = \exp \left\{ \int_t^s r_u du + \int_t^s \rho_u^\top dW_u \right\},$$

and the supremum runs over the set

$$\mathcal{U}_t^2(f^\#) := \left\{ (r_s, \rho_s)_{s \geq t} \text{ predictable; } \int_t^T E[|f^\#(s, r_s, \rho_s)|^2] ds < \infty \right\}.$$

This is the non-reflected continuous time analogue to the primal optimization problem in Theorem 3.4 in a Brownian environment.

The pathwise approach to the dual minimization problem in Theorem 3.1 requires the use of Malliavin calculus in continuous time. For the corresponding definitions and notations we refer to Nualart (2006). Given a stochastic process θ such that θ_t is Malliavin differentiable for a.e. $t \in [0, T]$, we denote by $D\theta$ the Malliavin derivative of θ . Notice that the field $(D_s \theta_t)_{s, t \in [0, T]^2}$ is only defined almost everywhere on $[0, T]^2$, and consequently the trace $D_t \theta_t$ of $D_s \theta_t$ is not well-defined. We shall therefore make use of the one-sided trace $(D^+ \theta)_t$, as introduced on p. 173 in Nualart (2006) for $p = 2$.

Now given a martingale M^0 such that $M_T^0 \in \mathbb{D}^{1,2}$, (i.e. the random variable M_T^0 is Malliavin differentiable with square-integrable Malliavin derivative), we say that a possibly non-adapted process θ is a M^0 -solution of

$$-d\theta_t = f(t, \theta_t, (D^+ \theta)_t) dt - dM_t^0, \quad \theta_T = \xi \quad (17)$$

if $E[\int_0^T |\theta_t|^2 dt] < \infty$, $(D^+ \theta)_t$ exists, $f(\cdot, \theta, (D^+ \theta)) \in \mathbb{L}^{1,2}$, and for every $t \in [0, T]$

$$\theta_t = \xi + \int_t^T f(s, \theta_s, (D^+ \theta)_s) ds - (M_T^0 - M_t^0).$$

Now suppose that θ is a M^0 -solution for some martingale M^0 such that $M_T^0 \in \mathbb{D}^{1,2}$. Define $\tilde{Y}_t = E[\theta_t | \mathcal{F}_t^W]$, $\tilde{Z}_t = E[(D^+ \theta)_t | \mathcal{F}_t^W]$, and

$$c_s = E[f(s, \theta_s, (D^+ \theta)_s) | \mathcal{F}_s^W] - f(s, E[\theta_s | \mathcal{F}_s^W], E[(D^+ \theta)_s | \mathcal{F}_s^W]).$$

Then,

$$\tilde{Y}_t + \int_0^t \left(f(s, \tilde{Y}_s, \tilde{Z}_s) + c_s \right) ds = E \left[\xi + \int_0^T E[f(s, \theta_s, (D^+\theta)_s) | \mathcal{F}_s^W] ds \middle| \mathcal{F}_t^W \right] =: \tilde{M}_t. \quad (18)$$

Assuming that $\xi \in \mathbb{D}^{1,2}$, we next note that

$$\tilde{Z}_t = E \left[D_t \xi + \int_t^T D_t f(s, \theta_s, (D^+\theta)_s) ds \middle| \mathcal{F}_t^W \right]. \quad (19)$$

Indeed, by the martingale representation theorem and Lemma 1.3.4 in Nualart (2006), there is an adapted process $u \in \mathbb{L}^{1,2}$ such that

$$M_t^0 = M_0^0 + \int_0^t u_s dW_s.$$

Then, by Proposition 1.3.8 and the same argument as in Proposition 3.1.1 in Nualart (2006),

$$(D^+\theta)_t = D_t \xi + \int_t^T D_t f(s, \theta_s, (D^+\theta)_s) ds - \int_t^T D_t u_s dW_s.$$

The last integral is a martingale increment by adaptedness and square-integrability of the integrand. Hence, taking conditional expectation yields (19). We are now in the position to link \tilde{Z} to the martingale \tilde{M} , which was defined in (18). By the Clark-Ocone formula (Nualart, 2006, Proposition 1.3.14), we obtain

$$\begin{aligned} \tilde{M}_t - \tilde{Y}_0 &= \int_0^t E \left[D_r \left(\xi + \int_0^T E[f(s, \theta_s, (D^+\theta)_s) | \mathcal{F}_s^W] ds \right) \middle| \mathcal{F}_r^W \right] dW_r \\ &= \int_0^t E \left[D_r \xi + \int_r^T E[D_r f(s, \theta_s, (D^+\theta)_s) | \mathcal{F}_s^W] ds \middle| \mathcal{F}_r^W \right] dW_r \\ &= \int_0^t E \left[D_r \xi + \int_r^T D_r f(s, \theta_s, (D^+\theta)_s) \middle| \mathcal{F}_r^W \right] dW_r \\ &= \int_0^t \tilde{Z}_r dW_r, \end{aligned}$$

where we used Proposition 1.2.8 from Nualart (2006) to interchange Malliavin derivative and conditional expectation, and (19). Since $\tilde{Y}_T = E[\theta_T | \mathcal{F}_T^W] = \xi$, we conclude, thanks to (18), that (\tilde{Y}, \tilde{Z}) solves the BSDE

$$\tilde{Y}_t = \xi + \int_t^T (f(s, \tilde{Y}_s, \tilde{Z}_s) + c_s) ds - \int_t^T \tilde{Z}_s dW_s.$$

By the convexity of f we observe that $c_s \geq 0$. Hence, by the comparison theorem (see El Karoui et al., 1997, Theorem 2.2), we end up with

$$E[\theta_t | \mathcal{F}_t^W] = \tilde{Y}_t \geq Y_t.$$

Finally, Proposition 5.3 in El Karoui et al. (1997) shows that the unique adapted solution (Y, Z) to BSDE (16) satisfies $Z_t = (D^+Y)_t$ under some technical conditions on f and ξ , which we assume from now on. In particular, Y is a M^0 -solution to (17) for $M^0 = \int_0^T Z_s dW_s$. Summarizing the above, we arrive at the following result:

Proposition A.1. *Suppose that the assumptions of Proposition 5.3 in El Karoui et al. (1997) on (f, ξ) are in force. Then,*

$$Y_t = \operatorname{ess\,inf}_{\theta} E[\theta_t | \mathcal{F}_t^W],$$

where the infimum runs over the set of those processes θ , which are M^0 -solutions of (17) for some martingale M^0 such that $M_T^0 \in \mathbb{D}^{1,2}$.

Comparing this result with the discrete time result in Theorem 3.1, we immediately observe a major difference: In continuous time only the choice of a one-dimensional martingale M^0 is required, while in discrete time one additionally needs to choose a D -dimensional martingale M . This phenomenon is easily explained. Notice first that, under at most technical conditions,

$$(D^+\theta)_t = \lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \int_t^{t+\epsilon} D_s \theta_{t+\epsilon} ds = \lim_{\epsilon \downarrow 0} \left(\frac{W_{t+\epsilon} - W_t}{\epsilon} \theta_{t+\epsilon} - \frac{W_{t+\epsilon} - W_t}{\epsilon} \diamond \theta_{t+\epsilon} \right),$$

where the diamond denotes the Wick product, see Theorem 6.8 in Di Nunno et al. (2009). The first term on the right hand side corresponds to the expression $\beta_{i+1} \theta_{i+1}^{up}$ in (5), when $\beta_{i+1}(t_{i+1} - t_i)$ equals the truncated Brownian increment over $[t_i, t_{i+1}]$. The second term on the right hand side has zero conditional expectation, because the Wick product interchanges with the conditional expectation, i.e.

$$E \left[\frac{W_{t+\epsilon} - W_t}{\epsilon} \diamond \theta_{t+\epsilon} \middle| \mathcal{F}_t^W \right] = E \left[\frac{W_{t+\epsilon} - W_t}{\epsilon} \middle| \mathcal{F}_t^W \right] \diamond E[\theta_{t+\epsilon} | \mathcal{F}_t^W] = 0,$$

see e.g. Lemma 6.20 in Di Nunno et al. (2009). As one cannot expect that the Wick product $\beta_{i+1} \diamond \theta_{i+1}$ can be computed in closed form, a generic term with zero conditional expectation, namely the martingale increment $M_{i+1} - M_i$, is subtracted in (5). Due to the convexity of f , subtracting this generic term with zero conditional expectation pushes the solution of the recursion (5) upwards.

References

- S. Alanko, M. Avellaneda. Reducing variance in the numerical solution of BSDEs. *C. R. Math. Acad. Sci. Paris* 351, 135–138, 2013.
- L. Andersen, M. Broadie. A primal-dual simulation algorithm for pricing multidimensional American options. *Management Sci.* 50, 1222–1234, 2004.
- V. Bally, G. Pagès. A quantization algorithm for solving multi-dimensional discrete-time optimal stopping problems. *Bernoulli* 9, 1003–1049, 2003.
- D. Belomestny, C. Bender, J. Schoenmakers. True upper bounds for Bermudan products via non-nested Monte Carlo. *Math. Finance* 19, 53–71, 2009.
- C. Bender, R. Denk. A forward scheme for backward SDEs. *Stochastic Process. Appl.* 117, 1793–1812, 2007.
- C. Bender, J. Steiner. Least-squares Monte Carlo for BSDEs. In: Carmona, R. A. et al. (eds.) *Numerical Methods in Finance*, 257–289, Springer, 2012.

- C. Bender, J. Steiner. A-posteriori estimates for backward SDEs. *SIAM/ASA J. Uncertainty Quantification* 1, 139–163, 2013.
- Y. Z. Bergman. Option pricing with differential interest rates. *Rev. Financ. Stud.* 8, 475–500, 1995.
- B. Bouchard, J.-F. Chassagneux. Discrete-time approximation for continuously and discretely reflected BSDEs. *Stochastic Process. Appl.* 118, 2269–2293, 2008.
- B. Bouchard, R. Elie. Discrete-time approximation of decoupled forward-backward SDE with jumps. *Stochastic Process. Appl.* 118, 53–75, 2008.
- B. Bouchard, N. Touzi. Discrete-time approximation and Monte Carlo simulation of backward stochastic differential equations. *Stochastic Process. Appl.* 111, 175–206, 2004.
- D. B. Brown, J. E. Smith, P. Sun. Information relaxations and duality in stochastic dynamic programs. *Oper. Res.* 58, 785–801, 2010.
- J.-F. Chassagneux, A. Richou. Numerical simulation of quadratic BSDEs. arXiv preprint 1307.5741, 2013.
- P. Cheridito, M. Stadje. BSΔEs and BSDEs with non-Lipschitz drivers: comparison, convergence and robustness. *Bernoulli* 19, 1047–1085, 2013.
- S. Crépey, R. Gerboud, Z. Grbac, N. Ngor. Counterparty risk and funding: the four wings of the TVA. *Int. J. Theor. Appl. Finance* 16, 1350006, 2013.
- D. Crisan, K. Manolarakis. Solving backward stochastic differential equations using the cubature method: Application to nonlinear pricing. *SIAM J. Financial Math.* 3, 534–571, 2012.
- V. V. Desai, V. F. Farias, C. C. Moallemi. Pathwise optimization for optimal stopping problems. *Management Sci.* 58, 2292–2308, 2012.
- D. Duffie, M. Schroder, C. Skidas. Recursive valuation of defaultable securities and the timing of resolution of uncertainty. *Ann. Appl. Probab.* 6, 1075–1090, 1996.
- N. El Karoui, S. Peng, M. C. Quenez. Backward stochastic differential equations in finance. *Math. Finance* 7, 1–71, 1997.
- A. Fahim, N. Touzi, X. Warin. A probabilistic numerical method for fully nonlinear parabolic PDEs. *Ann. Appl. Probab.* 21, 1322–1364, 2011.
- E. Gobet, C. Labart. Error expansion for the discretization of backward stochastic differential equations. *Stochastic Process. Appl.* 117, 803–829, 2007.
- E. Gobet, A. Makhlouf. L^2 -time regularity of BSDEs with irregular terminal functions. *Stochastic Process. Appl.* 120, 1105–1132, 2010.
- J. Guyon, P. Henry-Labordère. Uncertain volatility model: a Monte Carlo approach. SSRN preprint 1540043, 2010.
- M. Haugh, L. Kogan. Pricing American options: a duality approach. *Oper. Res.* 52, 258–270, 2004.

- P. Henry-Labordère. Counterparty risk valuation: A marked branching diffusion approach. SSRN preprint 1995503, 2012.
- H. Johnson. Options on the maximum or the minimum of Several Assets. *J. Financial Quant. Anal.* 22, 277–283, 1987.
- J.-P. Laurent, P. Amzelek, J. Bonnaud. An overview of the valuation of collateralized derivative contracts. Working Paper, Université Paris I, 2012.
- J.-P. Lemor, E. Gobet, X. Warin. Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations. *Bernoulli* 12, 889–916, 2006.
- J. Ma, J. Zhang. Representations and regularities for solutions to BSDEs with reflections. *Stochastic Process. Appl.* 115, 539 – 569, 2005.
- D. Nualart. *The Malliavin Calculus and Related Topics*, 2nd ed., Springer, 2006.
- G. Di Nunno, B. Øksendal, F. Proske. *Malliavin Calculus for Lévy Processes with Applications in Finance*, Springer, 2009.
- A. Pallavicini, D. Perini, D. Brigo. Funding, collateral and hedging: uncovering the mechanics and the subtleties of funding valuation adjustments. arXiv preprint 1210.3811, 2012.
- L.C.G. Rogers. Monte Carlo valuation of American options. *Math. Finance* 12, 271–286, 2002.
- J. Schoenmakers, J. Zhang, J. Huang. Optimal dual martingales, their analysis, and application to new algorithms for Bermudan products. *SIAM J. Financial Math.* 4, 86–116, 2013.
- J. Zhang. A numerical scheme for BSDEs. *Ann. Appl. Probab.* 14, 459–488, 2004.
- G. Zhang, M. Gunzburger, W. Zhao. A sparse-grid method for multi-dimensional backward stochastic differential equations. *J. Comput. Math.* 31, 221–248, 2013.