

Enhanced policy iteration for American options via scenario selection

Schoenmakers, John; Bender, Christian; Kolodko, Anastasia

Postprint / Postprint

Zeitschriftenartikel / journal article

Zur Verfügung gestellt in Kooperation mit / provided in cooperation with:

www.peerproject.eu

Empfohlene Zitierung / Suggested Citation:

Schoenmakers, J., Bender, C., & Kolodko, A. (2008). Enhanced policy iteration for American options via scenario selection. *Quantitative Finance*, 8(2), 135-146. <https://doi.org/10.1080/14697680701253013>

Nutzungsbedingungen:

Dieser Text wird unter dem "PEER Licence Agreement zur Verfügung" gestellt. Nähere Auskünfte zum PEER-Projekt finden Sie hier: <http://www.peerproject.eu>. Gewährt wird ein nicht exklusives, nicht übertragbares, persönliches und beschränktes Recht auf Nutzung dieses Dokuments. Dieses Dokument ist ausschließlich für den persönlichen, nicht-kommerziellen Gebrauch bestimmt. Auf sämtlichen Kopien dieses Dokuments müssen alle Urheberrechtshinweise und sonstigen Hinweise auf gesetzlichen Schutz beibehalten werden. Sie dürfen dieses Dokument nicht in irgendeiner Weise abändern, noch dürfen Sie dieses Dokument für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen.

Mit der Verwendung dieses Dokuments erkennen Sie die Nutzungsbedingungen an.

gesis
Leibniz-Institut
für Sozialwissenschaften

Terms of use:

This document is made available under the "PEER Licence Agreement". For more information regarding the PEER-project see: <http://www.peerproject.eu>. This document is solely intended for your personal, non-commercial use. All of the copies of this document must retain all copyright information and other information regarding legal protection. You are not allowed to alter this document in any way, to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public.

By using this particular document, you accept the above-stated conditions of use.

Mitglied der
Leibniz
Leibniz-Gemeinschaft



Enhanced policy iteration for American options via scenario selection

Journal:	<i>Quantitative Finance</i>
Manuscript ID:	RQUF-2006-0005.R2
Manuscript Category:	Research Paper
Date Submitted by the Author:	14-Dec-2006
Complete List of Authors:	Schoenmakers, John; Weierstrass Institute, Stochastic algorithms and nonparametric statistics Bender, Christian; Weierstrass Institute Kolodko, Anastasia; Weierstrass Institute
Keywords:	Monte Carlo Methods, American Style Derivative Securities, Pricing of Derivatives Securities, Optimal Policies
JEL Code:	G13 - Contingent Pricing Futures Pricing < G1 - General Financial Markets < G - Financial Economics
Note: The following files were submitted by the author for peer review, but cannot be converted to PDF. You must view these files (e.g. movies) online.	
BenderKolodkoSchoenmakers_Rev2.tex tab-bas-call.tex tab-bas-call-ns-cor.tex tab-np-ls.tex tab-np-ls-N1.tex tab-np-ls-N2.tex	

SCHOLARONE™
Manuscripts

Enhanced policy iteration for American options via scenario selection

Christian Bender¹, Anastasia Kolodko^{1,2}, and John Schoenmakers¹

December 22, 2006

Abstract

In Kolodko & Schoenmakers [9] and Bender & Schoenmakers [3] a policy iteration was introduced, which allows to achieve tight lower approximations of the price for early exercise options via a nested Monte-Carlo simulation in a Markovian setting. In this paper we enhance the algorithm by a scenario selection method. It is demonstrated by numerical examples that the scenario selection can significantly reduce the number of actually performed inner simulations, and thus can heavily speed up the method (up to factor 15 in some examples). Moreover, it is shown that the modified algorithm retains the desirable properties of the original one such as the monotone improvement property, termination after a finite number of iteration steps, and numerical stability.

Keywords: American options, Monte-Carlo simulation, optimal stopping, policy improvement.

2000 AMS subject classification: 60G40, 62L15, 91B28

1 Introduction

In recent years the pricing of American options on a high-dimensional system of underlyings via Monte Carlo has become an ever growing field of interest. While, in principle, the backward dynamic program provides a recursive representation of the (time-discretized) price process of an American option, it requires the evaluation of high order nestings of conditional expectations. Therefore Monte Carlo estimators for regression functions, which do not run into explosive cost when nested several times, have been proposed by several authors, see Longstaff & Schwartz [10], Tsitsiklis & Van Roy [15], Broadie & Glasserman [5], and Bouchard et al. [4].

An alternative to solving the backward dynamic program recursively are policy iterations for dynamic programming. The main advantage of policy iterations is, that they yield lower approximations of the price process for any

¹Weierstrass Institute for Applied Analysis and Stochastics, Mohrenstr. 39, D-10117 Berlin, Germany. [bender\[kolodko\]\[schoenma\]@wias-berlin.de](mailto:bender[kolodko][schoenma]@wias-berlin.de).

Supported by the DFG Research Center MATHEON ‘Mathematics for Key Technologies’ in Berlin.

²Institute of Computational Mathematics and Mathematical Geophysics, Russian Acad. Sci., Lavrentieva str. 6, 630090 Novosibirsk, Russia.

given order of nested conditional expectations, which are typically of increasing quality the higher the order. (The latter property is referred to as monotone improvement property.) In a Markovian setting, this methodology allows to apply the plain Monte Carlo estimator to evaluate the conditional expectations, at least for nestings of order one. As one can only obtain approximations corresponding to low order iterations this way, the quality of a single improvement step is of prime importance. A new policy improvement algorithm was developed in Kolodko & Schoenmakers [9] and Bender & Schoenmakers [3] which outperforms for instance the more classical Howard improvement (e.g. [13]). In principle, one could start the algorithm with a very simple input policy. For example, in case European options are given (quasi-)analytically, two iterations of the policy ‘exercise immediately’ can be obtained by a one degree nested Monte Carlo simulation, yielding surprisingly good results [9]. Given today’s computer power however, more than one degree of nesting is virtually impossible. This means that in practice only one step of the algorithm in [9] can be carried out. Therefore, the choice of the input stopping family is important. In this respect we found that, in particular for complex high-dimensional products, it turns out more effective to apply one improvement step to some input policy obtained by a popular standard method (for example Longstaff & Schwartz [10], Tsitsiklis & Van Roy [15], Piterbarg [12], or Andersen [1]), as exemplified in [2].

Since a one step version of the algorithm in [9] typically requires nested Monte Carlo simulation, it is still quite costly. In the present paper we enhance the policy improvement algorithm by a scenario selection method, while retaining the monotone improvement property of the original procedure. In this way the number of actually performed inner simulation can be reduced, which in some of our numerical examples speeds up the procedure by a factor 15. The basic idea is as follows: Suppose the holder of an American option has some pre-information, for example he knows good closed form approximations of the price for the corresponding European options. Such approximations are often available in the literature for practically relevant options. Given a trajectory of the underlying system, the investor rules out some time points, at which an optimal strategy cannot (or at least is very unlikely to) exercise, by the pre-information. Then the policy improvement is run only at the remaining time points. (Here, the set of remaining time points depends on the state of the underlying system. Hence we do not simply reduce to an other American option with a smaller set of exercise dates.)

We illustrate the enhancement procedure by numerical examples for American basket-call options on dividend paying stocks and an asset based cancelable swap. The latter example indicates that it may be more efficient to apply one enhanced improvement step to an input family obtained via the Longstaff-Schwartz algorithm with a low number of basis functions, than running the Longstaff-Schwartz algorithm alone with a huge number of basis functions. Hence, our algorithm can efficiently complement linear regression approaches such as [10, 15] in situations where the latter algorithms do not yield a fully satisfactory lower bound using standard bases. Moreover, the enhanced version of the policy improvement algorithm is compared to the original version showing that the scenario selection may drastically increase the efficiency.

After a short recap of American options and optimal stopping in discrete time (Section 2), we introduce the enhanced algorithm in Section 3.1 and verify the monotone improvement property. We also prove that the algorithm termi-

ates after a finite number of iteration steps. The latter result is of theoretical interest rather, since in practice only one or two iterations can be calculated. We also estimate the additional error when time points are ruled out which are only unlikely but not impossible to be in the range of an optimal policy. In Section 3.2 we provide a pseudo-code for the Monte Carlo implementation of one improvement step and prove the numerical stability in Section 3.3. The numerical examples are presented in Section 4. Some proofs are postponed to an Appendix.

2 Optimal stopping in discrete time

First we recall some facts about the optimal stopping problem in discrete time. Suppose $(Z(i) : i = 0, 1, \dots, k)$ is a nonnegative stochastic process in discrete time on a probability space (Ω, \mathcal{F}, P) adapted to some filtration $(\mathcal{F}_i : 0 \leq i \leq k)$ which satisfies

$$\sum_{i=1}^k E|Z(i)| < \infty.$$

We may think of the process Z as a cashflow, which an investor may exercise once. The investors' problem is to maximize his expected gain by choosing the optimal time for exercising. This problem is known as optimal stopping in discrete time. In the context of option pricing it is well known by the no arbitrage principle that the pricing of a (discrete time) American option is equivalent to the (discrete) optimal stopping problem, where Z is the discounted pay-off and P is a pricing measure corresponding to the discount numeraire.

To formalize the stopping problem we define \mathcal{S}_i as the set of \mathcal{F}_i stopping times taking values in $\{i, \dots, k\}$. The stopping problem can now be stated as follows: Find stopping times $\tau^*(i) \in \mathcal{S}_i$ such that for $0 \leq i \leq k$

$$E^{\mathcal{F}_i} [Z(\tau^*(i))] = \text{esssup}_{\tau \in \mathcal{S}_i} E^{\mathcal{F}_i} [Z(\tau)]. \quad (1)$$

The process on the right hand side is called the *Snell envelope* of Z and we denote it by $Y^*(i)$. We collect some facts, which can be found in Neveu [11] for example.

1. The Snell envelope Y^* of Z is the smallest supermartingale that dominates Z . It can be constructed recursively by backward dynamic programming:

$$\begin{aligned} Y^*(k) &= Z(k) \\ Y^*(i) &= \max\{Z(i), E^{\mathcal{F}_i}[Y^*(i+1)]\}. \end{aligned}$$

2. A family of optimal stopping times is given by

$$\hat{\tau}^*(i) = \inf\{j : i \leq j \leq k, Z(j) \geq E^{\mathcal{F}_j}[Y^*(j+1)]\}.$$

If several optimal stopping families exist, then the above family is the family of first optimal stopping times. In that case

$$\hat{\tau}^*(i) = \inf\{j : i \leq j \leq k, Z(j) > E^{\mathcal{F}_j}[Y^*(j+1)]\}$$

is the family of last optimal stopping times.

For the remainder of the paper we assume that

$$P(Z(k) > 0) > 0. \quad (2)$$

Clearly, this is no loss of generality: Let $\tilde{k} = \max\{i : 0 \leq i \leq k, P(Z(i) > 0) > 0\}$. Then exercising at $i > \tilde{k}$ cannot be optimal and hence the stopping problem is equivalent to the one with exercise set $\{0, \dots, \tilde{k}\}$.

3 Enhancing the policy iteration method

3.1 Definition and monotone improvement property

Suppose the buyer of the option chooses ad hoc a family of stopping times $(\tau(i) : 0 \leq i \leq k)$ taking values in the set $\{0, \dots, k\}$. We interpret $\tau(i)$ as the time, at which the buyer will exercise his option, provided he has not exercised prior to time i . This interpretation requires the following consistency condition:

Definition 3.1. A family of integer-valued stopping times $(\tau(i) : 0 \leq i \leq k)$ is said to be *consistent*, if

$$\begin{aligned} i &\leq \tau(i) \leq k, \quad \tau(k) \equiv k, \\ \tau(i) > i &\Rightarrow \tau(i) = \tau(i+1), \quad 0 \leq i < k. \end{aligned} \quad (3)$$

Indeed, suppose $\tau(i) > i$, i.e. according to our interpretation the investor has not exercised the first right prior to time $i+1$. Then he has not exercised the first right prior to time i , either. This means he will exercise the first right at times $\tau(i)$ and $\tau(i+1)$, which requires $\tau(i) = \tau(i+1)$. A typical example of a consistent stopping family can be obtained by comparison with the still-alive European options, i.e.

$$\tau(i) := \inf \left\{ j : i \leq j \leq k, Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(p)] \right\}. \quad (4)$$

In addition to the algorithm introduced in Kolodko & Schoenmakers [9] and further developed in Bender & Schoenmakers [3] we suppose that the investor has in some sense a-priori knowledge about an optimal exercise strategy. We consider a random set A , $A \subset \{0, \dots, k\}$, for which $\mathbf{1}_A(i)$ is \mathcal{F}_i -adapted, and $k \in A$ almost surely. Henceforth we will call such a set *an adapted random set*. Given some consistent stopping family τ we then consider a new stopping family by

$$\tilde{\tau}(i) := \inf \left\{ j : i \leq j \leq k, (Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(\tau(p))]) \wedge (j \in A) \right\}, \quad (5)$$

where \wedge denotes the logical ‘and’. Note that the stopping family $\tilde{\tau}$ is consistent. In particular $\tilde{\tau}(k) = k$, since $\max \emptyset = -\infty$ and $k \in A$. Moreover, by the definition of $\tilde{\tau}$, we have for all $0 \leq i \leq k$,

$$\tilde{\tau}(i) \in A. \quad (6)$$

In (5) the investor exploits his ‘a-priori knowledge’ by not exercising outside the set $A(\omega)$. If there exists some optimal stopping family τ^* such that

$$\tau^*(i) \in A, \quad i = 0, \dots, k, \quad P - a.s. \quad (7)$$

we call A an *a-priori set*. For instance, given any \mathcal{F}_i -adapted lower bound $L(i)$ of the Snell envelope $Y^*(i)$,

$$A(\omega) = \{i : 0 \leq i \leq k, Z(i, \omega) \geq L(i, \omega)\} \quad (8)$$

is an a-priori set. For an a-priori set A , (5) means that, due to the new family $\tilde{\tau}$, the investor will not exercise at a date j which is either suboptimal or for optimality not necessary to exercise since $\tau^*(j) \in A$.

We call $\tilde{\tau}$ a *one-step improvement* of τ for the following reason: Denote by $Y(i; \tau)$ the value process corresponding to the stopping family τ , namely

$$Y(i; \tau) = E^{\mathcal{F}_i} [Z(\tau(i))]. \quad (9)$$

Then due to the next theorem which is in fact a generalization of Theorem 3.5 in Bender & Schoenmakers [3], the one-step improvement yields a higher value than the given family, provided the input family $\tau(i)$ takes values in A .

Theorem 3.2. *Suppose A is an adapted random set, τ is a consistent input stopping family such that $\tau(i) \in A$ a.s. for all $0 \leq i \leq k$. Consider*

$$\hat{\tau}(i) := \inf \left\{ j : i \leq j \leq k, (Z(j) > \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(\tau(p))]) \wedge (j \in A) \right\}, \quad (10)$$

and let $\bar{\tau}$ be a consistent stopping family such that

$$\bar{\tau}(i) \leq \tau(i) \leq \hat{\tau}(i), \quad 0 \leq i \leq k. \quad (11)$$

Then,

$$Y(i; \bar{\tau}) \geq Y(i; \hat{\tau}) \geq \max_{i \leq p \leq k} E^{\mathcal{F}_i} [Z(\tau(p))] \geq Y(i; \tau), \quad 0 \leq i \leq k. \quad (12)$$

Moreover, $Y(i; \bar{\tau}) \geq Z(i)$ on $\{i \in A\}$.

Remark 3.1. (i) In the case of the trivial adapted random set $A \equiv \{1, \dots, k\}$, i.e. without scenario selection, the above Theorem is proved in [3].

(ii) It is interesting to note, that $\bar{\tau}(i)$ need not take values in A .

Proof. Define $Z_A(i) := \mathbf{1}_A(i)Z(i)$. Since $\tau(i) \in A$, we have by (2)

$$\max_{i \leq p \leq k} E^{\mathcal{F}_i} [Z_A(\tau(p))] = \max_{i \leq p \leq k} E^{\mathcal{F}_i} [Z(\tau(p))] > 0, \quad 0 \leq i \leq k.$$

Consequently,

$$\bar{\tau}(i) = \inf \left\{ j : i \leq j \leq k, Z_A(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z_A(\tau(p))] \right\}, \quad (13)$$

$$\hat{\tau}(i) = \inf \left\{ j : i \leq j \leq k, Z_A(j) > \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z_A(\tau(p))] \right\}. \quad (14)$$

We can now apply Remark 3.1, (i), to the cashflow Z_A and obtain

$$E^{\mathcal{F}_i} [Z_A(\bar{\tau}(i))] \geq E^{\mathcal{F}_i} [Z_A(\hat{\tau}(i))] \geq \max \left\{ Z_A(i), \max_{i \leq p \leq k} E^{\mathcal{F}_i} [Z_A(\tau(p))] \right\}.$$

So,

$$Y(i; \bar{\tau}) = E^{\mathcal{F}_i} [Z(\bar{\tau}(i))] \geq E^{\mathcal{F}_i} [Z_A(\bar{\tau}(i))] \geq E^{\mathcal{F}_i} [Z_A(\tilde{\tau}(i))] = Y(i; \tilde{\tau})$$

by (6), and

$$Y(i; \tilde{\tau}) \geq \max \left\{ Z_A(i), \max_{i \leq p \leq k} E^{\mathcal{F}_i} [Z(\tau(p))] \right\}$$

since $\tau(i) \in A$. \square

The following example shows that assumption $\tau(i) \in A$ cannot be dispensed with in Theorem 3.2.

Example 3.3. Suppose ξ is a binary trial with $P(\xi = 1) = P(\xi = -1)$. Define the process Z by $Z(0) = 9/4$, $Z(1) = Z(3) = 2$ and $Z(2) = 2 + \xi$. The filtration \mathcal{F}_i is assumed to be generated by Z . Then $\sigma = 5/2 - \xi/2$ is a stopping time which yields an expected payoff $E[Z(\sigma)] = 5/2$. Consequently, immediate exercise at $t = 0$ cannot be optimal. With this knowledge we define an a-priori set $A(\omega) \equiv \{1, 2, 3\}$. We want to improve upon the trivial starting family $\tau(i) = i$, which obviously violates the condition $\tau(0) \in A$. We define

$$\tilde{\tau}(0) = \inf \{j : 0 \leq j \leq 3, (Z(j) \geq \max_{p=1,2,3} E[Z(p)]) \wedge (j \in A)\}.$$

A simple calculation gives $\tilde{\tau}(0) = 1$ and hence $E[Z(\tilde{\tau}(0))] = 2 < 9/4 = E[Z(\tau(0))]$. Hence, $\tilde{\tau}(0)$ does not improve upon $\tau(0)$.

It is natural to iterate the policy improvement (5). Suppose A is an adapted random set, and τ_0 is some consistent stopping family satisfying $\tau_0(i) \in A$ for all $0 \leq i \leq k$. Define, recursively,

$$\begin{aligned} \tau_m &= \tilde{\tau}_{m-1}, \\ Y_m(i) &= Y(i; \tau_m), \quad m = 1, 2, \dots \end{aligned}$$

By Theorem 3.2, Y_m is an increasing sequence. Taking (13) and Proposition 4.4 in Kolodko & Schoenmakers [9] into account, we observe that the algorithm terminates after at most k steps.

Proposition 3.4. Suppose $m \geq k - i$. Then,

$$\tau_m(i) = \tilde{\tau}_A^*(i), \quad Y_m(i) = Y_A^*(i),$$

where $\tilde{\tau}_A^*$ and Y_A^* denote the first optimal stopping family and the Snell envelope for the stopping problem with cashflow $Z_A(i) = \mathbf{1}_A(i)Z(i)$. In particular, it follows that

$$Y_m(i) = Y^*(i)$$

for $m \geq k - i$, if A is an a-priori set.

If A is an a-priori set, the proposition states that the policy improvement algorithm terminates at the Snell envelope as fast as backward dynamic programming does. Most importantly, in every iteration step we obtain increased lower approximations of the Snell envelope, simultaneously at all exercise dates. If A is only an adapted random set, but not an a-priori set, the algorithm terminates at the Snell envelope Y_A^* of the cashflow $Z_A(i) = \mathbf{1}_A(i)Z(i)$ and not

at Y^* . As we will demonstrate by numerical examples in Section 4, it can be numerically more efficient to choose an adapted random set which contains the image of an optimal stopping family with only high probability. The following theorem estimates the difference between the two Snell envelopes in such a situation. The proof is postponed to the Appendix.

Theorem 3.5. *Let A be an adapted random set containing k a.s., and suppose that for some $q > 1$, $E[|Z(i)|^q] < \infty$ for all $0 \leq i \leq k$. Then for every consistent stopping family τ^* which is optimal for the cashflow Z the following estimate holds:*

$$E[Y^*(i) - Y_A^*(i)] \leq K_{q,i} P(\{\tau^*(i) \notin A\})^{1-1/q},$$

where

$$K_{q,i} = (k-i)^{1/q} \max_{i \leq j \leq k-1} (E[|Z(j)1_{\{0,\dots,k\} \setminus A}(j)|^q])^{1/q}.$$

In the case $A(\omega) \equiv \{1, \dots, k\}$ the policy iteration presented in this subsection coincides with the one suggested in [9]. As will be explained in Section 3.2 and exemplified in Section 4, an appropriate choice of A can significantly reduce the computational cost for a Monte Carlo implementation of one improvement step. In this respect the following corollary, which follows directly from Proposition A.1, is interesting.

Corollary 3.6. *Suppose τ is a consistent stopping family and $A_1 \subset A_2$ are adapted random sets. Define*

$$\begin{aligned} \tilde{\tau}(i) &:= \inf \left\{ j : i \leq j \leq k, (Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(\tau(p))]) \wedge (j \in A_1) \right\}, \\ \tilde{\sigma}(i) &:= \inf \left\{ j : i \leq j \leq k, (Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(\tau(p))]) \wedge (j \in A_2) \right\}. \end{aligned}$$

Then obviously $\tilde{\tau}$ and $\tilde{\sigma}$ are consistent and $\tilde{\tau}(i) \geq \tilde{\sigma}(i)$, $0 \leq i \leq k$. So by Proposition A.1 and Jensen's inequality, we have

$$\begin{aligned} (Y(i; \tilde{\tau}) - Y(i; \tilde{\sigma}))_- &\leq \sum_{j=i}^{k-1} E^{\mathcal{F}_i} [\mathbf{1}_{\{\tilde{\tau}(i) > j\}} \mathbf{1}_{\{\tilde{\sigma}(i) = j\}} (Y(j; \tilde{\tau}) - Z(j))_-], \quad (15) \\ (Y(i; \tilde{\tau}) - Y(i; \tilde{\sigma}))_+ &\leq \sum_{j=i}^{k-1} E^{\mathcal{F}_i} [\mathbf{1}_{\{\tilde{\tau}(i) > j\}} \mathbf{1}_{\{\tilde{\sigma}(i) = j\}} (Y(j; \tilde{\tau}) - Z(j))_+], \end{aligned}$$

where for any real x , $x =: x_+ - x_-$, with minimal $x_{\pm} \geq 0$.

As a special case we may compare one step of the plain version of the algorithm, $\tilde{\sigma}$ with $A_2(\omega) = \{0, \dots, k\}$, with a modified version (5) due to a non-trivial adapted random set A_1 containing k a.s. Obviously, constructing $\tilde{\tau}$ is generally cheaper than constructing $\tilde{\sigma}$, and the quality loss with respect to $\tilde{\sigma}$, due to $\tilde{\tau}$ may be estimated by (15). In fact (15) means that $\tilde{\tau}$ may be worse than $\tilde{\sigma}$ only if $Y(i; \tilde{\tau})$ can be below the cashflow at a time where $\tilde{\sigma}$ says 'exercise' but $\tilde{\tau}$ refuses to do so.

3.2 On the implementation

We now give some comments on the practical Monte Carlo implementation of an improvement step. Henceforth, we suppose that the cashflow Z is of the form $Z(i) = f(i, X(i))$ where $f(i, x)$ is a deterministic function and $(X(i), \mathcal{F}_i)$ is a – possibly high-dimensional – Markovian chain. Note that one improvement step of an initial lower bound $Y(0; \tau_0)$ requires a nested Monte Carlo simulation provided there are no closed form expressions for the conditional expectations in (5). The introduction of an adapted random set can significantly reduce the number of actually performed inner simulations and consequently increases the efficiency of the method. We suggest to implement an improvement step as follows.

Step 1: Choose an adapted random set A such that $(j \in A)$ can be checked in closed form given the state $X(j)$. This means, there are Borel sets $B_j \subset \mathbb{R}$, $B_k = \mathbb{R}$, which are explicitly known to the investor, such that $(j \in A)$ if and only if $X(j) \in B_j$. For instance good closed form approximations of the price processes of still alive Europeans are often available for practically relevant products. In such situation let $L(i)$ be a closed form approximation of their maximum, $\max_{i+1 \leq p \leq k} E^{\mathcal{F}_i}[Z(p)]$. Typically $L(i) = g(i, X(i))$ for some explicitly known deterministic function $g(i, x)$. Define

$$\begin{aligned} A(\omega) &= \{i : 0 \leq i \leq k, Z(i, \omega) \geq L(i, \omega)\} \\ &= \{i : 0 \leq i \leq k, f(i, X(i, \omega)) \geq g(i, X(i, \omega))\}. \end{aligned} \quad (16)$$

Obviously, A is in closed form in the above sense. One can just define $B_j = (f(j, \cdot) - g(j, \cdot))^{-1}([0, \infty))$. Clearly, A is an a-priori set, if L is a lower approximation of the maximum of still alive Europeans, but only an adapted random set in general. An alternative construction of adapted random sets is given in Remark 3.3 below.

Step 2: Choose an initial stopping family τ_0 such that $\tau_0(i) \in A$ for all $0 \leq i \leq k$. A natural choice is

$$\tau_0(i) = \inf\{j : i \leq j \leq k, j \in A\}.$$

Notice that, by Step 1, $\{\tau_0(i) = i\} = \{X(i) \in B_i\}$ belongs to the σ -field generated by $X(i)$. This guarantees that conditional expectations of the form $E^{\mathcal{F}_i}[Z(\tau_0(j))]$ can be replaced by regressions $E^{X(i)}[Z(\tau_0(j))]$ in the present Markovian setting.

Step 3: Construct a lower bound $Y(0; \tau_1)$ due to the improved policy $\tau_1 := \tilde{\tau}_0$ using the following pseudo-code.

Simulate M trajectories $X^{(m)}$, $m = 1, \dots, M$, starting at $X(0)$;

Along each trajectory $X^{(m)}$ we compute $\eta^{(m)} \approx \tau_1^{(m)}(0)$ as follows:

$i := 0$;

A: Search the first exercise date $\eta \geq i$ such that η belongs to A , which formally means $X^{(m)}(\eta) \in B_\eta$, since A is in closed form.

If $\eta = k$ (i.e. $\tau_1^{(m)}(0) = k$) then set $\eta^{(m)} := k$, else:

Consider η as a candidate for $\tau_1^{(m)}(0)$.

To decide whether $\eta \approx \tau_1^{(m)}(0)$ or not we do the following:

Simulate M_1 trajectories $(X^{(m,p)}(q), q = \eta, \dots, k), p = 1, \dots, M_1$, under the conditional measure $P^{X^{(m)}(\eta)}$ (hence $X^{(m,p)}(\eta) = X^{(m)}(\eta)$);

Along each trajectory (m, p) search all exercise dates $\geq \eta$ where the policy τ_0 says ‘exercise’. From these dates we can detect easily (an approximation of) the family $(\tau_0^{(m,p)}(q), q \geq \eta)$ along the path (m, p) ;

Then, for $q = \eta, \dots, k$ compute

$$\begin{aligned} Dummy[q] &:= \frac{1}{M_1} \sum_{p=1}^{M_1} f(\tau_0^{(m,p)}(q), X^{(m,p)}(\tau_0^{(m,p)}(q))) \\ &\approx E^{X^{(m)}(\eta)} [Z(\tau_0(q))]; \end{aligned}$$

Next determine

$$Max_Dummy := \max_{\eta \leq q \leq k} Dummy[q] \approx \max_{\eta \leq q \leq k} E^{X^{(m)}(\eta)} [Z(\tau_0(q))];$$

Check whether $f(\eta, X^{(m)}(\eta)) \geq Max_Dummy$:

If yes, set $\eta^{(m)} := \eta \approx \tau_1^{(m)}(0)$;

If no, do $i := \eta + 1$ and go to (A);

We so end up with $\eta^{(m)} \approx \tau_1^{(m)}(0)$;

Finally compute $\frac{1}{M} \sum_{m=1}^M f(\eta^{(m)}, X^{(m)}(\eta^{(m)})) \approx E[Z(\tau_1(0))] = Y(0; \tau_1)$.

Step 4: Given a consistent stopping family τ_0 as in Step 2, $Y(0; \tau_0)$ is a lower bound of $Y^*(0)$. We recommend to construct an upper bound from this lower bound by the duality method developed by Rogers [14] and Haugh & Kogan [6]. Define,

$$Y_{up}(0; \tau) = E \left[\max_{0 \leq j \leq k} (Z(j) - M(j)) \right], \quad (17)$$

where $M(0) = 0$ and, for $1 \leq i \leq k$,

$$M(i) = \sum_{p=1}^i (Y(p; \tau) - E^{\mathcal{F}_{p-1}} [Y(p; \tau)]) .$$

Approximation of this upper bound due to $Y(\cdot; \tau_0)$ by Monte Carlo also requires nested simulation. It is, thus, roughly as expensive as the improvement of τ_0 described in Step 3. Note, however, that calculating $Y_{up}(0; \tau_1)$ would require another layer of simulations within simulations which is infeasible with today’s computer power. For a detailed treatment of efficient computation of dual upper bounds see for example Kolodko & Schoenmakers [8]. A multiplicative analogon of the duality method is due to Jamshidian [7].

Remark 3.2. The suggested implementation primarily is a procedure for computing the option price. It does not yield an improved exercise boundary in

terms of functions of the underlying process at different exercise dates. However, at every exercise date the improved exercise decision, ‘exercise’ or ‘do not exercise’, can be obtained via a standard (non-nested) Monte Carlo simulation, hence relatively fast.

Remark 3.3. The algorithm described above can be generically combined with other methods, such as [1], [10], and [15]. Suppose an approximative exercise boundary $h(i, X(i))$ is at hand which was pre-computed by another method. Then one can construct adapted random sets by shifting this exercise boundary

$$A_\gamma = \{i : 0 \leq i \leq k, f(i, X(i)) \geq h(i, X(i)) - \gamma\}, \quad \gamma \geq 0. \quad (18)$$

We can then define the input stopping time based on the approximative exercise boundary

$$\tau_0(i) = \inf\{j : i \leq j \leq k, j \in A_0\},$$

which clearly satisfies $\tau_0(i) \in A_\gamma$ for all $\gamma \geq 0$, since the sets A_γ are increasing in γ . The choice of the parameter γ for the adapted random set can be seen as a trade off between speed and accuracy. The closer it is to 0 the bigger the gain due to the scenario selection. However, when γ is away from 0, then the adapted random set A_γ becomes ‘closer’ to an a priori set.

3.3 Stability

As described in the previous section, for practical implementation one typically has to approximate the conditional expectations in the exercise criterion. We now extend a stability result from Bender & Schoenmakers [3] to the case of a nontrivial adapted random set. Let A be an adapted random set and τ be a consistent stopping family which satisfies $\tau(i) \in A$ for all $0 \leq i \leq k$. Further suppose $\epsilon^{(N)}(i)$ is a sequence of \mathcal{F}_i -adapted processes such that

$$\lim_{N \rightarrow \infty} \epsilon^{(N)}(i) = 0, \quad P - a.s.$$

A perturbed version of the one step improvement is then defined by

$$\begin{aligned} \tilde{\tau}^{(N)}(i) := \inf \Big\{ j : i \leq j \leq k, (Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(\tau(p))] + \epsilon^{(N)}(j)) \\ \wedge (j \in A) \Big\}. \end{aligned} \quad (19)$$

The sequence $\epsilon^{(N)}$ accounts for the errors when approximating the conditional expectation. We may and will assume that $\epsilon^{(N)}(k) = 0$, since no conditional expectation is to be evaluated at $j = k$. In accordance with the previous section we suppose that the criterion $j \in A$ can be checked in closed form. We first recall that even with a trivial a-priori set we can neither expect

$$\tilde{\tau}^{(N)}(i) \rightarrow \tilde{\tau}(i) \quad \text{in probability}$$

nor

$$Y(0; \tilde{\tau}^{(N)}) \rightarrow Y(0; \tilde{\tau})$$

in general. For corresponding counterexamples we refer to [3]. However, we can generalize a stability result from [3] where the error is measured in terms of the shortfall instead of the absolute value. As emphasized in [3], preventing shortfall (viz. change to the worse) is the relevant criterion to look at since our goal is improvement.

Theorem 3.7. For all $0 \leq i \leq k$,

$$\lim_{N \rightarrow \infty} \left(Y(i; \tilde{\tau}^{(N)}) - Y(i; \tilde{\tau}) \right)_- = 0,$$

where the limit is P -almost surely and in $L^1(P)$.

Proof. Since $\epsilon^{(N)}(k) = 0$ we may write as in (13),

$$\tilde{\tau}^{(N)}(i) = \inf \left\{ j : i \leq j \leq k, Z_A(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z_A(\tau(p))] + \epsilon^{(N)}(j) \right\},$$

where $Z_A(i) = \mathbf{1}_A(i)Z(i)$. Hence, the claim follows from the corresponding stability result for the plain version of the algorithm (without scenario selection) applied to the cashflow Z_A . Such stability result is proved in [3]. \square

Remark 3.4. Theorem 3.7 provides stability of one improvement step. More generally, one can prove that the shortfall of the expected gain corresponding to m perturbed steps of the algorithm below the expected gain corresponding to m theoretical steps converges to zero. In the case of the trivial a-priori set $A = \{0, \dots, k\}$ this statement is made precise and proved in [3], Section 4.2. This result carries over to the case of a general adapted random set. Here it is crucial that the criterion ($j \in A$) involves no approximation so that it is guaranteed that e.g. $\tilde{\tau}^{(N)}(i) \in A$ for all $0 \leq i \leq k$.

4 Numerical examples

We now test our algorithm with two examples. The first example is a Bermudan call on a basket of dividend paying assets. In this comparably simple example we construct an a priori set based on a lower approximation of European option prices to illustrate the efficiency gain stemming from the scenario selection. The second example is a complex structured asset based cancelable swap. It turns out that the duality gap due to the Longstaff-Schwartz algorithm [10] is not fully satisfactory for this product. We hence utilize the approximative exercise boundary of the Longstaff-Schwartz algorithm to run one enhanced improvement step as described in Remark 3.3. In particular we demonstrate the trade-off between accuracy and computational time by shifting the approximative exercise boundary for the construction of different adapted random sets.

Both examples are treated in the context of a multi-dimensional Black-Scholes model, i.e. we consider an underlying system of n assets where each asset is governed under the risk-neutral measure by the following SDE:

$$dS_l(t) = (r - \delta)S_l(t)dt + \sigma_l S_l(t)dW_l(t), \quad 1 \leq l \leq n. \quad (20)$$

Here $W_1(t), \dots, W_n(t)$ are correlated n -dimensional Brownian motions with time independent correlations $\rho_{lm} := E[W_l(t)W_m(t)]$, $1 \leq l, m \leq n$. The continuously compounded interest rate r and δ , a dividend rate, are assumed to be constant.

4.1 Basket-call

The price of a Bermudan basket-call option on n assets is given by (1) with

$$Z(i) = e^{-rT_i} \left(\frac{S_1(T_i) + \cdots + S_n(T_i)}{n} - K \right)^+, \quad 0 \leq i \leq k.$$

For our experiments we assume, that this option can be exercised at 9 dates, $0 < T_1 < \cdots < T_9 = 3$, which are equidistant on $[0, 3]$. We take the following parameter values,

$$n = 5, \quad r = 0.05, \quad \delta = 0.1, \quad S_1(0) = \cdots = S_5(0) = S_0, \quad K = 100.$$

Further, we consider two different asset systems:

Case 1: uncorrelated assets with identical volatilities,

$$\sigma_l = 0.2, \quad \rho_{lm} = \delta_{lm}, \quad 1 \leq l, m \leq 5.$$

Case 2: assets with dispersed volatilities and exponential correlation structure,

$$\sigma_l = 0.16 + 2(l-1), \quad \rho_{lm} = \exp(-0.4|l-m|), \quad 1 \leq l, m \leq 5.$$

In this example we make use of the a-priori set A given by (16) with

$$L(j) := \max_{j+1 \leq p \leq k} e^{-rT_p} E^{\mathcal{F}_j} [(S_1(T_p) \cdots S_5(T_p))^{1/5} - K]^+, \quad 0 \leq j \leq k, \quad (21)$$

which clearly is a lower approximation of the Snell envelope since

$$L(j) \leq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(p)] \leq Y^*(j), \quad 0 \leq j \leq k.$$

Notice that the process $G_n(t) := (S_1(t) \cdots S_n(t))^{1/n}$ has a log-normal distribution and can be represented as

$$G_n(t) = e^{((r-\delta) - \frac{1}{2n} \sum_{l=1}^n \sigma_l^2)t + \frac{1}{n} \sum_{l=1}^n \sigma_l W_l(t)}.$$

Then

$$\widetilde{W}(t) := \frac{\sum_{l=1}^n \sigma_l W_l(t)}{\sqrt{\sum_{l=1}^n \sigma_l^2 + 2 \sum_{l=1}^n \sum_{m=l+1}^n \sigma_l \sigma_m \rho_{lm}}}$$

is a standard Brownian motion and we thus have

$$G_n(t) = e^{(r - \frac{1}{2} \tilde{\sigma}^2)t + \tilde{\sigma} \widetilde{W}(t)} \cdot e^{\varphi(t)}$$

with

$$\tilde{\sigma} := \frac{1}{n} \sqrt{\sum_{l=1}^n \sigma_l^2 + 2 \sum_{l=1}^n \sum_{m=l+1}^n \sigma_l \sigma_m \rho_{lm}}$$

and

$$\varphi(t) := \frac{1}{2} t (\tilde{\sigma}^2 - \frac{1}{n} \sum_{l=1}^n \sigma_l^2) = \frac{t}{n^2} \left(\sum_{l=1}^n \sum_{m=l+1}^n \sigma_l \sigma_m \rho_{lm} - \frac{n-1}{2} \sum_{l=1}^n \sigma_l^2 \right).$$

So the right hand side of (21) can be given in closed form by the well-known Black-Scholes formula (*BS*),

$$\begin{aligned} e^{-rT_p} E^{\mathcal{F}_j} [(G_5(T_p) - K)^+] &= \\ e^{-rT_p} E^{\mathcal{F}_j} [(e^{((r-\delta)-\frac{1}{2}\tilde{\sigma}^2)T_p + \tilde{\sigma}\tilde{W}(T_p)} - Ke^{-\varphi(T_p)})^+ e^{\varphi(T_p)}] &= \\ e^{\varphi(T_p)-rT_j} BS(G_5(T_j), r, \delta, \tilde{\sigma}, Ke^{-\varphi(T_p)}, T_p - T_j). \end{aligned}$$

Based on the a-priori set A we consider the simple initial stopping family $\tau(i) = \inf\{j : i \leq j \leq k, j \in A\}$. We construct the lower bound $Y(0; \tau) = E[Z(\tau(0))]$ of the zero Bermudan price, the improved lower bound $Y(0; \tilde{\tau}) = EZ[(\tilde{\tau}(0))]$ with $\tilde{\tau}$ given by (5), and the dual upper bound $Y_{up}(0; \tau)$ given by (17). For comparison, we also compute the one-step improvement of the lower bound without scenario selection, $Y(0; \check{\tau}) = E[Z(\check{\tau})]$, where

$$\check{\tau}(i) := \inf \left\{ j : i \leq j \leq k, Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(\tau(p))] \right\}.$$

In Table 1, we present the results for the uncorrelated assets with identical volatilities (case 1), while we deal with the positively correlated assets having dispersed volatilities (case 2) in Table 2.

We first simulate $Y(0; \tau)$ by 10^7 Monte Carlo trajectories. To reduce the variance we then compute $Y(0; \tilde{\tau})$ and $Y(0; \check{\tau})$ via the representations

$$\begin{aligned} Y(0; \tilde{\tau}) &= Y(0; \tau) + E[Z(\tilde{\tau}(0)) - Z(\tau(0))] \quad \text{and} \\ Y(0; \check{\tau}) &= Y(0; \tau) + E[Z(\check{\tau}(0)) - Z(\tau(0))] \end{aligned} \quad (22)$$

The second term in the respective representation is approximated using $2 \cdot 10^5$ outer and 1000 inner Monte Carlo trajectories for Table 1 (case 1) and $7 \cdot 10^4$ outer and 500 inner Monte Carlo trajectories for Table 2 (case 2). Further, $Y_{up}(0; \tau) - Y(0; \tau)$ in the Tables 1-2 are simulated correspondingly by 20 000 and 5000 outer with 1000 inner trajectories. We can see that, although the initial stopping family gives a rather crude lower bound (the gap between $Y(0; \tau)$ and its dual upper bound $Y_{up}(0; \tau)$ is 8%-17% relative to the value), the improvements $Y(0; \tilde{\tau})$ and $Y(0; \check{\tau})$ are pretty close to the Bermudan price (the error is typically less than 1.5% relative to the value).

The example nicely illustrates the efficiency gain due to the scenario selection. For simulating $Y(0; \tilde{\tau})$ we need to estimate the conditional expectations by nested Monte Carlo simulation at each exercise date until the decision to exercise is made. However, since a closed form expression of the process L is available, we can avoid the nested Monte Carlo simulation at many exercise dates by rejecting the dates, which are not in A , see Section 3.2. In columns 7 and 8, we display the average number of points (per trajectory), where the nested Monte Carlo simulation has been carried out for constructing $Y(0; \tilde{\tau})$ and $Y(0; \check{\tau})$ respectively. We see, that pre-selecting exercise dates by checking $Z(i) < L(i)$ for each i reduces the number of nested Monte Carlo simulations up to 15 times. However, the values of $Y(0; \tilde{\tau})$ and $Y(0; \check{\tau})$ are the same within one standard deviation.

Table 1.

S_0	$Y(0; \tau)$ (SD)	$Y(0; \tilde{\tau})$ (SD)	$Y(0; \bar{\tau})$ (SD)	$Y_{up}(0; \tau)$ (SD)	\tilde{N}	\bar{N}
90	0.369(0.000)	0.427(0.002)	0.425(0.002)	0.431(0.002)	0.4	6.3
95	0.916(0.001)	1.052(0.003)	1.058(0.003)	1.064(0.003)	0.5	6.4
100	2.136(0.001)	2.364(0.004)	2.370(0.004)	2.395(0.004)	0.8	5.6
103	3.430(0.001)	3.668(0.005)	3.677(0.005)	3.716(0.004)	1.0	4.6

Table 2.

S_0	$Y(0; \tau)$ (SD)	$Y(0; \tilde{\tau})$ (SD)	$Y(0; \bar{\tau})$ (SD)	$Y_{up}(0; \tau)$ (SD)	\tilde{N}	\bar{N}
90	2.229(0.001)	2.421(0.011)	2.411(0.012)	2.462(0.009)	0.4	6.7
95	3.428(0.002)	3.709(0.015)	3.723(0.015)	3.762(0.016)	0.6	6.5
100	5.087(0.002)	5.495(0.017)	5.496(0.017)	5.523(0.017)	0.8	6.0
103	6.371(0.002)	6.810(0.019)	6.829(0.019)	6.890(0.018)	0.9	5.5

4.2 Asset based cancelable coupon swap

Nowadays bonds which bear coupons depending on the performance of a basket of stocks have become popular. Our second example is a stylized version of such product.

With respect to the system (20) we consider an exotic structured product specified as follows: Let T_1, \dots, T_k be a sequence of exercise dates. The option holder (a bank for example) pays a sequence of coupons on a \$1 nominal account at the exercise dates T_i up to a cancellation time (index) τ , according to the following scheme: Fix a quantile α , $0 < \alpha < 1$, numbers $1 \leq n_1 < n_2 \leq n$ (we assume $n \geq 2$), and three rates s_1, s_2, s_3 . Let

$$M(i) := \#\{l : 1 \leq l \leq n, S_l(T_i) \leq (1 - \alpha)S_l(0)\},$$

i.e. $M(i)$ is the number of assets which at T_i have fallen more than α with respect to their starting value at time zero. We then introduce the random rate

$$a(i) = s_1 1_{M(i) \leq n_1} + s_2 1_{n_1 < M(i) \leq n_2} + s_3 1_{n_2 < M(i)}$$

and specify the T_i -coupon to be

$$C(i) := a(i)(T_i - T_{i-1}).$$

For pricing this structure we need to compare the coupons $C(i)$ with risk free coupons over the period $[T_{i-1}, T_i]$ and thus consider the discounted net cash-flow process

$$\mathcal{Z}(i) := e^{-rT_i}(e^{r(T_i - T_{i-1})} - 1 - C(i)) = e^{-rT_{i-1}} - e^{-rT_i} - a(i)e^{-rT_i}(T_i - T_{i-1}).$$

The product value at time zero may be represented as an optimal stopping problem,

$$V_0 = \sup_{\tau \in \{1, \dots, k\}} E[Z(\tau)] := \sup_{\tau \in \{1, \dots, k\}} E\left[\sum_{i=1}^{\tau} \mathcal{Z}(i)\right].$$

Indeed, the discounted net cash-flows may be regarded as investments in the numeraire, which aggregate up to a virtual pay-off $Z(\tau)$ at the cancellation date τ .

Remark 4.1. Note that the cashflow $Z(i)$ can become negative in this example. It is, however, bounded from below. Hence, one can shift the cashflow to obtain an equivalent stopping problem with nonnegative cashflow and therefore all theoretical results of this paper still hold true.

For our experiments, we choose a five-year option with semiannual exercise possibility,

$$k = 10; \quad T_i - T_{i-1} = 0.5, \quad 1 \leq i \leq k.$$

For simplicity we consider a basket of 20 independent and identically distributed assets. Precisely, we take the following values of the parameters,

$$n = 20, \quad r = 0.05, \quad \delta = 0, \quad \sigma_l = 0.2, \quad \rho_{lm} = \delta_{lm}, \quad S_l(0) = 100, \quad 1 \leq l, m \leq 20,$$

$$n_1 = 5, \quad n_2 = 10, \quad \alpha = 0.05, \quad s_1 = 0.09, \quad s_2 = 0.03, \quad s_3 = 0.$$

Since no good closed form approximation of still-alive Europeans is available for this product, we estimate an initial exercise boundary by the method of Longstaff & Schwartz [10] and construct an adapted random set A_γ by

$$A_\gamma = \{i : 1 \leq i \leq k-1, 0 \geq \sum_{q \geq 1} c_{iq} \beta_q(S(T_i)) - \gamma\} \cup \{k\}, \quad \gamma \geq 0.$$

Here, $(\beta_q)_{q \geq 1}$ is a system of basis functions on the state space. The coefficients $(c_{iq})_{q \geq 1}$, $1 \leq i \leq k-1$ are pre-computed via a backward least square regression on pre-simulated asset trajectories, see, e.g., [10]. The criterion

$$0 \geq \sum_{q \geq 1} c_{iq} \beta_q(S(T_i))$$

approximates, backwards in time, the true exercise criterion

$$0 \geq E^{\mathcal{F}_i} \left[\sum_{p=i+1}^{\tilde{\tau}^{*(i+1)}} \mathcal{Z}(p) \right]$$

at time i and is shifted by γ as motivated in Remark 3.3. In all cases below we pre-compute the coefficients c_{iq} by simulating $2 \cdot 10^5$ trajectories.

Starting from different systems of basis functions, we compute a lower bound $Y(0; \tau^{LS})$ of the Bermudan price with $\tau^{LS}(i) = \inf\{j : i \leq j \leq k, j \in A_0\}$, an improved lower bound $Y(0; \tilde{\tau}_{[\gamma]}^{LS})$ with

$$\tilde{\tau}_{[\gamma]}^{LS}(i) = \inf\{j : i \leq j \leq k, (0 \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} \sum_{q=j+1}^{\tau^{LS}(p)} \mathcal{Z}(q)) \wedge (j \in A_\gamma)\} \quad (23)$$

and a dual upper bound $Y_{up}(0; \tau^{LS})$ given by (17). We also run, for comparison, the one-step improvement $Y(0; \tilde{\tau}^{LS})$ without scenario selection, i.e.

$$\tilde{\tau}^{LS}(i) = \inf\{j : i \leq j \leq k, 0 \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} \sum_{q=j+1}^{\tau^{LS}(p)} \mathcal{Z}(q)\}. \quad (24)$$

The corresponding numerical price bounds are presented in Table 3.

To obtain these results, we simulate $Y(0; \tau^{LS})$ with 10^7 trajectories. For $Y(0; \tilde{\tau}_{[\gamma]}^{LS})$ and $Y_{up}(0; \tau^{LS})$ we use a variance reducing representation, analogous to (22), where the second term is simulated using $3.5 \cdot 10^4$ outer and 500 inner Monte Carlo trajectories. Further, we compute an approximation of $Y_{up}(0; \tau^{LS}) - Y(0; \tau^{LS})$ by 2500 outer and 1000 inner trajectories.

It turns out that the choice of the basis functions is a nontrivial issue. Recall that the ‘classical’ system of basis functions consists of polynomials on the underlying assets S_1, \dots, S_{20} and the net cashflow \mathcal{Z} . For polynomials of first degree, this basis provides a 12% gap between the Longstaff-Schwartz lower bound $Y(0; \tau^{LS})$ and the corresponding dual upper bound $Y_{up}(0; \tau^{LS})$ relative to the value (an absolute error of 19 b.p.), see Table 3, row 1. This system of basis functions can be naturally extended as follows. Let us order the assets at time T_i for a fixed trajectory ω from the worst to the best: $S_{p_1(\omega)}(T_i, \omega) \leq \dots \leq S_{p_{20}(\omega)}(T_i, \omega)$. It then holds,

$$\begin{aligned} M(i) \leq n_1 &\Leftrightarrow S_{p_{n_1+1}}(T_i) > (1 - \alpha)S_{p_{n_1+1}}(0), \\ n_1 < M(i) \leq n_2 &\Leftrightarrow (S_{p_{n_1+1}}(T_i) \leq (1 - \alpha)S_{p_{n_1+1}}(0)) \wedge \\ &\quad (S_{p_{n_2+1}}(T_i) > (1 - \alpha)S_{p_{n_2+1}}(0)), \\ n_2 < M(i) &\Leftrightarrow S_{p_{n_2+1}}(T_i) \leq (1 - \alpha)S_{p_{n_2+1}}(0). \end{aligned}$$

So, it is promising to add the processes $S_{p_{n_1+1}}(T_i)$ and $S_{p_{n_2+1}}(T_i)$ to the basis, because they trigger the coupon level. This extension decreases the gap between $Y(0; \tau^{LS})$ and $Y_{up}(0; \tau^{LS})$ to 7.5% relative to the value (12 b.p.), see Table 3, row 5. An increase of the polynomial degree leads to a rapid growth of the number of basis functions and thus the computing time. In fact, the set of second order polynomials on the extended asset system $S_1, \dots, S_{20}, \mathcal{Z}, S_{p_1+1}, S_{p_2+1}$ consists of 300 functions. Pre-computation of the regression coefficients in this case requires in our implementation more than 80 minutes (for comparison, in all the other cases pre-computation takes less than 1 minute). However, the gap between $Y(0; \tau^{LS})$ and $Y_{up}(0; \tau^{LS})$ is still not satisfactory (3.8% relative to the value, 6.5 b.p.), see Table 3, row 5. Another way to define the basis functions, suggested in Piterbarg [12], is to construct polynomials on a small set of functions, called explanatory variables, which contain important information about the product. In our example, S_{p_1+1} , S_{p_2+1} , and \mathcal{Z} constitute a straightforward choice of explanatory variables. In Table 3, row 2-4, we present the values for polynomials on this system of explanatory variables up to 4th degree. We see, that the gap between $Y(0; \tau^{LS})$ and $Y_{up}(0; \tau^{LS})$ is approximately 9% relative to the value (15 b.p.) and decreases only slightly with the growth of the basis.

In contrast, the improved stopping family $\tilde{\tau}^{LS}$ provides a very close approximation of the Bermudan price for all considered basis functions. The gap between $Y(0; \tilde{\tau}^{LS})$ and $Y_{up}(0; \tau^{LS})$ does not exceed 1.5% relative to the value (2.5 b.p.) in almost all cases, see Table 3, column 4. As in the previous example, we can essentially speed up the procedure by computing conditional expectations in (23) only when the date j belongs to A_γ . Note that the bigger γ , the closer is $\tilde{\tau}_{[\gamma]}^{LS}$ to the standard improved policy $\tilde{\tau}$, but the smaller is the gain in the computing time. We illustrate this phenomenon by computing $Y(0; \tilde{\tau}_{[\gamma]}^{LS})$ for $\gamma = 0$ and $\gamma = 0.01$, see Table 3, columns 2-3, respectively. In Table 4, we display the average (per trajectory) number of nested Monte Carlo simulations for computing $Y(0; \tilde{\tau}_{[0]}^{LS})$, $Y(0; \tilde{\tau}_{[0.01]}^{LS})$ and $Y(0; \tilde{\tau}^{LS})$, denoted

by $\tilde{N}_{[0]}$, $\tilde{N}_{[0.01]}$ and \tilde{N} , respectively. As we see, $Y(0; \tilde{\tau}_{[0.01]}^{LS})$ coincides with $Y(0; \tilde{\tau}^{LS})$ within 1 standard deviation in almost all cases, while the number of inner simulations for computing $Y(0; \tilde{\tau}_{[0.01]}^{LS})$ is 4.5-5 times smaller. We also report that computing $Y(0; \tilde{\tau}_{[0.01]}^{LS})$ (provided the regression coefficients and $Y(0; \tau^{LS})$ are pre-computed) takes in all cases 10-15 minutes on our Pentium-III computer. The computing times for $Y(0; \tilde{\tau}_{[0]}^{LS})$ are roughly two times smaller, but the lower bound due to $\tilde{\tau}_{[0]}^{LS}$ does not close the duality gap as satisfactory as $\tilde{\tau}_{[0.01]}^{LS}$ does.

Table 3. (all values in base points)

Basis functions	$Y(0; \tau^{LS})$	$Y(0; \tilde{\tau}_{[0]}^{LS})$	$Y(0; \tilde{\tau}_{[0.01]}^{LS})$	$Y(0; \tilde{\tau}^{LS})$	$Y_{up}(0; \tau^{LS})$
$1, \mathcal{Z}, S_i$	159.8(0.2)	166.5(0.5)	173.0(0.8)	174.2(0.8)	179.0(0.8)
$1, \mathcal{Z}, S_{p_{1,2}+1}$	161.4(0.2)	167.0(0.5)	172.1(0.8)	174.2(0.8)	176.5(0.6)
2-polynomial $\mathcal{Z}, S_{p_{1,2}+1}$	161.4(0.2)	167.9(0.5)	174.4(0.8)	175.6(0.8)	176.1(0.7)
4-polynomial $\mathcal{Z}, S_{p_{1,2}+1}$	162.0(0.2)	170.2(0.6)	174.3(0.8)	175.5(0.8)	177.0(0.7)
$1, \mathcal{Z}, S_i, S_{p_{1,2}+1}$	164.8(0.2)	168.4(0.5)	174.3(0.7)	174.5(0.7)	177.0(0.6)
2-polynomial $\mathcal{Z}, S_i, S_{p_{1,2}+1}$	169.7(0.2)	171.4(0.4)	174.5(0.6)	175.3(0.6)	176.2(0.4)

Table 4.

Basis functions	$\tilde{N}_{[0]}$	$\tilde{N}_{[0.01]}$	\tilde{N}	Basis functions	$\tilde{N}_{[0]}$	$\tilde{N}_{[0.01]}$	\tilde{N}
$1, \mathcal{Z}, S_i$	0.62	1.41	6.53	$1, \mathcal{Z}, S_{p_{1,2}+1}$	0.61	1.39	6.57
$1, \mathcal{Z}, S_i, S_{p_{1,2}+1}$	0.58	1.35	6.60	2-polynomial $\mathcal{Z}, S_{p_{1,2}+1}$	0.65	1.49	6.58
2-polynomial $\mathcal{Z}, S_i, S_{p_{1,2}+1}$	0.59	1.46	6.67	4-polynomial $\mathcal{Z}, S_{p_{1,2}+1}$	0.67	1.51	6.61

Remark 4.2. We finally remark that often very crude input families already yield surprisingly good improved lower bounds. In this respect we report simulation results starting from the initial exercise policy "exercise, when a cashflow becomes negative or zero", i.e.

$$\tau(i) = \inf\{j : i \leq j \leq k-1, 0 \geq \mathcal{Z}(j)\} \cup \{k\}.$$

While this initial stopping family gives a very crude lower bound $Y(0; \tau)$, namely 85.7(0.2) b.p., which is roughly 50% off the true price, its one step improvement $Y(0; \tilde{\tau})$ is at 172.8(0.8) b.p. already better than the lower bounds provided by the Longstaff-Schwartz algorithm in our experiments. Here, we use 10^7 Monte Carlo trajectories for $Y(0; \tau)$ and $1.5 \cdot 10^5$ outer (with 500 inner) trajectories for

$Y(0; \tilde{\tau})$. Moreover, when we enhance the improvement procedure by computing

$$\max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} \sum_{q=j+1}^{\tau(p)} Z(q)$$

only at those exercise dates j , which are in the set $A = \{i : 1 \leq i \leq k-1, 0 \geq Z(i)\} \cup \{k\}$, the value of the improved lower bound becomes equal to 171.6(0.8) which is slightly higher than the best Longstaff-Schwartz lower bound and requires approximately the same computing time. Of course, it is more efficient with respect to computing time and quality of the lower bound to run the enhanced improvement step starting from one of the Longstaff-Schwartz policies (computed with a ‘small’ basis).

A Proof of Theorem 3.5

We first prove the next Proposition.

Proposition A.1. *Let τ and σ be two consistent stopping families, such that $\sigma(i) \leq \tau(i)$, $0 \leq i \leq k$. Then,*

$$Y(i; \tau) - Y(i; \sigma) = \sum_{j=i}^{k-1} E^{\mathcal{F}_i} [1_{\tau(i) > j} 1_{\sigma(i)=j} (Y(j; \tau) - Z(j))].$$

Proof. We have,

$$\begin{aligned} Y(i; \tau) - Y(i; \sigma) &= E^{\mathcal{F}_i} [1_{\tau(i) > \sigma(i)} (Z(\tau(i)) - Z(\sigma(i)))] \\ &= \sum_{j=i}^{k-1} E^{\mathcal{F}_i} [1_{\tau(i) > j} 1_{\sigma(i)=j} (Z(\tau(i)) - Z(j))] \\ &= \sum_{j=i}^{k-1} E^{\mathcal{F}_i} [1_{\tau(i) > j} 1_{\sigma(i)=j} (Z(\tau(j)) - Z(j))] \\ &= \sum_{j=i}^{k-1} E^{\mathcal{F}_i} [1_{\tau(i) > j} 1_{\sigma(i)=j} (E^{\mathcal{F}_j} Z(\tau(j)) - Z(j))] \\ &= \sum_{j=i}^{k-1} E^{\mathcal{F}_i} [1_{\tau(i) > j} 1_{\sigma(i)=j} (Y(j; \tau) - Z(j))]. \end{aligned}$$

Here we use that, due to the consistency, if $i \leq j \leq \tau(i)$, then $\tau(j) = \tau(i)$. \square

As a second preliminary result for the proof of Theorem 3.5 we have the following lemma.

Lemma A.2. *Suppose τ^* is some optimal stopping family for the cashflow Z . Then $A^*(\omega) = \{\tau^*(i, \omega), 0 \leq i \leq k\}$ is an a-priori set. Moreover,*

$$\tau^*(i) = \inf \left\{ j : i \leq j \leq k, (Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j} [Z(\tau^*(p))]) \wedge (j \in A^*) \right\}$$

provided τ^ is consistent.*

Proof. Since $\{i \in A^*\} = \bigcup_{0 \leq j \leq i} \{\tau^*(j) = i\} \in \mathcal{F}_i$ and $\tau^*(i) \in A^*$, A^* is an a-priori set. Suppose now that additionally τ^* is consistent. Then, by consistency and optimality of τ^* , and by the supermartingale property of the Snell envelope, we have

$$\begin{aligned} & \inf \left\{ j : i \leq j \leq k, (Z(j) \geq \max_{j+1 \leq p \leq k} E^{\mathcal{F}_j}[Z(\tau^*(p))]) \wedge (j \in A^*) \right\} \\ &= \inf \left\{ j : i \leq j \leq k, (Z(j) \geq \max_{j \leq p \leq k} E^{\mathcal{F}_j}[Z(\tau^*(p))]) \wedge (j \in A^*) \right\} \\ &= \inf \left\{ j : i \leq j \leq k, (Z(j) \geq \max_{j \leq p \leq k} E^{\mathcal{F}_j}[Y^*(p)]) \wedge (j \in A^*) \right\} \\ &= \inf \{ j : i \leq j \leq k, (Z(j) \geq Y^*(j)) \wedge (j \in A^*) \}. \end{aligned}$$

Moreover, by consistency, $j \in A^*$ if and only if $\tau^*(j) = j$. Hence,

$$\begin{aligned} & \inf \{ j : i \leq j \leq k, (Z(j) \geq Y^*(j)) \wedge (j \in A^*) \} \\ &= \inf \{ j : i \leq j \leq k, (Z(\tau^*(j)) \geq Y^*(\tau^*(j))) \wedge (\tau^*(j) = j) \} \\ &= \inf \{ j : i \leq j \leq k, \tau^*(j) = j \} = \tau^*(i). \end{aligned}$$

Note, for the second identity we applied the well-known fact, that evaluated at any optimal stopping time the Snell envelope Y^* equals the cashflow Z . Thus, $(Z(\tau^*(j)) \geq Y^*(\tau^*(j)))$ is always satisfied for all $0 \leq j \leq k$. \square

After these preparations we prove Theorem 3.5.

Proof of Theorem 3.5. Let τ be a consistent and optimal stopping family for the cashflow Z . We define $\tilde{\tau}$ and $\tilde{\sigma}$ as in Corollary 3.6 for $A_2(\omega) = \{\tau(i, \omega), 0 \leq i \leq k\}$ and $A_1(\omega) = A_2(\omega) \cap A(\omega)$. Then $\tilde{\tau} \geq \tilde{\sigma}$ and $\tilde{\sigma} = \tau$ is optimal due to Lemma A.2. Hence by Proposition A.1 we have,

$$E[Y^*(i) - Y(i; \tilde{\tau})] = \sum_{j=i}^{k-1} E[1_{\tilde{\tau}(i) > j} 1_{\tau(i)=j} (Z(j) - Y(j; \tilde{\tau}))]$$

As $\tilde{\tau}$ takes values in A and is possibly suboptimal for the cashflow Z_A , we obtain

$$Y(i; \tilde{\tau}) = E^{\mathcal{F}_i}[Z_A(\tilde{\tau}(i))] \leq Y_A^*(i).$$

Consequently, by Hölder's inequality,

$$\begin{aligned} & E[Y^*(i) - Y_A^*(i)] \leq E[Y^*(i) - Y(i; \tilde{\tau})] \\ & \leq \sum_{j=i}^{k-1} E[1_{\{j \notin A\}} 1_{\{\tau(i)=j\}} Z(j)] \\ & \leq \max_{i \leq j \leq k-1} (E[|Z(j) 1_{\{0, \dots, k\} \setminus A}(j)|^q])^{1/q} \sum_{j=i}^{k-1} P(\{\tau(i) = j\} \cap \{j \notin A\})^{1-1/q} \\ & \leq \max_{i \leq j \leq k-1} (E[|Z(j) 1_{\{0, \dots, k\} \setminus A}(j)|^q])^{1/q} (k-i)^{1/q} \\ & \quad \times \left(\sum_{j=i}^{k-1} P(\{\tau(i) = j\} \cap \{j \notin A\}) \right)^{1-1/q}. \end{aligned}$$

The obvious equation

$$\sum_{j=i}^{k-1} P(\{\tau(i) = j\} \cap \{j \notin A\}) = P(\{\tau(i) \notin A\})$$

concludes. \square

References

- [1] Andersen, L. (1999) A simple approach to the pricing of Bermudan swaptions in a multifactor LIBOR market model. *Journal of Computational Finance*, **3**, 5-32.
- [2] Bender, C., Kolodko, A., Schoenmakers, J. (2006) Iterating cancelable snowballs and related exotics in a many-factor Libor model. *Risk*, to appear.
- [3] Bender, C., Schoenmakers, J. (2006) An iterative algorithm for multiple stopping: Convergence and stability. *Advances in Appl. Prob.*, to appear.
- [4] Bouchard, B., Ekeland, I., Touzi, N. (2004) On the Malliavin approach to Monte Carlo approximation of conditional expectations. *Finance and Stochastics*, **8**, 45-71.
- [5] Broadie, M., Glasserman, P. (2004) A stochastic mesh method for pricing high-dimensional American options. *Journal of Computational Finance*, **7**(4), 35-72.
- [6] Haugh, M. B., Kogan, L. (2004) Pricing american options: a duality approach. *Operations Research*, **52**, 258-270.
- [7] Jamshidian, F. (2006) Duality of optimal stopping times and domineering claims: A Doob-Meyer decomposition approach to American options, Working paper.
- [8] Kolodko, A., Schoenmakers, J. (2004) Upper Bounds for Bermudan Style Derivatives. *Monte Carlo Methods and Applications*, **10**(3-4), 331-343.
- [9] Kolodko, A., Schoenmakers, J. (2006) Iterative construction of the optimal Bermudan stopping time. *Finance and Stochastics*, **10**(1), 27-49.
- [10] Longstaff, F. A., Schwartz, R. S. (2001) Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies*, **14**, 113-147.
- [11] Neveu, J. (1975) *Discrete parameter martingales*. North-Holland: Amsterdam.
- [12] Piterbarg, V. V. (2004) Pricing and hedging callable Libor exotics in forward Libor models, *Journal of Computational Finance*, **8** (2), 65-117.
- [13] Puterman, M. (1994) *Markov decision processes*. New York: Wiley.
- [14] Rogers, L. C. G. (2002) Monte Carlo valuation of American options. *Math. Finance*, **12**, 271-286.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

[15] Tsitsiklis, J. N. and Van Roy, B. (2001) Regression methods for pricing complex American-style options. *IEEE Trans. on Neural Networks*, **12** (4), 694–703.

For Peer Review Only