

Dual pricing of American options by Wiener chaos expansion*

Jérôme Lelong[†]

March 3, 2020

Premia 22

Abstract

In this work, we propose an algorithm to price American options by directly solving the dual minimization problem introduced by Rogers [2002]. Our approach relies on approximating the set of uniformly square integrable martingales by a finite dimensional Wiener chaos expansion. Then, we use a sample average approximation technique to efficiently solve the optimization problem. Unlike all the regression based methods, our method can transparently deal with path dependent options without extra computations and a parallel implementation writes easily with very little communication and no centralized work. We test our approach on several multi-dimensional options with up to 40 assets and show the impressive scalability of the parallel implementation.

Key words: American option, duality, Snell envelope, stochastic optimization, sample average approximation, high performance computing, Wiener chaos expansion.

AMS subject classification: 62L20, 62L15, 91G60, 65Y05, 60H07

1 Introduction

The pricing of American options quickly becomes challenging as the dimension increases and the payoff gets complex. Many people have contributed to this problem usually by considering its dynamic programming principle formulation Tilley [1993], Carriere [1996], Tsitsiklis and Roy [2001], Longstaff and Schwartz [2001], Broadie and Glasserman [2004] and Bally and Pages [2003]. Among this so extensive literature, the practitioners seem to prefer the iterative optimal policy approach proposed by Longstaff and Schwartz [2001], which proves to be quite efficient in many situations. However, *true* path-dependent options cannot be handled by this approach. Solving the dynamic programming principle requires the computation of a conditional expectation, which is eventually dealt with regression

*This project was supported by the Finance for Energy Market Research Centre, www.fime-lab.org. The High Performance Computations presented in this paper were performed using the Froggy platform of the CIMENT infrastructure (<https://ciment.ujf-grenoble.fr>), which is supported by the Rhône-Alpes region (GRANT CPER07_13 CIRA) and the Equip@Meso project (reference ANR-10-EQPX-29-01) of the programme Investissements d'Avenir supervised by the Agence Nationale pour la Recherche.

[†]Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France.
jerome.lelong@univ-grenoble-alpes.fr

techniques. These techniques are known to suffer from the curse of dimensionality: global regression methods lead to high dimensional linear algebra problems, whereas local methods see the number of domains blow up with the dimension. Despite the numerous parallel implementation of this techniques (see for instance Dung Doan et al. [2010], Abbas-Turki et al. [2014]), we cannot expect to obtain a fully scalable algorithm. In this work, we follow the dual approach initiated by Rogers [2002] and Haugh and Kogan [2004], which can naturally handle path dependent options. To make it implementable, we need a smart and finite dimensional approximation of the set of uniformly integrable martingales. We chose the set of truncated Wiener chaos expansions, which have some magic features in our problem: its density makes the optimization differentiable almost everywhere and computing its conditional expectation exactly is straightforward. Then, the pricing problem boils down to a finite dimensional and convex optimization problem. The optimization problem is solved using a *Sample Average Approximation* (see Rubinstein and Shapiro [1993]), which can be easily and efficiently implemented using parallel computing.

We fix some finite time horizon $T > 0$ and a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P})$, where $(\mathcal{F}_t)_{0 \leq t \leq T}$ is supposed to be the natural augmented filtration of a d -dimensional Brownian motion B . On this space, we consider an adapted process $(S_t)_{0 \leq t \leq T}$ with values in $\mathbb{R}^{d'}$ modeling a d' -dimensional underlying asset. The number of assets d' can be smaller than the dimension d of the Brownian motion to encompass the case of stochastic volatility models or stochastic interest rate. We assume that the short interest rate is modeled by an adapted process $(r_t)_{0 \leq t \leq T}$ with values in \mathbb{R}_+ and that \mathbb{P} is an associated risk neutral measure. We consider an adapted payoff process \tilde{Z} and introduce its discounted value process $\left(Z_t = e^{-\int_0^t r_s ds} \tilde{Z}_t \right)_{0 \leq t \leq T}$. We assume that the paths of Z are right continuous and that $\sup_{t \in [0, T]} |Z_t| \in L^2$. The process \tilde{Z} can obviously take the simple form $(\phi(S_t))_{t \leq T}$ but it can also depend on the whole path of the underlying asset S up to the current time. So, our framework transparently deals with path-dependent option, which are far more difficult to handle using regression techniques.

We consider the American option paying \tilde{Z}_t to its holder if exercised at time t . Standard arbitrage pricing theory defines the discounted time- t value of the American option to be

$$U_t = \text{esssup}_{\tau \in \mathcal{T}_t} \mathbb{E}[Z_\tau | \mathcal{F}_{t_k}]$$

where \mathcal{T}_t denotes the set of \mathcal{F} -stopping times with values in $[t, T]$. The integrability properties of Z ensure that U is a supermartingale of class (D) and hence has a Doob–Meyer decomposition

$$U_t = U_0 + M_t^* - A_t^*$$

where M^* is a martingale vanishing at zero and A^* is a predictable integrable increasing process also vanishing at zero. With our assumptions on Z , M^* is square integrable. Following Davis and Karatzas [1994], Rogers [2002] and Haugh and Kogan [2004] found an alternative representation of the price at time-0 of the American option as the minimum value of the following optimization problem

$$U_0 = \inf_{M \in H_0^2} \mathbb{E} \left[\sup_{t \leq T} (Z_t - M_t) \right] = \mathbb{E} \left[\sup_{t \leq T} (Z_t - M_t^*) \right] \quad (1)$$

where H_0^2 denotes the set of square integrable martingales vanishing at zero. A martingale reaching the infimum is called an *optimal* martingale. As the dual price problem writes as a convex minimisation problem, the set of all optimal martingales is a convex subset of H_0^2 . Among the martingales reaching the infimum in (1), some of them actually satisfy the pathwise equality $\sup_{t \leq T} Z_t - M_t = U_0$. These martingales are called *surely optimal*. Any surely optimal martingale reaches the lower bound in (1) but not all optimal martingales are surely optimal. We refer to Schoenmakers et al. [2013] for a detailed characterisation of optimal martingales. Anyway, Jamshidian [2007] proved the uniqueness of surely optimal martingales within the continuing region, ie. for any surely optimal martingale M and any optimal strategy τ , $(M_{t \wedge \tau})_t = (M_{t \wedge \tau}^*)_t$ a.s.

The most famous method using the dual representation (1) is probably the primal–dual approach of Andersen and Broadie [2004], which heavily relies on the knowledge of an optimal exercising policy. The a priori knowledge may take the form of nested Monte Carlo simulations as in Schoenmakers [2005], and Kolodko and Schoenmakers [2004]. To circumvent this difficulty, Rogers [2010] explained how to construct a *good* martingale. In a Wiener framework, Belomestny et al. [2009] investigated this approach by relying on the martingale representation theorem to build *good* martingales. When trying to practically use the dual formulation (1), the first difficulty is to find a rich enough but finite dimensional approximation of H_0^2 (Belomestny [2013] suggests to use the martingale representation theorem) and then we face a finite although potentially high–dimensional minimization problem, which can be written as a linear programming problem with as many constraints as the number of sample paths used (see Desai et al. [2012]). Belomestny [2013] considered a penalized version of (1) of the form

$$\inf_{M \in H_0^1} \left(\mathbb{E} \left[\sup_{0 \leq t \leq T} (Z_t - M_t) \right] + \kappa \text{Var} \left(\sup_{0 \leq t \leq T} (Z_t - M_t) \right) \right)$$

with $\kappa > 0$. This criteria naturally selects surely optimal martingales (see Schoenmakers et al. [2013]). Although adding such a penalization looks interesting from a theoretical point of view, it breaks the convexity of the minimization criterion, which makes the problem harder to solve from a practical point of view. We refer to Belomestny et al. [2017] for a detailed analysis of the properties of the penalized problem.

The minimization problem (1) can be equivalently formulated as

$$U_0 = \inf_{X \in L_0^2(\Omega, \mathcal{F}_T, \mathbb{P})} \mathbb{E} \left[\sup_{0 \leq t \leq T} (Z_t - \mathbb{E}[X | \mathcal{F}_t]) \right] \quad (2)$$

where $L_0^2(\Omega, \mathcal{F}_T, \mathbb{P})$ is the set of square integrable \mathcal{F}_T – random variables with zero mean. In this work, we suggest to use the truncated Wiener chaos expansion as a finite dimensional approximation of $L^2(\Omega, \mathcal{F}_T, \mathbb{P})$. Since Wiener chaos are orthogonal for the L^2 inner product, the computations of the conditional expectations $\mathbb{E}[X | \mathcal{F}_t]$ become straightforward and boil down to dropping some terms in the chaos expansion, which makes our approach very convenient. Based on this approximation, we propose a scalable algorithm and study its convergence.

The paper starts with the presentation of the Wiener chaos expansion and some of its useful properties in Section 2. Then, we can develop the core of our work in Section 3 in which we explain how the price of the American option can be approximated by the solution of a finite dimensional optimization problem. First, we analyze the properties of the optimization problem in order to prove the convergence of its solution to the American option price. Second,

we study its sample average approximation, which makes the problem tractable, and prove its convergence. Based on all these theoretical results, we present our algorithm in Section 4 and discuss its parallel implementation on distributed memory architectures. Finally, some numerical examples are presented in Section 5.

Notation

- For $\alpha \in \mathbb{N}^q$, $|\alpha|_1 = \sum_{i=1}^q \alpha_i$.
- For $\alpha \in (\mathbb{N}^q)^d$, $|\alpha|_1 = \sum_{j=1}^d \sum_{i=1}^q \alpha_i^j$.
- For $n \geq 1$, $0 = t_0 < t_1 < \dots < t_n = T$ is a time grid of $[0, T]$ satisfying $\lim_{n \rightarrow \infty} \sup_{0 \leq k \leq n-1} |t_{k+1} - t_k| = 0$.
- For $n \geq 1$, the discrete time filtration \mathcal{G} is defined by $\mathcal{G}_k = \sigma(B_{t_{i+1}} - B_{t_i}, i = 0, \dots, k-1)$ for all $1 \leq k \leq n$, while \mathcal{G}_0 is the trivial sigma algebra. Obviously, $\mathcal{G}_k \subset \mathcal{F}_{t_k}$ for all $0 \leq k \leq n$.
- For $1 \leq q \leq d$, $\mathbb{I}(r) \in \{0, 1\}^n$ denotes the vector $(\underbrace{0, \dots, 0}_{r-1}, \underbrace{1, 0, \dots, 0}_{n-r})$.
- For $1 \leq q \leq d$, and $1 \leq r \leq n$, $\mathbb{I}(r, q) \in \mathbb{N}^{n \times d}$ with all components equal to 0 except the component with index (r, q) which is equal to 1.

We recall some useful definitions related to Malliavin calculus using the notation of Nualart [1998]. These notations are mainly used to study the regularity properties of the minimization problem (2). The reader more interested by the algorithm can skip this part.

- Let \mathcal{S} denote the class of smooth random variables of the form $F = f(W(h_1), \dots, W(h_\ell))$ where $m \geq 1$, $f \in C_p^\infty(\mathbb{R}^{\ell \times d}, \mathbb{R})$, for all $j \leq \ell$, $h_j = (h_j^1, \dots, h_j^d) \in L^2([0, T], \mathbb{R}^d)$ and for all $i \leq d$, $W^i(h_j^i) = \int_0^T h_j^i(t) dW_t^i$.
- For $F \in \mathcal{S}$, the Malliavin derivative of F denoted by $DF = (D^1, \dots, D^d)$ is a stochastic process with values in \mathbb{R}^d . For $t \leq T$ and $1 \leq i \leq d$, D_t^i is defined by

$$D_t^i F = \sum_{j=1}^{\ell} \partial_j f(W(h_1), \dots, W(h_m)) h_j^i(t).$$

With this notation, D_t is a gradient operator.

- For $m \geq 1$, a multi-index $\alpha \in \{1, \dots, d\}^m$ and a tuple of dates (t_1, \dots, t_m) , we write

$$D_{t_1, \dots, t_m}^\alpha F = D_{t_1}^{\alpha_1} (\dots (D_{t_m}^{\alpha_m} F)).$$

$D^{(m)}F = \{D_{t_1, \dots, t_m}^\alpha F : \alpha \in \{1, \dots, d\}^m, (t_1, \dots, t_m) \in [0, T]^m\}$ can be seen as a measurable function defined on $\Omega \times [0, T]^m$. When $d = 1$, we drop the multi-index and just write D_{t_1, \dots, t_m} .

- Let $\mathbb{D}^{m,2}$ be the closure of \mathcal{S} w.r.t. the following norm

$$\|F\|_{\mathbb{D}^{m,2}}^2 = \mathbb{E}[|F|^2] + \sum_{r=1}^m \sum_{|\alpha|_1=r} \mathbb{E} \left[\int_{[0,T]^r} |D_{t_1, \dots, t_r}^\alpha F|^2 dt_1 \dots dt_r \right].$$

2 Wiener chaos expansion

In this section, we briefly recall the principles of Wiener chaos expansion and its basic properties. We refer to Nualart [1998] for theoretical details.

Let H_i be the i -th Hermite polynomial defined by

$$H_0(x) = 1; \quad H_i(x) = (-1)^i e^{x^2/2} \frac{d^i}{dx^i} (e^{-x^2/2}), \text{ for } i \geq 1.$$

They satisfy for all integer i , $H'_i = H_{i-1}$ with the convention $H_{-1} = 0$. We recall that if (X, Y) is a standard random normal vector $\mathbb{E}[H_i(X)H_j(Y)] = i! (\mathbb{E}[XY])^i \mathbf{1}_{i=j}$.

It is well-known that every square integrable \mathcal{F}_T -measurable random variable F admits the following orthonormal decomposition

$$F = \mathbb{E}[F] + \sum_{k=1}^{\infty} \sum_{\alpha \in (\mathbb{N}^{\mathbb{N}})^d, |\alpha|_1 = k} \lambda_{\alpha} \prod_{j=1}^d \prod_{i \geq 1} H_{\alpha_i^j} \left(\int_0^T \eta_i^j(t) dB_t^j \right)$$

where $(\eta_i^j)_{i \geq 1, 1 \leq j \leq d}$ is an orthonormal basis of $L^2([0, T], \mathbb{R}^d)$. For all $p \geq 0$, we define the Wiener chaos of order p by

$$\mathcal{H}_p = \overline{\text{span}}^{L^2(\Omega, \mathcal{F}_T)} \left\{ \prod_{j=1}^d H_{p_j} \left(\int_0^T f_t^j dB_t^j \right) : f \in L^2([0, T], \mathbb{R}^d), \sum_{j=1}^d p_j = p \right\}.$$

We denote the projection of a random variable $F \in L^2(\mathcal{F}_T)$ on to $\bigoplus_{\ell=0}^p \mathcal{H}_{\ell}$ by $C_p(F)$.

Consider the indicator functions of the grid defined by $0 = t_0 < t_1 < \dots < t_n = T$ with values in \mathbb{R}^d defined by

$$f_i^j(t) = \frac{\mathbf{1}_{[t_{i-1}, t_i]}(t)}{\sqrt{t_i - t_{i-1}}} \mathbf{e}_j, \quad i = 1, \dots, n, \quad j = 1, \dots, d$$

where $(\mathbf{e}_1, \dots, \mathbf{e}_d)$ denotes the canonical basis of \mathbb{R}^d . Based on the definition of the \mathcal{H}_p , we introduce the truncated Wiener chaos of order up to p

$$\mathcal{C}_{p,n} = \text{span} \left\{ \prod_{j=1}^d \hat{H}_{\alpha^j} (G_1^j, \dots, G_n^j) : \alpha \in (\mathbb{N}^n)^d, |\alpha|_1 \leq p \right\}$$

where $G_i^j = \frac{B_{t_i}^j - B_{t_{i-1}}^j}{\sqrt{t_i - t_{i-1}}}$ and $\hat{H}_{\alpha}(x) = \prod_{i \geq 1} H_{\alpha_i}(x_i)$, for $x \in \mathbb{R}^{\mathbb{N}}$. The truncated chaos expansion of order p of a square integrable random variable F is defined as the L^2 projection onto $\mathcal{C}_{p,n}$ and writes

$$C_{p,n}(F) = \sum_{\alpha \in A_{p,n}^{\otimes d}} \lambda_{\alpha} \hat{H}_{\alpha}^{\otimes d}(G_1, \dots, G_n)$$

where for $\alpha \in (\mathbb{N}^n)^d$, $\hat{H}_{\alpha}^{\otimes d}(G_1, \dots, G_n) = \prod_{j=1}^d \hat{H}_{\alpha^j}(G_1^j, \dots, G_n^j)$ and

$$A_{p,n}^{\otimes d} = \left\{ \alpha \in (\mathbb{N}^n)^d : |\alpha|_1 \leq p \right\}.$$

With an obvious abuse of notation, we write, for $\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}$,

$$C_{p,n}(\lambda) = \sum_{\alpha \in A_{p,n}^{\otimes d}} \lambda_\alpha \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n). \quad (3)$$

The space of truncated Wiener chaos $\mathcal{C}_{p,n}$ has the key property to be stable by the conditional expectation operator. More precisely, the following result explains how to compute, in a closed form, the conditional expectation of an element of $\mathcal{C}_{p,n}$. The proof is postponed to Section A.

Proposition 2.1 *Let F be a real valued random variable in $L^2(\Omega, \mathcal{F}_T, \mathbb{P})$ and let $k \in \{1, \dots, n\}$ and $p \geq 0$*

$$\mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}] = \sum_{\alpha \in A_{p,n}^{\otimes d,k}} \lambda_\alpha \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n)$$

where $A_{p,n}^{\otimes d,k}$ is the set of multi-indices vanishing after time t_k

$$A_{p,n}^{\otimes d,k} = \left\{ \alpha \in A_{p,n}^{\otimes d} : \forall j \in \{1, \dots, d\}, \forall \ell > k, \alpha_\ell^j = 0 \right\}.$$

Remark 2.2 *Since the sum appearing in $\mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}]$ is reduced to a sum over the set of multi-indices $\alpha \in A_{p,n}^{\otimes d,k}$, it actually only depends on the first k increments (G_1, \dots, G_k) . One can easily check that $\mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}]$ is actually given by the chaos expansion of F on the first k Brownian increments. Hence, computing a conditional expectation simply boils down to dropping non measurable terms. While it may look like a naive way to proceed, it is indeed correct in our setting.*

Remark 2.3 *The discrete time sequence $(\mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}])_{0 \leq k \leq n}$ is of course adapted to the filtration $(\mathcal{F}_{t_k})_k$ but also to the smaller filtration $(\mathcal{G}_k)_k$. This property plays a crucial role when approximating a random variable $F \in L^2(\Omega, \mathcal{G}_n, \mathbb{P})$ as we know from [Nualart, 1998, Theorem 1.1.1] that in such a case $\lim_{p \rightarrow \infty} C_{p,n}(F) = F$ in the L^2 -sense. This result holds for the fixed value n . If F were only \mathcal{F}_T -measurable and not \mathcal{G}_n -measurable, we would need to let both p and n go to infinity, $\lim_{p \rightarrow \infty, n \rightarrow \infty} C_{p,n}(F) = F$.*

3 Pricing American options using Wiener chaos expansion and sample average approximation

In this section, we aim at approximating the dual price (2) by a tractable optimization problem. This involves two kinds of approximations: first, approximate the space $L_0^2(\Omega, \mathcal{F}_T, \mathbb{P})$ by a finite dimensional vector space; second, replace the expectation by a sample average approximation.

3.1 A stochastic optimization problem

The dual price writes

$$\inf_{X \in L_0^2(\Omega, \mathcal{F}_T, \mathbb{P})} \mathbb{E} \left[\sup_{0 \leq t \leq T} (Z_t - \mathbb{E}[X|\mathcal{F}_t]) \right].$$

In this optimization problem, we replace X by its chaos expansion $C_{p,n}(X)$, which has no constant term as $\mathbb{E}[X] = 0$ and we approximate the supremum by a discrete time maximum. Then, we face a finite dimensional minimization problem to determine the optimal solution in the subset $\mathcal{C}_{p,n}$

$$\inf_{\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}, \lambda_0=0} \mathbb{E} \left[\max_{0 \leq k \leq n} (Z_{t_k} - \mathbb{E}[C_{p,n}(\lambda) | \mathcal{F}_{t_k}]) \right]. \quad (4)$$

We introduce the random functions $v_{p,n}$

$$v_{p,n}(\lambda, k, Z, G) = Z_{t_k} - \sum_{\alpha \in A_{p,n}^{\otimes d,k}} \lambda_{\alpha} \hat{H}_{\alpha}^{\otimes d}(G_1, \dots, G_n)$$

and consider the cost function $V_{p,n} : \lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}} \mapsto V_{p,n}(\lambda) = \mathbb{E} [\max_{0 \leq k \leq n} v_{p,n}(\lambda, k; Z, G)]$.

By using Proposition 2.1, we deduce that the minimization problem (4) is equivalent to

$$\inf_{\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}, \lambda_0=0} V_{p,n}(\lambda). \quad (5)$$

As the function $V_{p,n}$ is barely tractable in a closed form, we need some stochastic optimization techniques to solve the minimization problem. Two different approaches are commonly used: either, one uses a stochastic algorithm or replaces the expectation by a sample average approximation. In this work, we target large problems, which puts scalability as a primary requirement. The intrinsic sequential nature of stochastic algorithms has led us to prefer the sample average approximation approach.

We introduce the sample average approximation of $V_{p,n}$ defined by

$$V_{p,n}^m(\lambda) = \frac{1}{m} \sum_{i=1}^m \max_{0 \leq k \leq n} v_{p,n}(\lambda, k; Z^{(i)}, G^{(i)})$$

where $(Z^{(i)}, G^{(i)})_{1 \leq i \leq m}$ are i.i.d samples from the distribution of (Z, G) . Then, we aim at solving the empirical minimization problem

$$\inf_{\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}, \lambda_0=0} V_{p,n}^m(\lambda). \quad (6)$$

The solution of this problem is typically computed by a gradient descent algorithm whose baselines are summarized in Algorithm 3.1. We refer the reader to Section 4 for a detailed description of the algorithm.

In Section 3.2, we prove that this optimization problem is convex and has a solution (see Proposition 3.1) and converges to the price of the American option (see Proposition 3.2). Then, we prove in Section 3.3 that the solution of the sample average approximation converges to the solution of (4) when the number of samples goes to infinity (Proposition 3.6).

3.2 Properties of the finite dimensional problem

Proposition 3.1 *The minimization problem (5) has at least one solution.*

Proof. As the supremum of linear functions is convex, the random function $\lambda \mapsto \max_{k \leq n} v_{p,n}(\lambda, t_k, Z, G)$ is almost surely convex. The convexity of $V_{p,n}$ ensues from the linearity of the expectation.

```

1 Generate  $(G^{(1)}, Z^{(1)}), \dots, (G^{(m)}, Z^{(m)})$   $m$  i.i.d. samples following the law of  $(Z, G)$ 
2  $x_0 \leftarrow 0 \in \mathbb{R}^{A_{p,n}^{\otimes d}}$ 
3  $\ell \leftarrow 0, d_0 \leftarrow 0, v_0 \leftarrow \infty$ 
4 while True do
5    $d_{\ell+1} \leftarrow \nabla V_{p,n}^m(x_{\ell+1})$ 
6    $x_{\ell+1} \leftarrow x_\ell - \gamma \alpha_\ell d_{\ell+1}$ 
7   if the desired accuracy is reached then return
8 end

```

Algorithm 3.1: Sample Average Approximation of the dual price

Let us prove that $V_{p,n}(\lambda) \rightarrow \infty$ when $|\lambda| \rightarrow \infty$. Note that $V_{p,n}(\lambda) \geq \mathbb{E}[(C_{p,n}(\lambda))_-] \geq \frac{1}{2} \mathbb{E}[|C_{p,n}(\lambda)|]$, where we have used that $|x| = 2x_- + x$ and $\mathbb{E}[C_{p,n}(\lambda)] = 0$.

$$\mathbb{E}[|C_{p,n}(\lambda)|] = |\lambda| \mathbb{E}[|C_{p,n}(\lambda/|\lambda|)|] \geq |\lambda| \inf_{\mu \in \mathbb{R}^{A_{p,n}^{\otimes d}}, |\mu|=1} \mathbb{E}[|C_{p,n}(\mu)|].$$

By a standard continuity argument, the infimum is attained. Moreover, it is strictly positive as otherwise there would exist $\mu \in \mathbb{R}^{A_{p,n}^{\otimes d}}$ with $|\mu| = 1$ s.t. $\mathbb{E}[|C_{p,n}(\mu)|] = 0$. Using the orthogonality of the family $(\hat{H}_\alpha^{\otimes d})_{\alpha \in A_{p,n}^{\otimes d}}$, we would immediately deduce that $\mu = 0$. Hence, we show that $V_{p,n}(\lambda) \rightarrow \infty$ when $|\lambda| \rightarrow \infty$. The growth at infinity of $V_{p,n}$ combined with its convexity yields the existence of a solution to the minimization problem (5). ■

Proposition 3.1 ensures the existence of $\lambda_{p,n}^\sharp$ solving (5), ie.

$$V_{p,n}(\lambda_{p,n}^\sharp) = \inf_{\lambda \text{ s.t. } \lambda_0=0} V_{p,n}(\lambda).$$

To study the convergence of the $V_{p,n}(\lambda_{p,n}^\sharp)$, we introduce the Bermudan option with exercising dates t_0, \dots, t_n and with discounted payoff $(Z_{t_k})_k$. Let \hat{U}_k^n be its time- t_k price. The sequence $(\hat{U}_k^n)_{0 \leq k \leq n}$ is a supermartingale admitting the Doob–Meyer decomposition $\hat{U}_k^n = \hat{U}_0^n + \hat{M}_k^{*,n} - \hat{A}_k^{*,n}$ where \hat{M}^n is a square integrable $(\mathcal{F}_{t_k})_k$ -martingale and \hat{A}^n a predictable increasing process for the filtration $(\mathcal{F}_{t_k})_k$. The time-0 price can be expressed as

$$\hat{U}_0^n = \inf_{X \in L_0^2(\Omega, \mathcal{F}_T, \mathbb{P})} \mathbb{E} \left[\max_{0 \leq k \leq n} (Z_{t_k} - \mathbb{E}[X | \mathcal{F}_{t_k}]) \right] = \mathbb{E} \left[\max_{0 \leq k \leq n} (Z_{t_k} - M_{t_k}^{*,n}) \right]. \quad (7)$$

Note that $V_{p,n}(\lambda_{p,n}^\sharp) \geq \hat{U}_0^n$.

Proposition 3.2 *We have*

$$|V_{p,n}(\lambda_{p,n}^\sharp) - U_0| \leq U_0 - \hat{U}_0^n + 2 \|M_T^* - C_{p,n}(M_T^*)\|_2. \quad (8)$$

Moreover, assume \hat{U}_0^n converges to U_0 with n . Then, $V_{p,n}(\lambda_{p,n}^\sharp)$, converges to U_0 when both p and n go to infinity.

We refer to Carverhill and Webber [1990], Lamberton [2002] for results on the convergence of \hat{U}_0^n to U_0 .

Proof. We introduce the truncated chaos expansion of M_T^* and denote its coefficients by $\lambda_{p,n}^*$, ie. $C_{p,n}(M_T^*) = C_{p,n}(\lambda_{p,n}^*)$. We write that

$$\begin{aligned}
V_{p,n}(\lambda_{p,n}^\#) - U_0 &= V_{p,n}(\lambda_{p,n}^\#) - \mathbb{E} \left[\sup_t (Z_t - M_t^*) \right] \\
&= V_{p,n}(\lambda_{p,n}^\#) - \mathbb{E} \left[\max_k (Z_{t_k} - M_{t_k}^*) \right] \\
&\quad + \mathbb{E} \left[\max_k (Z_{t_k} - M_{t_k}^*) \right] - \mathbb{E} \left[\sup_t (Z_t - M_t^*) \right] \\
|V_{p,n}(\lambda_{p,n}^\#) - U_0| &\leq V_{p,n}(\lambda_{p,n}^\#) - \mathbb{E} \left[\max_k (Z_{t_k} - M_{t_k}^*) \right] \\
&\quad + \mathbb{E} \left[\sup_t (Z_t - M_t^*) \right] - \mathbb{E} \left[\max_k (Z_{t_k} - M_{t_k}^*) \right]. \tag{9}
\end{aligned}$$

From the definition of $\lambda_{p,n}^\#$, we know that $V_{p,n}(\lambda_{p,n}^\#) \leq V_{p,n}(\lambda_{p,n}^*)$. Then,

$$\begin{aligned}
0 \leq V_{p,n}(\lambda_{p,n}^\#) - \mathbb{E} \left[\max_k (Z_{t_k} - M_{t_k}^*) \right] &\leq V_{p,n}(\lambda_{p,n}^*) - \mathbb{E} \left[\max_k (Z_{t_k} - M_{t_k}^*) \right] \\
&= \mathbb{E} \left[\max_k (Z_{t_k} - \mathbb{E}[C_{p,n}(\lambda_{p,n}^*) | \mathcal{F}_{t_k}]) - \max_k (Z_{t_k} - M_{t_k}^*) \right] \\
&\leq \mathbb{E} \left[\max_k |M_{t_k}^* - \mathbb{E}[C_{p,n}(\lambda_{p,n}^*) | \mathcal{F}_{t_k}]| \right] \\
&\leq \mathbb{E} \left[\max_k \mathbb{E} \left[|M_T^* - C_{p,n}(\lambda_{p,n}^*)| | \mathcal{F}_{t_k} \right] \right] \\
&\leq \sqrt{\mathbb{E} \left[\max_k \mathbb{E} \left[|M_T^* - C_{p,n}(\lambda_{p,n}^*)|^2 | \mathcal{F}_{t_k} \right] \right]} \\
&\leq 2 \|M_T^* - C_{p,n}(M_T^*)\|_2
\end{aligned}$$

where the last upper-bound ensues from Doob's inequality. We combine this inequality with (9) and note that $\mathbb{E} [\max_k (Z_{t_k} - M_{t_k}^*)] \geq \hat{U}_0^n$ to obtain

$$|V_{p,n}(\lambda_{p,n}^\#) - U_0| \leq 2 \|M_T^* - C_{p,n}(M_T^*)\|_2 + U_0 - \hat{U}_0^n.$$

The convergence of $\|M_T^* - C_{p,n}(M_T^*)\|_2$ to 0 when p, n go to infinity ensues from [Nualart, 1998, Theorem 1.1.1, Proposition 1.1.1]. ■

From this proof, we can deduce the following corollary for the case of Bermudan options.

Corollary 3.3 *Assume the discounted payoff $(Z_{t_k})_k$ of the Bermudan option is \mathcal{G} -adapted. Then, $V_{p,n}(\lambda_{p,n}^\#)$ converges to the price \hat{U}_0^n of the Bermudan option when p goes to infinity.*

By writing that $\|M_T^* - C_{p,n}(M_T^*)\|_2 \leq \|M_T^* - C_p(M_T^*)\|_2 + \|C_p(M_T^*) - C_{p,n}(M_T^*)\|_2$, we can split the effects of the order of the chaos approximation from the impact of the truncation of the basis of $L^2([0, T])$. We apply [Geiss and Labart, 2017, Lemma 2.4] to handle the error w.r.t p and [Briand and Labart, 2014, Lemma 4.14] to handle the error due to n . Then, we obtain the following convergence rate.

Proposition 3.4 Assume $M_T^* \in \mathbb{D}^{m,2}$ for some $1 \leq m \leq p+1$, and for all $\ell \leq m$ and all $(t_1, \dots, t_\ell) \in [0, T]^\ell$ and $(s_1, \dots, s_\ell) \in [0, T]^\ell$

$$\forall |\alpha| = \ell, \quad |\mathbb{E}[D_{t_1, \dots, t_\ell} M_T^*] - \mathbb{E}[D_{s_1, \dots, s_\ell} M_T^*]| \leq K_\ell(|t_1 - s_1|^{\beta^*} + \dots + |t_\ell - s_\ell|^{\beta^*})$$

where $(K_\ell)_\ell$ is an increasing sequence of positive real numbers and β^* is a positive real constant. Then,

$$\|M_T^* - C_{p,n}(M_T^*)\|_2 \leq \frac{\|M_T^*\|_{\mathbb{D}^{m,2}}}{\sqrt{(p+1) \dots (p-m+2)}} + 2\sqrt{T(1+T)} e^T \left(\frac{T}{n}\right)^{\beta^*} K_p.$$

Proposition 3.5 Let $p \geq 1$. Assume that

$$\begin{aligned} \forall 1 \leq r \leq k \leq n, \forall F \mathcal{F}_{t_k} - \text{measurable}, F \in \mathcal{C}_{p-1,n}, F \neq 0, \exists q' \in \{1, \dots, d\} \text{ s.t.} \\ \mathbb{P}\left(\forall t \in]t_{r-1}, t_r], D_t^{q'} Z_{t_k} + F = 0 \mid Z_{t_k} > 0\right) = 0. \end{aligned} \quad (10)$$

Define the open set

$$\Lambda = \{(\lambda_\alpha)_\alpha \in \mathbb{R}^{A_{p,n}^{\otimes d}} : \forall r \in \{1, \dots, n\}, \exists (\alpha, q \in \{1, \dots, d\}) \text{ s.t. } \alpha_r^q \geq 1 \text{ and } \lambda_\alpha \neq 0\}.$$

Then, the function $V_{p,n}$ is differentiable on the set Λ and the gradient $\nabla V_{p,n}$ is given by

$$\nabla V_{p,n}(\lambda) = \mathbb{E} \left[\mathbb{E} \left[\hat{H}^{\otimes d}(G_1, \dots, G_n) \mid \mathcal{F}_{t_i} \right]_{\{i\}=\mathcal{I}(\lambda, Z, G)} \right].$$

with

$$\mathcal{I}(\lambda, Z, G) = \left\{ 0 \leq k \leq n : v_{p,n}(\lambda, k; Z, G) = \max_{\ell \leq n} v_{p,n}(\lambda, \ell; Z, G) \right\}.$$

The proof of this result being quite technical, we defer it to Appendix B.

We refer the reader to section 5.1 for a detailed discussion on which kinds of models and payoffs satisfy (10). Note that the complementary of the set Λ has Lebesgue measure 0.

3.3 Convergence of the Sample Average Approximation

For large enough m , $V_{p,n}^m$ inherits from the smoothness of $V_{p,n}$ and is in particular convex and a.s. differentiable at any point with no zero component. Then, we easily deduce from Proposition 3.1 that there exists $\lambda_{p,n}^m$ such that

$$V_{p,n}^m(\lambda_{p,n}^m) = \inf_{\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}, \lambda_0=0} V_{p,n}^m(\lambda).$$

The main difficulty in studying the convergence of $V_{p,n}^m(\lambda_{p,n}^m)$ when m goes to infinity comes from the non compactness of the set $\mathbb{R}^{A_{p,n}^{\otimes d}}$. To circumvent this difficulty, we adapt to non strictly convex problems the technique used in Jourdain and Lelong [2009].

Proposition 3.6 The sequence $(V_{p,n}^m(\lambda_{p,n}^m))_m$ converges a.s. to $V_{p,n}(\lambda_{p,n}^\#)$ when $m \rightarrow \infty$. Moreover, the distance between $\lambda_{p,n}^m$ and the convex set of minimizers in (5) converges to zero as m goes to infinity.

Remark 3.7 We saw in Section 3.2 that $V_{p,n}(\lambda_{p,n}^\sharp)$ is always greater than the price of the Bermudan option \hat{U}_0^n . However, this does not ensure that the sample average estimator $V_{p,n}^m(\lambda_{p,n}^m)$ is upper biased. If one only aims at computing an approximation of the option price, it is fine to directly use $V_{p,n}(\lambda_{p,n}^\sharp)$, but if the goal is to provide an upper-bound of the price then a second stage Monte Carlo is needed. From a fixed set of samples $(G^{(1)}, \dots, G^{(m)})$, compute $\lambda_{p,n}^m$ minimizing $\lambda \mapsto V_{p,n}^m$. Then draw a second set of samples $(\bar{G}^{(1)}, \dots, \bar{G}^{(m)})$ independent of the first one and compute

$$\bar{V}_{p,n}^m(\lambda_{p,n}^m) = \frac{1}{m} \sum_{i=1}^m \max_{0 \leq k \leq n} v_{p,n}(\lambda_{p,n}^m, k; \bar{Z}^{(i)}, \bar{G}^{(i)}).$$

Proof. The random function $\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}} \mapsto \max_{0 \leq k \leq n} v_{p,n}(\lambda, k; Z, G)$ is a.s. continuous. For $\Lambda > 0$,

$$\begin{aligned} & \sup_{|\lambda| \leq \Lambda} \max_{0 \leq k \leq n} v_{p,n}(\lambda, k; Z, G) \\ & \leq \max_{0 \leq k \leq n} Z_{t_k} + \sup_{|\lambda| \leq \Lambda} \max_{0 \leq k \leq n} \sum_{\alpha \in A_{p,n}^{\otimes d}} \lambda_\alpha \mathbb{E} \left[\hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) \middle| \mathcal{F}_{t_k} \right] \\ & \leq \max_{0 \leq k \leq n} Z_{t_k} + \Lambda \sup_{|\lambda| \leq 1} \max_{0 \leq k \leq n} \sum_{\alpha \in A_{p,n}^{\otimes d}} \lambda_\alpha \mathbb{E} \left[\hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) \middle| \mathcal{F}_{t_k} \right] \\ & \leq \max_{0 \leq k \leq n} Z_{t_k} + \Lambda \max_{0 \leq k \leq n} \sum_{\alpha \in A_{p,n}^{\otimes d}} \left| \mathbb{E} \left[\hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) \middle| \mathcal{F}_{t_k} \right] \right| \\ & \leq \max_{0 \leq k \leq n} Z_{t_k} + \Lambda \sum_{\alpha \in A_{p,n}^{\otimes d}} \max_{0 \leq k \leq n} \mathbb{E} \left[\left| \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) \right| \middle| \mathcal{F}_{t_k} \right] \\ & \leq \max_{0 \leq k \leq n} Z_{t_k} + \Lambda \sum_{\alpha \in A_{p,n}^{\otimes d}} \sum_{k=0}^n \mathbb{E} \left[\left| \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) \right| \middle| \mathcal{F}_{t_k} \right]. \end{aligned}$$

The right hand side of the above inequality is integrable. We apply [Rubinstein and Shapiro, 1993, Lemma A1 Chapter 2] to deduce that a.s. $V_{p,n}^m$ converges locally uniformly to $V_{p,n}$. From the proof of the Proposition 3.1, there exists $\Lambda > 0$ such that

$$\gamma = \inf_{|\lambda - \lambda_{p,n}^\sharp| \geq \Lambda} V_{p,n}(\lambda) - V_{p,n}(\lambda_{p,n}^\sharp) > 0.$$

The local uniform convergence of $V_{p,n}^m$ to $V_{p,n}$ ensures that

$$\exists m_\gamma \in \mathbb{N}^*, \forall m \geq m_\gamma, \forall \lambda \text{ s.t. } |\lambda - \lambda_{p,n}^\sharp| \leq \Lambda, \quad \left| V_{p,n}^m(\lambda) - V_{p,n}(\lambda) \right| \leq \frac{\gamma}{3}.$$

For $m \geq m_\gamma$ and λ such that $|\lambda - \lambda_{p,n}^\sharp| \geq \Lambda$, we deduce, using the convexity of $V_{p,n}^m$, that

$$\begin{aligned} & V_{p,n}^m(\lambda) - V_{p,n}^m(\lambda_{p,n}^\sharp) \\ & \geq \frac{|\lambda - \lambda_{p,n}^\sharp|}{\Lambda} \left\{ V_{p,n}^m \left(\lambda_{p,n}^\sharp + \Lambda \frac{\lambda - \lambda_{p,n}^\sharp}{|\lambda - \lambda_{p,n}^\sharp|} \right) - V_{p,n}^m(\lambda_{p,n}^\sharp) \right\} \\ & \geq \frac{|\lambda - \lambda_{p,n}^\sharp|}{\Lambda} \left\{ V_{p,n} \left(\lambda_{p,n}^\sharp + \Lambda \frac{\lambda - \lambda_{p,n}^\sharp}{|\lambda - \lambda_{p,n}^\sharp|} \right) - V_{p,n}(\lambda_{p,n}^\sharp) - \frac{2\gamma}{3} \right\} \geq \frac{\gamma}{3}. \end{aligned}$$

Since $V_{p,n}^m(\lambda_{p,n}^m) - V_{p,n}^m(\lambda_{p,n}^\#) \leq 0$, we conclude that the above inequality does not hold for $\lambda_{p,n}^m$, which proves that $|\lambda_{p,n}^m - \lambda_{p,n}^\#| < \Lambda$ for $m \geq m_\gamma$.

Hence, for $m \geq m_\gamma$, it is sufficient to minimize $V_{p,n}^m$ on the compact set $\{\lambda : |\lambda - \lambda_{p,n}^\#| \leq \Lambda\}$. Now, we can apply [Rubinstein and Shapiro, 1993, Theorem A1 of Chapter 2] to prove that $V_{p,n}^m(\lambda_{p,n}^m)$ converges to $V_{p,n}(\lambda_{p,n}^\#)$ a.s. when m goes to infinity. The second assertion of our proposition is discussed right after the proof of Theorem A1 in Rubinstein and Shapiro [1993]. \blacksquare

Although $V_{p,n}^m$ is not twice differentiable and the classical central limit theorem for sample average approximations cannot be applied, we can study the variance of $V_{p,n}^m(\lambda_{p,n}^m)$ and obtain some asymptotic bounds. Before stating our result, we introduce, for $\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}$, the notation $M_k(\lambda) = \mathbb{E}[C_{p,n}(\lambda) | \mathcal{F}_{t_k}]$ for $0 \leq k \leq n$. We write $M_k^{(i)}(\lambda)$ for the value computed using the sample $G^{(i)}$.

Proposition 3.8 *Assume $\lambda_{p,n}^\#$ is unique. Then,*

$$\frac{1}{m} \sum_{i=1}^m \left(\max_{0 \leq k \leq n} Z_{t_k}^{(i)} - M_k^{(i)}(\lambda_{p,n}^m) \right)^2 - V_{p,n}^m(\lambda_{p,n}^m)^2$$

is a convergent estimator of $\text{Var}(\max_{k \leq 0 \leq n} Z_{t_k} - M_k(\lambda_{p,n}^\#))$ and moreover if $\lambda_{p,n}^m$ is bounded, $\lim_{m \rightarrow \infty} m \text{Var}(V_{p,n}^m(\lambda_{p,n}^m)) = \text{Var}(\max_{k \leq 0 \leq n} Z_{t_k} - M_k(\lambda_{p,n}^\#))$.

Proof. We know that $V_{p,n}^m(\lambda_{p,n}^m)$ converges a.s. to $V_{p,n}(\lambda_{p,n}^\#)$. Following the beginning of the proof of Proposition 3.6, one can easily prove that a.s. the sequence of random functions $\zeta^m : \lambda \mapsto \zeta^m(\lambda) = \frac{1}{m} \sum_{i=1}^m \left(\max_{0 \leq k \leq n} Z_{t_k}^{(i)} - M_k^{(i)}(\lambda) \right)^2$ converges locally uniformly to the function $\lambda \mapsto \mathbb{E}[(\max_{0 \leq k \leq n} Z_{t_k} - M_k(\lambda))^2]$. We have already seen that for large enough m , we can assume to have solved the optimization problem under a compact constraint. Hence, we deduce that $\frac{1}{m} \sum_{i=1}^m \left(\max_{0 \leq k \leq n} Z_{t_k}^{(i)} - M_k^{(i)}(\lambda_{p,n}^m) \right)^2$ converges a.s. to $\mathbb{E}[(\max_{0 \leq k \leq n} Z_{t_k} - M_k(\lambda_{p,n}^\#))^2]$. This proves the first statement of the proposition.

As $\text{Var}(V_{p,n}^m(\lambda_{p,n}^\#)) = m^{-1} \text{Var}(\max_{k \leq 0 \leq n} Z_{t_k} - M_k(\lambda_{p,n}^\#))$, it is sufficient to compute

$$\begin{aligned} & \mathbb{E} \left[\left(V_{p,n}^m(\lambda_{p,n}^m) - V_{p,n}^m(\lambda_{p,n}^\#) \right)^2 \right] \\ & \leq \frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[\max_k \left| M_k^{(i)}(\lambda_{p,n}^m) - M_k^{(i)}(\lambda_{p,n}^\#) \right|^2 \right] \\ & \leq \frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[\left| \lambda_{p,n}^m - \lambda_{p,n}^\# \right|^2 \max_k \left| \mathbb{E} \left[\hat{H}^{\otimes d}(G_1^{(i)}, \dots, G_n^{(i)}) \mid \mathcal{F}_{t_k} \right] \right|^2 \right] \\ & \leq \frac{16}{9} \mathbb{E} \left[\left| \lambda_{p,n}^m - \lambda_{p,n}^\# \right|^4 \right]^{1/2} \mathbb{E} \left[\left| \hat{H}^{\otimes d}(G_1^{(i)}, \dots, G_n^{(i)}) \right|^4 \right]^{1/2} \end{aligned}$$

where we have used Cauchy Schwartz' inequality and Doob's maximal inequality. Then, we easily conclude that $V_{p,n}^m(\lambda_{p,n}^m) - V_{p,n}^m(\lambda_{p,n}^\#)$ converges to 0 in \mathbb{L}^2 if $\lambda_{p,n}^m$ is bounded. Hence, $\lim_{m \rightarrow \infty} \text{Var}(V_{p,n}^m(\lambda_{p,n}^m)) - \text{Var}(V_{p,n}^m(\lambda_{p,n}^\#)) = 0$. \blacksquare

Proposition 3.8 enables us to monitor the variance of our estimator online as for a standard Monte Carlo estimator. Even though the terms involved in $V_{p,n}^m(\lambda_{p,n}^m)$ are not independent, the classical variance estimator gives the right result. In practice, one should not feel concerned with the boundedness condition used in the proposition as we know from the proof of Proposition 3.6 that for large enough m we can impose a compactness constraint to the optimization problem without changing its result. Hence, one can pragmatically rely on the proposed variance estimator.

4 The algorithm

Any optimization algorithm requires to repeatedly compute $V_{p,n}^m$ and therefore the truncated chaos expansion, which becomes the most time consuming part of our approach as the dimension and/or p increase. A lot of computational time can be saved by considering slightly modified martingales, which only start the first time the option goes in the money.

4.1 An improved set of martingales

We define the first time the option goes in the money by

$$\tau_0 = \inf\{k \geq 0 : Z_{t_k} > 0\} \wedge n,$$

which is a \mathcal{F} -stopping time and becomes a \mathcal{G} -stopping time when the sequence $(Z_{t_k})_k$ is \mathcal{G} -adapted. To consider martingales only starting once the option has been in the money, we define

$$N_k(\lambda) = \sum_{\ell=1}^k (M_\ell(\lambda) - M_{\ell-1}(\lambda)) \mathbf{1}_{\ell-1 \geq \tau_0} = (M_k(\lambda) - M_{\tau_0}(\lambda)) \mathbf{1}_{k > \tau_0} = M_k(\lambda) - M_{k \wedge \tau_0}(\lambda)$$

We easily check that $N(\lambda)$ is a $(\mathcal{F}_{t_k})_{0 \leq k \leq n}$ -martingale. It is clear from the proof proposed by Rogers [2002] that in the dual price of a Bermudan option (see (1)) the maximum can be shrunk to the random interval $[\tau_0, n]$. Hence, it is sufficient to consider

$$\inf_{\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}, \lambda_0=0} \mathbb{E} \left[\max_{\tau_0 \leq k \leq n} (Z_{t_k} - M_k(\lambda)) \right].$$

Using Doob's stopping theorem, we have, for any fixed λ ,

$$\mathbb{E} \left[\max_{\tau_0 \leq k \leq n} (Z_{t_k} - M_k(\lambda)) \right] = \mathbb{E} \left[\max_{\tau_0 \leq k \leq n} (Z_{t_k} - (M_k(\lambda) - M_{\tau_0}(\lambda))) \right] = \mathbb{E} \left[\max_{\tau_0 \leq k \leq n} (Z_{t_k} - N_k(\lambda)) \right].$$

We deduce from this equality that minimizing over either set of martingales $M(\lambda)$ or $N(\lambda)$ leads to the same minimum value and that both problems share the same properties, which justifies why we did not take into account the in-the-money condition for the theoretical study. However, considering the set of martingales N^λ is far more efficient from a practical point of view.

In our numerical examples, we modify $V_{p,n}$ and $V_{p,n}^m$ to take into account this improvement and consider instead

$$\tilde{V}_{p,n}(\lambda) = \mathbb{E} \left[\max_{\tau_0 \leq k \leq n} (Z_{t_k} - N_k(\lambda)) \right] \quad \text{and} \quad \tilde{V}_{p,n}^m(\lambda) = \frac{1}{m} \sum_{i=1}^m \max_{\tau_0 \leq k \leq n} (Z_{t_k}^{(i)} - N_k^{(i)}(\lambda)).$$

The idea of using martingales starting from the first time the option goes in the money is actually owed to Rogers [2002]. Although he did not discuss it much, this was his choice in the examples he treated.

4.2 Our implementation of the algorithm

To practically compute the infimum of $\tilde{V}_{p,n}^m$, we advise to use a gradient descent algorithm, see Algorithm 4.1. The efficiency of such an approach mainly depends on the computation of the descend direction. When the problem is not twice differentiable, the gradient at the current point is used as a descent direction but it often needs to be scaled, which makes the choice of the step size α_ℓ a burning issue to ensure a fast numerical convergence. We refer to Boyd et al. [2003] for a comprehensive survey of several step size rules. After many tests, we found that the step size rule proposed by Polyak [1987] was the best in our context

$$\alpha_\ell = \frac{\tilde{V}_{p,n}^m(x_\ell) - v^\sharp}{\|\nabla \tilde{V}_{p,n}^m(x_\ell)\|^2}$$

where v^\sharp is the price of the American option we are looking for. In practice, we use the price of the associated European option instead of v^\sharp , which makes α_ℓ too large and explains the need of the magnitude factor γ . The value of the European price does not need to be very accurate. A decent and fast approximation can be computed with a few thousand samples within few seconds no matter the dimension of the problem.

```

1 Generate  $(G^{(1)}, Z^{(1)}), \dots, (G^{(m)}, Z^{(m)})$   $m$  i.i.d. samples following the law of  $(Z, G)$ 
2  $x_0 \leftarrow 0 \in \mathbb{R}^{A_{p,n}^{\otimes d}}$ 
3  $\ell \leftarrow 0, \gamma \leftarrow 1, d_0 \leftarrow 0, v_0 \leftarrow \infty$ 
4 while True do
5   Compute  $v_{\ell+1/2} \leftarrow \tilde{V}_{p,n}^m(x_\ell - \gamma \alpha_\ell d_\ell)$ 
6   if  $v_{\ell+1/2} < v_\ell$  then
7      $x_{\ell+1} \leftarrow x_\ell - \gamma \alpha_\ell d_\ell$ 
8      $v_{\ell+1} \leftarrow v_{\ell+1/2}$ 
9      $d_{\ell+1} \leftarrow \nabla \tilde{V}_{p,n}^m(x_{\ell+1})$ 
10    if  $\frac{|v_{\ell+1} - v_\ell|}{v_\ell} \leq \varepsilon$  then return
11  else
12     $\gamma \leftarrow \gamma/2$ 
13  end
14 end

```

Algorithm 4.1: Sample Average Approximation of the dual price

To better understand how this algorithm works, it is important to note that as $N(\lambda)$ linearly depends on λ , $N(\lambda) = \lambda \cdot \nabla_\lambda N(\lambda)$ and therefore both the value function and its gradient are computed at the same time without extra cost. So, $\nabla \tilde{V}_{p,n}^m(x_{\ell+1})$ is not actually computed on line 9 but at the same time as $v_{\ell+1/2}$ on line 5.

The HPC approach. Our method targets large problems with as many as several thousands of components for λ . This requires to design a scalable algorithm capable of making

```

1 In parallel do
2   | Generate  $(G^{(1)}, Z^{(1)}), \dots, (G^{(m)}, Z^{(m)})$   $m$  i.i.d. samples following the law of  $(Z, G)$ 
3 end
4  $x_0 \leftarrow 0 \in \mathbb{R}^{A_{p,n}^{\otimes d}}$ 
5  $\ell \leftarrow 0, \gamma \leftarrow 1, d_0 \leftarrow 0, v_0 \leftarrow \infty$ 
6 while True do
7   | Broadcast  $x_\ell, d_\ell, \gamma, \alpha_\ell$ 
8   | In parallel do
9     | Compute  $\max_{\tau_0 \leq k \leq n} (Z_{t_k}^{(i)} - N_k^{(i)}(x_\ell - \gamma \alpha_\ell d_\ell))$  for  $i = 1, \dots, m$ 
10  | end
11  | Make a reduction of the above contributions to obtain  $\tilde{V}_{p,n}^m(x_\ell - \gamma \alpha_\ell d_\ell)$  and
    |  $\nabla \tilde{V}_{p,n}^m(x_\ell - \gamma \alpha_\ell d_\ell)$ 
12  |  $v_{\ell+1/2} \leftarrow \tilde{V}_{p,n}^m(x_\ell - \gamma \alpha_\ell d_\ell)$ 
13  | if  $v_{\ell+1/2} < v_\ell$  then
14    |  $x_{\ell+1} \leftarrow x_\ell - \gamma \alpha_\ell d_\ell$ 
15    |  $v_{\ell+1} \leftarrow v_{\ell+1/2}$ 
16    |  $d_{\ell+1} \leftarrow \nabla \tilde{V}_{p,n}^m(x_{\ell+1})$ 
17    | if  $\frac{|v_{\ell+1} - v_\ell|}{v_\ell} \leq \varepsilon$  then return
18  | else
19    |  $\gamma \leftarrow \gamma/2$ 
20  | end
21 end

```

Algorithm 4.2: Parallel implementation of the Sample Average Approximation of the dual price

the most of cluster architectures with hundreds of nodes. At each iteration, the computation of $\tilde{V}_{p,n}^m$ and $\nabla \tilde{V}_{p,n}^m$ is nothing but a standard Monte Carlo method and it inherits from its embarrassingly parallel nature.

A parallel algorithm for distributed memory systems based on the master/slave paradigm is proposed in Algorithm 4.2. At the beginning, each process samples a bunch of the m paths (lines 1–3). Then, at each iteration the master process broadcasts the value of $d_\ell, x_\ell, \alpha_\ell$ and γ (line 7 of Algorithm 4.1). With these new values, each process computes its contribution to $\tilde{V}_{p,n}^m(x_\ell - \gamma \alpha_\ell d_\ell)$ and $\nabla \tilde{V}_{p,n}^m(x_\ell - \gamma \alpha_\ell d_\ell)$ (lines 8–9) and the Monte Carlo summations are obtained by two simple reductions (line 11). Then, the master process tests whether the move is admissible and updates the parameter for the next iteration or returns the solution if the algorithm is not moving enough anymore. This part carried out by the master process is very fast compared to the rest of the code and we dare say that there is no centralized computation in our algorithm. Moreover the communications are reduced to four broadcasts, which guarantees an almost perfect very good scalability. The number of communications is monitored by the number of function evaluations, which remains quite small (between 10 and 20). We study the efficiency of our algorithm on a few examples at the end of Section 5.

Study of the complexity. Most of the computational time is spent computing the martingale part; remember that the cardinality of $\mathcal{C}_{p,n}$ is given by $\binom{nd+p}{nd} = \frac{(nd+p) \dots (nd+1)}{p!}$. Using

martingales only starting once the option has been in the money enables us to only compute the martingale part on paths going in the money strictly before maturity time. Depending on the product, this may allow for saving a lot of computational time. The complexity of one iteration of the loop line 3 in Algorithm 4.1 is proportional to

$$\#\{\text{paths in the money strictly before } T\} \times \binom{nd+p}{nd}.$$

The payoffs are computed once and for all before starting the descent algorithm. It is worth noting that its computational cost becomes negligible compared to the optimization part when the dimension of the model or the number of dates increase, the most demanding computation being the evaluation of the martingale decomposition.

5 Applications

5.1 Some frameworks satisfying the assumption of Proposition 3.5

Let $(r_t)_t$ be the instantaneous interest rate supposed to be deterministic.

5.1.1 A put basket option in the multi-dimensional Black Scholes model

The d -dimensional Black Scholes model writes for $j \in \{1, \dots, d\}$

$$dS_t^j = S_t^j((r_t - \delta^j)dt + \sigma^j L_j dB_t)$$

where W is a Brownian motion with values in \mathbb{R}^d , $\sigma_t = (\sigma_t^1, \dots, \sigma_t^d)$ is the vector of volatilities, assumed to be deterministic and positive at all times, $\delta = (\delta^1, \dots, \delta^d)$ is the vector of instantaneous dividend rates and L_j is the j -th row of the matrix L defined as a square root of the correlation matrix Γ , ie. $\Gamma = LL'$. Moreover, we assume that L is lower triangular. Clearly, for every t , the random vector S_t is an element of $\mathbb{D}^{1,2}$.

The payoff of the put basket option writes as $\phi(S_t) = \left(K - \sum_{i=1}^d \omega^i S_t^i\right)_+$ where $\omega = (\omega^1, \dots, \omega^d)$ is a vector of real valued weights. The function ϕ is Lipschitz continuous and hence $\phi(S_t) \in \mathbb{D}^{1,2}$ for all t . Moreover, for $s \leq t$ and $q \in \{1, \dots, d\}$, we have on the set $\{\phi(S_t) > 0\}$

$$D_s^q \phi(S_t) = \sum_{j=1}^d \omega^j S_t^j \sigma^j L_{j,q}.$$

In particular for $q = d$, we get $D_s^d \phi(S_t) = \omega^d S_t^d \sigma^d L_{d,d}$.

Let $1 \leq k \leq n$ and F be a non zero and \mathcal{F}_{t_k} -measurable element of $\mathcal{C}_{p-1,n}$, ie.

$$F = \sum_{\alpha \in A_{p-1,n}^{\otimes d,k}} \lambda_\alpha \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n)$$

for some $\lambda \in \mathbb{R}^{A_{p,n}^{\otimes d}}$. Let $1 \leq r \leq k$.

$$\begin{aligned}
& \mathbb{P} \left(\forall t \in]t_{r-1}, t_r], D_t^d \phi(S_{t_k}) + F = 0 \mid \phi(S_{t_k}) > 0 \right) \\
&= \mathbb{P} \left(\forall t \in]t_{r-1}, t_r], \omega^d S_{t_k}^d \sigma_t^d L_{d,d} + F = 0 \mid \phi(S_{t_k}) > 0 \right) \\
&\leq \frac{\mathbb{P} \left(\forall t \in]t_{r-1}, t_r], \omega^d S_{t_k}^d \sigma_t^d L_{d,d} + F = 0 \right)}{\mathbb{P}(\phi(S_{t_k}) > 0)}. \tag{11}
\end{aligned}$$

If $p = 1$, then F is a deterministic non zero constant. In this case, the numerator vanishes because $S_{t_k}^d$ has a density. Assume $p \geq 2$, then F is a multivariate polynomial with global degree $p - 1 \geq 1$. Then we can find $\ell \in \{1, \dots, k\}$, $q \in \{1, \dots, d\}$ and α such that $\alpha_\ell^q \geq 1$ and $\lambda_\alpha \neq 0$. Let $\hat{\mathcal{G}}$ be the sigma algebra generated by $(G_i^j, 1 \leq i \leq k, 1 \leq j \leq d, (i, j) \neq (\ell, q))$.

$$\mathbb{P} \left(\forall t \in]t_{r-1}, t_r], \omega^d S_t^d \sigma_t^d L_{d,d} + F = 0 \right) = \mathbb{E} \left[\mathbb{P} \left(\forall t \in]t_{r-1}, t_r], \omega^d S_t^d \sigma_t^d L_{d,d} + F = 0 \mid \hat{\mathcal{G}} \right) \right].$$

Conditioning on $\hat{\mathcal{G}}$, the random variable $\omega^d S_t^d \sigma_t^d L_{d,d} + F$ only depends on G_ℓ^q . Consider the algebraic equation for $x \in \mathbb{R}$

$$a e^{bx+c} = P(x) \tag{12}$$

where $(a, b, c) \in \mathbb{R}^3, a \neq 0, b \neq 0$ and P is polynomial with degree $p - 1 \geq 1$. Let $f(x) = a e^{bx+c} - P(x)$, $f^{(p)}(x) = ab^p e^{bx+c}$. Clearly, $f^{(p)}$ never vanishes, which ensures that f has at most p different roots. Hence, we deduce that for any $t \in]t_{r-1}, t_r]$, $\mathbb{P} \left(\omega^d S_t^d \sigma_t^d L_{d,d} + F = 0 \mid \hat{\mathcal{G}} \right) = 0$. Combining this result along with (11) proves that Equation (10) holds in this setting.

5.1.2 A put option on the minimum of a basket in the multi-dimensional Black Scholes model

We use the notation of the previous example. The payoff of the put option on the minimum of d assets write $\phi(S_t) = (K - \min_j(S_t^j))_+$. One can prove by induction on d that the function $x \in \mathbb{R}^d \mapsto \min_j(x^j)$ is 1-Lipschitz for the 1-norm on \mathbb{R}^d . Hence, as the positive part function is also Lipschitz, the payoff function ϕ is Lipschitz. Then, [Nualart, 1998, Proposition 1.2.4] yields that for all $t \in [0, T]$, $\phi(S_t) \in \mathbb{D}^{1,2}$ and for all $q \in \{1, \dots, d\}$,

$$D^q(\phi(S_t)) = \sum_{j=1}^d \partial_{x^j} \phi(S_t) D^q(S_t^j) = \sum_{j=1}^d \partial_{x^j} \phi(S_t) S_t^j \sigma^j L_{j,q}.$$

With our choice for the matrix L ,

$$D^d(\phi(S_t)) = \partial_{x^d} \phi(S_t) S_t^d \sigma^d L_{d,d} = -S_t^d \sigma^d L_{d,d} \mathbf{1}_{\phi(S_t) > 0} \mathbf{1}_{\min_j(S_t^j) = S_t^d}.$$

Let $1 \leq k \leq n$ and F be a non zero and \mathcal{F}_{t_k} -measurable element of $\mathcal{C}_{p-1,n}$. For $1 \leq r \leq k$,

$$\begin{aligned}
& \mathbb{P} \left(\forall t \in]t_{r-1}, t_r], D_t^d \phi(S_{t_k}) + F = 0 \mid \phi(S_{t_k}) > 0 \right) \\
&= \mathbb{P} \left(\forall t \in]t_{r-1}, t_r], -S_t^d \sigma^d L_{d,d} + F = 0 \mid \phi(S_{t_k}) > 0, \min_j(S_t^j) = S_t^d \right) \mathbb{P} \left(\min_j(S_t^j) = S_t^d \right) \\
&\quad + \mathbb{P} \left(\forall t \in]t_{r-1}, t_r], F = 0 \mid \phi(S_{t_k}) > 0, \min_j(S_t^j) \neq S_t^d \right) \mathbb{P} \left(\min_j(S_t^j) \neq S_t^d \right)
\end{aligned}$$

Clearly, the second term in the above sum is zero as F has a density. Hence,

$$\mathbb{P}\left(\forall t \in]t_{r-1}, t_r], D_t^d \phi(S_{t_k}) + F = 0 \mid \phi(S_{t_k}) > 0\right) \leq \frac{\mathbb{P}\left(\forall t \in]t_{r-1}, t_r], -S_t^d \sigma^d L_{d,d} + F = 0\right)}{\mathbb{P}(\phi(S_{t_k}) > 0)}.$$

We conclude as in the case of the put basket option.

5.1.3 A put option in the Heston model

The Heston model can be written

$$\begin{aligned} dS_t &= S_t(r_t dt + \sqrt{\sigma_t}(\rho dW_t^1 + \sqrt{1 - \rho^2} dW_t^2)) \\ d\sigma_t &= \kappa(\theta - \sigma_t)dt + \xi \sqrt{\sigma_t} dW_t^1. \end{aligned}$$

For $s \leq t$, $D_s^2 S_t = S_t \sqrt{1 - \rho^2} \sqrt{\sigma_t}$. Conditionally on W^1 , $D_s^2 S_t$ writes as $a e^{bW_t^2 + c}$ and we can unfold the same reasoning as after (12).

5.2 Numerical experiments

In this part, we present results obtained from a sequential implementation of our approach as described in Algorithm 4.1. The algorithm has been implemented in C++ and crucially relies on the PNL library (see Lelong [2007-2017]) especially for its generic and very efficient implementation of multivariate polynomials. The computations are run on a standard laptop with an Intel Core i5 processor 2.9 Ghz. For each experiment, we report the estimator $\tilde{V}_{p,n}^m(\lambda_{p,n}^\#)$ obtained using Algorithm 4.1 along with its computational time and standard deviation. For people focusing on providing an upper biased estimator of the price, we also report the value of the second stage independent Monte Carlo $\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$ along with its standard deviation and computational time. The computational time reported for $\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$ takes into account both the resolution of the optimization problem to get $\lambda_{p,n}^\#$ and the independent second stage Monte Carlo. In the analyses of the numerical experiments, we refer to the estimator $V_{p,n}^m(\lambda_{p,n}^\#)$ as the *price* and to the estimator $\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$ as the *upper price*.

5.2.1 Examples in the Black Scholes models

We consider the d -dimensional Black Scholes model as presented in Section 5.1.1. For the sake of simplicity in choosing the parameters, we have decided to use the same correlation between all the assets, which amounts to considering the following simple structure for Γ .

$$\Gamma = \begin{pmatrix} 1 & \rho & \dots & \rho \\ \rho & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \rho \\ \rho & \dots & \rho & 1 \end{pmatrix}$$

where $\rho \in]-1/(d-1), 1]$ to ensure that Γ is positive definite.

A basket option in the Black–Scholes model. We consider a put option on several assets as presented in Section 5.1.1. We report in Table 1 the price obtained with our approach for $m = 20,000$. The last column *reference price* corresponds to the prices reported in Schoenmakers et al. [2013] on the same examples. These *reference* prices were obtained within a few minutes according to the authors whereas here we manage to get similar values within a few seconds. We can see that a second order chaos expansion, $p = 2$, already gives very accurate prices within a few tenths of a second for a 5–dimensional problem with 6 dates, which proves the impressive efficiency of our approach. As expected, the upper price obtained by an independent second stage Monte Carlo is always larger than the price directly obtained from the SAA estimator and can be taken for granted as an upper bound of the true price. If we compare the computational times reported for both prices, we note that the independent second stage Monte Carlo does not increase the computational times by a great deal.

p	n	S_0	$\tilde{V}_{p,n}^m(\lambda_{p,n}^\#)$	Stdev	time	$\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$	Stdev	time	ref. price
2	3	100	2.27	0.029	0.23	2.29	0.02	0.33	2.17
3	3	100	2.23	0.025	0.9	2.25	0.02	1.5	2.17
2	3	110	0.56	0.014	0.07	0.57	0.01	0.2	0.55
3	3	110	0.53	0.012	0.5	0.55	0.01	1	0.55
2	6	100	2.60	0.021	1.28	2.62	0.017	1.7	2.43
3	6	100	2.42	0.021	14	2.52	0.01	17.5	2.43
2	6	110	0.61	0.012	0.33	0.64	0.01	0.6	0.61
3	6	110	0.57	0.008	10	0.64	0.01	13.6	0.61

Table 1: Prices for the put basket option with parameters $T = 3$, $r = 0.05$, $K = 100$, $\rho = 0$, $\sigma^j = 0.2$, $\delta^j = 0$, $d = 5$, $\omega^j = 1/d$.

A call on the maximum of d assets in the Black–Scholes model. We consider a call option on the maximum of d assets in the Black Scholes model. As in the previous example, the last column *reference price* corresponds to the prices reported in Schoenmakers et al. [2013] on the same examples. With no surprise, the computational times reported in Table 2

d	p	m	S_0	$\tilde{V}_{p,n}^m(\lambda_{p,n}^\#)$	Stdev	time	$\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$	Stdev	time	ref. price
2	2	20,000	90	9.98	0.07	0.4	10.05	0.05	0.5	8.15
2	3	20,000	90	8.5	0.05	4.1	8.6	0.02	7.1	8.15
2	2	20,000	100	16.2	0.06	0.54	16.3	0.05	0.7	14.01
2	3	20,000	100	14.4	0.06	5.6	15	0.05	6.5	14.01
5	2	20,000	90	20.2	0.09	4	21.2	0.07	5	16.77
5	3	40,000	90	16.3	0.05	210	20.13	0.1	225	16.77
5	2	20,000	100	30.7	0.09	3.4	31.8	0.07	4.5	26.34
5	3	40,000	100	26.0	0.05	207	29	0.1	230	26.34

Table 2: Prices for the call option on the maximum of d assets with parameters $T = 3$, $r = 0.05$, $K = 100$, $\rho = 0$, $\sigma^j = 0.2$, $\delta^j = 0.1$, $n = 9$.

increase exponentially with the dimension $n \times d$ and the degree p . Whereas a second order expansion provides very accurate results for the basket option, it only gives a rough upper estimation for the call option on the maximum of d assets. Considering a third order expansion $p = 3$ takes far longer but enables us to get much tighter prices. For the cases $p = 3$ and $n = 6$, the prices obtained are smaller than the reference prices, which might highlight some over fitting phenomenon. One way of avoiding this is to consider the upper price given the independent second stage Monte Carlo, which always has a positive bias. Actually, we notice from the experiments that when over fitting occurs in the minimization problem, the upper price shows an even larger bias.

A geometric basket option in the Black–Scholes model Benchmarking a new method on high dimensional products becomes hardly feasible as almost no high dimensional American options can be priced accurately in a reasonable time. An exception to this is the geometric option with payoff $(K - (\prod_{j=1}^d S_t^j)^{1/d})_+$ for the put option. Easy calculations show that the price of this d –dimensional option equals the one of the 1–dimensional option with parameters

$$\hat{S}_0 = \left(\prod_{j=1}^d S_0^j \right)^{1/d} ; \quad \hat{\sigma} = \frac{1}{d} \sqrt{\sum_{i,j} \sigma^i \sigma^j \Gamma_{ij}}; \quad \hat{\delta} = \frac{1}{d} \sum_{j=1}^d \left(\delta^j + \frac{1}{2} (\sigma^j)^2 \right) - \frac{1}{2} (\hat{\sigma})^2.$$

Table 3 summarizes the corresponding values used in the examples.

d	S_0	σ	ρ	\hat{S}_0	$\hat{\sigma}$	$\hat{\delta}$
2	100	0.2	0	100	0.14	0.01
10	100	0.3	0.1	100	0.131	0.036
40	100	0.3	0.1	100	0.105	0.039

Table 3: Correspondence table for the parameters of the geometric options with $\delta^j = 0$.

d	σ^j	ρ	p	m	$\tilde{V}_{p,n}^m(\lambda_{p,n}^\sharp)$	Stdev	time	$\bar{V}_{p,n}^m(\lambda_{p,n}^\sharp)$	Stdev	time	1d price
2	0.2	0	2	5000	4.32	0.04	0.18	4.42	0.03	0.2	4.20
2	0.2	0	3	15000	4.29	0.02	4.23	4.40	0.03	4.7	4.20
10	0.3	0.1	1	5000	5.50	0.06	0.12	5.55	0.04	0.2	4.60
10	0.3	0.1	2	20000	4.55	0.02	17	4.87	0.03	18	4.60
40	0.3	0.1	1	10000	4.4	0.03	1.4	4.51	0.03	1.7	3.69
40	0.3	0.1	2	20000	3.61	0.02	170	5.37	0.04	190	3.69

Table 4: Prices for the geometric basket put option with parameters $T = 1$, $r = 0.0488$ (it corresponds to a 5% annual interest rate), $K = 100$, $\delta^j = 0$, $n = 9$.

The 1– d price is computed using a tree method with several thousand steps. We can see in Table 4 that a second order approximation gives very accurate result within a few seconds for an option with 10 underlying assets, which proves the efficiency of our approach. We cannot beat the curse of dimensionality, which significantly slows down the algorithm for very large

problems. For an option on 40 assets, we obtain a price up to a 3% relative error within 3 minutes which is already very fast for such a high dimensional problem. The number of terms involved in the chaos expansion can become very large: for $d = 40$ and $p = 2$, there are 65340 elements in $\mathcal{C}_{p,n}$. Even though we are not working in a linear algebra framework, it is advisable to ensure that the number of samples m used in the sample average approximation is larger than the number of free parameters in the optimization problem. When m becomes too small, we may face an over-fitting phenomenon as the number of parameters is far too large compared to the information contained in the sample average approximation. This probably explains why the price obtained for $p = 2$, $d = 40$ and $m = 40$ is slightly smaller than the true price. This is even confirmed by the upper price estimator, which is far too large. We have run more experiments for this 40 dimensional problem to analyse the influence of the number of samples m on the sample average approximation $\tilde{V}_{p,n}^m(\lambda_{p,n}^\#)$ but also on the independent resimulation $\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$. Figure 1 shows the evolution of the two estimators when m increases has been obtained from running the parallel implementation on 128 cores. As one may have expected, the gap between the two estimators significantly decreases when m increases. Actually a large gap between $\tilde{V}_{p,n}^m(\lambda_{p,n}^\#)$ and $\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$ is a typical sign of over-fitting caused because the number of samples m is too small compared to the number of degrees of freedom in the chaos expansion. We can also from this experiment that when m is large enough, the sample average approximation is larger than the true price computed using the $1d$ simplification. To avoid any over-fitting issues, one definitely needs to ensure that enough samples are used compared to the number of elements in $\mathcal{C}_{p,n}$ given by $\binom{nd+p}{nd}$. In the example of Figure 1, $\binom{362}{360} = 65341$ and the sample average approximation price becomes larger than the true price for values of m greater than 5 times the size of $\mathcal{C}_{p,n}$.

5.2.2 Scalability of the parallel algorithm

We consider the 40-dimensional geometric put option studied in Table 4 with $p = 2$ and test the scalability of our parallel implementation for $m = 200,000$. The tests are run on a BullX DLC supercomputer containing 190 nodes for a total of 3204 CPU cores. We report in Table 5 the results of our scalability study using from 1 to 512 cores. Despite the two levels of parallelism available on this supercomputer, we have used a pure MPI implementation without any reference to multithread programming. We could probably have improved the efficiency a bit using two levels of parallelism, but the results are already so much convincing and do not justify the need of a two level approach, which makes the implementation more delicate. The sequential Algorithm runs within one hour and a quarter whereas using 512 cores we manage to get the computational time down to a dozen of seconds, which corresponds to a 0.6 efficiency. Considering the so short wall time required by the run on 512 cores, keeping the efficiency at this level represents a great achievement. Note that with 128 cores, the code runs within a minute with an efficiency of three quarters. These experiments prove the impressive scalability of our algorithm.

6 Conclusion

We have proposed a purely dual algorithm to compute the price of American or Bermudan options using some stochastic optimization tools. The starting point of our algorithm is the use of Wiener chaos expansion to build a finite dimensional vector space of martingales.




Figure 1: Influence of the number of samples m on the prices $\hat{V}_{p,n}^m(\lambda_{p,n}^\#)$ and $\bar{V}_{p,n}^m(\lambda_{p,n}^\#)$ for the geometric put option studied in Table 4 with $d = 40$, $\sigma^j = 0.3$, $\rho = 0.1$ and $p = 2$.

Then, we rely on a sample average approximation to effectively optimize the coefficients of the expansion. Our algorithm is very fast: for problems up to dimension 5, a price is obtained within a few seconds, which is a tremendous improvement compared to existing purely dual methods. For higher dimensional problems, we can use a very scalable parallel algorithm to tackle very high dimensional problems (40 underlying assets). We can transparently deal with complex path-dependent payoffs without any extra computational cost. Even though our sample average estimator may not yield an upper bound as we cannot control its bias, it is easy and relatively cheap to run an independent second stage Monte Carlo using the optimizer from the first stage to obtain an upper biased estimator, which can indeed be used as an upper bound of the price. We restricted to a Brownian setting in this work, but our approach could easily be extended to jump diffusion models by introducing Poisson chaos expansion, which is linked to Charlier polynomials (see Geiss and Labart [2017]). We believe that our approach could be improved by cleverly reducing the number of terms in the chaos expansion, the computation of which centralizes most of the effort.

#processes	time (sec.)	efficiency
1	4365	1
2	2481	0.99
4	1362	0.90
16	282	0.84
32	272	0.75
64	87	0.78
128	52	0.73
256	34	0.69
512	10.7	0.59

Table 5: Scalability of Algorithm 4.2 on the 40-dimensional geometric put option described above with $T = 1$, $r = 0.0488$, $K = 100$, $\sigma^j = 0.3$, $\rho = 0.1$, $\delta^j = 0$, $n = 9$, $p = 2$.

A Wiener chaos expansion

A.1 Proof of Proposition 2.1

Taking the conditional expectation in Eq. (3) leads to

$$\mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}] = \sum_{\alpha \in A_{p,n}^{\otimes d}} \lambda_{\alpha} \left(\prod_{i=1}^k \prod_{j=1}^d H_{\alpha_i^j}(G_i^j) \right) \mathbb{E} \left[\prod_{i=k+1}^n \prod_{j=1}^d H_{\alpha_i^j}(G_i^j) | \mathcal{F}_{t_k} \right]. \quad (13)$$

Since the Brownian increments after time t_k are independent of \mathcal{F}_{t_k} and are independent of one another, $\mathbb{E} \left[\prod_{i=k+1}^n \prod_{j=1}^d H_{\alpha_i^j}(G_i^j) | \mathcal{F}_{t_k} \right] = \prod_{i=k+1}^n \prod_{j=1}^d \mathbb{E} \left[H_{\alpha_i^j}(G_i^j) \right]$, which is zero as soon as $\sum_{i=k+1}^n \sum_{j=1}^d \alpha_i^j > 0$. Hence, the sum in (13) is reduced to the sum over the set of multi-indices $\alpha \in A_{p,n}^{\otimes d}$ such that $\alpha_i^j = 0$ for all $i > k$ and $1 \leq j \leq d$, which is exactly the definition of the set $A_{p,n}^{\otimes d,k}$.

A.2 The Malliavin derivative of a truncated chaos expansion

Proposition A.1 *Let F be a real valued random variable in $L^2(\Omega, \mathcal{F}_T, \mathbb{P})$ and let $k \in \{1, \dots, n\}$ and $p \geq 1$. For $t > t_k$, $D_t \mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}] = 0$. For all $t \in [t_{r-1}, t_r]$ with $1 \leq r \leq k$, and $q = 1, \dots, d$,*

$$D_t^q \mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}] = \frac{1}{\sqrt{t_r - t_{r-1}}} \sum_{\alpha \in A_{p,n}^{\otimes d,k}, \alpha_r^q \geq 1} \lambda_{\alpha} \hat{H}_{\alpha - \mathbb{I}(r,q)}^{\otimes d}(G_1, \dots, G_n)$$

where $(\alpha - \mathbb{I}(r,q))_i^j = \alpha_i^j - \mathbf{1}_{j=q, i=r}$.

Remark A.2 *The conditional expectation preserves the nature of a chaos expansion. Similarly, the Malliavin derivative of a chaos expansion still writes as a chaos expansion and hence is a Hermite polynomial of Brownian increments. The roots of a non zero polynomial being a zero measure set and since the Brownian increments have a joined density, the Malliavin derivative of a chaos expansion is almost surely non zero as soon as one of the coefficients λ_{α} is non zero for $\alpha \in A_{p,n}^{\otimes d,k}$ such that $\alpha_r^j \geq 1$ for some $j \in \{1, \dots, d\}$.*

Proof. From Proposition 2.1, we know that for all $1 \leq k \leq n$

$$\mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}] = \sum_{\alpha \in A_{p,n}^{\otimes d,k}} \lambda_\alpha \prod_{j=1}^d \hat{H}_{\alpha^j}(G_1^j, \dots, G_n^j).$$

Let $r \leq k$ and $t \in]t_{r-1}, t_r]$. Let $1 \leq q \leq d$. The chain rule for the Malliavin derivative yields

$$\begin{aligned} & D_t^q \mathbb{E}[C_{p,n}(F)|\mathcal{F}_{t_k}] \\ &= \sum_{\alpha \in A_{p,n}^{\otimes d,k}} \lambda_\alpha D_t^q \left(\prod_{j=1}^d \hat{H}_{\alpha^j}(G_1^j, \dots, G_n^j) \right) \\ &= \sum_{\alpha \in A_{p,n}^{\otimes d,k}} \lambda_\alpha \left(\prod_{j=1, j \neq q}^d \hat{H}_{\alpha^j}(G_1^j, \dots, G_n^j) \right) D_t^q \left(\hat{H}_{\alpha^q}(G_1^q, \dots, G_n^q) \right) \\ &= \frac{1}{\sqrt{t_r - t_{r-1}}} \sum_{\alpha \in A_{p,n}^{\otimes d,k}} \lambda_\alpha \left(\prod_{j=1, j \neq q}^d \hat{H}_{\alpha^j}(G_1^j, \dots, G_n^j) \right) \hat{H}_{\alpha^q - \mathbb{I}(r)}(G_1^q, \dots, G_n^q) \\ &= \frac{1}{\sqrt{t_r - t_{r-1}}} \sum_{\alpha \in A_{p,n}^{\otimes d,k}, \alpha_r^q \geq 1} \lambda_\alpha \hat{H}_{\alpha - \mathbb{I}(r,q)}(G_1, \dots, G_n). \quad \blacksquare \end{aligned}$$

B Differentiability of the optimization problem: proof of Proposition 3.5

We already know that the function $V_{p,n}$ is convex. Moreover, for all Z and G , the function $\lambda \mapsto \max_{k \leq n} v_{p,n}(\lambda, k, Z, G)$ has a subdifferential given by

$$\left\{ \sum_{i \in \mathcal{I}(\lambda, Z, G)} \beta_i \mathbb{E}[\hat{H}^{\otimes d}(G_1, \dots, G_n)|\mathcal{F}_{t_i}] : \beta_i \geq 0, \beta_i \mathcal{F}_T - \text{measurable s.t.} \sum_{i \in \mathcal{I}(\lambda, Z, G)} \beta_i = 1 \right\}$$

Then, the expression of the subdifferential $\partial V_{p,n}(\lambda)$ ensues from Bertsekas [1973].

$$\partial V_{p,n}(\lambda) = \left\{ \mathbb{E} \left[\sum_{i \in \mathcal{I}(\lambda, Z, G)} \beta_i \mathbb{E}[\hat{H}^{\otimes d}(G_1, \dots, G_n)|\mathcal{F}_{t_i}] \right] : \beta_i \geq 0, \beta_i \mathcal{F}_T - \text{meas.}, \sum_i \beta_i = 1 \right\}.$$

It is sufficient to prove for any $\lambda \in \Lambda$, the set $\mathcal{I}(\lambda, Z, G)$ is almost surely reduced to a single value as in this case the subdifferential $\partial V_{p,n}(\lambda)$ contains a unique element, which is then the gradient.

By the equality

$$\{\exists t_i \neq t_k ; v_{p,n}(\lambda, i, Z, G) = v_{p,n}(\lambda, k, Z, G)\} = \bigcup_{i < k \leq n} \{v_{p,n}(\lambda, i, Z, G) = v_{p,n}(\lambda, k, Z, G)\},$$

it is sufficient to prove that for any $i < k \leq n$, $\mathbb{P}(v_{p,n}(\lambda, i, Z, G) = v_{p,n}(\lambda, k, Z, G)) = 0$. Fix $i < k$ and set $X_\lambda = v_{p,n}(\lambda, k, Z, G) - v_{p,n}(\lambda, i, Z, G)$. According to [Nualart, 1998, Theorem 2.1.3], proving that $\|DX_\lambda\|_{L^2([0,T])} > 0$.a.s ensures that X_λ is absolutely continuous with

respect to the Lebesgue measure on \mathbb{R} and hence is almost surely non zero. Note that $\|DX_\lambda\|_{L^2([0,T])}^2 = \int_0^T |D_t X_\lambda|^2 dt \geq \int_{t_i}^{t_k} |D_t X_\lambda|^2 dt$.

For $t \in [0, T]$, and $1 \leq q \leq d$, the Malliavin derivative of X_λ is given by

$$\begin{aligned} D_t^q X_\lambda &= D_t^q(Z_{t_k} - Z_{t_i}) - D_t^q \left(\sum_{\alpha \in A_{p;n}^{\otimes d,k}} \lambda_\alpha \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) - \sum_{\alpha \in A_{p;n}^{\otimes d,i}} \lambda_\alpha \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) \right) \\ &= D_t^q(Z_{t_k} - Z_{t_i}) - D_t^q \left(\sum_{\alpha \in A_{p;n}^{\otimes d,i:k}} \lambda_\alpha \hat{H}_\alpha^{\otimes d}(G_1, \dots, G_n) \right) \end{aligned}$$

where for $i, k \in \{1, \dots, n\}$, with $i < k$, the set $A_{p,n}^{\otimes d,i:k} = A_{p,n}^{\otimes d,k} \setminus A_{p,n}^{\otimes d,i}$ writes

$$A_{p,n}^{\otimes d,i:k} = \left\{ \alpha \in (\mathbb{N}^n)^d : |\alpha|_1 \leq p, \text{ and } \forall 1 \leq j \leq d, \forall \ell \notin \{i+1, \dots, k\}, \alpha_\ell^j = 0 \right\}.$$

Clearly, w.p.1. $D_t^q X_\lambda = 0$ for all $t > t_k$. Hence,

$$\{D_t^q X_\lambda = 0 \forall t \in [0, T] \text{ a.e.}\} \subset \bigcap_{i < r \leq k} \{D_t^q X_\lambda = 0 \forall t \in [t_{r-1}, t_r] \text{ a.e.}\}.$$

From Proposition A.1, we can deduce that for $i < r \leq k$, and $t \in]t_{r-1}, t_r]$

$$D_t^q X_\lambda = D_t^q(Z_{t_k}) + \frac{1}{\sqrt{h}} \sum_{\alpha \in A_{p;n}^{\otimes d,i:k}, \alpha_r^q \geq 1} \lambda_\alpha \hat{H}_{\alpha - \mathbb{I}(r,q)}^{\otimes d}(G_1, \dots, G_n).$$

Using the locality of the operator D , we know that a.s $D_t^q(Z_{t_k}) = 0$ for all $t \in]t_{r-1}, t_r]$ on the set $\{Z_{t_k} = 0\}$. Hence, we can write for any pair $(q, q') \in \{1, \dots, d\}^2$

$$\begin{aligned} \mathbb{P}(\forall t \in]t_{r-1}, t_r], D_t X_\lambda = 0) &\leq \mathbb{P}\left(\frac{1}{\sqrt{h}} \sum_{\alpha \in A_{p;n}^{\otimes d,i:k}, \alpha_r^q \geq 1} \lambda_\alpha \hat{H}_{\alpha - \mathbb{I}(r,q)}^{\otimes d}(G_1, \dots, G_n) = 0, Z_{t_k} = 0\right) \\ &\quad + \mathbb{P}(\forall t \in]t_{r-1}, t_r], D_t^{q'} X_\lambda = 0 \mid Z_{t_k} > 0) \mathbb{P}(Z_{t_k} > 0). \end{aligned} \quad (14)$$

Assume $\lambda \in \Lambda$ and pick the corresponding value of q , then the chaos polynomial $\frac{1}{\sqrt{h}} \sum_{\alpha \in A_{p;n}^{\otimes d,i:k}, \alpha_r^q \geq 1} \lambda_\alpha \hat{H}_{\alpha - \mathbb{I}(r,q)}^{\otimes d}(G_1, \dots, G_n)$ is either a non zero constant if $p = 1$ or it has an absolutely continuous density thanks to Remark A.2. In both cases, it is a non zero element of $\mathcal{C}_{p-1,n}$ and

$$\mathbb{P}\left(\frac{1}{\sqrt{h}} \sum_{\alpha \in A_{p;n}^{\otimes d,i:k}, \alpha_r^q \geq 1} \lambda_\alpha \hat{H}_{\alpha - \mathbb{I}(r,q)}^{\otimes d}(G_1, \dots, G_n) = 0\right) = 0.$$

To treat the other term, we pick a $q' \in \{1, \dots, d\}$ as in (10) and it yields that

$$\mathbb{P}(\forall t \in]t_{r-1}, t_r], D_t^{q'} X_\lambda = 0 \mid Z_{t_k} > 0) = 0.$$

Hence, we deduce from the last two results and (14) that $\|DX_\lambda\|_{L^2([0,T])}^2 > 0$ a.s..

References

- L. Abbas-Turki, S. Vialle, B. Lapeyre, and P. Mercier. Pricing derivatives on graphics processing units using monte carlo simulation. *Concurrency and Computation: Practice and Experience*, 26(9):1679–1697, 2014.
- L. Andersen and M. Broadie. Primal-dual simulation algorithm for pricing multidimensional american options. *Management Science*, 50(9):1222–1234, 2004.
- V. Bally and G. Pages. A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli*, 9(6):1003–1049, 2003.
- D. Belomestny. Solving optimal stopping problems via empirical dual optimization. *Ann. Appl. Probab.*, 23(5):1988–2019, 2013.
- D. Belomestny, C. Bender, and J. Schoenmakers. True upper bounds for Bermudan products via non-nested Monte Carlo. *Math. Finance*, 19(1):53–71, 2009.
- D. Belomestny, R. Hildebrand, and J. Schoenmakers. Optimal stopping via pathwise dual empirical maximisation. *Applied Mathematics & Optimization*, Nov 2017.
- D. P. Bertsekas. Stochastic optimization problems with nondifferentiable cost functionals. *J. Optimization Theory Appl.*, 12:218–231, 1973.
- S. Boyd, L. Xiao, and A. Mutapcic. Subgradient methods. *lecture notes of EE392o, Stanford University, Autumn Quarter*, 2004:2004–2005, 2003.
- P. Briand and C. Labart. Simulation of BSDEs by Wiener Chaos Expansion. *Annals of Applied Probability*, 24(3):1129–1171, 2014.
- M. Broadie and P. Glasserman. A stochastic mesh method for pricing high-dimensional american options. *Journal of Computational Finance*, 7:35–72, 2004.
- J. F. Carriere. Valuation of the early-exercise price for options using simulations and non-parametric regression. *Insurance: mathematics and Economics*, 19(1):19–30, 1996.
- A. Carverhill and N. Webber. American options: theory and numerical analysis. *Options: recent advances in theory and practice*, pages 80–94, 1990.
- M. H. A. Davis and I. Karatzas. A deterministic approach to optimal stopping. In *Probability, statistics and optimisation*, Wiley Ser. Probab. Math. Statist. Probab. Math. Statist., pages 455–466. Wiley, Chichester, 1994.
- V. V. Desai, V. F. Farias, and C. C. Moallemi. Pathwise optimization for optimal stopping problems. *Management Science*, 58(12):2292–2308, 2012.
- V. Dung Doan, A. Gaiwad, M. Bossy, F. Baude, and I. Stokes-Rees. Parallel pricing algorithms for multidimensional bermudan/american options using Monte Carlo methods. *Mathematics and Computers in Simulation*, 81(3):568–577, 2010.
- C. Geiss and C. Labart. Simulation of BSDEs with jumps by wiener chaos expansion. *Stochastic Processes and their Applications*, 127(3), 2017.

- M. B. Haugh and L. Kogan. Pricing american options: a duality approach. *Operations Research*, 52(2):258–270, 2004.
- F. Jamshidian. The duality of optimal exercise and domineering claims: a Doob-Meyer decomposition approach to the Snell envelope. *Stochastics*, 79(1-2):27–60, 2007.
- B. Jourdain and J. Lelong. Robust Adaptive Importance Sampling for Normal Random Vectors. *Ann. Appl. Probab.*, 19(5):1687–1718, 2009.
- A. Kolodko and J. Schoenmakers. Upper bounds for bermudan style derivatives. *Monte Carlo Methods and Applications mcma*, 10(3-4):331–343, 2004.
- Lamberton. Brownian optimal stopping and random walks. *Applied Mathematics & Optimization*, 45(3):283–324, 2002.
- J. Lelong. Pnl : a free scientific library. <https://pnlnum.github.io/pnl>, 2007-2017.
- F. Longstaff and R. Schwartz. Valuing American options by simulation : A simple least-square approach. *Review of Financial Studies*, 14:113–147, 2001.
- D. Nualart. Analysis on Wiener space and anticipating stochastic calculus. In B. Springer-Verlag, editor, *Lectures on Probability Theory and Statistics (Saint-Flour, 1995)*, pages 123–227. 1998.
- B. T. Polyak. Introduction to optimization. *Optimization Software*, 1987.
- L. C. G. Rogers. Monte Carlo valuation of American options. *Math. Finance*, 12(3):271–286, 2002.
- L. C. G. Rogers. Dual valuation and hedging of Bermudan options. *SIAM J. Financial Math.*, 1:604–608, 2010.
- R. Y. Rubinstein and A. Shapiro. *Discrete event systems*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., Chichester, 1993. ISBN 0-471-93419-4. Sensitivity analysis and stochastic optimization by the score function method.
- J. Schoenmakers. *Robust Libor modelling and pricing of derivative products*. CRC Press, 2005.
- J. Schoenmakers, J. Zhang, and J. Huang. Optimal dual martingales, their analysis, and application to new algorithms for bermudan products. *SIAM Journal on Financial Mathematics*, 4(1):86–116, 2013.
- J. A. Tilley. Valuing american options in a path simulation model. *Transactions of the Society of Actuaries*, 45(83):104, 1993.
- J. Tsitsiklis and B. V. Roy. Regression methods for pricing complex American-style options. *IEEE Trans. Neural Netw.*, 12(4):694–703, 2001.