



HAL
open science

Pricing Bermudan options using regression trees/random forests

Zineb El Filali Ech-Chafiq, Pierre Henry Labordère, Jérôme Lelong

► **To cite this version:**

Zineb El Filali Ech-Chafiq, Pierre Henry Labordère, Jérôme Lelong. Pricing Bermudan options using regression trees/random forests. *SIAM Journal on Financial Mathematics*, 2023, 14 (4), pp.1113-1139. 10.1137/21M1460648 . hal-03436046v2

HAL Id: hal-03436046

<https://hal.science/hal-03436046v2>

Submitted on 22 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pricing Bermudan options using Regression Trees / Random Forests

Zineb El Filali Ech-chafiq * Pierre Henry Labordère † Jérôme Lelong ‡

June 22, 2023

Abstract

The value of an American option is the maximized value of the discounted cash flows from the option. At each time step, one needs to compare the immediate exercise value with the continuation value and decide to exercise as soon as the exercise value is strictly greater than the continuation value. We can formulate this problem as a dynamic programming equation, where the main difficulty comes from the computation of the conditional expectations representing the continuation values at each time step. In (Longstaff and Schwartz, 2001), these conditional expectations were estimated using regressions on a finite-dimensional vector space (typically a polynomial basis). In this paper, we follow the same algorithm; only the conditional expectations are estimated using Regression trees or Random forests. We discuss the convergence of the LS algorithm when the standard least squares regression is replaced by regression trees. Finally, we expose some numerical results with regression trees and random forests. The random forest algorithm gives excellent results in high dimensions.

keywords: Regression trees, Random forests, Bermudan options, Optimal stopping

1 Introduction

Bermudan options are very widespread in financial markets. Their valuation adds a challenge of optimal stopping determination in comparison to European options. Bermudan options offer the investor the possibility to exercise his option at any date of his choice among a certain number of dates prior to the option expiry, called exercise dates. Naturally, the option holder will have to find the most optimal date to exercise. To do so, at each exercise date, he will compare the payoff of the immediate exercise to the expected value of continuation of the option and decide to exercise only if the immediate exercise value is the highest. We can formulate this problem as a dynamic programming equation, where the main difficulty comes from the computation of the conditional expectation representing the expected continuation value of the option. Many papers have discussed this issue, starting with regression-based algorithms; see for example (Tsitsiklis and Van Roy, 1999) and (Carriere, 1996). Also, in this category falls the most commonly used method for pricing Bermudan options which is the Least Squares Method (LSM) presented by Longstaff and Schwarz in (Longstaff and Schwartz, 2001) where the conditional expectation is estimated by a least squares regression of the post realized payoffs from continuation on some basis functions of the state variables (usually polynomial functions). Alternatively, an approximately optimal but interpretable policy can be obtained as the solution to a sample average approximation problem, see Ciocan and Mišić (2022). Another class of algorithms focuses on quantization approaches, see

*Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France. Quantitative analyst at Natixis, Paris. zinebelfilaliechchafique@gmail.com

†Head of Quantitative Research cross asset, Natixis, 47 quai d'austerlitz 75013, Paris. CMAP, Ecole Polytechnique

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

for example (Bally et al., 2005). The algorithm consists in computing the conditional expectations by projecting the diffusion on some optimal grid. We also have a class of duality based methods that give an upper bound on the option value for a given exercise policy by adding a quantity that penalizes the incorrect exercise decisions made by the sub-optimal policy, see for example (Rogers, 2002), (Andersen and Broadie, 2004) and (Lelong, 2018). The last class of algorithms is based on machine learning techniques. Deep learning has been widely used to solve stochastic control problems Han and E (2016) or their PDE formulation Han et al. (2018). Neural networks can also be used to estimate the continuation values as in (Kohler et al., 2010; Han et al., 2018; Lapeyre and Lelong, 2021) or the boundary (Reppen et al., 2022). The continuation value can also be approximated using Gaussian processes as in (Ludkovski, 2018). Our solution falls in this last category of algorithms. We examine Bermudan options' prices when the continuation values' estimation is done using regression trees or random forests.

Let X, Y be two random variables with values in $[0, 1]^d$ and \mathbb{R} respectively. A regression tree approximates the conditional expectation $\mathbb{E}[Y|X]$ with a piecewise constant function. The tree is built recursively, generating a sequence of partitions of $[0, 1]^d$ that are finer and finer. The approximation value on each set in the partition can be seen as a terminal leaf of the tree. This algorithm is very simple and efficient. However, it can easily over-fit the data, which results in high generalization errors. To solve this issue, we use ensemble methods to aggregate multiple trees, which means that we create multiple trees and then combine them to produce improved results. We suggest using random forests (see (Breiman, 2001)). This method consists in averaging a combination of trees where each tree depends on a random vector sampled independently and identically for each tree in the forest. This vector will allow to differentiate the trees in the random forest and can be chosen in different ways. For example, one can draw for each tree a sub-sample of training from the global training data without replacement (this method is called bagging and is thoroughly studied in (Breiman, 1999)). A second method is random split selection, where at each node, the split is selected at random from among the K best splits, see (Dietterich, 2000). Other methods for aggregating regression trees into random forests can be found in the literature, see for example (Breiman, 2001) or (Ho, 1998).

The structure of the paper will be as follows. First, we present the regression trees algorithm and the algorithm of least squares using regression trees. Then, we proceed to present some convergence results for regression trees and study the convergence of the LS algorithm when regression trees are used to estimate the continuation values. We prove two results: the convergence of the regression solution for infinite sample size as the depth of the regression tree goes to infinity and the convergence of the empirical price as the number of samples go to infinity for a fixed depth. Then, we briefly talk about Random Forests before we finally study some numerical examples.

2 Regression trees

Let X be a r.v with values in $[0, 1]^d$ and Y a squared integrable real-valued r.v. We want to approximate the conditional expectation $\mathbb{E}[Y|X]$. Throughout this paper, we will consider for computational convenience that X has a density f_X in $[0, 1]^d$ w.r.t the Lebesgue measure. We assume given a training sample $D_M = \{(X_1, Y_1), \dots, (X_M, Y_M)\} \in [0, 1]^d \times \mathbb{R}$ where the (X_i, Y_i) 's are i.i.d random variables following the law of (X, Y) . An approximation using a regression tree consists in writing the conditional expectation as a piecewise constant function of X . Each domain where the function is constant can be seen as a terminal leaf of a tree. Let us start with the one-dimensional case ($d = 1$). Consider the function

$$m_M : y_l \in \mathbb{R}, y_r \in \mathbb{R}, x \in [0, 1] \mapsto \frac{1}{M} \sum_{i=1}^M (Y_i - y_l \mathbb{1}_{\{X_i \leq x\}} - y_r \mathbb{1}_{\{X_i > x\}})^2.$$

With probability $q > 0$, choose x as the midpoint $x = \frac{1}{2}$ and solve the minimization problem $\inf_{y_l, y_r} m_M(y_l, y_r, \frac{1}{2})$. With probability $0 < 1 - q < 1$, solve the minimization problem

$\inf_{y_l, y_r, x} m_M(y_l, y_r, x)$. For a fixed x^* , the optimal values of y_l and y_r are given by

$$y_r^* = \frac{\sum_{i=1}^M Y_i \mathbb{1}_{\{X_i > x^*\}}}{\sum_{i=1}^M \mathbb{1}_{\{X_i > x^*\}}}; \quad y_l^* = \frac{\sum_{i=1}^M Y_i \mathbb{1}_{\{X_i \leq x^*\}}}{\sum_{i=1}^M \mathbb{1}_{\{X_i \leq x^*\}}}. \quad (1)$$

The problem $\inf_{y_l, y_r, x} m_M(y_l, y_r, x)$ may admit more than one minimizer. To make the tree unique, we choose the minimizer (y_l^*, y_r^*, x^*) with the minimal third component, ie any solution (y_l, y_r, x) to $\inf_{y_l, y_r, x} m_M(y_l, y_r, x)$ satisfies $x > x^*$.

Once the threshold x^* is determined, we split the samples into two groups following the sign of $X_i - x^*$ and repeat the process for each group. We stop the process if introducing a new leaf does not improve the MSE (ie when x^* is on the boundary of the interval on which we solve the minimization problem) or when enough iterations have been made. In the end, we have a tree that approximates the conditional expectation with a piecewise constant function. The regression trees are an algorithmic tool to find an adapted partition and the corresponding weights of this piecewise constant function.

In the multi-dimensional case, we choose the direction (the index along which the optimization is performed) uniformly for each new split. Then, the process is iterated as in the one-dimensional case.

We denote the resulting tree by $\hat{\mathcal{T}}_p^M(X)$ where p represents the depth of the tree, i.e., the number of iterations done in the process described above. A tree of depth p has 2^p leaves. When the size of the training data is infinite, we use the same procedure as above but with expectations rather than empirical means. Consider the function

$$m : y_l \in \mathbb{R}, y_r \in \mathbb{R}, x \in [0, 1] \mapsto \mathbb{E} [(Y - y_l \mathbb{1}_{\{X \leq x\}} - y_r \mathbb{1}_{\{X > x\}})^2].$$

Equation (1) becomes

$$y_r^* = \mathbb{E}[Y|X > x^*]; \quad y_l^* = \mathbb{E}[Y|X \leq x^*].$$

Again, when we optimize the value of x^* , we choose the minimizer (y_l^*, y_r^*, x^*) with the minimal third component. We denote the tree of depth p obtained with an infinite data set by $\mathcal{T}_p(X)$.

3 LS algorithm with regression trees

3.1 Notation

For $p \in \mathbb{N}$, let $\mathcal{A} = \left(\prod_{j=1}^d [a_p^{i-1}(j), a_p^i(j)] \right)_{1 \leq i \leq 2^p}$ be a partition of $[0, 1]^d$ with 2^p elements. We write

$$[a_p^{i-1}, a_p^i[:= \prod_{j=1}^d [a_p^{i-1}(j), a_p^i(j)[$$

For $(\alpha_p^i)_{1 \leq i \leq 2^p} \in \mathbb{R}^{2^p}$, we define \mathcal{P}_p as the piecewise constant function on the partition \mathcal{A} with values α_p^i . For $x \in [0, 1]^d$,

$$\mathcal{P}_p(x, (\alpha_p^i)_{0 \leq i \leq 2^p}, (\alpha_p^i)_{1 \leq i \leq 2^p}) = \sum_{i=1}^{2^p} \alpha_p^i \mathbb{1}_{\{x \in [a_p^{i-1}, a_p^i)\}}$$

If \mathcal{A} denotes the partition obtained in the regression tree \mathcal{T}_p , and we choose

$$\alpha_p^i = \mathbb{E}[Y|X \in [a_p^{i-1}, a_p^i)],$$

then the regression tree $\mathcal{T}_p(X)$ can be written as follows

$$\mathcal{T}_p(X) = \mathcal{P}_p(X, (\alpha_p^i)_{0 \leq i \leq 2^p}, (\alpha_p^i)_{1 \leq i \leq 2^p}) \quad (2)$$

Similarly, if $\mathcal{A}^M = ([a_p^{i-1,M}, a_p^{i,M}])_{1 \leq i \leq 2^p}$ denotes the partition obtained in the regression tree $\hat{\mathcal{T}}_p^M$ and

$$\alpha_p^{i,M} = \frac{\sum_{m=1}^M Y_m \mathbb{1}_{\{X_m \in [a_p^{i-1,M}, a_p^{i,M}]\}}}{\sum_{m=1}^M \mathbb{1}_{\{X_m \in [a_p^{i-1,M}, a_p^{i,M}]\}}},$$

then the regression tree $\hat{\mathcal{T}}_p^M(X)$ can be written as follows

$$\hat{\mathcal{T}}_p^M(X) = \mathcal{P}_p(X, (a_p^{i,M})_{0 \leq i \leq 2^p}, (\alpha_p^{i,M})_{1 \leq i \leq 2^p}) \quad (3)$$

3.2 Description of the algorithm

Let T be a fixed maturity, and consider the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P})$ where \mathbb{P} is the risk neutral measure. Consider a Bermudan option that can be exercised at dates $0 = t_0 < t_1 < t_2 < \dots < t_N = T$. When exercised at time t_j , the option's discounted payoff is given by $Z_{t_j} = h_j(X_{t_j})$ with $(X_{t_j})_j$ being an adapted Markov process taking values in $[0, 1]^d$ such that for every j , X_{t_j} has a density f_X^j on $[0, 1]^d$ uniformly bounded from below on any compact set of $[0, 1]^d$ by a strictly positive number. The discounted value $(U_{t_j})_{0 \leq j \leq N}$ of this option is given by

$$U_{t_j} = \operatorname{esssup}_{\tau \in \mathcal{S}_{t_j, T}} \mathbb{E}[Z_\tau | \mathcal{F}_{t_j}], \quad (4)$$

where $\mathcal{S}_{t_j, T}$ is the set of all stopping times taking values in $\{t_j, \dots, T\}$. Using the Snell envelope theory, we know that U solves the dynamic programming equation

$$\begin{cases} U_{t_N} &= Z_{t_N} \\ U_{t_j} &= \max(Z_{t_j}, \mathbb{E}[U_{t_{j+1}} | \mathcal{F}_{t_j}]) \text{ for } 1 \leq j \leq N-1. \end{cases} \quad (5)$$

This equation can be rewritten in terms of optimal policy as follows

$$\begin{cases} \tau_N &= t_N = T \\ \tau_j &= t_j \mathbb{1}_{\{Z_{t_j} \geq \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} + \tau_{j+1} \mathbb{1}_{\{Z_{t_j} < \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} \text{ for } 1 \leq j \leq N-1 \end{cases} \quad (6)$$

where τ_j is the smallest optimal stopping time after t_j . As we are in a Markovian setting, we can write $\mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] = \mathbb{E}[Z_{\tau_{j+1}} | X_{t_j}]$. The main difficulty in solving this equation comes from the computation of the continuation value $\mathbb{E}[Z_{\tau_{j+1}} | X_{t_j}]$. In the Least Squares approach presented by (Longstaff and Schwartz, 2001), this conditional expectation is estimated by a linear regression on a countable set of basis functions of X_{t_j} . In our approach, we suggest to estimate it using a regression tree of depth p . Let $\mathcal{T}_p^j(X_{t_j})$ be the regression tree approximating $\mathbb{E}[Z_{\tau_{j+1}} | X_{t_j}]$ with an infinite data set. The algorithm solves the following policy

$$\begin{cases} \tau_N^p &= t_N = T \\ \tau_j^p &= t_j \mathbb{1}_{\{Z_{t_j} \geq \mathcal{T}_p^j(X_{t_j})\}} + \tau_{j+1} \mathbb{1}_{\{Z_{t_j} < \mathcal{T}_p^j(X_{t_j})\}} \text{ for } 1 \leq j \leq N-1. \end{cases} \quad (7)$$

We sample M paths of the model $X_{t_0}^{(m)}, \dots, X_{t_N}^{(m)}$ along with the corresponding payoff paths $Z_{t_0}^{(m)}, \dots, Z_{t_N}^{(m)}$, $m = 1, \dots, M$. For each date $j = 1, \dots, N-1$ we approximate the conditional expectations $\mathbb{E}[Z_{\tau_{j+1}} | X_{t_j}]$ on the path m using the regression tree $\hat{\mathcal{T}}_p^{j,M}(X_{t_j}^{(m)})$ built with the samples $(X_{t_j}^{(m)}, Z_{\hat{\tau}_{j+1}^p}^m)_{1 \leq m \leq M}$ where

$$\begin{cases} \hat{\tau}_N^{p,(m)} &= t_N = T \\ \hat{\tau}_j^{p,(m)} &= t_j \mathbb{1}_{\{Z_{t_j}^{(m)} \geq \hat{\mathcal{T}}_p^{j,M}(X_{t_j}^{(m)})\}} + \hat{\tau}_{j+1}^{p,(m)} \mathbb{1}_{\{Z_{t_j}^{(m)} < \hat{\mathcal{T}}_p^{j,M}(X_{t_j}^{(m)})\}} \text{ for } 1 \leq j \leq N-1. \end{cases} \quad (8)$$

Finally, the time-0 price of the option is approximated by

$$U_0^{p,M} = \max \left(Z_0, \frac{1}{M} \sum_{m=1}^M Z_{\hat{\tau}_1^{p,(m)}}^{(m)} \right). \quad (9)$$

4 Convergence of the algorithm

4.1 Notation

For each time step $1 \leq j \leq N-1$, let $\theta_j^p = ((a_{i,j}^p)_{0 \leq i \leq 2^p}, (\alpha_{i,j}^p)_{1 \leq i \leq 2^p})$ be the parameters of \mathcal{T}_p^j and $\hat{\theta}_j^{p,M} = ((a_{i,j}^{p,M})_{0 \leq i \leq 2^p}, (\alpha_{i,j}^{p,M})_{1 \leq i \leq 2^p})$ be the parameters of $\hat{\mathcal{T}}_p^{j,M}$. Note that with this notation and Equations (2) and (3), we have

$$\mathcal{T}_p^j(X_{t_j}) = \mathcal{P}_p(X_{t_j}, \theta_j^p) \quad (10)$$

$$\hat{\mathcal{T}}_p^{j,M}(X_{t_j}^{(m)}) = \mathcal{P}_p(X_{t_j}^{(m)}, \hat{\theta}_j^{p,M}) \quad (11)$$

where for the sake of conciseness we have shrunk the partition and weight parameters of the tree into a single multi-dimensional parameter.

We introduce the vector ϑ of the coefficients of the successive expansions $\vartheta^p = (\theta_1^p, \dots, \theta_{N-1}^p)$ and $\hat{\vartheta}^{p,M} = (\hat{\theta}_1^{p,M}, \dots, \hat{\theta}_{N-1}^{p,M})$.

Let $t^p = (t_1^p, \dots, t_{N-1}^p) \in ([0, 1]^d)^{2^p+1}^{N-1}$ be a deterministic parameter, $z = (z_1, \dots, z_N) \in \mathbb{R}^N$ and $x = (x_1, \dots, x_N) \in ([0, 1]^d)^N$ be deterministic vectors. Following the notation of (Clément et al., 2002), we define the function $F = (F_1, \dots, F_N)$ taking values in \mathbb{R}^N by

$$\begin{cases} F_N(t^p, z, x) &= z_N \\ F_j(t^p, z, x) &= z_j \mathbb{1}_{\{z_j \geq \mathcal{P}_p(x_j, t_j^p)\}} + F_{j+1}(t^p, z, x) \mathbb{1}_{\{z_j < \mathcal{P}_p(x_j, t_j^p)\}}, \text{ for } 1 \leq j \leq N-1. \end{cases}$$

$F_j(t^p, z, x)$ only depends on t_j^p, \dots, t_{N-1}^p and not on the first $j-1$ components of t^p . Using (10), we have

$$\begin{aligned} F_j(\vartheta^p, Z, X) &= Z_{\tau_j^p}, \\ F_j(\hat{\vartheta}^{p,M}, Z^{(m)}, X^{(m)}) &= Z_{\hat{\tau}_j^{p,(m)}}^{(m)}. \end{aligned}$$

Moreover, we clearly have that for all $t^p \in ([0, 1]^d)^{2^p+1}^{N-1}$:

$$|F_j(t^p, Z, X)| \leq \max_{k \geq j} |Z_{t_k}|. \quad (12)$$

4.2 Convergence of the conditional expectations

4.2.1 Preliminary results

We define

$$J_p(a_p^0, \dots, a_p^p) := \mathbb{E} \left[\left| \sum_{i=1}^p \alpha_p^i \mathbb{1}_{\{X \in [a_p^{i-1}, a_p^i]\}} - \mathbb{E}[Y|X] \right|^2 \right].$$

with $\alpha_p^i = \mathbb{E}[Y|X \in [a_p^{i-1}, a_p^i]]$. We will use the following standard lemma to prove the convergence of random tree estimators.

Lemma 4.1. *Let X be a r.v with a density f_X w.r.t to the Lebesgue measure on $[0, 1]^d$ and Y be a real valued square integrable random variable. Let $\left(([a_p^{i-1}, a_p^i])_{1 \leq i \leq p} \right)_{p \in \mathbb{N}}$ be a sequence of partitions of $[0, 1]^d$ such that $\lim_{p \rightarrow \infty} \max_{1 \leq i \leq p} \max_{1 \leq j \leq d} |a_p^i(j) - a_p^{i-1}(j)| = 0$. Then,*

$$\lim_{p \rightarrow \infty} J_p(a_p^0, \dots, a_p^p) = 0.$$

Theorem 4.2.

$$\lim_{p \rightarrow \infty} \mathbb{E} \left[|\mathcal{T}_p(X) - \mathbb{E}[Y|X]|^2 \right] = 0.$$

Proof. With probability $\frac{1}{d}$, the index j is chosen for optimisation. In the other $d - 1$ cases, we do not even cut along that direction, in which case the interval length is at most equal to the length of the largest interval at step $p - 1$. When the index j is chosen: with probability q , the length of the interval is cut in two, and with probability $1 - q$, it is cut to optimize the MSE. In that case, the interval length is at most equal to the length of the largest interval at step $p - 1$.

For all $1 \leq j \leq d$

$$\begin{aligned} & \mathbb{E} \left[\max_{1 \leq i \leq 2^p} |a_p^i(j) - a_p^{i-1}(j)| \right] \\ & \leq \frac{1}{d} \left[q \frac{1}{2} \mathbb{E} \left[\max_{1 \leq i \leq 2^{p-1}} |a_{p-1}^i(j) - a_{p-1}^{i-1}(j)| \right] + (1 - q) \mathbb{E} \left[\max_{1 \leq i \leq 2^{p-1}} |a_{p-1}^i(j) - a_{p-1}^{i-1}(j)| \right] \right] \\ & \quad + \frac{d-1}{d} \mathbb{E} \left[\max_{1 \leq i \leq 2^{p-1}} |a_{p-1}^i(j) - a_{p-1}^{i-1}(j)| \right] \\ & \leq \left(1 - \frac{q}{2d}\right) \mathbb{E} \left[\max_{1 \leq i \leq 2^{p-1}} |a_{p-1}^i(j) - a_{p-1}^{i-1}(j)| \right] \\ & \leq \left(1 - \frac{q}{2d}\right)^{2^p} \end{aligned}$$

Since $\sum_{p=0}^{\infty} \left(1 - \frac{q}{2d}\right)^{2^p} < \infty$, so is $\sum_{p=0}^{\infty} \mathbb{E} \left[\max_{1 \leq i \leq 2^p} |a_p^i(j) - a_p^{i-1}(j)| \right]$. As $\max_{1 \leq i \leq 2^p} |a_p^i(j) - a_p^{i-1}(j)|$ is non negative for all p , using Tonelli's theorem we conclude that $\mathbb{E} \left[\sum_{p=0}^{\infty} \max_{1 \leq i \leq 2^p} |a_p^i(j) - a_p^{i-1}(j)| \right] < \infty$.

As a result, the series $\sum_{p=0}^{\infty} \max_{1 \leq i \leq 2^p} |a_p^i(j) - a_p^{i-1}(j)|$ converges a.s. Then, $\lim_{p \rightarrow \infty} \max_{1 \leq i \leq 2^p} |a_p^i(j) - a_p^{i-1}(j)| = 0$ a.s for all j and $\lim_{p \rightarrow \infty} \max_{1 \leq i \leq 2^p} \max_{1 \leq j \leq d} |a_p^i(j) - a_p^{i-1}(j)| = 0$.

Let \mathcal{G} be the σ -field generated by the splitting strategy (direction choice and threshold strategy). Conditioning by \mathcal{G} allows us to consider the partition $([a_p^{i-1}, a_p^i])_{1 \leq i \leq 2^p}$ deterministic and we can apply Lemma 4.1 to prove

$$\lim_{p \rightarrow \infty} \mathbb{E} \left[|\mathcal{T}_p(X) - \mathbb{E}[Y|X]|^2 | \mathcal{G} \right] = 0 \text{ a.s.}$$

Note that

$$\begin{aligned} \mathbb{E} \left[|\mathcal{T}_p(X)|^2 | \mathcal{G} \right] & \leq \mathbb{E} \left[\sum_{i=1}^{2^p} \mathbb{E} [Y^2 | X \in [a_p^{i-1}, a_p^i]] \mathbb{1}_{\{X \in [a_p^{i-1}, a_p^i]\}} | \mathcal{G} \right] \\ & \leq \sum_{i=1}^{2^p} \mathbb{E} \left[\mathbb{E} [Y^2 \mathbb{1}_{\{X \in [a_p^{i-1}, a_p^i]\}} | X \in [a_p^{i-1}, a_p^i]] | \mathcal{G} \right] \\ & \leq \sum_{i=1}^{2^p} \mathbb{E} [Y^2 \mathbb{1}_{\{X \in [a_p^{i-1}, a_p^i]\}} | \mathcal{G}] \\ & \leq \mathbb{E} [Y^2 | \mathcal{G}] \end{aligned}$$

Then,

$$\begin{aligned} \mathbb{E} \left[|\mathcal{T}_p(X) - \mathbb{E}[Y|X]|^2 | \mathcal{G} \right] & \leq 2 \mathbb{E} \left[|\mathcal{T}_p(X)|^2 | \mathcal{G} \right] + 2 \mathbb{E} \left[\mathbb{E}[Y|X]^2 \right] \\ & \leq 2 \left(\mathbb{E} [Y^2 | \mathcal{G}] + \mathbb{E} \left[\mathbb{E}[Y|X]^2 \right] \right) \end{aligned}$$

Using Lebesgue's bounded convergence theorem,

$$\lim_{p \rightarrow \infty} \mathbb{E} \left[|\mathcal{T}_p(X) - \mathbb{E}[Y|X]|^2 \right] = \lim_{p \rightarrow \infty} \mathbb{E} \left[\mathbb{E} \left[|\mathcal{T}_p(X) - \mathbb{E}[Y|X]|^2 | \mathcal{G} \right] \right] = 0. \quad \square$$

Once the partition of the tree is computed, the values associated to each element of the partition are uniquely determined by a global optimization problem

Proposition 4.3. *Let $\mathcal{A} = \left(\prod_{j=1}^d [a_p^{i-1}(j), a_p^i(j)] \right)_{1 \leq i \leq 2^p}$ denote the partition associated to \mathcal{T}_p . Then,*

$$\mathbb{E} \left[(\mathcal{T}_p(X) - \mathbb{E}[Y|X])^2 \right] = \inf_{(\alpha_p^i)_{1 \leq i \leq 2^p}} \mathbb{E} \left[(\mathcal{P}_p(X, (a_p^i)_{0 \leq i \leq 2^p}, (\alpha_p^i)_{1 \leq i \leq 2^p}) - \mathbb{E}[Y|X])^2 \right]$$

and moreover the right hand side admits a unique minimizer given by

$$\alpha_p^i = \mathbb{E} [Y|X \in [a_p^{i-1}, a_p^i]].$$

Proof. Since the function \mathcal{P}_p is linear w.r.t the parameter α , the function

$$\alpha \in \mathbb{R}^{2^p+1} \longmapsto \mathbb{E} \left[(\mathcal{P}_p(X, (a_p^i)_{0 \leq i \leq 2^p}, (\alpha_p^i)_{1 \leq i \leq 2^p}) - \mathbb{E}[Y|X])^2 \right]$$

is strongly convex. Hence, it admits a unique minimizer defined by the first order optimality condition stating that for all $0 \leq i \leq 2^p + 1$,

$$\begin{aligned} \mathbb{E} \left[(\mathcal{P}_p(X, (a_p^i)_{0 \leq i \leq 2^p}, (\alpha_p^i)_{1 \leq i \leq 2^p}) - \mathbb{E}[Y|X]) \mathbb{1}_{\{X \in [a_p^{i-1}, a_p^i]\}} \right] &= 0 \\ \mathbb{E} \left[(\alpha_p^i - \mathbb{E}[Y|X]) \mathbb{1}_{\{X \in [a_p^{i-1}, a_p^i]\}} \right] &= 0 \\ \alpha_p^i &= \mathbb{E} [Y|X \in [a_p^{i-1}, a_p^i]]. \end{aligned} \quad \square$$

A similar result holds for $\hat{\mathcal{T}}_p^M$ by replacing the conditional expectations by empirical conditional expectations.

4.2.2 Approximation of the conditional expectations with regression trees

The following result proves that the approximation of the Bermudan price produced by our algorithm converges to the true Bermudan price.

Proposition 4.4.

$$\lim_{p \rightarrow \infty} \mathbb{E} \left[Z_{\tau_j^p} | \mathcal{F}_{t_j} \right] = \mathbb{E} \left[Z_{\tau_j} | \mathcal{F}_{t_j} \right] \text{ in } \mathbb{L}^2(\Omega) \text{ for } 1 \leq j \leq N \quad (13)$$

Proof. We proceed by induction.

For $j = N$, the proposition is true since $\tau_N^p = \tau_N = T$. Assume that the result holds for $j + 1$, Let us prove that it still holds for j :

$$\begin{aligned} \mathbb{E} \left[Z_{\tau_j^p} - Z_{\tau_j} | \mathcal{F}_{t_j} \right] &= Z_{t_j} \left(\mathbb{1}_{\{Z_{t_j} \geq \tau_p^j(X_{t_j})\}} - \mathbb{1}_{\{Z_{t_j} \geq \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} \right) \\ &\quad + \mathbb{E} \left[Z_{\tau_{j+1}^p} \mathbb{1}_{\{Z_{t_j} < \tau_p^j(X_{t_j})\}} - Z_{\tau_{j+1}} \mathbb{1}_{\{Z_{t_j} < \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} | \mathcal{F}_{t_j} \right] \\ &= (Z_{t_j} - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]) \left(\mathbb{1}_{\{Z_{t_j} \geq \tau_p^j(X_{t_j})\}} - \mathbb{1}_{\{Z_{t_j} \geq \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} \right) \\ &\quad + \mathbb{E}[Z_{\tau_{j+1}^p} - Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] \mathbb{1}_{\{Z_{t_j} < \tau_p^j(X_{t_j})\}} \\ &= A_j^p + \mathbb{E}[Z_{\tau_{j+1}^p} - Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] \mathbb{1}_{\{Z_{t_j} < \tau_p^j(X_{t_j})\}} \end{aligned}$$

Where A_j^p is defined by

$$A_j^p = (Z_{t_j} - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]) \left(\mathbb{1}_{\{Z_{t_j} \geq \tau_p^j(X_{t_j})\}} - \mathbb{1}_{\{Z_{t_j} \geq \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} \right)$$

On one hand, since the conditional expectation is an orthogonal projection, we have

$$\mathbb{E} \left[\left| \mathbb{E} \left[Z_{\tau_{j+1}^p} - Z_{\tau_{j+1}} | \mathcal{F}_{t_j} \right] \right|^2 \right] \leq \mathbb{E} \left[\left| \mathbb{E} \left[Z_{\tau_{j+1}^p} - Z_{\tau_{j+1}} | \mathcal{F}_{t_{j+1}} \right] \right|^2 \right]$$

and using the induction assumption $\mathbb{E} \left[Z_{\tau_{j+1}^p} - Z_{\tau_{j+1}} | \mathcal{F}_{t_{j+1}} \right] \rightarrow 0$ in $\mathbb{L}^2(\Omega)$ when $p \rightarrow \infty$. On the other hand,

$$\begin{aligned} |A_j^p| &= \left| Z_{t_j} - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] \right| \left| \mathbb{1}_{\{Z_{t_j} \geq \mathcal{T}_p^j(X_{t_j})\}} - \mathbb{1}_{\{Z_{t_j} \geq \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} \right| \\ &\leq \left| Z_{t_j} - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] \right| \left| \mathbb{1}_{\{\mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] > Z_{t_j} \geq \mathcal{T}_p^j(X_{t_j})\}} - \mathbb{1}_{\{\mathcal{T}_p^j(X_{t_j}) > Z_{t_j} \geq \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]\}} \right| \\ &\leq \left| Z_{t_j} - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] \right| \left| \mathbb{1}_{\{|Z_{t_j} - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]| \leq |\mathcal{T}_p^j(X_{t_j}) - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]|\}} \right| \\ &\leq |\mathcal{T}_p^j(X_{t_j}) - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}]| \\ &\leq \left| \mathcal{T}_p^j(X_{t_j}) - \mathbb{E}[Z_{\tau_{j+1}^p} | \mathcal{F}_{t_j}] \right| + \left| \mathbb{E}[Z_{\tau_{j+1}^p} | \mathcal{F}_{t_j}] - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] \right|. \end{aligned}$$

Using the induction assumption, the second term goes to zero in $\mathbb{L}^2(\Omega)$ when $p \rightarrow \infty$. Let $([a_{i-1}(p), a_i(p)])_{1 \leq i \leq 2^p}$ be the partition generated by \mathcal{T}_p^j . We define

$$\bar{\mathcal{T}}_p^j(X_{t_j}) = \sum_{i=1}^{2^p} \mathbb{E}[Z_{\tau_{j+1}} | X_{t_j} \in [a_{i-1}(p), a_i(p)]] \mathbb{1}_{\{X_{t_j} \in [a_{i-1}(p), a_i(p)]\}}.$$

Note that $\bar{\mathcal{T}}_p^j$ uses the partition given by $\mathcal{T}_p^j(X_{t_j})$ but the coefficients $\alpha_i(p)$ are given by the conditional expectations of $Z_{\tau_{j+1}}$ w.r.t X_{t_j} and not those of $Z_{\tau_{j+1}^p}$. Using Proposition 4.3, we have the following inequality stating that $\bar{\mathcal{T}}_p^j(X_{t_j})$ is sub-optimal compared to $\mathcal{T}_p^j(X_{t_j})$

$$\begin{aligned} &\mathbb{E} \left[\left| \mathcal{T}_p^j(X_{t_j}) - \mathbb{E}[Z_{\tau_{j+1}^p} | \mathcal{F}_{t_j}] \right|^2 \right] \\ &\leq \mathbb{E} \left[\left| \bar{\mathcal{T}}_p^j(X_{t_j}) - \mathbb{E}[Z_{\tau_{j+1}^p} | \mathcal{F}_{t_j}] \right|^2 \right] \\ &\leq 2\mathbb{E} \left[\left| \bar{\mathcal{T}}_p^j(X_{t_j}) - \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] \right|^2 \right] + 2\mathbb{E} \left[\left| \mathbb{E}[Z_{\tau_{j+1}} | \mathcal{F}_{t_j}] - \mathbb{E}[Z_{\tau_{j+1}^p} | \mathcal{F}_{t_j}] \right|^2 \right]. \end{aligned}$$

The second term goes to 0 using the induction assumption. As for the first term, note that the partition obtained with $\bar{\mathcal{T}}_p^j$ verifies the conditions of Lemma 4.1. Then, using the same arguments as in the proof of Theorem 4.2, we can show that the first term also goes to 0. \square

4.3 Convergence of the Monte Carlo approximation

In this section, the depth p of the trees is fixed and we study the convergence with respect to the number of samples M . We assume that, for all dates j , the trees $\mathcal{T}_p^j(X_{t_j})$ and $\hat{\mathcal{T}}_p^{j,M}(X_{t_j})$ are built using the same splitting strategy (ie the same splitting directions and threshold strategies).

4.3.1 Convergence of optimisation problems

We present three major results to study the convergence of stochastic optimization problems related to regression trees.

The first result is a uniform strong law of large numbers, see (Leake et al., 1994, Chap. 2, Lemma. A1), which can be seen as a particular case of the strong law of large numbers in Banach spaces (Ledoux and Talagrand, 1991, Corollary 7.10, page 189).

Lemma 4.5. *Let $(\xi_i)_{i \geq 1}$ be a sequence of i.i.d \mathbb{R}^n -valued random vectors and $h : \mathbb{R}^d \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a measurable function. Assume that*

(i) For all $\bar{\theta} \in \mathbb{R}^d$, the function $\theta \in \mathbb{R}^d \mapsto h(\theta, \xi_1)$ is continuous at $\bar{\theta}$ a.s.,

(ii) $\forall C > 0, \mathbb{E} \left[\sup_{|\theta| \leq C} |h(\theta, \xi_1)| \right] < \infty$.

Then, a.s $\theta \in \mathbb{R}^d \mapsto \frac{1}{n} \sum_{i=1}^n h(\theta, \xi_i)$ converges locally uniformly to the continuous function $\theta \in \mathbb{R}^d \mapsto \mathbb{E} [h(\theta, \xi_1)]$, i.e

$$\forall C > 0, \lim_{n \rightarrow \infty} \sup_{|\theta| \leq C} \left| \frac{1}{n} \sum_{i=1}^n h(\theta, \xi_i) - \mathbb{E} [h(\theta, \xi_1)] \right| = 0 \text{ a.s.}$$

This lemma is a slight improvement of (Leake et al., 1994, Chap. 2, Lemma. A1), which was formulated under the assumption that the function $\theta \in \mathbb{R}^d \mapsto h(\theta, \xi_1)$ is almost surely continuous. However, looking closely at their proof, it turns that it sufficient to assume (i) for the conclusion to hold. Condition (i) allows the \mathbb{P} -null sets on which the continuity at $\bar{\theta}$ does not hold to depend on $\bar{\theta}$.

Consider a sequence of real valued functions $(f_n)_n$ defined on a compact set $K \subset \mathbb{R}^d$ such that there exists a sequence of $(x_n)_n$ satisfying

$$f_n(x_n) = \inf_{x \in K} f_n(x).$$

From (Leake et al., 1994, Chap. 2, Theorem A1), we can derive the following lemma

Lemma 4.6. *Assume that the sequence $(f_n)_n$ converges uniformly on K to a continuous function f . Let $v^* = \inf_{x \in K} f(x)$ and $\mathcal{S}^* = \{x \in K : f(x) = v^*\}$. Then, $f_n(x_n) \rightarrow \inf_{x \in K} f(x)$ and $d(x_n, \mathcal{S}^*) \rightarrow 0$.*

Now, we focus on a canonical minimization problem appearing at each node of the regression tree. Let $[\underline{a}, \bar{a}]$ be a rectangle in $[0, 1]$. For some index $1 \leq \delta \leq d$, and a real number a in $]\underline{a}^\delta, \bar{a}^\delta[$, we define the two new rectangles

$$\mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(a) = \{x \in [\underline{a}, \bar{a}] : x^\delta \leq a\}; \quad \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, r}(a) = \{x^\delta \in [\underline{a}, \bar{a}] : x > a\}. \quad (14)$$

We consider the cost functions

$$m_M : y_l \in \mathbb{R}, y_r \in \mathbb{R}, x \in [\underline{a}^\delta, \bar{a}^\delta] \mapsto \frac{1}{M} \sum_{i=1}^M \left(Y_i - y_l \mathbb{1}_{\{X_i \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(x)\}} - y_r \mathbb{1}_{\{X_i \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, r}(x)\}} \right)^2 \quad (15)$$

$$m : y_l \in \mathbb{R}, y_r \in \mathbb{R}, x \in [\underline{a}^\delta, \bar{a}^\delta] \mapsto \mathbb{E} \left[\left(Y - y_l \mathbb{1}_{\{X \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(x)\}} - y_r \mathbb{1}_{\{X \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, r}(x)\}} \right)^2 \right]. \quad (16)$$

Let $(\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M)$ be the solution to

$$\inf_{y_l, y_r, x \in [\underline{a}^\delta + \varepsilon_0, \bar{a}^\delta - \varepsilon_0]} m_M(y_l, y_r, x) \quad (17)$$

with the smallest third component for some arbitrary small $\varepsilon_0 > 0$. It is clear that (17) may not have a unique solution. Choosing the solution with the minimal third component is a standard way to get a unique minimizer (see for instance Seijo and Sen (2011)).

Lemma 4.7. *Assume that the density f_X of X satisfies $f_X(x) \geq \underline{f} > 0$ for all x in any compact set of $]0, 1[^d$ and that the minimization problem*

$$\inf_{y_l, y_r, x \in [\underline{a}^\delta + \varepsilon_0, \bar{a}^\delta - \varepsilon_0]} m(y_l, y_r, x) \quad (18)$$

has a unique minimizer (y_l^*, y_r^*, x^*) .

Then, $(\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M)$ converges almost surely to (y_l^*, y_r^*, x^*) and $m_M(\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M)$ converges almost surely to $m(y_l^*, y_r^*, x^*)$ when M goes to infinity.

Proof. Consider the function

$$h(y_l, y_r, x, y, \xi) = \left(y - y_l \mathbb{1}_{\{\xi \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(x)\}} - y_r \mathbb{1}_{\{\xi \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, r}(x)\}} \right)^2$$

Since X has a density on $[0, 1]^d$, the function $(y_l, y_r, x) \mapsto h(y_l, y_r, x, Y, X)$ satisfies Condition (i) of Lemma 4.5. Moreover, for $C > 0$

$$\mathbb{E} \left[\sup_{|y_l| \leq C, |y_r| \leq C, x \in [\underline{a}^\delta + \varepsilon_0, \bar{a}^\delta - \varepsilon_0]} |h(y_l, y_r, x, Y, X)| \right] \leq 2\mathbb{E}[Z^2] + 2C$$

Hence, we deduce from Lemma 4.5 that m_M converges a.s. locally uniformly m . Note that for a fixed $x \in [\underline{a}^\delta + \varepsilon_0, \bar{a}^\delta - \varepsilon_0]$, the optimal values of y_l and y_r are given by

$$\begin{aligned} y_l &= \mathbb{E}[Y | X \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(x)] \\ y_r &= \mathbb{E}[Y | X \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, r}(x)] \end{aligned}$$

Observe that for $x \in [\underline{a}^\delta + \varepsilon_0, \bar{a}^\delta - \varepsilon_0]$

$$|\mathbb{E}[Y | X \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(x)]| \leq \frac{\sqrt{\mathbb{E}[Y^2]}}{\sqrt{\mathbb{P}(X \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(x))}} \leq \frac{\sqrt{\mathbb{E}[Y^2]}}{\sqrt{\mathbb{P}(X \in \mathcal{R}_{\underline{a}, \bar{a}}^{\delta, l}(\underline{a}^\delta - \varepsilon_0))}}$$

The uniform upper bound is uniform is finite thanks to the assumption on the density f_X of f . Hence, as a function of x , y_l is uniformly bounded. A similar results holds for y_r . Then, we deduce from Lemma 4.6, that $m_M(\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M)$ converges almost surely to $m(y_l^*, y_r^*, x^*)$ and $(\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M)$ converges almost surely to (y_l^*, y_r^*, x^*) . \square

To the best of our knowledge, without the uniqueness assumption on (y_l^*, y_r^*, x^*) , we can only prove that $(\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M)$ converges in probability to (y_l^*, y_r^*, x^*) . See for instance (Ferguson, 2004, Corollary 1), which yields that for any $\eta > 0$, $\mathbb{P}(\hat{x}^M < x^* + \eta) \rightarrow 1$. Combining this result with $d((\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M), S^*) \rightarrow 0$ a.s. yields the convergence in probability of $(\hat{y}_l^M, \hat{y}_r^M, \hat{x}^M)$ to (y_l^*, y_r^*, x^*) .

4.3.2 Strong law of large numbers

In order to prove the almost sure convergence of $\hat{\mathcal{T}}_{j,p}^M(X)$, we slightly modify the design of our regression trees. At every node, when computing the optimal splitting point x^* , we do not perform the optimization on the entire cell but we actually a margin to make sure that y_l^* and y_r^* are uniformly bounded w.r.t x^* . Consider a node at depth p , the optimal splitting is obtained by minimizing m_M defined in (15) where \underline{a} and \bar{a} were computed at depth $p - 1$. As in (17), we search for the optimal value of x^* in $[\underline{a}^\delta - \varepsilon_p, \bar{a}^\delta - \varepsilon_p]$ where δ is the coordinate chosen for the splitting and $\varepsilon_p > 0$ is a technical security padding decreasing to 0. When there is no room left for this padding ($|\underline{a}^\delta - \bar{a}^\delta| < 2\varepsilon_p$), we stop the splitting procedure.

From (Clément et al., 2002), we have the following result

Lemma 4.8. *For every $j = 1, \dots, N - 1$,*

$$|F_j(a, Z, X) - F_j(b, Z, X)| \leq \left(\sum_{i=j}^N |Z_{t_i}| \right) \left(\sum_{i=j}^{N-1} \mathbb{1}_{\{|Z_{t_i} - \mathcal{P}_p^i(X_{t_i}, b_i)| \leq |\mathcal{P}_p^i(X_{t_i}, a_i) - \mathcal{P}_p^i(X_{t_i}, b_i)|\}} \right).$$

Proposition 4.9. *Assume that for all $p \in \mathbb{N}^*$, and all $1 \leq j \leq N - 1$, $\mathbb{P}(Z_{t_j} = \mathcal{P}_p^j(X_{t_j}, \theta_p^j)) = 0$. For all $1 \leq j \leq N - 1$ and all $1 \leq i \leq 2^p$, the optimization problems*

$$\inf_{y_l, y_r, x} \mathbb{E} \left[\left(Z_{\tau_j^p} - y_l \mathbb{1}_{\left\{ X_{t_j} \in \mathcal{R}_{a_p^{i-1}, a_p^i}^{\delta, l}(x) \right\}} - y_r \mathbb{1}_{\left\{ X_{t_j} \in \mathcal{R}_{a_p^{i-1}, a_p^i}^{\delta, r}(x) \right\}} \right)^2 \right]$$

with $x \in [a_p^{i-1,\delta} + \varepsilon_p, a_p^{i,\delta} - \varepsilon_p]$ admit a unique solution where δ is the coordinate selected by the splitting strategy.

Then, for all $j = 1, \dots, N-1$, $\hat{\theta}_j^{p,M}$ converges a.s. to θ_j^p and $\mathcal{P}_p^j(X_{t_j}, \hat{\theta}_j^{p,M})$ converges a.s. to $\mathcal{P}_p^j(X_{t_j}, \theta_j^p)$ as $M \rightarrow \infty$.

Proof. We proceed by backward induction on j with a nested forward induction on p .

Step 1: $j = N-1$

- For $p = 1$, conditionally on splitting at the best point (and not at the midpoint) the new nodes are obtained by solving

$$\begin{aligned} \mathcal{P}_1^{N-1}(X_{t_{N-1}}, \hat{\theta}_{N-1}^{1,M}) &= \inf_{\alpha, \beta, a} \frac{1}{M} \sum_{m=1}^M \left| Z_{t_N}^{(m)} - \alpha \mathbb{1}_{\{X_{t_{N-1}}^{(m)} \in \mathcal{R}_{0,1}^{\delta,l}(a)\}} - \beta \mathbb{1}_{\{X_{t_{N-1}}^{(m)} \in [a,1]\}} \right|^2 \\ \mathcal{P}_1^{N-1}(X_{t_{N-1}}, \theta_{N-1}^1) &= \inf_{\alpha, \beta, a} \mathbb{E} \left[\left| Z_{t_N} - \alpha \mathbb{1}_{\{X_{t_{N-1}} \in \mathcal{R}_{0,1}^{\delta,l}(a)\}} - \beta \mathbb{1}_{\{X_{t_{N-1}} \in [a,1]\}} \right|^2 \right] \end{aligned}$$

where $a \in [\varepsilon_1, 1 - \varepsilon_1]$. We can apply Lemma 4.7 to obtain $\mathcal{P}_1^{N-1}(X_{t_{N-1}}, \hat{\theta}_{N-1}^{1,M})$ converges to $\mathcal{P}_1^{N-1}(X_{t_{N-1}}, \theta_{N-1}^1)$ a.s as $M \rightarrow \infty$ and so does $\hat{\theta}_{N-1}^{1,M}$ converge a.s. to θ_{N-1}^1 . If the split is at the midpoint, the conclusion is even easier to obtain as the infimums are only computed w.r.t α and β but not a . The same situation will occur repeatedly and for the sake of clearness, we will only treat the case of splitting at the best point which is harder to handle.

- Assume that $\mathcal{P}_p^{N-1}(X_{t_{N-1}}, \hat{\theta}_{N-1}^{p,M})$ converges to $\mathcal{P}_p^{N-1}(X_{t_{N-1}}, \theta_{N-1}^p)$ a.s and that $\hat{\theta}_{N-1}^{p,M}$ converges a.s. to θ_{N-1}^p as $M \rightarrow \infty$ for $p \geq 1$, we will prove it for $p+1$. For $i \in \{1, \dots, 2^p\}$, we consider the computation of the i -th node in $\hat{\mathcal{T}}_{p+1}^{j,M}$ at depth $p+1$.

$$\hat{\nu}_{p,N-1}^M(\alpha, \beta, a) = \frac{1}{M} \sum_{m=1}^M \left| Z_{t_N}^{(m)} - \alpha \mathbb{1}_{\{X_{t_{N-1}}^{(m)} \in \mathcal{R}_{a_{i-1}^{p,M}, N-1, a_{i-1}^{p,M}}^{\delta,l}(a)\}} - \beta \mathbb{1}_{\{X_{t_{N-1}}^{(m)} \in \mathcal{R}_{a_{i-1}^{p,M}, N-1, a_{i-1}^{p,M}}^{\delta,r}(a)\}} \right|^2.$$

The parameters $\hat{\theta}_{N-1}^{p+1,M}$ are obtained by computing all the nodes. We also introduce a modified version of $\hat{\nu}_{p,N-1}^M$ in which we use the splits computed in \mathcal{T}_p^j .

$$\nu_{p,N-1}^M(\alpha, \beta, a) = \frac{1}{M} \sum_{m=1}^M \left| Z_{t_N}^{(m)} - \alpha \mathbb{1}_{\{X_{t_{N-1}}^{(m)} \in \mathcal{R}_{a_{i-1}^p, N-1, a_{i-1}^p}^{\delta,l}(a)\}} - \beta \mathbb{1}_{\{X_{t_{N-1}}^{(m)} \in \mathcal{R}_{a_{i-1}^p, N-1, a_{i-1}^p}^{\delta,r}(a)\}} \right|^2.$$

The random functions $\nu_{p,N-1}^M$ write as standard empirical means and we will prove they are close to $\hat{\nu}_{p,N-1}^M$ for M large.

Using Lemma 4.5 along with the arguments of the proof of Lemma 4.7, it is easy to see that the random function $\alpha, \beta, a \mapsto \nu_{p,N-1}^M(\alpha, \beta, a)$ converges a.s locally uniformly to the function

$$\alpha, \beta, a \mapsto \mathbb{E} \left[\left| Z_{t_N} - \alpha \mathbb{1}_{\{X_{t_{N-1}} \in \mathcal{R}_{a_{i-1}^p, N-1, a_{i-1}^p}^{\delta,l}(a)\}} - \beta \mathbb{1}_{\{X_{t_{N-1}} \in \mathcal{R}_{a_{i-1}^p, N-1, a_{i-1}^p}^{\delta,r}(a)\}} \right|^2 \right].$$

$$\begin{aligned} & \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} \left| \hat{\nu}_{p,N-1}^M(\alpha, \beta, a) - \nu_{p,N-1}^M(\alpha, \beta, a) \right| \\ & \leq \frac{1}{M} \sum_{m=1}^M \left[\left| 2Z_{t_N}^{(m)} \right| + 4C \right] C \left(\mathbb{1}_{\{X_{t_{N-1}}^{(m),\delta} \in (a_{i-1}^{p,M,\delta}, a_{i-1}^{p,M,\delta})\}} + \mathbb{1}_{\{X_{t_{N-1}}^{(m),\delta} \in (a_i^{p,M,\delta}, a_i^{p,M,\delta})\}} \right). \end{aligned}$$

Using the induction assumption on p , $a_i^{p,M}$ (resp. $a_{i-1}^{p,M}$) converges a.s. to a_i^p (resp. a_{i-1}^p). Let $\epsilon > 0$,

$$\begin{aligned} & \limsup_M |\hat{\nu}_{p,N-1}^M(\alpha, \beta, a) - \nu_{p,N-1}^M(\alpha, \beta, a)| \\ & \leq \limsup_M \frac{1}{M} \sum_{m=1}^M \left[|2Z_{t_N}^{(m)}| + 4C \right] C \left(\left| \mathbb{1}_{\{|X_{t_{N-1}}^{(m),\delta} - a_i^{p,\delta}| \leq \epsilon\}} \right| + \left| \mathbb{1}_{\{|X_{t_{N-1}}^{(m),\delta} - a_{i+1}^{p,\delta}| \leq \epsilon\}} \right| \right) \\ & \leq C(4C + 2\mathbb{E}[|2Z_{t_N}|]) \left(\mathbb{P}(|X_{t_{N-1}}^\delta - a_i^{p,\delta}| \leq \epsilon) + \mathbb{P}(|X_{t_{N-1}}^\delta - a_{i+1}^{p,\delta}| \leq \epsilon) \right). \end{aligned}$$

Since $X_{t_{N-1}}$ has a density, $\lim_{\epsilon \rightarrow 0} \mathbb{P}(|X_{t_{N-1}}^\delta - a_i^{p,\delta}| \leq \epsilon) = \mathbb{P}(X_{t_{N-1}}^\delta = a_i^{i,\delta}) = 0$ and $\lim_{\epsilon \rightarrow 0} \mathbb{P}(|X_{t_{N-1}}^\delta - a_{i+1}^{p,\delta}| \leq \epsilon) = \mathbb{P}(X_{t_{N-1}}^\delta = a_{i+1}^{i,\delta}) = 0$. As a result, $|\hat{\nu}_{p,N-1}^M - \nu_{p,N-1}^M| \rightarrow 0$ a.s. locally uniformly when $M \rightarrow \infty$. Thus, the random function $\hat{\nu}_{p,N-1}^M$ converges a.s. locally uniformly to the function

$$\alpha, \beta, a \mapsto \mathbb{E} \left[\left| Z_{t_N} - \alpha \mathbb{1}_{\{X_{t_{N-1}} \in \mathcal{R}_{a_{i-1}^p, a_i^p}^{\delta,l}\}} - \beta \mathbb{1}_{\{X_{t_{N-1}} \in \mathcal{R}_{a_{i-1}^p, a_i^p}^{\delta,r}\}} \right|^2 \right].$$

Now, we apply Lemma 4.6 along with the same arguments of Lemma 4.7 to conclude that $\mathcal{P}_{p+1}^{N-1}(X_{t_{N-1}}, \hat{\theta}_{N-1}^{p+1,M})$ converges to $\mathcal{P}_{p+1}^{N-1}(X_{t_{N-1}}, \theta_{N-1}^{p+1})$ a.s as $M \rightarrow \infty$ and that $\hat{\theta}_{N-1}^{p+1,M}$ converges a.s. to θ_{N-1}^{p+1} .

Step 2: $j < N - 1$.

So far, we have proved that for all p , $\mathcal{P}_p^{N-1}(X_{t_{N-1}}, \hat{\theta}_{N-1}^{p,M})$ converges to $\mathcal{P}_p^{N-1}(X_{t_{N-1}}, \theta_j^p)$ a.s and that $\hat{\theta}_{N-1}^{p,M}$ converges a.s. to θ_j^p as $M \rightarrow \infty$. Now, suppose that $\mathcal{P}_p^k(X_{t_k}, \hat{\theta}_k^{p,M})$ (resp. $\hat{\theta}_k^{p,M}$) converges to $\mathcal{P}_p^k(X_{t_k}, \theta_k^p)$ (resp. θ_k^p) a.s as $M \rightarrow \infty$ for all p and for $k = N - 1, \dots, j + 1$. We should prove that these convergence results hold for j and this will be done by induction on p . Now that we have understood that considering multidimensional random variables X_{t_j} does not play any role in the proof, we will make the rest of the proof as if X were having in \mathbb{R} in order to use a little lighter notation

- For $p = 1$, consider

$$\begin{aligned} \hat{\nu}_{1,j}^M(\alpha, \beta, a) &= \frac{1}{M} \sum_{m=1}^M \left| F_{j+1}(\hat{\vartheta}^{1,M}, Z^{(m)}, X^{(m)}) - \alpha \mathbb{1}_{\{X_{t_j}^{(m)} \in [0,a)\}} - \beta \mathbb{1}_{\{X_{t_j}^{(m)} \in [a,1]\}} \right|^2 \\ \nu_{1,j}^M(\alpha, \beta, a) &= \frac{1}{M} \sum_{m=1}^M \left| F_{j+1}(\vartheta^1, Z^{(m)}, X^{(m)}) - \alpha \mathbb{1}_{\{X_{t_j}^{(m)} \in [0,a)\}} - \beta \mathbb{1}_{\{X_{t_j}^{(m)} \in [a,1]\}} \right|^2. \end{aligned}$$

The function $\nu_{1,j}^M$ writes as the sum of i.i.d random variables. Let $C \geq 0$, using Equation (12)

$$\begin{aligned} & \mathbb{E} \left[\sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} \left| F_{j+1}(\vartheta^1, Z, X) - \alpha \mathbb{1}_{\{X_{t_j} \in [0,a)\}} - \beta \mathbb{1}_{\{X_{t_j} \in [a,1]\}} \right|^2 \right] \\ & \leq 2\mathbb{E} \left[|F_{j+1}(\vartheta^1, Z, X)|^2 \right] + 2\mathbb{E} \left[\sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} \left| \alpha \mathbb{1}_{\{X_{t_j} \in [0,a)\}} + \beta \mathbb{1}_{\{X_{t_j} \in [a,1]\}} \right|^2 \right] \\ & \leq 2\mathbb{E} \left[\max_{l \geq j+1} (Z_{t_l})^2 \right] + 2C^2 < \infty. \end{aligned}$$

Using Lemma 4.5, $\alpha, \beta, a \mapsto \nu_{1,j}^M(\alpha, \beta, a)$ converges a.s locally uniformly to the function $\alpha, \beta, a \mapsto \mathbb{E} \left[\left| F_{j+1}(\vartheta^1, Z, X) - \alpha \mathbb{1}_{\{X_{t_j} \in [0,a)\}} - \beta \mathbb{1}_{\{X_{t_j} \in [a,1]\}} \right|^2 \right]$. It remains to prove that

$\forall C > 0 \quad \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} |\hat{\nu}_{1,j}^M(a, \alpha, \beta) - \nu_{1,j}^M(a, \alpha, \beta)| \rightarrow 0$ a.s. when $M \rightarrow \infty$.

Then, using Equation (12) and Lemma 4.8

$$\begin{aligned} & \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} |\hat{\nu}_{1,j}^M(a, \alpha, \beta) - \nu_{1,j}^M(a, \alpha, \beta)| \\ & \leq \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} \frac{1}{M} \sum_{m=1}^M \left| F_{j+1} \left(\hat{\vartheta}^{1,M}, Z^{(m)}, X^{(m)} \right) - F_{j+1} \left(\vartheta^1, Z^{(m)}, X^{(m)} \right) \right| \\ & \left| F_{j+1} \left(\hat{\vartheta}^{1,M}, Z^{(m)}, X^{(m)} \right) + F_{j+1} \left(\vartheta^1, Z^{(m)}, X^{(m)} \right) - 2\alpha \mathbb{1}_{\{X_{t_j}^{(m)} \in [0,a]\}} - 2\beta \mathbb{1}_{\{X_{t_j}^{(m)} \in [a,1]\}} \right| \\ & \leq \sum_{m=1}^M 2 \left(\max_{l \geq j+1} |Z_{t_l}^{(m)}| + 2C \right) \left| F_{j+1} \left(\hat{\vartheta}^{1,M}, Z^{(m)}, X^{(m)} \right) - F_{j+1} \left(\vartheta^1, Z^{(m)}, X^{(m)} \right) \right| \\ & \leq \frac{1}{M} \sum_{m=1}^M 2 \left(\max_{l \geq j+1} |Z_{t_l}^{(m)}| + 2C \right) \left(\sum_{i=j+1}^N |Z_{t_i}^{(m)}| \sum_{i=j+1}^{N-1} \mathbb{1}_{\{|Z_{t_i}^{(m)} - \mathcal{T}_1^i(X_{t_i}^{(m)})| \leq |\hat{\tau}_1^{i,M}(X_{t_i}^{(m)}) - \mathcal{T}_1^i(X_{t_i}^{(m)})|\}} \right). \end{aligned}$$

Using the induction assumption on j , $\hat{\tau}_1^{i,M}(X_{t_i}^{(m)})$ converges a.s. to $\mathcal{T}_1^i(X_{t_i}^{(m)})$ for all $N-1 \geq i \geq j+1$. Let $\epsilon < 0$.

$$\begin{aligned} & \limsup_M \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} |\hat{\nu}_{1,j}^M(a, \alpha, \beta) - \nu_{1,j}^M(a, \alpha, \beta)| \\ & \leq \frac{1}{M} \sum_{m=1}^M 2 \left(\max_{l \geq j+1} |Z_{t_l}^{(m)}| + 2C \right) \left(\sum_{i=j+1}^N |Z_{t_i}^{(m)}| \sum_{i=j+1}^{N-1} \mathbb{1}_{\{|Z_{t_i}^{(m)} - \mathcal{T}_1^i(X_{t_i}^{(m)})| \leq \epsilon\}} \right). \end{aligned}$$

Since $\mathbb{P}(Z_{t_j}^{(m)} = \mathcal{T}_p^j(X_{t_j}^{(m)})) = 0$, then $\lim_{\epsilon \rightarrow 0} \mathbb{1}_{\{|Z_{t_i}^{(m)} - \mathcal{T}_1^i(X_{t_i}^{(m)})| \leq \epsilon\}} = 0$ a.s. and we conclude that a.s. $|\hat{\nu}_{1,j}^M(a, \alpha, \beta) - \nu_{1,j}^M(a, \alpha, \beta)|$ converges to zero uniformly. Thus, $\hat{\nu}_{1,j}^M$ converges a.s. uniformly to the function $a, \alpha, \beta \mapsto \mathbb{E} \left[\left| F_{j+1}(\vartheta^1, Z, X) - \alpha \mathbb{1}_{\{X_{t_j} \in [0,a]\}} - \beta \mathbb{1}_{\{X_{t_j} \in [a,1]\}} \right|^2 \right]$.

Then, we apply Lemma 4.5 and the arguments of the proof of Lemma 4.7 to deduce that $\mathcal{P}_1^j(X_{t_j}, \hat{\theta}_j^{1,M})$ converges to $\mathcal{P}_j^j(X_{t_j}, \theta_j^j)$ a.s. and that $\hat{\theta}_j^{j,M}$ converges a.s. to θ_j^j as $M \rightarrow \infty$.

- We assume that $\mathcal{P}_p^k(X_{t_k}, \hat{\theta}_k^{p,M})$ (resp. $\hat{\theta}_k^{p,M}$) converges to $\mathcal{P}_p^k(X_{t_k}, \theta_k^p)$ (resp. θ_k^p) a.s. as $M \rightarrow \infty$ for some p and for $k = N-1, \dots, j$. We will prove that it holds for $p+1$.

Let $i \in \{1, \dots, 2^p\}$ and consider

$$\begin{aligned} \hat{\nu}_{p,j}^M(\alpha, \beta, a) &= \frac{1}{M} \sum_{m=1}^M \left| F_{j+1}(\hat{\vartheta}^{p,M}, Z^{(m)}, X^{(m)}) - \alpha \mathbb{1}_{\{X_{t_j}^{(m)} \in [a_{i-1,j}^p, a]\}} - \beta \mathbb{1}_{\{X_{t_j}^{(m)} \in [a, a_{i,j}^p]\}} \right|^2. \\ \nu_{p,j}^M(\alpha, \beta, a) &= \frac{1}{M} \sum_{m=1}^M \left| F_{j+1}(\vartheta^p, Z^{(m)}, X^{(m)}) - \alpha \mathbb{1}_{\{X_{t_j}^{(m)} \in [a_{i-1,j}^p, a]\}} - \beta \mathbb{1}_{\{X_{t_j}^{(m)} \in [a, a_{i,j}^p]\}} \right|^2 \end{aligned}$$

The function $\nu_{p,j}^M$ writes as the sum of i.i.d random variables. Let $C \geq 0$,

$$\begin{aligned} & \mathbb{E} \left[\sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} \left| F_{j+1} \left(\vartheta^p, Z^{(m)}, X^{(m)} \right) - \alpha \mathbb{1}_{\{X_{t_j}^{(m)} \in [a_{i-1,j}^p, a]\}} - \beta \mathbb{1}_{\{X_{t_j}^{(m)} \in [a, a_{i,j}^p]\}} \right|^2 \right] \\ & \leq 2\mathbb{E} \left[\left| F_{j+1} \left(\vartheta^p, Z^{(m)}, X^{(m)} \right) \right|^2 \right] + 2C^2 \\ & \leq 2\mathbb{E} \left[\max_{l \geq j+1} (Z_{t_l})^2 \right] + 2C^2 < \infty. \end{aligned}$$

We conclude using Lemma 4.5 that a.s $\nu_{p,j}^M$ converges locally uniformly to the function

$$\alpha, \beta, a \mapsto \mathbb{E} \left[\left| F_{j+1}(\vartheta^p, Z, X) - \alpha \mathbb{1}_{\{X_{t_j} \in [a_{i-1,j}^p, a]\}} - \beta \mathbb{1}_{\{X_{t_j} \in (a, a_{i,j}^p]\}} \right|^2 \right].$$

Let $C > 0$,

$$\begin{aligned} & \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} |\hat{\nu}_{p,j}^M(\alpha, \beta, a) - \nu_{p,j}^M(\alpha, \beta, a)| \\ & \leq \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} \frac{1}{M} \sum_{m=1}^M \left(2 \max_{l \geq j+1} |Z_{t_l}^{(m)}| + 4C \right) \\ & \quad \left[\left(\left(\sum_{i=j+1}^N |Z_{t_i}^{(m)}| \right) \sum_{i=j+1}^{N-1} \mathbb{1}_{\{|Z_{t_i} - \mathcal{T}_p^i(X_{t_i}^{(m)})| \leq |\hat{\tau}_p^{i,M}(X_{t_i}^{(m)}) - \mathcal{T}_p^i(X_{t_i}^{(m)})|\}} \right) \right. \\ & \quad \left. + \left| \alpha \mathbb{1}_{\{X_{t_j}^{(m)} \in [a_{i-1,j}^p, a_{i-1,j}^{p,M}]\}} + \beta \mathbb{1}_{\{X_{t_j}^{(m)} \in [a_{i,j}^{p,M}, a_{i,j}^p]\}} \right| \right]. \end{aligned}$$

Using the induction assumption on p , $\lim_{M \rightarrow \infty} a_{i-1,j}^{p,M} = a_{i-1,j}^p$ and $\lim_{M \rightarrow \infty} a_{i,j}^{p,M} = a_{i,j}^p$ a.s and using the induction assumption on j , $\lim_{M \rightarrow \infty} \hat{\tau}_p^{i,M}(X_{t_i}^{(m)}) = \mathcal{T}_p^i(X_{t_i}^{(m)}) \forall i \geq j+1$. Let $\epsilon < 0$,

$$\begin{aligned} & \limsup_M \sup_{a \in [0,1]^d, |\alpha| \leq C, |\beta| \leq C} |\hat{\nu}_{p,j}^M(\alpha, \beta, a) - \nu_{p,j}^M(\alpha, \beta, a)| \\ & \leq \limsup_M \frac{1}{M} \sum_{m=1}^M \left(2 \max_{l \geq j+1} |Z_{t_l}^{(m)}| + 4C \right) \\ & \quad \left[C \left(\mathbb{1}_{\{|X_{t_j}^{(m)} - a_{i-1,j}^p| \leq \epsilon\}} + \beta \mathbb{1}_{\{|X_{t_j}^{(m)} - a_{i,j}^p| \leq \epsilon\}} \right) + \left(\left(\sum_{k=j+1}^N |Z_{t_k}^{(m)}| \right) \sum_{k=j+1}^{N-1} \mathbb{1}_{\{|Z_{t_k}^{(m)} - \mathcal{T}_p^k(X_{t_k}^{(m)})| \leq \epsilon\}} \right) \right]. \end{aligned}$$

Since $\lim_{\epsilon \rightarrow 0} \mathbb{1}_{\{|Z_{t_i} - \mathcal{T}_p^i(X_{t_i}^{(m)})| \leq \epsilon\}} = \lim_{\epsilon \rightarrow 0} \left| \alpha \mathbb{1}_{\{|X_{t_j}^{(m)} - a_{i-1,j}^p| \leq \epsilon\}} + \beta \mathbb{1}_{\{|X_{t_j}^{(m)} - a_{i,j}^p| \leq \epsilon\}} \right| = 0$, we conclude that a.s $|\hat{\nu}_{p,j}^M(\alpha, \beta, a) - \nu_{p,j}^M(\alpha, \beta, a)| \rightarrow 0$ locally uniformly when $M \rightarrow \infty$, and thus the random function $\hat{\nu}_{p,j}^M$ converges a.s locally uniformly to the function $\alpha, \beta, a \mapsto \mathbb{E} \left[\left| F_{j+1}(\vartheta^p, Z, X) - \alpha \mathbb{1}_{\{X_{t_j} \in [a_{i-1,j}^p, a]\}} - \beta \mathbb{1}_{\{X_{t_j} \in (a, a_{i,j}^p]\}} \right|^2 \right]$. Then we apply Lemma 4.6 along with the arguments of the proof of Lemma 4.7 to conclude that $\mathcal{P}_{p+1}^k(X_{t_k}, \hat{\theta}_k^{p+1,M})$ (resp. $\hat{\theta}_k^{p+1,M}$) converges to $\mathcal{P}_{p+1}^k(X_{t_k}, \theta_k^p)$ (resp. θ_k^{p+1}) a.s as $M \rightarrow \infty$ for $k = N-1, \dots, j$. This completes the induction. \square

Theorem 4.10. Assume that for all $p \in \mathbb{N}^*$, and all $1 \leq j \leq N-1$, $\mathbb{P}(Z_{t_j} = \mathcal{T}_p^j(X_{t_j})) = 0$. Then, for $\alpha = 1, 2$ and for every $j = 1, \dots, N$,

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M \left(Z_{\tau_j^{p,(m)}}^{(m)} \right)^\alpha = \mathbb{E} \left[(Z_{\tau_j^p})^\alpha \right] \text{ a.s.}$$

Proof. Note that $\mathbb{E} \left[(Z_{\tau_j^p})^\alpha \right] = \mathbb{E} [F_j(\vartheta^p, Z, X)^\alpha]$ and by the strong law of large numbers

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M F_j \left((\vartheta^p, Z^{(m)}, X^{(m)})^\alpha \right) = \mathbb{E} [F_j(\vartheta^p, Z, X)^\alpha] \text{ a.s.}$$

It remains to prove that

$$\Delta F_M = \frac{1}{M} \sum_{m=1}^M F_j(\hat{\vartheta}^{p,M}, Z^{(m)}, X^{(m)})^\alpha - F_j(\vartheta^p, Z^{(m)}, X^{(m)})^\alpha \xrightarrow[M \rightarrow \infty]{a.s.} 0.$$

For any $x, y \in \mathbb{R}$, and $\alpha = 1, 2$, $|x^\alpha - y^\alpha| \leq |x - y| |x^{\alpha-1} + y^{\alpha-1}|$. Using Lemma 4.8 and Equation (12), we have

$$\begin{aligned} |\Delta F_M| &\leq \frac{1}{M} \sum_{m=1}^M \left| F_j(\hat{\vartheta}^{p,M}, Z^{(m)}, X^{(m)})^\alpha - F_j(\vartheta^p, Z^{(m)}, X^{(m)})^\alpha \right| \\ &\leq 2 \frac{1}{M} \sum_{m=1}^M \sum_{i=j}^N \max_{k \geq j} |Z_{t_k}^{(m)}|^{\alpha-1} |Z_{t_i}^{(m)}| \sum_{i=j}^{N-1} \mathbb{1}_{\{|Z_{t_i}^{(m)} - \mathcal{T}_p^i(X_{t_i}^{(m)})| \leq |\hat{\mathcal{T}}_p^{i,M}(X_{t_i}^{(m)}) - \mathcal{T}_p^i(X_{t_i}^{(m)})|\}}. \end{aligned}$$

Using Proposition 4.9, for all $i = j, \dots, N-1$, $|\hat{\mathcal{T}}_p^{i,M}(X_{t_i}) - \mathcal{T}_p^i(X_{t_i})| \rightarrow 0$ a.s. when $M \rightarrow \infty$. Then for any $\epsilon > 0$,

$$\begin{aligned} &\limsup_M |\Delta F_M| \\ &\leq 2 \limsup_M \frac{1}{M} \sum_{m=1}^M \sum_{i=j}^N \max_{k \geq j} |Z_{t_k}^{(m)}|^{\alpha-1} |Z_{t_i}^{(m)}| \sum_{i=j}^{N-1} \mathbb{1}_{\{|Z_{t_i}^{(m)} - \mathcal{T}_p^i(X_{t_i}^{(m)})| \leq \epsilon\}} \\ &\leq 2 \mathbb{E} \left[\sum_{i=j}^N \max_{k \geq j} |Z_{t_k}|^{\alpha-1} |Z_{t_i}| \sum_{i=j}^{N-1} \mathbb{1}_{\{|Z_{t_i} - \mathcal{T}_p^i(X_{t_i})| \leq \epsilon\}} \right]. \end{aligned}$$

We conclude that $\limsup_M |\Delta F_M| = 0$ by letting ϵ go to 0 which ends the proof. \square

Proving the global convergence of our algorithm amounts to studying the difference

$$\frac{1}{M} \sum_{i=1}^M Z_{\tau_j^{p,(m)}}^{(m)} - \mathbb{E}[(Z_{\tau_j})] = \left(\frac{1}{M} \sum_{i=1}^M Z_{\tau_j^{p,(m)}}^{(m)} - \mathbb{E}[(Z_{\tau_j^p})] \right) + \left(\mathbb{E}[Z_{\tau_j^p}] - \mathbb{E}[(Z_{\tau_j})] \right)$$

By Proposition 4.4, the $\mathbb{E}[(Z_{\tau_j^p})] - \mathbb{E}[(Z_{\tau_j})]$ goes to zero when $p \rightarrow \infty$ independently of M . Even if Theorem 4.10 proves the a.s. convergence of the first term for any fixed p when $M \rightarrow \infty$, its behaviour seems less clear when both M and p go to infinity. Note that the same difficult already occurred in Clément et al. (2002) for the more standard least square approach. The main difficulty stems for the inability to compute the variance of the limiting distribution appearing in the central limit theorem governing the convergence of $\frac{1}{M} \sum_{i=1}^M \left(Z_{\tau_j^{p,(m)}}^{(m)} \right)$ even in the least square framework. In particular, we do seem to control the behaviour of the limiting variance with respect to p .

5 Random forests

Regression trees are barely used as standalone estimators but are often aggregated to obtain a random forest. Considering our training sample $\mathcal{D}_M = \{(X_i, Y_i)_{1 \leq i \leq M}\}$ of size M . Let Θ be a random variable independent of \mathcal{D}_M and used to resample the training set without replacement. Note that the size of the resampled set is usually smaller than M . Let $\Theta_1, \dots, \Theta_B$ B iid samples of Θ . For every $k = 1, \dots, B$, $\hat{\mathcal{T}}_p^M(X, \Theta_k)$ denotes the regression tree of depth p computed on the resampled set obtained from \mathcal{D}_M by using Θ_k . Then, the random forest estimator is defined by

$$\mathcal{H}_{B,p}(X) = \sum_{k=1}^B \frac{1}{B} \hat{\mathcal{T}}_{p,\Theta_k}^M(X).$$

We also introduce $\mathcal{H}_p(X) = \mathbb{E}_\Theta \left[\hat{\mathcal{T}}_{p,\Theta}^M(X) \right] = \lim_{B \rightarrow \infty} \mathcal{H}_{B,p}(X)$.

Theorem 5.1.

$$\lim_{B \rightarrow \infty} \mathbb{E} \left[|Y - \mathcal{H}_{B,p}(X)|^2 \right] = \mathbb{E} \left[|Y - \mathcal{H}_p(X)|^2 \right]$$

See Theorem 11.1 in (Breiman, 2001).

Theorem 5.2.

$$\mathbb{E} \left[|Y - \mathcal{H}_p(X)|^2 \right] \leq \bar{\rho} \mathbb{E}_{\Theta} \left[\mathbb{E} \left[|Y - \hat{\mathcal{T}}_{p,\Theta}^M(X)|^2 \right] \right]$$

where $\bar{\rho}$ is the weighted correlation between the residuals $Y - \hat{\mathcal{T}}_{p,\Theta}^M(X)$ and $Y - \hat{\mathcal{T}}_{p,\Theta'}^M(X)$ and Θ and Θ' are independent. See Theorem 11.2 in (Breiman, 2001)

Theorem 5.2 says that to have a good generalization error in the random forest, one should have small generalization errors in the basis trees, and the basis trees should not be highly correlated.

6 Numerical results

6.1 Description

This section studies the price of some Bermudan options using regression trees or random forests to approximate the conditional expectations. We compare the results to some reference prices and those given by the standard Longstaff Schwarz method with regression on polynomial functions, which is a basic LSM algorithms. More sophisticated representations can be used to mildly temper the curse of dimensionality. We use the Scikit-Learn library in Python, (Pedregosa et al., 2011). For regression trees, this library offers two methods of splitting: “best” to choose the best split, meaning that the split threshold is the one that minimizes the MSE and the direction for splitting is the one that gives the lowest MSE among all directions; “random” to choose the best random split, meaning that the split threshold is the one that minimizes the MSE and the direction for splitting is chosen randomly. For the following tests, we will use the latter method, which is just slightly different from what we presented in Section 2 in the way that no mid-point cuts will be considered. We also use the feature `min_samples_leaf` which allows us to set a minimum number of samples in each node. This will allow us to avoid over-fitting. For random forests, we will use the bootstrapping method (`Bootstrap=True`), meaning that for each tree in the forest, we will use a sub-sample drawn randomly and with replacement from the training data. We will also use the feature `max_samples` which allows having a specific number of data points or a percentage of the training data attributed to each tree. Having the trees trained on different data as much as possible allows us to have a low correlation between the trees which, using Theorem 5.2, should make the random forest more robust.

Following the work of (Longstaff and Schwartz, 2001), we only use the in-the-money paths to learn the continuations values, which significantly improves the numerical computations. All the prices that we show are obtained after resimulation, meaning that the paths used in the estimation of the conditional expectations are not the same ones used by the Monte Carlo which means that the prices we show are unbiased.

For small dimensional problems, our algorithm takes approximately the same computation time as the LSM approach for a comparable accuracy. To measure the computational time of our algorithm, we take into account both the training and running times. However, note that the training step over weights the prediction step by a great deal. Once the trees are trained, evaluating them is straightforward, which makes the resimulation price almost free. The computational time of the random forest approach linearly depends on the number of regression trees in the forest.

6.2 Black and Scholes

Consider the Black and Scholes model

$$\begin{cases} dS_t^i & = rS_t^i dt + \sigma_i S_t^i dB_t^i, \\ d \langle B^i, B^j \rangle_t & = \rho_{ij} dt. \end{cases}$$

where σ_i is the volatility of the underlying S^i , assumed to be deterministic, r is the interest rate, assumed constant, and ρ_{ij} , represents the correlation between the underlyings S^i and S^j , assumed constant.

6.2.1 One-dimensional put option

We consider the Bermudan put option with payoff $(K - S_\tau)^+$ with maturity $T = 1$ year, $K = 110$, $S_0 = 100$, $\sigma = 0.25$, exercisable at $N = 10$ different dates. We consider $r = 0.1$. We have a reference price for this option of 11.987 computed by a convolution method in (Lord et al., 2007). The LSM algorithm converges to the correct price with only a polynomial of degree 3. Figure 1, shows the price of the option when we use regression trees with a random split strategy (continuous line) or a best split strategy (dotted line) to estimate the conditional expectations. With the random strategy, the best price we get is 11.89. The case `min_samples_leaf=1` and `max_depth=20` gives a price of 10.5, which is far from the reference price. This result is due to over-fitting. In fact, for this case, the number of degrees of freedom is too big. The tree fits the training data too well, but it cannot generalize when confronted with new data. For the best split strategy, we obtain a slightly better price of 11.94. However, depending on the tree parameters, the price fluctuates, and we can see that the best split strategy is not necessarily better than the random split strategy. Thus, for the following, we will keep using the random split strategy. Random forests with basis trees of maximum depth 5 and minimum 100 samples in each leaf converge to the correct price with only ten trees.




Figure 1: one dimensional put with regression trees, true price=11.987

6.2.2 Call option on the maximum of two assets

We consider a call option on the maximum of 2 assets with payoff $(\max(S_\tau^1, S_\tau^2) - K)^+$, we use the same set of parameters as in (Glasserman, 2004), for which we have reference prices of 13.90, 8.08 and 21.34 for $S_0^i = 100, 90$ and 110 respectively. The LSM algorithm using a polynomial of degree 5 converges to a price of 13.90, 8.06, 21.34 for the cases $K = 100, 90, 110$ respectively. This is a small dimensional problem, so the convergence of the LSM is expected. With regressions trees we have slightly less satisfying results as shown in Figure 2. We can still see the case of over-fitting when giving the regression trees too many degrees of freedom. Aggregating the regression trees into random forests immediately improves the results as shown in Figure 3. Note that the lower the percentage of data in each basis tree, the better the results. This confirms the results of Theorem 5.2 .




Figure 2: Call on the maximum of two assets with regression trees, $K = 100, T = 3$ years, $\sigma^i = 0.2, r = 0.05, \rho_{ij} = 0, \delta_i = 0.1, N = 9, M = 100,000$

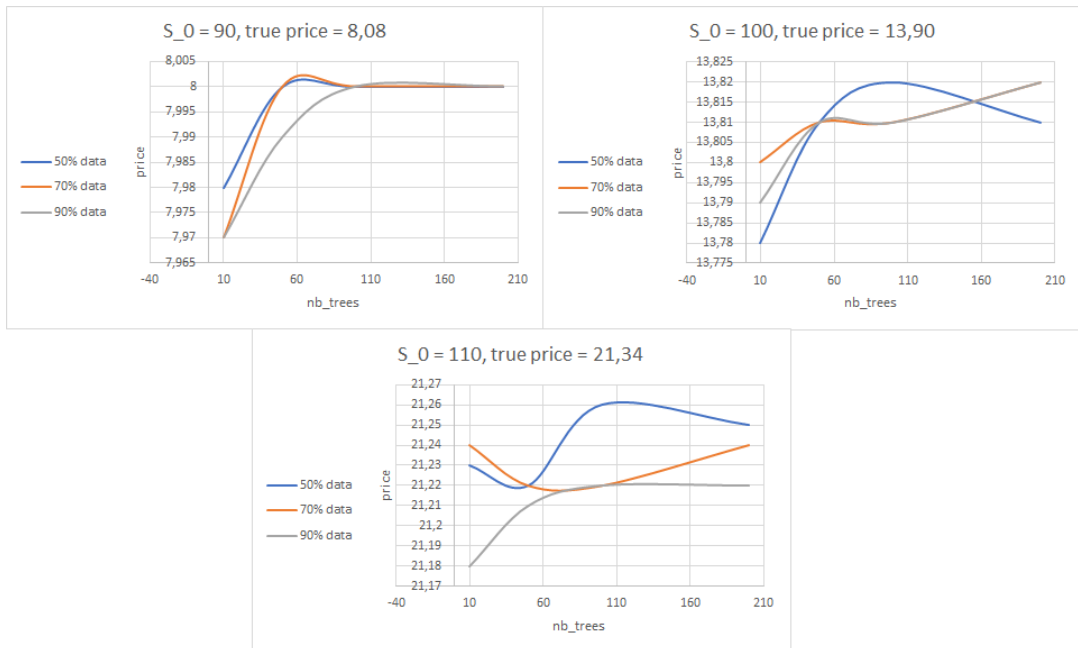


Figure 3: Call on the maximum of two assets with random forests, $K = 100, T = 3$ years, $\sigma^i = 0.2, r = 0.05, \rho_{ij} = 0, \delta_i = 0.1, N = 9, M = 100,000$

6.2.3 Geometric basket option

We consider a Bermudan Put option on a geometric basket of d underlying with payoff

$$\left(K - \left(\prod_{i=1}^d S_{\tau}^i \right)^{\frac{1}{d}} \right)^+.$$

We test the following option for $d = 2, 10, 40$ for which we have reference prices from (Cox et al., 1979) using the CRR tree method. With the LSM algorithm, we converge to the correct price 4.57 for the case $d = 2$, using only a polynomial of degree 3. For the case $d = 10$, we can at most use a polynomial of degree 3 due to the curse of dimensionality. With this parametrization, we obtain a price of 2.90 for a true price of 2.92. For the case $d = 40$, we cannot go further than a polynomial of degree 1, which yields a price of 2.48 for a reference price of 2.52. Figure 4 shows the results obtained with regression trees. For the case $d = 2$, the best price we get is 4.47 and, as expected, the LSM algorithm has a better performance. This is also the case for the cases $d = 10$ and $d = 40$ where the best prices we obtain are 2.84 and 2.46 respectively. Notice that even though these are high dimensional cases, the trees converge with only a depth of 5 or 8. For $d = 10$, our algorithm with `depth` equal 8 and `min_samples_leaf` equal 200 runs in 3 seconds compared to the 11 seconds needed by the LSM with polynomial degree 3. The running times of the two approaches are pretty comparable. For $d = 40$, our algorithm takes the same computational time as the LSM with degree 1 polynomials. We also notice the importance of the parameter `min_samples_leaf`. In fact, letting the trees grow without managing this parameter (case `leaf1`) leads to a problem of over-fitting. The results get better when we use random forests as shown in Figure 5. For these random forests, we used basis trees of `max_depth=8` and `min_samples_leaf=100`. Notice for the case $d = 2$, the curve where only 50% of the data is used gives much better results as in this case the basis trees are the less correlated. For the cases $d = 10$ and $d = 40$, the best choice is not necessarily to use 50% of the data in each tree. As these are larger dimensions, having the trees trained on a small percentage of the training data maybe not enough. One may consider extending the size of the training data itself. Furthermore, we notice that once the percentage of data to use in each tree is chosen, the price of the option converges as the number of trees in the forest grows. However, note the computational time of the random forest method linearly depends on the number of trees inside the forest. Although the regression tree approach runs within roughly the same computational time as the LSM for a similar accuracy, the random forests approach may take much longer as the number of trees increases.




Figure 4: Geometric put option with regression trees

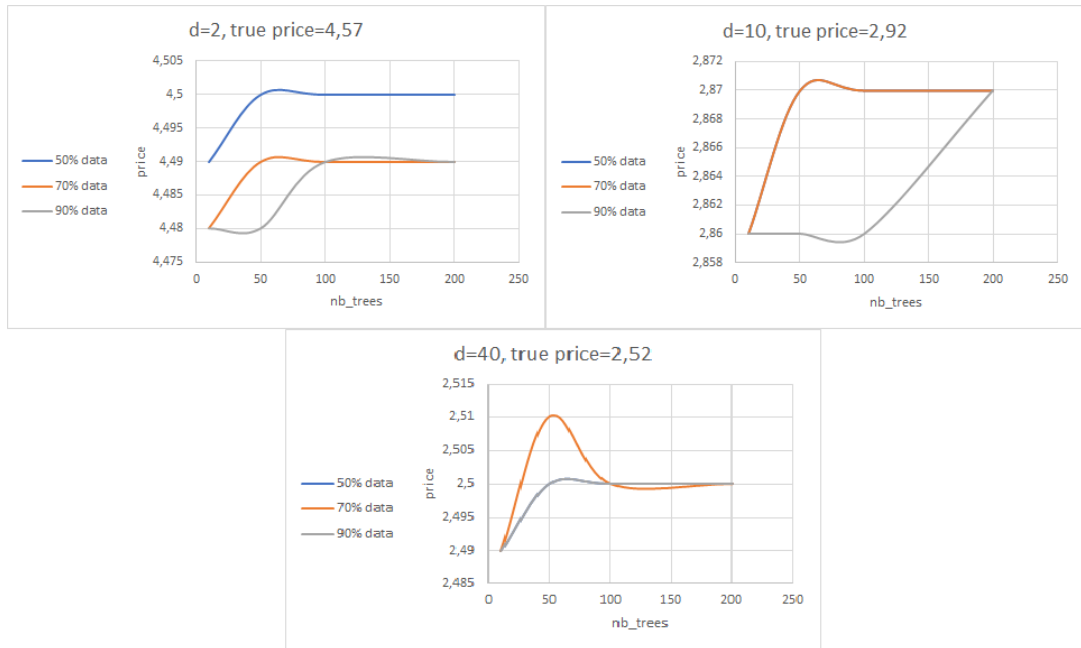


Figure 5: Geometric put option with random forests

6.2.4 A put basket option

We consider a put option on the basket of $d = 40$ asset with payoff $\left(K - \sum_{i=1}^d \omega_i S_T^i\right)^+$. We test this payoff for $d = 40$ for which we have a reference price from (Goudenège et al., 2019) between 2.15 and 2.22 using the following set of parameters: $T = 1, S_i = 100, K = 100, r = 0.05, \sigma_i =$

0.2, $\rho_{ij} = 0.2$, $\omega_i = \frac{1}{d}$ and $N = 10$. With a polynomial of degree 1, we obtain a price of 2.15 using the LSM algorithm. The results obtained with regression trees are shown in Figure 6.




Figure 6: Put on a basket of 40 asset with regression trees

Even though this example is high dimensional, we do not need a lot of parameters to estimate the conditional expectations (the trees converge for very small depths). This will not be the case for the next example which is very non linear. The aggregation into random forests leads to a price of 2.16 using only 50 trees.

6.2.5 A call on the max of 50 asset

We consider a call option on the maximum of $d = 50$ asset with payoff $\left(\max_{1 \leq i \leq d} (S_T^i) - K\right)^+$ with the following characteristics: $K = 100$, $T = 3$ years, $S_0^i = 100$, $\sigma_i = 0.2$, $\delta_i = 0.1$, $\rho_{ij} = 0 \forall i, j$, $r = 0.05$, $N = 9$, $M = 100000$. (Becker et al., 2019) report [69.56, 69.95] as the 95% confidence interval for the option price. With the LSM algorithm we find a price of 67.88 with a polynomial of degree 1. This is a difficult example for which bigger trees are needed to approach the conditional expectations. At maturity, the payoff depends only on one direction (corresponding to the best performance), if the cuts in the tree never consider that direction, the estimation will not be correct. As a result, we consider a number of cuts big enough to ensure that each direction is taken into consideration. We allow the depth to grow while monitoring the `min_samples_leaf` in order to have a significant number of samples in each leaf. Table 1 shows the results obtained with regression trees. As the best price we obtain is given by `depth=100` and `min_samples_leaf = 100`, we use this set of parameters for the random forest part. Table 2 shows the results that we obtain with this method.

depth	min_samples_leaf	price
50	50	66,89
50	100	66.88
100	50	67.13
100	100	67.31
200	50	67.16
200	100	67.28

Table 1: A call option on the maximum of 50 asset with regression trees

nb_trees	max_samples	price
10	50%	68,32
10	70%	68,32
10	90%	68,29

Table 2: A call option on the maximum of 50 asset with random forests

Using only regression trees is not enough to have acceptable results. However, as soon as we aggregate the regressor into random forests, we obtain very satisfying results and with just 10 trees we converge to a good price. We can also notice in this example that using uncorrelated trees leads to better results (see the case `max_samples= 50%` or `70%` against the case `max_samples = 90%`).

6.3 A put in the Heston model

We consider the Heston model defined by

$$dS_t = S_t(r_t dt + \sqrt{\sigma_t}(\rho dW_t^1 + \sqrt{1 - \rho^2} dW_t^2))$$

$$d\sigma_t = \kappa(\theta - \sigma_t)dt + \xi\sqrt{\sigma_t}dW_t^1$$

and we consider a put option with payoff $(K - S_T)^+$. we have no reference price for this option, so we will just compare the results of regression trees and random forests to the LSM method. We use the following set of parameters: $K = 100, S_0 = 100, T = 1, \sigma_0 = 0.01, \xi = 0.2, \kappa = 2, \rho = -0.3, r = 0.1, N = 10$ and $M = 100,000$. The LSM method yields a price of 1.70. Figures 7 and 8 show the results obtained with regression trees and random forests. Both methods converge to the same price of LSM. We notice for this example the occurrence of the over-fitting phenomenon for regression trees with `max_depth=15` and `min_sample_leaf=1`. We also have the same behavior for random forests in function of the percentage of data given to each basis tree.




Figure 7: A put option in the Heston model with regression trees




Figure 8: A put option in the Heston model with random forests

7 Conclusion

Pricing Bermudan options comes down to solving a dynamic programming equation where the main trouble comes from the computation of the conditional expectations representing the conditional expectations. We have explored the usage of regression trees and random forests for the computations of these quantities. We have proved in two steps the convergence of the algorithm when regression trees are used: first, the convergence of the conditional expectations; Then, the convergence of the Monte Carlo approximation. This problem was particularly hard to solve given that the regression trees do not solve a global optimization problem as does the functional regression used in the LSM algorithm. We have shown through numerical experiments that we obtain good prices for some classical examples using regression trees. The aggregation of regression trees into random forests yields even better results. We came to the conclusion that for small dimensional problems, a simpler algorithm like the LSM is efficient enough. However, for high dimensional problems, it is interesting to consider using random forests. Instead of using all the features of the problem, the basis trees in the forest only use a subset of the features which can help combat the problem of the curse of dimensionality.

References

- Leif Andersen and Mark Broadie. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science*, 50(9), 2004. ISSN 00251909. doi: 10.1287/mnsc.1040.0258.
- Vlad Bally, Gilles Pagès, and Jacques Printems. A quantization tree method for pricing and hedging multidimensional american options. *Mathematical Finance*, 15(1), 2005. ISSN 09601627. doi: 10.1111/j.0960-1627.2005.00213.x.
- Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Deep optimal stopping. *The Journal of Machine Learning Research*, 20(1), January 2019. doi: 10.5555/3322706.3362015.
- Leo Breiman. Using Adaptive Bagging to Debias Regressions. *Technical Report 547*, 1(October), 1999. ISSN 1098-6596.
- Leo Breiman. Random forests. *Machine Learning*, 45(1), 2001. ISSN 08856125. doi: 10.1023/A:1010933404324.
- Jacques F. Carriere. Valuation of the early-exercise price for options using simulations and non-parametric regression. *Insurance: Mathematics and Economics*, 19(1), 1996. ISSN 01676687. doi: 10.1016/S0167-6687(96)00004-2.
- Dragos Florin Ciocan and Velibor V Mišić. Interpretable optimal stopping. *Management Science*, 68(3):1616–1638, 2022.
- Emmanuelle Clément, Damien Lamberton, and Philip Protter. An analysis of a least squares regression method for American option pricing. *Finance and Stochastics*, 6(4), 2002. ISSN 09492984. doi: 10.1007/s007800200071.
- John C. Cox, Stephen A. Ross, and Mark Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3), 1979. ISSN 0304405X. doi: 10.1016/0304-405X(79)90015-1.
- Thomas G. Dietterich. Experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization. *Machine Learning*, 40(2), 2000. ISSN 08856125. doi: 10.1023/A:1007607513941.
- Dietmar Ferger. A continuous mapping theorem for the argmax-functional in the non-unique case. *Statistica Neerlandica*, 58(1):83–96, February 2004. doi: 10.1046/j.0039-0402.2003.

- Paul Glasserman. Monte Carlo method in financial engineering, 2004. ISSN 14697688.
- Ludovic Goudenège, Andrea Molent, and Antonino Zanette. Variance reduction applied to machine learning for pricing bermudan/american options in high dimension, 2019.
- Jiequn Han and Weinan E. Deep learning approximation for stochastic control problems. *CoRR*, abs/1611.07422, 2016. URL <http://arxiv.org/abs/1611.07422>.
- Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- Tin Kam Ho. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8), 1998. ISSN 01628828. doi: 10.1109/34.709601.
- Michael Kohler, Adam Krzyzak, and Nebojsa Todorovic. Pricing of high-dimensional american options by neural networks. *Mathematical Finance*, 20(3), 2010. ISSN 09601627. doi: 10.1111/j.1467-9965.2010.00404.x.
- Bernard Lapeyre and Jérôme Lelong. Neural network regression for Bermudan option pricing. *Monte Carlo Methods and Applications*, 27(3), 2021. ISSN 15693961. doi: 10.1515/mcma-2021-2091.
- Charles Leake, Reuven Y. Rubinstein, and Alexander Shapiro. Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method. *The Journal of the Operational Research Society*, 45(8), 1994. ISSN 01605682. doi: 10.2307/2584023.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces*. Springer, 1991. doi: 10.1007/978-3-642-20212-4.
- Jérôme Lelong. Dual pricing of American options by wiener chaos expansion. *SIAM Journal on Financial Mathematics*, 9(2), 2018. ISSN 1945497X. doi: 10.1137/16M1102161.
- Francis A. Longstaff and Eduardo S. Schwartz. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1), 2001. ISSN 08939454. doi: 10.1093/rfs/14.1.113.
- R. Lord, F. Fang, F. Bervoets, and C. W. Oosterlee. A fast and accurate FFT-based method for pricing early-exercise options under levy processes. *SIAM Journal on Scientific Computing*, 30(4), 2007. ISSN 10648275. doi: 10.1137/070683878.
- Mike Ludkovski. Kriging metamodels and experimental design for bermudan option pricing. *Journal of Computational Finance*, 22(1), 2018. ISSN 17552850. doi: 10.21314/JCF.2018.347.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- A Max Reppen, H Mete Soner, and Valentin Tissot-Daguette. Neural optimal stopping boundary. *arXiv preprint arXiv:2205.04595*, 2022.
- L. C.G. Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3), 2002. ISSN 09601627. doi: 10.1111/1467-9965.02010.
- Emilio Seijo and Bodhisattva Sen. A continuous mapping theorem for the smallest argmax functional. *Electronic Journal of Statistics*, 5(none):421 – 439, 2011.
- John N. Tsitsiklis and Benjamin Van Roy. Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control*, 44(10), 1999. ISSN 00189286. doi: 10.1109/9.793723.