René Aïd, Luciano Campi, Nicolas Langrené, Huyên Pham

A probabilistic numerical method for optimal multiple switching problems in high dimension

Article (Accepted version) (Refereed)

Original citation: Aïd, René, Campi, Luciano, Langrené, Nicolas and Pham, Huyên (2014) *A probabilistic numerical method for optimal multiple switching problems in high dimension.* SIAM Journal on Financial Mathematics, 5 (1). pp. 191-231. ISSN 1945-497X

DOI: 10.1137/120897298

A probabilistic numerical method for optimal multiple switching problem in high dimension *

René Aïd[†] , Luciano Campi[‡] , Nicolas Langrené[§] , Huyên Pham[¶]

First version: October 30, 2012
This revised version: December 13, 2013

Abstract

In this paper, we present a probabilistic numerical algorithm combining dynamic programming, Monte Carlo simulations and local basis regressions to solve non-stationary optimal multiple switching problems in infinite horizon. We provide the rate of convergence of the method in terms of the time step used to discretize the problem, of the regression basis used to approximate conditional expectations, and of the truncating time horizon. To make the method viable for problems in high dimension and long time horizon, we extend a memory reduction method to the general Euler scheme, so that, when performing the numerical resolution, the storage of the Monte Carlo simulation paths is not needed. Then, we apply this algorithm to a model of optimal investment in power plants in dimension eight, i.e. with two different technologies and six random factors.

1 Introduction

This paper presents a probabilistic numerical method for multiple switching problem. Our approach in this paper takes advantage of the considerable progress made in the last ten years by numerical methods for high-dimensional American options valuation problems. For an up-to-date state of the art on this subject, the reader is referred to the recent book [9].

In this paper, we first adapt the resolution of American options problems by Monte-Carlo methods and regression ([28, 35]), to the more general class of optimal switching problems. The crucial choice of regression basis is done here in the light of the work of [7], so as to obtain a stable algorithm suited to high-dimensional problems, aiming at the best possible numerical complexity. The memory complexity, often acknowledged as the major drawback of this Monte Carlo approach (see [10]), is drastically slashed by generalizing the memory reduction method from [12, 13, 14] to any stochastic differential equation. We provide a rigorous and comprehensive analysis of the rate of convergence of our algorithm, taking advantage of the works of, most notably, [6], [33] and [18]. Note that such unusual features as infinite horizon and non-stationarity are encompassed here.

Finally, we apply our algorithm to a long-term investment model for electricity generation based on a structural model for the spot price of electricity developed in [3] and [1]. This model has been shown to suitably reproduce the statistical and dynamical properties of the spot price of electricity. Nevertheless, to suit the purpose of long-term electricity price modeling, it has been adapted

*The authors would like to thank Thomas Vareschi, Xavier Warin and the participants of the FiME seminar and the Energy Finance conference for their helpful remarks. The authors are also grateful to the two anonymous referees for all their comments, which helped to improve the first version of this paper.

[†]EDF R&D and FiME (Finance for Energy Market Research Centre, www.fime-lab.org) ; rene.aid@edf.fr

[‡]Univ Paris Nord, Sorbonne Paris Cité, LAGA (Laboratoire Analyse, Géométrie et Applications), FiME and CREST ; campi@math.univ-paris13.fr

[§]Univ Paris Diderot, Sorbonne Paris Cité, LPMA (Laboratoire de Probabilités et Modèles Aléatoires, CNRS, UMR 7599) ; langrene@math.univ-paris-diderot.fr

[¶]Univ Paris Diderot, Sorbonne Paris Cité, LPMA and CREST-ENSAE, pham@math.univ-paris-diderot.fr

and extended in several directions (cointegrated fuels and CO_2 prices, stochastic availability rate of production capacities, new scarcity function). The resolution of this problem using our algorithm is illustrated on a simple numerical example with two different technologies, leading to an eight-dimensional problem (demand, CO_2 price, and, for each technology, fuel price, random outages and the controlled installed capacity). The time evolution of the distribution of power prices and of the generation mix is illustrated on a forty-year time horizon. To the knowledge of the authors, the highest dimension considered so far in the case of long-term investment models in electricity generation was three ([29, 5]).

The contribution of the paper is twofold. Firstly, it provides a comprehensive analysis of convergence of a regression-based Monte-Carlo algorithm for a class of infinite horizon optimal multiple switching problems, large enough to handle realistic short term profit functions and investment cost structures with possible seasonality patterns. Secondly, we implement successfully our algorithm to a new stylized investment model for electricity generation, by adapting and generalizing a memory reduction method. A numerical resolution of this investment problem with our algorithm is illustrated on a specific example, providing, among many other outputs, an electricity spot price dynamics consistent with the investment decision process in power generation.

The outline of the paper is the following. Section 2 presents the class of optimal switching problems studied here, including the detailed list of assumptions considered. Section 3 describes the resolution algorithm and analyzes its rate of convergence, in terms of the discretization step, of the choice of regression basis, and of the truncating time horizon. Section 4 details the computational complexity of the algorithm, as well as its memory complexity, along with the construction of the memory reduction method. In Section 5, we implement and illustrate numerically our algorithm on an investment model in electricity generation based on an extended structural model of power spot price. Finally, Section 6 concludes the paper.

Notation

Here are some notation that will be used throughout the paper:

- The notation $\mathbf{1}\{.\}$ stands for the indicator function.
- Throughout the paper, $C > 0$ denotes a generic constant whose value may differ from line to line, but which does not depend on any parameter of our scheme.
- For any stochastic process $X = (X_s)_{s\geq 0}$ taking values in a given set $\mathcal{X}$, and any $(t, x) \in \mathbb{R}_+ \times \mathcal{X}$, we denote as $X^{t,x} = (X_s^{t,x})_{s\geq t}$ the stochastic process with the same dynamics as X, but starting from x at time t: $X_t^{t,x} = x$.
- For any $(a, b) \in \mathbb{R} \times \mathbb{R}$, $a \wedge b := \min(a, b)$ and $a \vee b := \max(a, b)$.
- $\forall p \geq 1$, the norms $\|.\|_p$ and $\|.\|_{L_p}$ denote respectively the $p-$norm and the L_p- norm: $\forall x \in \mathbb{R}^n$ and any $\mathbb{R}$-valued random variable X such that $\mathbb{E}[|X|^p] < \infty$:

$$\|x\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{\frac{1}{p}} \quad , \quad \|X\|_{L_p} = \mathbb{E}[|X|^p]^{\frac{1}{p}}$$

  We recall that $\forall p \geq 1$, $\forall x \in \mathbb{R}^n$, $\|x\|_p \leq \|x\|_1 \leq n^{\frac{p-1}{p}} \|x\|_p$

2 Optimal switching problem

2.1 Formulation

Fix a filtered probability space $\left(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}\right)$, where $\mathbb{F}$ satisfies the usual conditions of right-continuity and $\mathbb{P}$-completeness. We consider the following general class of (non-stationary) optimal

switching problems:

$$v(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}} \mathbb{E}\left[\int_t^\infty f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{\tau_n \geq t} k\left(\tau_n, \zeta_n\right)\right] \tag{2.1}$$

where:

- $X^{t,x} = (X_s^{t,x})_{s \geq t}$ is an $\mathbb{R}^d$-valued, $\mathbb{F}$-adapted Markovian diffusion starting from $X_t = x \in \mathbb{R}^d$, with generator $\mathcal{L}$.
- $I^\alpha = (I_s^\alpha)_{s \geq 0}$ is a càd-làg, $\mathbb{R}^{d'}$-valued, $\mathbb{F}$-adapted piecewise constant process. It is controlled by a strategy α, described below. We suppose it can only take values into a fixed finite set $\mathbb{I}_q = \{i_1, i_2, \ldots, i_q\}$, $q \in \mathbb{N}^*$ with $i_1 = 0$ $\left(\in \mathbb{R}^{d'}\right)$, which means that equation (2.1) corresponds to an optimal switching problem.
- An impulse control strategy α corresponds to a sequence $(\tau_n, \iota_n)_{n \in \mathbb{N}}$ of increasing stopping times $\tau_n \geq 0$, and $\mathcal{F}_{\tau_n}$-measurable random variables ι_n valued in $\mathbb{I}_q$. Using this sequence, $I^\alpha = (I_s^\alpha)_{s \geq 0}$ is defined as follows:

  $$I_s^\alpha = \iota_0 \mathbf{1}\{0 \leq s < \tau_0\} + \sum_{n \in \mathbb{N}} \iota_n \mathbf{1}\{\tau_n \leq s < \tau_{n+1}\} \in \mathbb{I}_q$$

  Alternatively, α can be described by the sequence $(\tau_n, \zeta_n)_{n \in \mathbb{N}}$, where $\zeta_n := \iota_n - \iota_{n-1}$ (and $\zeta_0 := 0$). Using this alternative sequence, I^α can be written as follows:

  $$I_s^\alpha = \iota_0 + \sum_{\tau_n \leq s} \zeta_n \in \mathbb{I}_q$$

- $\mathcal{A}$ is the set of admissible strategies: a strategy α belongs to $\mathcal{A}$ if $\tau_n \to +\infty$ a.s. as $n \to \infty$.
- For any $(t,i) \in \mathbb{R}_+ \times \mathbb{I}_q$, the set $\mathcal{A}_{t,i} \subset \mathcal{A}$ is defined as the subset of admissible strategies α such that $I_t^\alpha = i$.
- f and k are $\mathbb{R}$-valued measurable functions.

2.2 Assumptions

We complete the above formulation with the following relevant assumptions.

Assumption 1. *[Diffusion] The $\mathbb{R}^d$-valued uncontrolled process X is a diffusion process, governed by the dynamics*

$$\begin{aligned} dX_s &= b(s, X_s)\, ds + \sigma(s, X_s)\, dW_s \\ X_0 &= x_0 \in \mathbb{R}^d \end{aligned} \tag{2.2}$$

where W is a d-dimensional Brownian motion, and b and σ are respectively $\mathbb{R}^d$-valued and $\mathbb{R}^{d \times d}$-valued functions.

Assumption 2. *[Lipschitz] The functions $b : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}^{d \times d}$ are Lipschitz-continuous (uniformly in t) with linear growth: $\exists C_b, C_\sigma > 0$ s.t. $\forall t \in \mathbb{R}_+$, $\forall (x, x') \in \left(\mathbb{R}^d\right)^2$:*

$$\begin{aligned} |b(t,x) - b(t,x')| &\leq C_b |x - x'| \\ |b(t,x)| &\leq C_b (1 + |x|) \\ |\sigma(t,x) - \sigma(t,x')| &\leq C_\sigma |x - x'| \\ |\sigma(t,x)| &\leq C_\sigma (1 + |x|) \end{aligned}$$

Remark 2.1. Assumption 2 is sufficient to prove the existence and uniqueness of a strong solution to the SDE (2.2) (see for instance Theorem 4.5.3 in [23]).

Remark 2.2. Under Assumption 2, there exist, for every $p \geq 1$, positive constants C_p and ρ_p such that $\forall s \geq t \geq 0$ and $\forall x \in \mathbb{R}^d$:

$$\mathbb{E}\left[\left|X_s^{t,x}\right|^p\right] \leq C_p\left(1+|x|^p\right)\exp\left(\rho_p\left(s-t\right)\right) \tag{2.3}$$

(use Burkholder-Davis-Gundy inequality and Gronwall's Lemma, see for instance [23] Theorem 4.5.4 for the even power case).

Assumption 3. *[Lipschitz&Discount] The functions f and k decrease exponentially in time: $\exists \rho > 0$ s.t. $\forall\left(t,x,i,j\right) \in \mathbb{R}_+ \times \mathbb{R}^d \times \left(\mathbb{I}_q\right)^2$:*

$$\begin{aligned} f\left(t,x,i\right) &= e^{-\rho t}\tilde{f}\left(t,x,i\right) \\ k\left(t,j-i\right) &= e^{-\rho t}\tilde{k}\left(t,j-i\right) \end{aligned}$$

where the functions $\tilde{f}$ and $\tilde{k}$ are Lipschitz continuous with linear growth:

$\exists C_f, C_k > 0$ *s.t.* $\forall\left\{\left(t,x,i,j\right),\left(t',x',i',j'\right)\right\} \in \left\{\mathbb{R}_+ \times \mathbb{R}^d \times \left(\mathbb{I}_q\right)^2\right\}^2$*:*

$$\begin{aligned} \left|\tilde{f}\left(t,x,i\right)-\tilde{f}\left(t',x',i'\right)\right| &\leq C_f\left(|t-t'|+|x-x'|+|i-i'|\right) \\ \left|\tilde{f}\left(t,x,i\right)\right| &\leq C_f\left(1+|x|\right) \\ \left|\tilde{k}\left(t,j-i\right)-\tilde{k}\left(t',j'-i'\right)\right| &\leq C_k\left(|t-t'|+|(j-i)-(j'-i')|\right) \end{aligned}$$

Moreover, we assume in the following that $\rho > \rho_1$ where ρ_1 is defined in equation (2.3).

Assumption 4. *[Fixed costs] The cost function $k : \mathbb{R}_+ \times \mathbb{R}^{d'} \to \mathbb{R}_+$ is such that:*

- $\forall t \in \mathbb{R}_+,\ k\left(t,0\right) = 0.$
- $\exists \kappa > 0$ s.t. $\forall t \in \mathbb{R}_+,\ \forall\left(i,j\right) \in \left(\mathbb{I}_q\right)^2,\ \{i \neq j\} \Rightarrow \left\{\tilde{k}(t,j-i) \geq \kappa\right\}.$
- *(triangular inequality)* $\forall t \in \mathbb{R}_+,\ \forall\left(i,j,k\right) \in \left(\mathbb{I}_q\right)^3$ *with* $i \neq j$ *and* $j \neq k$*:*

$$k(t,k-i) < k(t,j-i) + k(t,k-j)\ .$$

Remark 2.3. The economic interpretations of Assumption 4 are the following:

1. There is no cost for not switching, but any switch incurs at least a positive fixed cost.
2. At any given date, it is always cheaper to switch directly from i to k than to switch first from i to j and then from j to k.

Remark 2.4. Under those standard assumptions, the value function v from equation (2.1) is well-defined and finite. Indeed, using equation (2.3), $\forall\left(t_0,t,x,i\right) \in \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{R}^{d'}$ with $t_0 \leq t$ and $\forall \alpha \in \mathcal{A}_{t_0,i}$:

$$\begin{aligned} \mathbb{E}\left[\int_t^\infty \left|f\left(s,X_s^{t_0,x},I_s^\alpha\right)\right| ds\right] &\leq C_f\int_t^\infty e^{-\rho s}\left(1+\mathbb{E}\left[\left|X_s^{t_0,x}\right|\right]\right) ds \\ &\leq C_f\left(e^{-\rho t}+\left(1+|x|\right)\int_t^\infty e^{-\rho s}e^{\rho_1(s-t_0)}ds\right) \\ &\leq C_f\left(1+|x|\right)e^{-\bar{\rho}t-\rho_1 t_0} \end{aligned} \tag{2.4}$$

where $\bar{\rho} := \rho - \rho_1 > 0$ (Assumption 3). In particular, the costs being positive (Assumption 4), and recalling (2.1), it holds that:

$$v\left(t,x,i\right) \leq C_f\left(1+|x|\right)e^{-\rho t} \tag{2.5}$$

2.3 Outline of the solution

From a theoretical point of view, the value functions $v_i := v(.,.,i)$, $i \in \mathbb{I}_q$ from equation (2.1) are known to satisfy (under suitable conditions on $f_i(.,.) := f(.,.,i)$ and k, see for instance [32] in a much more general setting) the following Hamilton-Jacobi-Bellman Quasi-Variational Inequalities (HJBQVI): $\forall (t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{I}_q$

$$\min\left\{-\frac{\partial v_i}{\partial t}(t,x) - \mathcal{L}v_i(t,x) - f_i(t,x)\ ,\ v_i(t,x) - \max_{j\in\mathbb{I}_P,\, j\neq i}\left(v_j(t,x) - k(t,j-i)\right)\right\} = 0 \tag{2.6}$$

together with suitable limit condition, which ensure existence and unicity of the solution to this system (cf. [20] for instance).

Alternatively, the process $v(t,X_t,i)$, $t \geq 0$ can be characterized as the solution of a particular Reflected Backward Stochastic Differential Equation ([21, 16]).

Moreover, the value function (2.1) satisfies the well-known dynamic programming principle, i.e., for any stopping time $\tau \geq t$:

$$v(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}} \mathbb{E}\left[\int_t^\tau f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t\leq\tau_n\leq\tau} k(\tau_n,\zeta_n) + v\left(\tau, X_\tau^{t,x}, I_\tau^\alpha\right)\right]. \tag{2.7}$$

From a practical point of view, apart from a few simple examples in low-dimension, finding directly the solution of the HJBQVI (2.6) is usually infeasible, and the numerical PDE tools become cumbersome and inefficient in the multi-dimensional setting. Instead, probabilistic methods based on (2.7), in the spirit of [10], are usually more practical and versatile.

Indeed, as the diffusion X is not controlled, this optimal switching problem can be seen as an extended American option problem. This suggests that, up to some adjustments, the probabilistic numerical tools developed in this context (see [7] for instance) may be adapted to solve (2.1).

To be more specific, define a finite time grid $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$ for a fixed $T > 0$, and consider the function v^Π defined as v (equation (2.1)) but with the strategy set $\mathcal{A}$ replaced by $\mathcal{A}^\Pi \subset \mathcal{A}$, defined as the subset of strategies that can be modified only at the dates $t \in \Pi$. In other words, the switching decisions can now only take place on the time grid Π. Suppose, moreover, that the cost function k is such that at most one switch can occur on a given date t_k (triangular condition). Then $\forall i \in \mathbb{I}_q$, $\forall x \in \mathbb{R}^d$, and $\forall t_k \in \Pi$, the dynamic programming principle (2.7) becomes:

$$v^\Pi(t_k,x,i) = \max_{j\in\mathbb{I}_q}\left\{E_j(t_k,x) - k(t_k,j-i)\,\mathbf{1}_{\{j\neq i\}}\right\} \tag{2.8}$$

where:

$$E_i(T,x) := \mathbb{E}\left[\int_T^\infty f_i\left(s, X_s^{T,x}\right) ds\right] \tag{2.9}$$

$$E_i(t_k,x) := \mathbb{E}\left[\int_{t_k}^{t_{k+1}} f_i\left(s, X_s^{t_k,x}\right) ds + v^\Pi\left(t_{k+1}, X_{t_{k+1}}^{t_k,x}, i\right)\right],\ k = N-1,\ldots,0 \tag{2.10}$$

which is explicit in the sense that $v^\Pi(t_k,.,.)$ directly depends on $v^\Pi(t_{k+1},.,.)$.

In practice, apart from the potential approximation of the stochastic process X and of the final values (2.9), the difficulty lies in the efficient computation of the conditional expectations (2.10).

In the American option literature, various approaches have been developed to solve (2.8) efficiently. Notable examples are the least-squares regression approach ([28, 35]), the quantization approach and the Malliavin calculus based formulation (see [7] for a thorough comparison and improvements of these techniques). In the spirit of [11], one may also consider non-parametric regression (see [24] and [34]) combined with speeding up techniques like Kd-trees or the Fast Gauss Transform in the case of kernel regression.

Here, we intend to solve (2.1) on numerical applications which bears the particularity of handling stochastic processes in high dimension ($\dim(X) = d \gg 3$, with however $\dim(I) = d' \approx 3$, see Section 5). For such problems, the most adequate technique so far seems to be the local regression method developed in [7]. We are thus going to make use of this specific method to solve (2.8) in practice.

In the following, we provide a detailed analysis of the above suggested computational method.

3 Numerical approximation and convergence analysis

This section is devoted to the precise description of the resolution of (2.1), along the lines of the discussions from Subsection 2.3. Moreover, the convergence rate of the proposed algorithm will be precisely assessed.

3.1 Approximations

Recall equation (2.1) defining the value function $v(t,x,i)$:

$$v(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}} \mathbb{E}\left[\int_t^\infty f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{\tau_n \geq t} k\left(\tau_n, \zeta_n\right)\right] \tag{3.1}$$

We are going to consider the following sequence of approximations:

- *[Finite time horizon]* The time horizon will be truncated to a finite horizon T.
- *[Time discretization]* The continuous state process X and investment process I will be discretized with a time step h.
- *[Space localization]* The $\mathbb{R}^d$- valued process X will be projected into a bounded domain $\mathcal{D}_\varepsilon$, parameterized by ε.
- *[Conditional expectation approximation]* The conditional expectation involved in the dynamic programming equation will be replaced by an empirical least-squares regression, computed on a bundle of M Monte Carlo trajectories, on a finite basis of local hypercubes with edges of size δ.

The rate of convergence of the algorithm will then be provided, as a function of these five numerical parameters: T, h, ε, M and δ.

3.1.1 Finite time horizon

The first step is to reduce the set of strategies to a finite horizon:

$$v_T(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}^T} \mathbb{E}\left[\int_t^T f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t \leq \tau_n \leq T} k\left(\tau_n, \zeta_n\right) + g_f\left(T, X_T^{t,x}, I_T^\alpha\right)\right] \tag{3.2}$$

$$g_f(T,x,i) := \mathbb{E}\left[\int_T^\infty f\left(s, X_s^{T,x}, i\right) ds\right] \tag{3.3}$$

where $0 \leq t \leq T < +\infty$, and $\mathcal{A}_{t,i}^T \subset \mathcal{A}_{t,i}$ is the subset of strategies without switches strictly after time T. Hence the final value g_f corresponds to the remaining gain after T.

Alternatively, one may choose, for convenience, another final value g instead of g_f, as long as it is Lipschitz-continuous and satisfies a suitable condition (cf. equation (3.20)). The set of such functions will be denoted as Θ_{g_f}. The difference between the two value functions is quantified in Proposition 3.1.

This freedom on the final values will be used in practice to avoid a computation on an infinite interval $[T, \infty[$ as in the definition of g_f.

From now on, we choose and fix one such $g \in \Theta_{g_f}$.

3.1.2 Time discretization

Then, we discretize the time segment $[0,T]$. Introduce a time grid $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$ with constant mesh h. Consider the following approximation:

$$v_\Pi(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}^\Pi} \mathbb{E}\left[\int_t^T f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t \leq \tau_n \leq T} k\left(\tau_n, \zeta_n\right) + g\left(T, X_T^{t,x}, I_T^\alpha\right)\right] \tag{3.4}$$

where $\mathcal{A}_{t,i}^{\Pi} \subset \mathcal{A}_{t,i}^{T}$ is the subset of strategies such that switches can only occur at dates $\tau_n \in \Pi \cap [t,T]$.

Now, with a slight abuse of notation, we can safely switch from the notation $\alpha = (\tau_n, \zeta_n)_{n\geq 0}$ to the notation $\alpha = (\tau_n, \iota_n)_{n\geq 0}$ (remember Subsection 2.1), replacing the quantity $\sum_{t\leq\tau_n\leq T} k(\tau_n, \zeta_n)$ by $\sum_{t\leq\tau_n\leq T} k\left(\tau_n, I^{\alpha}_{\tau_n - h}, I^{\alpha}_{\tau_n}\right)$ or by $\sum_{t\leq\tau_n\leq T} k(\tau_n, \iota_{n-1}, \iota_n)$, where $k(t,i,j) = k(t, j-i)$. The error between v_T and v_{Π} is quantified in Proposition 3.2.

Next we also approximate the stochastic process X by its Euler scheme $\bar{X} = \left(\bar{X}_s\right)_{0\leq s\leq T}$, with dynamics:

$$\begin{aligned} d\bar{X}_s &= b\left(\pi(s), \bar{X}_{\pi(s)}\right) ds + \sigma\left(\pi(s), \bar{X}_{\pi(s)}\right) dW_s\,,\ 0 \leq s \leq T \\ \bar{X}_0 &= x_0 \in \mathbb{R}^d \end{aligned} \tag{3.5}$$

where $\forall s \in [0,T]$, $\pi(s) := \max\{t \in \Pi; t \leq s\}$. More precisely, we substitute the piecewise constant $\bar{X}_{\pi(t)}$ for X_t. (Note that at this stage the process I^{α} is already piecewise constant). The new value function reads:

$$\bar{v}_{\Pi}(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}^{\Pi}} \mathbb{E}\left[\int_t^T f\left(\pi(s), \bar{X}^{t,x}_{\pi(s)}, I^{\alpha}_s\right) ds - \sum_{t\leq\tau_n\leq T} k(\tau_n, \iota_{n-1}, \iota_n) + g\left(T, \bar{X}^{t,x}_{\pi(T)}, I^{\alpha}_T\right)\right] \tag{3.6}$$

The error between v_{Π} and $\bar{v}_{\Pi}$ is computed in Proposition 3.3.

3.1.3 Space localization

In order to derive a rigorous convergence analysis, our subsequent choices in terms of conditional expectation approximation (Subsection 3.1.4 below) and specific choice of basis (Assumption 5) will require the underlying state process X to lie into a bounded set (cf. equation (3.16)). Thus, we explicitly build such an approximation and assess the associated error. Remark, though, that the usefulness of this step is more theoretical (for a proper convergence speed to hold) than practical (on a finite sample, this localization step would be somewhat redundant, and may safely be omitted).

Hence, let $\mathcal{D} = [-R,R]^d$, $R > 0$, be a bounded convex domain of $\mathbb{R}^d$ that contains x_0. For every $i = 1,\dots,d$, define the stopping time τ_i and the killed process $\bar{X}^{i,\mathcal{D}} = \left(\bar{X}^{i,\mathcal{D}}_t\right)_{0\leq t\leq T}$ as follows:

$$\begin{aligned} \tau_i &= \inf\left\{t \in [0,T]\,; \bar{X}^i_t \notin [-R,R]\right\} \\ \bar{X}^{i,\mathcal{D}}_t &= \bar{X}^i_{t\wedge\tau_i}\,,\ t \in [0,T] \end{aligned}$$

In other words the $d-$dimensional process $\bar{X}^{\mathcal{D}} = \left(\bar{X}^{i,\mathcal{D}}_t\right)^{1\leq i\leq d}_{0\leq t\leq T}$ is equal to $\bar{X}$ most of the time (i.e. when $\bar{X}_t \in \mathcal{D}$), except when one component of $\bar{X}_t$ gets outside $\mathcal{D}$, in which case the corresponding component of $\bar{X}^{\mathcal{D}}_t$ is killed and remains on the border of the domain $\mathcal{D}$ (the other components being unaffected). In particular, the killed process $\bar{X}^{\mathcal{D}}$ is bounded and Markovian.

Finally, one can choose R sufficiently large such that

$$\sup_{t\in[0,T]} \mathbb{E}\left[\left|\bar{X}_t - \bar{X}^{\mathcal{D}}_t\right|\right] \leq \varepsilon \tag{3.7}$$

for some $\varepsilon > 0$ (in which case $R = R(T,\varepsilon)$). This is the parameterization of the domain $\mathcal{D} = \mathcal{D}_{\varepsilon}$ that we adopt in the following.

Define $\bar{v}^{\varepsilon}_{\Pi}$ as the value function $\bar{v}_{\Pi}$ from equation (3.6) with $\left(\bar{X}_{\pi(t)}\right)_{0\leq t\leq T}$ replaced by $\left(\bar{X}^{\mathcal{D}_{\varepsilon}}_{\pi(t)}\right)_{0\leq t\leq T}$. The error between those two value functions is computed in Proposition 3.4.

Example 3.1. To clarify this construction of space localization, we explicit it on the very simple example of a d-dimensional standard brownian motion $(W_t)_{t\in[0,T]}$. In this case, $\bar{X}_t = X_t = W_t$. In this example, equation (3.7) can be shown to hold by choosing for instance $R(T,\varepsilon) = \sqrt{2T\log\left(\frac{2dT}{\pi\varepsilon}\right)}$.

3.1.4 Conditional expectation approximation

Now that the problem has been localized, and in order to prevent the notation from becoming too cumbersome and clumsy, we are going to drop the ε index in the following final approximation step, i.e. $\bar{X}_t$ will stand for $\bar{X}_t^{\mathcal{D}_\varepsilon}$, and $\bar{v}_\Pi$ for $\bar{v}_\Pi^\varepsilon$.

Remark that the discrete process $\left(\bar{X}_{t_n}\right)_{n=0,\dots,N}$ is a Markov chain. Therefore, the dynamic programming principle applied to $\bar{v}_\Pi$ yields:

$$\begin{aligned}
&\bar{v}_\Pi(T,x,i) = g(T,x,i)\\
&\bar{v}_\Pi(t_n,x,i) = \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j) - k(t_n,i,j) + \mathbb{E}\left[\bar{v}_\Pi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},j\right)\right]\right\}, n = N-1,\dots,0
\end{aligned} \tag{3.8}$$

The last step is to approximate the conditional expectation appearing in equation (3.8). As discussed in Subsection 2.3, we choose to approximate it by least-squares regression. Consider basis functions $(e_k(x))_{1\le k\le K}$, $K\in\mathbb{N}\cup\{+\infty\}$, $x\in\mathbb{R}^d$. For suitable functions $\varphi:\Pi\times\mathbb{R}^d\times\mathbb{I}_q\to\mathbb{R}$, define:

$$\tilde{\lambda} = \tilde{\lambda}_i^{t_n}(\varphi) := \arg\min_{\lambda\in\mathbb{R}^K}\mathbb{E}\left[\left(\varphi\left(t_{n+1},\bar{X}_{t_{n+1}},i\right) - \sum_{k=1}^K \lambda_k e_k\left(\bar{X}_{t_n}\right)\right)^2\right] \tag{3.9}$$

Now, before using this projection, it is more cautious to truncate it within known bounds (see [6, 19, 33]). Hence, suppose that there exist known bounds $\underline{\Gamma}^{t_n,x}(\varphi)$ and $\overline{\Gamma}^{t_n,x}(\varphi)$ around $\mathbb{E}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right]$:

$$\underline{\Gamma}^{t_n,x}(\varphi) \le \mathbb{E}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right] \le \overline{\Gamma}^{t_n,x}(\varphi) \tag{3.10}$$

Then, $\forall i\in\mathbb{I}_q$ the quantity $\mathbb{E}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right]$ is approximated by:

$$\tilde{\mathbb{E}}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right] := \underline{\Gamma}^{t_n,x}(\varphi) \vee \sum_{k=1}^K \tilde{\lambda}_k e_k(x) \wedge \overline{\Gamma}^{t_n,x}(\varphi) \tag{3.11}$$

which is used to define the next approximation $\tilde{v}_\Pi$ of the value function:

$$\begin{aligned}
&\tilde{v}_\Pi(T,x,i) = g(T,x,i)\\
&\tilde{v}_\Pi(t_n,x,i) = \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j) - k(t_n,i,j) + \tilde{\mathbb{E}}\left[\tilde{v}_\Pi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},j\right)\right]\right\},\ n = N-1,\dots,0
\end{aligned} \tag{3.12}$$

Interesting discussions on the choice of function basis can be found in [7]. In particular they advocate bases of local polynomials, which is numerically efficient and well-suited to tackle large-dimensional problems (see Subsection 4.1). However, for the sake of simplicity, we will restrict our study in this section to a basis of indicator functions on local hypercubes (cf. [33] and the numerical experiments of [19]) (which is the simplest example of local polynomials). Assumption 5 below states this specific choice.

Assumption 5. *The regression basis is set to a basis of indicator function on disjoint local hypercubes, as described in Definition 3.1 below.*

Definition 3.1. For every $t_n\in\Pi$, consider a partition of the domain $\mathcal{D}_\varepsilon$ into hypercubes $\left(B_{t_n}^k\right)_{k=1,\dots,K_\varepsilon}$, i.e., $\cup_{k=1,\dots,K_\varepsilon} B_{t_n}^k = \mathcal{D}_\varepsilon$ and $B_{t_n}^i\cap B_{t_n}^j = \emptyset\ \forall i\neq j$. It may be deterministic, or computed from a sample of $\bar{X}$. We only assume that there exists $(\underline{\delta},\delta)\in\mathbb{R}_+^2$ with $\underline{\delta}\le\delta$ such that the lengths of the edges of the hypercubes, in each dimension, belong to $[\underline{\delta},\delta]$ (in particular, the volume of each hypercube $B_{t_n}^k$ belongs to $\left[\underline{\delta}^d,\delta^d\right]$). This liberty over the definition of the partition enables to encompass to some extend the kind of adaptative partition described in [7]. Then, the basis functions considered here are defined by $e_{t_n}^k(x) := \mathbf{1}\left\{x\in B_{t_n}^k\right\}$, $x\in\mathbb{R}^d$, $1\le k\le K_\varepsilon$.

Under Assumption 5, the error between $\bar{v}_\Pi$ and $\tilde{v}_\Pi$ is computed in Proposition 3.5.

Finally, let $\left(\bar{X}_{t_n}^m\right)_{1\leq n\leq N}^{1\leq m\leq M}$ be a finite sample of size M of paths of the process $\bar{X}$. The final step is to replace the regression (3.9) by a regression on this sample:

$$\hat{\lambda} = \hat{\lambda}_i^{t_n}(\varphi) := \arg\min_{\lambda\in\mathbb{R}^K} \frac{1}{M}\sum_{m=1}^{M}\left[\left(\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^m, i\right) - \sum_{k=1}^{K}\lambda_k e_k\left(\bar{X}_{t_n}^m\right)\right)^2\right]. \tag{3.13}$$

Then $\forall i \in \mathbb{I}_q$ the quantity $\mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, i\right)\right]$ is approximated by:

$$\hat{\mathbb{E}}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, i\right)\right] := \underline{\Gamma}^{t_n,x}(\varphi) \vee \sum_{k=1}^{K}\hat{\lambda}_k e_k(x) \wedge \overline{\Gamma}^{t_n,x}(\varphi) \tag{3.14}$$

leading to the final, computable approximation $\hat{v}_\Pi$ of the value function:

$$\begin{aligned}\hat{v}_\Pi(T,x,i) &= g(T,x,i)\\ \hat{v}_\Pi(t_n,x,i) &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j) - k(t_n,i,j) + \hat{\mathbb{E}}\left[\hat{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right]\right\},\ n = N-1,\ldots,0\end{aligned} \tag{3.15}$$

Under Assumption 5, the error between $\tilde{v}_\Pi$ and $\hat{v}_\Pi$ is given in Proposition 3.6. This proposition will make use of the following quantity:

$$p(T,\delta,\varepsilon) := \min_{t\in\Pi}\min_{B_t^k\subset\mathcal{D}^\varepsilon}\mathbb{P}\left(\bar{X}_t \in B_t^k\right) \tag{3.16}$$

which is strictly positive, as the domain $\mathcal{D}_\varepsilon$ is (purposely) bounded.

Example 3.2. Carrying on with Example 3.1 of a d-dimensional Brownian motion, we explicit a lower bound for $p(T,\delta,\varepsilon)$ in this simple case. First, $\mathbb{P}\left(W_T \in B_T^k\right) = \int_{B_T^k} f_{W_T}(x)\,dx$ where f_{W_T} is the density of W_T. As $\forall k$, $B_T^k \subset \mathcal{D}^\varepsilon$, with $R(T,\varepsilon) = \sqrt{2T\log\left(\frac{2dT}{\pi\varepsilon}\right)}$, it holds that $\forall x \in \mathcal{D}^\varepsilon$, $f_{W_T}(x) \geq \left(f_{W_T^1}(R(T,\varepsilon))\right)^d = \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}$. Hence $\mathbb{P}\left(W_t \in B_t^k\right) \geq \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}\mathrm{Vol}\left(B_t^k\right) \geq \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}\underline{\delta}^d$. As a conclusion, $p(T,\delta,\varepsilon) \geq \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}\underline{\delta}^d$. Remark however that this lower bound is very crude, and that it can be very far below $p(T,\delta,\varepsilon)$ for large δ.

Combining all these results, we obtain a rate of convergence of $\hat{v}_\Pi$ towards v:

Theorem 3.1. *$\forall p \geq 1$, $\exists C_p > 0$ such that:*

$$\begin{aligned}&\left\|\max_{i\in\mathbb{I}_q}\left|v(t_0,x_0,i) - \hat{v}_\Pi(t_0,x_0,i)\right|\right\|_{L_p}\\ \leq\ & C_p\left\{(1+|x_0|)e^{-\bar{\rho}T} + (1+|x_0|)^{\frac{3}{2}}\sqrt{h} + \varepsilon + \frac{\delta}{h} + \frac{1+R(T,\varepsilon)}{h\sqrt{M}p(T,\delta,\varepsilon)^{1-\frac{1}{p\vee 2}}} + \frac{1+R(T,\varepsilon)}{hMp(T,\delta,\varepsilon)}\right\}\end{aligned}$$

In particular, $\hat{v}_\Pi(0,x_0,i) \to_{L_p} v(0,x_0,i)$ uniformly in $i \in \mathbb{I}_q$ when $T\to\infty$, $h\to 0$, $\varepsilon\to 0$, $\delta\to 0$ and $M\to\infty$ with $\frac{\delta}{h}\to 0$, $\frac{1+R(T,\varepsilon)}{h\sqrt{M}p(T,\delta,\varepsilon)^{1-\frac{1}{p\vee 2}}}\to 0$ and $\frac{1+R(T,\varepsilon)}{hMp(T,\delta,\varepsilon)}\to 0$.

The proof of Theorem 3.1 will be given at the end of the next Subsection 3.2.

Remark 3.1. If the cost function k (recall Assumption 3) were to depend on x, then, under a usual Lipschitz condition on k (similar to that of f), Theorem 3.1 would still hold, replacing only the term $(1+|x_0|)^{\frac{3}{2}}\sqrt{h}$ by $\left(1+|x_0|^{\frac{5}{2}}\right)\sqrt{h\log\left(\frac{2T}{h}\right)}$ (recalling Remark 3.4).

Remark 3.2. The adaptative local basis can be such that each hypercube contains approximately the same number of Monte Carlo trajectories (see [7]). This means that $\frac{1}{p(T,\delta,\varepsilon)} \sim b$ where b is the number of functions in the regression basis. With this remark in mind, the leading error term in Theorem 3.1 behaves like $\frac{\sqrt{b}}{h\sqrt{M}}$ for $p=2$. This is close to the corresponding statistical error term in [27] ($\sqrt{\frac{b\log(M)}{hM}}$) in the context of BSDEs. The advantage of their approach is that it can handle any (orthonormal) regression basis, while our approach (in the context of optimal switching) provides a bound on the L_p error for every $p \geq 1$.

Example 3.3. In the case of a d-dimensional Brownian motion, the rate of convergence of Theorem 3.1 can be explicited further, using the upper bound on $R(T,\varepsilon)$ from Example 3.1 and the lower bound on $p(T,\delta,\varepsilon)$ from Example 3.2. Moreover, one can express the rate of convergence as a function of only one parameter, choosing the five numerical parameters T, h, ε, δ and M accordingly. For instance, assuming $\underline{\delta} = \delta$, and minimizing over δ, h, ε and T, one can get a convergence rate upper bounded by $C_p(1+|x|)^{\frac{3}{2}} z$ by choosing $M \sim z^{-\frac{1}{2}[6(d+1)]^2}$. This is admittedly highly demanding in terms of sample size M, but remember that this expression suffers from the crude lower bound on $p(T,\delta,\varepsilon)$ we chose previously.

3.2 Convergence analysis

From now on, we suppose that all the assumptions from Subsection 2.2 are in force.

3.2.1 Finite time horizon

Lemma 3.1. *There exists $C>0$ such that $\forall (t,x,i) \in [0,T]\times\mathbb{R}^d\times\mathbb{R}^{d'}$:*

$$0 \leq v(t,x,i) - v_T(t,x,i) \leq C(1+|x|)\, e^{-\bar{\rho}t\vee T - \rho_1 t}.$$

Proof. First, we introduce the following notations:

$$H(t,T,x,\alpha) := \int_t^T f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t<\tau_n\leq T} k(\tau_n,\zeta_n) \tag{3.17}$$

$$J(t,T,x,\alpha) := \mathbb{E}\left[H(t,T,x,\alpha)\right] \tag{3.18}$$

for any admissible strategy $\alpha \in \mathcal{A}_{t,i}$. In particular:

$$v(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}} J(t,+\infty,x,\alpha) \quad , \quad v_T(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}^T} J(t,+\infty,x,\alpha)\,. \tag{3.19}$$

Fix $(t,x,i) \in \mathbb{R}_+\times\mathbb{R}^d\times\mathbb{R}^{d'}$. Using equation (3.19):

$$v_T(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}^T} J(t,\infty,x,\alpha) \leq \sup_{\alpha\in\mathcal{A}_{t,i}} J(t,\infty,x,\alpha) = v(t,x,i)$$

which provides the first inequality. Consider now the second inequality. Choose $\varepsilon>0$. From the definition of v (equation (3.1)) there exists a strategy $\alpha^\varepsilon \in \mathcal{A}_{t,i}$ such that:

$$v(t,x,i) - \varepsilon \leq J(t,\infty,x,\alpha^\varepsilon) \leq v(t,x,i)$$

Define the truncated strategy $\alpha_T^\varepsilon \in \mathcal{A}_{t,i}^T$ such that $\forall s\in[t,T]$, $I_s^{\alpha_T^\varepsilon} = I_s^{\alpha^\varepsilon}$ and $\forall s>T$, $I_s^{\alpha_T^\varepsilon} = I_T^{\alpha^\varepsilon}$. In order not to mix up the variables τ_n and ζ_n from different strategies, we add the name of the strategy

in index when needed. Then:

$$
\begin{aligned}
&H\left(t,\infty,x,\alpha^{\varepsilon}\right)-H\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\\
&=\left\{\int_{t}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\right\}-\left\{\int_{t}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha_{T}^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha_{T}^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha_{T}^{\epsilon}},\zeta_{n}^{\alpha_{T}^{\epsilon}}\right)\right\}\\
&=\left\{\int_{t}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\right\}\\
&-\left\{\int_{t}^{t\vee T}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds+\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)ds-\sum_{t\vee T\geq\tau_{n}^{\alpha^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\right\}\\
&=\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha^{\epsilon}}\geq t\vee T}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\\
&\leq\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)ds
\end{aligned}
$$

as $k\left(s,0\right)=0$ and $k\geq0$ (Assumption 4). Hence, using Jensen's inequality and equation (2.4), $\exists C>0$ such that

$$
\begin{aligned}
\left|J\left(t,\infty,x,\alpha^{\varepsilon}\right)-J\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\right| &\leq \mathbb{E}\left[\left|H\left(t,\infty,x,\alpha^{\varepsilon}\right)-H\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\right|\right]\\
&\leq \mathbb{E}\left[\int_{t\vee T}^{\infty}\left|f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)\right|ds\right]+\mathbb{E}\left[\int_{t\vee T}^{\infty}\left|f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)\right|ds\right]\\
&\leq C\left(1+\left|x\right|\right)e^{-\bar{\rho}t\vee T-\rho_{1}t}
\end{aligned}
$$

Finally, given that $v\left(t,x,i\right)\leq\varepsilon+J\left(t,\infty,x,\alpha^{\varepsilon}\right)$ and $v_{T}\left(t,x,i\right)\geq J\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)$, the following holds:

$$
\begin{aligned}
v\left(t,x,i\right)-v_{T}\left(t,x,i\right) &\leq \varepsilon+J\left(t,\infty,x,\alpha^{\varepsilon}\right)-J\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\\
&\leq \varepsilon+C\left(1+\left|x\right|\right)e^{-\bar{\rho}t\vee T-\rho_{1}t}.
\end{aligned}
$$

Since this is true for any $\varepsilon>0$, and that C, ρ and ρ_{1} do not depend on ε, the proposition is proved. □

Now, we focus on the final boundary g_f. For the time being, denote the value function (3.2) as $v_T^{g_f}$ to emphasize the dependence of v on the terminal condition. As a consequence of equation (2.4), $\forall\left(x,i\right)\in\mathbb{R}^{d}\times\mathbb{I}_{q}$:

$$\left|g_{f}(T,x,i)\right|\leq C\left(1+\left|x\right|\right)e^{-\rho T} \tag{3.20}$$

Hence, define the class Θ_{g_f} of Lipschitz functions from $\mathbb{R}_{+}\times\mathbb{R}^{d}\times\mathbb{I}_{q}$ into $\mathbb{R}$ such that $\forall g\in\Theta_{g_f}$, $\forall\left(T,x,x',i\right)\in\mathbb{R}_{+}\times\mathbb{R}^{d}\times\mathbb{R}^{d}\times\mathbb{I}_{q}$:

$$
\begin{aligned}
\left|g(T,x,i)-g(T,x',i)\right| &\leq C_{g}e^{-\rho T}\left|x-x'\right| &(3.21)\\
\left|g(T,0,i)\right| &\leq C_{g}e^{-\rho T} &(3.22)
\end{aligned}
$$

for some $C_g>0$. In particular, the growth rate of such functions is at most linear in x:

$$\left|g(T,x,i)\right|\leq C_{g}e^{-\rho T}\left(1+\left|x\right|\right). \tag{3.23}$$

Obviously $g_f\in\Theta_{g_f}$. Now, for any $g\in\Theta_{g_f}$, denote as v_T^g the value function defined as in equation (3.2) with g instead of g_f. We are going to show that the precise approximation error due to the choice of final value g does not matter much as long as g is chosen in this class Θ_{g_f}.

Lemma 3.2. *There exists $C>0$ such that $\forall\left(t,x,i\right)\in\mathbb{R}_{+}\times\mathbb{R}^{d}\times\mathbb{I}_{q}$:*

$$\left|v_{T}^{g_{f}}\left(t,x,i\right)-v_{T}^{g}\left(t,x,i\right)\right|\leq C\left(1+\left|x\right|\right)e^{-\bar{\rho}t\vee T-\rho_{1}t}$$

Proof. Fix $(t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{I}_q$. To shorten the proof, we assume that $v_T^{g_f}$ (resp. v_T^g) admits an optimal strategy $\alpha_f^* \in \mathcal{A}_{t,i}^T$ (resp. $\alpha^* \in \mathcal{A}_{t,i}^T$) (this assumption can then be relaxed using ε-optimal strategies as in the proof of Proposition 3.1)[1]. Therefore, recalling the notations H (equation (3.17)) and J (equation (3.18)) introduced in the proof of Lemma 3.1:

$$\begin{aligned}
v_T^{g_f}(t,x,i) - v_T^g(t,x,i) &= J\left(t,T,x,\alpha_f^*\right) + \mathbb{E}\left[g_f\left(T, X_T^{t,x}, I_T^{\alpha_f^*}\right)\right] - J\left(t,T,x,\alpha^*\right) - \mathbb{E}\left[g\left(T, X_T^{t,x}, I_T^{\alpha^*}\right)\right] \\
&= J\left(t,T,x,\alpha_f^*\right) + \mathbb{E}\left[g\left(T, X_T^{t,x}, I_T^{\alpha_f^*}\right)\right] - J\left(t,T,x,\alpha^*\right) - \mathbb{E}\left[g\left(T, X_T^{t,x}, I_T^{\alpha^*}\right)\right] \\
&+ \mathbb{E}\left[g_f\left(T, X_T^{t,x}, I_T^{\alpha_f^*}\right) - g\left(T, X_T^{t,x}, I_T^{\alpha^*}\right)\right] \\
&\leq \mathbb{E}\left[g_f\left(T, X_T^{t,x}, I_T^{\alpha_f^*}\right) - g\left(T, X_T^{t,x}, I_T^{\alpha^*}\right)\right] \\
&\leq C\left(1 + \mathbb{E}\left[\left|X_T^{t,x}\right|\right]\right) e^{-\rho T} \;\leq\; C\left(1+|x|\right) e^{-\bar{\rho} t \vee T - \rho_1 t}
\end{aligned}$$

Symmetrically, the same inequality holds for $v_T^g(t,x,i) - v_T^{g_f}(t,x,i)$, ending the proof. □

Proposition 3.1. *There exists $C > 0$ independent of T such that $\forall (t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{I}_q$ and $\forall g \in \Theta_{g_f}$:*

$$\left|v(t,x,i) - v_T^g(t,x,i)\right| \leq C\left(1+|x|\right) e^{-\bar{\rho} t \vee T - \rho_1 t}$$

Proof. Combine Lemmas 3.1 and 3.2. □

From now on, we choose and keep one final value function $g \in \Theta_{g_f}$, and remove the index g from the notation of v and its subsequent approximations.

3.2.2 Time Discretization

Proposition 3.2. *There exists a positive constant C such that for any $(t,x,i) \in \Pi \times \mathbb{R}^d \times \mathbb{I}_q$:*

$$\left|v_T(t,x,i) - v_\Pi(t,x,i)\right| \leq C e^{-\rho t}\left(1 + |x|^{\frac{3}{2}}\right) h^{\frac{1}{2}} \tag{3.24}$$

Proof. Under the assumptions from Subsection 2.2, one can apply Theorem 3.1 in [18] to prove (3.24), noticing that the cost function k does not depend on the state variable x.

Use the discounting factor in the definition of f to factor the $e^{-\rho t}$ term and to get that C does not depend on T. □

Remark 3.3. Another alternative to get this rate of $h^{\frac{1}{2}}$ is to work with the reflected BSDE representation of v_Π, as in [10] (adapting [6]) or [15].

Remark 3.4. Were the cost function k to depend on the state variable, the upper bound in Proposition 3.2 would only be $Ce^{-\rho t}\left(1+|x|^{\frac{5}{2}}\right)\left(h \log\left(\frac{2T}{h}\right)\right)^{\frac{1}{2}}$, as stated in [18] (making use of results from [17]).

Proposition 3.3. *There exists $C > 0$ such that for any $(t,x,i) \in \Pi \times \mathbb{R}^d \times \mathbb{I}_q$:*

$$\left|v_\Pi(t,x,i) - \bar{v}_\Pi(t,x,i)\right| \leq C e^{-\rho t} h^{\frac{1}{2}}$$

[1] Note that under the assumptions from Subsection 2.2, one may use Theorem 3.1 from [22] to get the existence of a unique optimal strategy α^* for the value function (3.2), satisfying $\mathbb{E}\left[\left|\sum_{0 \leq \tau_n^{\alpha^*} \leq T} k\left(\tau_n^{\alpha^*}, \zeta_n^{\alpha^*}\right)\right|^2\right] < \infty$

Proof. T and g being fixed, we can define, in the spirit of equations (3.17) and (3.18), the following quantities:

$$H(t,x,\alpha) := \int_t^T f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t\le\tau_n\le T} k\left(\tau_n, \iota_{n-1}, \iota_n\right) + g\left(T, X_T^{t,x}, I_T^\alpha\right) \tag{3.25}$$

$$J(t,x,\alpha) := \mathbb{E}\left[H(t,x,\alpha)\right] \tag{3.26}$$

$$\bar{H}(t,x,\alpha) := \int_t^T f\left(\pi(s), \bar{X}_{\pi(s)}^{t,x}, I_s^\alpha\right) ds - \sum_{t\le\tau_n\le T} k\left(\tau_n, \iota_{n-1}, \iota_n\right) + g\left(T, \bar{X}_T^{t,x}, I_T^\alpha\right) \tag{3.27}$$

$$\bar{J}(t,x,\alpha) := \mathbb{E}\left[\bar{H}(t,x,\alpha)\right] \tag{3.28}$$

for any admissible strategy $\alpha \in \mathcal{A}_{t,i}^\Pi$. For these discretized problems, the existence of optimal controls α^* and $\bar{\alpha}^*$ is granted. Hence:

$$\begin{aligned}
v_\Pi(t,x,i) - \bar{v}_\Pi(t,x,i) &= J(t,x,\alpha^*) - \bar{J}(t,x,\bar{\alpha}^*) \\
&= J(t,x,\alpha^*) - \bar{J}(t,x,\alpha^*) + \left\{\bar{J}(t,x,\alpha^*) - \bar{J}(t,x,\bar{\alpha}^*)\right\} \\
&\le J(t,x,\alpha^*) - \bar{J}(t,x,\alpha^*) \\
&= \int_t^T e^{-\rho s}\mathbb{E}\left[\tilde{f}\left(s, X_s^{t,x}, I_s^{\alpha^*}\right) - \tilde{f}\left(\pi(s), \bar{X}_{\pi(s)}^{t,x}, I_s^{\alpha^*}\right)\right] ds \\
&+ \mathbb{E}\left[g\left(T, X_T^{t,x}, I_T^{\alpha^*}\right) - g\left(T, \bar{X}_T^{t,x}, I_T^{\alpha^*}\right)\right] \\
&\le C_f \int_t^T e^{-\rho s}\mathbb{E}\left[\left|X_s^{t,x} - \bar{X}_{\pi(s)}^{t,x}\right|\right] ds + C_g e^{-\rho T}\mathbb{E}\left[\left|X_T^{t,x} - \bar{X}_T^{t,x}\right|\right] \\
&\le C e^{-\rho t}\mathbb{E}\left[\sup_{t\le s\le T}\left|X_s^{t,x} - \bar{X}_{\pi(s)}^{t,x}\right|\right] \le C e^{-\rho t} h^{\frac{1}{2}}
\end{aligned}$$

using the strong convergence speed of the Euler scheme on $[t,T]$. Symmetrically, the same inequality holds for $\bar{v}_\Pi(t,x,i) - v_\Pi(t,x,i)$, ending the proof. □

3.2.3 Space localization

Recall from Subsection 3.1.3 the definition of the bounded domain $\mathcal{D}^\varepsilon$.

Proposition 3.4. *$\exists C > 0$ such that $\forall \varepsilon > 0$ and $\forall (x,i) \in \mathbb{R}^d \times \mathbb{I}_q$:*

$$\left|\bar{v}_\Pi(0,x,i) - \bar{v}_\Pi^\varepsilon(0,x,i)\right| \le C\varepsilon$$

Proof. Recall the definitions of $\bar{H}(t,x,\alpha)$ (equation (3.27)) and $\bar{J}(t,x,\alpha)$ (equation (3.28)), and define the quantities $\bar{H}^\varepsilon(t,x,\alpha)$ and $\bar{J}^\varepsilon(t,x,\alpha)$ like $\bar{H}(t,x,\alpha)$ and $\bar{J}(t,x,\alpha)$ but with $\bar{X}_{\pi(.)}$ replaced by $\bar{X}_{\pi(.)}^{\mathcal{D}_\varepsilon}$. Then, for every $(t,x,i) \in \Pi \times \mathbb{R}^d \times \mathbb{I}_q$ and $\alpha \in \mathcal{A}_{t,i}^\Pi$:

$$\begin{aligned}
\bar{J}(t,x,\alpha) = \bar{J}^\varepsilon(t,x,\alpha) &+ \int_t^T \mathbb{E}\left[f\left(\pi(s), \bar{X}_{\pi(s)}^{t,x}, I_s^\alpha\right) - f\left(\pi(s), \bar{X}_{\pi(s)}^{\mathcal{D}_\varepsilon,t,x}, I_s^\alpha\right)\right] ds \\
&+ \mathbb{E}\left[g\left(T, \bar{X}_T^{t,x}, I_T^\alpha\right) - g\left(T, \bar{X}_T^{\mathcal{D}_\varepsilon,t,x}, I_T^\alpha\right)\right]
\end{aligned}$$

But:

$$\begin{aligned}
&\left|\int_t^T \mathbb{E}\left[f\left(\pi(s), \bar{X}_{\pi(s)}^{t,x}, I_s^\alpha\right) - f\left(\pi(s), \bar{X}_{\pi(s)}^{\mathcal{D}_\varepsilon,t,x}, I_s^\alpha\right)\right] ds + \mathbb{E}\left[g\left(T, \bar{X}_T^{t,x}, I_T^\alpha\right) - g\left(T, \bar{X}_T^{\mathcal{D}_\varepsilon,t,x}, I_T^\alpha\right)\right]\right| \\
&\le C_f \int_t^T e^{-\rho s}\mathbb{E}\left[\left|\bar{X}_{\pi(s)}^{t,x} - \bar{X}_{\pi(s)}^{\mathcal{D}_\varepsilon,t,x}\right|\right] ds + C_g e^{-\rho T}\mathbb{E}\left[\left|\bar{X}_T^{t,x} - \bar{X}_T^{\mathcal{D}_\varepsilon,t,x}\right|\right]
\end{aligned}$$

It follows that:

$$\left|\bar{v}_\Pi(t,x,i) - \bar{v}_\Pi^\varepsilon(t,x,i)\right| \le C_f \int_t^T e^{-\rho s}\mathbb{E}\left[\left|\bar{X}_{\pi(s)}^{t,x} - \bar{X}_{\pi(s)}^{\mathcal{D}_\varepsilon,t,x}\right|\right] ds + C_g e^{-\rho T}\mathbb{E}\left[\left|\bar{X}_T^{t,x} - \bar{X}_T^{\mathcal{D}_\varepsilon,t,x}\right|\right]$$

In particular, at $t = 0$, using equation (3.7), $\exists C > 0$ such that:

$$|\bar{v}_\Pi(0, x, i) - \bar{v}_\Pi^\varepsilon(0, x, i)| \leq C\varepsilon$$

□

3.2.4 Conditional expectation approximation

From now on the domain $\mathcal{D}_\varepsilon$ is fixed once and for all, and, with a slight abuse of notation, we will drop ε from the subsequent notations.

We start with preliminary remarks. First, regarding the choice of regression basis, Assumption 5 is now supposed to hold. Then, recalling Subsection 3.1.4, and taking advantage of the orthogonality of the basis, one can easily compute the explicit solution of the minimisation equations that define the regression coefficients $\tilde{\lambda}_i^{t_n}(\varphi) = \left(\tilde{\lambda}_{i,k}^{t_n}(\varphi)\right)_{1\leq k\leq K}$ (equation (3.9)) and $\hat{\lambda}_i^{t_n}(\varphi) = \left(\hat{\lambda}_{i,k}^{t_n}(\varphi)\right)_{1\leq k\leq K_\varepsilon}$ (equation (3.13)). Namely:

$$\tilde{\lambda}_{i,k}^{t_n}(\varphi) = \frac{\mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, i\right)\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}^k\right\}\right]}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}^k\right)} = \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, i\right)\middle|\, \bar{X}_{t_n} \in B_{t_n}^k\right],\ 1 \leq k \leq K$$

$$\hat{\lambda}_{i,k}^{t_n}(\varphi) = \frac{\frac{1}{M}\sum_{m=1}^M \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^m, i\right)\mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}^k\right\}}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}^k\right\}},\ 1 \leq k \leq K$$

Extending these equations, define

$$\tilde{\lambda}_i^{t_n,x}(\varphi) := \frac{\mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, i\right)\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}(x)\right\}\right]}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}(x)\right)} = \mathbb{E}\left[\varphi\big(t_{n+1}, \bar{X}_{t_{n+1}}, i\big)\middle|\, \bar{X}_{t_n} \in B_{t_n}(x)\right] \quad (3.29)$$

$$\hat{\lambda}_i^{t_n,x}(\varphi) := \frac{\frac{1}{M}\sum_{m=1}^M \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^m, i\right)\mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}(x)\right\}}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}(x)\right\}} = \frac{1}{M_{t_n}^x}\sum_{m\in\mathcal{M}_{t_n}^x} \varphi\Big(t_{n+1}, \bar{X}_{t_{n+1}}^m, i\Big) \quad (3.30)$$

for every $(t_n, x, i) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$, where $\forall x \in \mathcal{D}$, $B_{t_n}(x)$ is the unique hypercube in the partition which contains x at time t_n, $\mathcal{M}_{t_n}^x := \left\{m \in [1, M],\ \bar{X}_{t_n}^m \in B_{t_n}(x)\right\}$ and $M_{t_n}^x := \#\mathcal{M}_{t_n}^x$.

Finally, recalling the approximated conditional expectations (3.11) and (3.14),

define for any $(t_n, x, j) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$ and any measurable function $\varphi : \Pi \times \mathbb{R}^d \times \mathbb{I}_q \to \mathbb{R}$, the following quantities:

$$\Phi_j^{t_n,x}(\varphi) := \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right] \quad (3.31)$$

$$\tilde{\Phi}_j^{t_n,x}(\varphi) := \tilde{\mathbb{E}}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right] = \underline{\Gamma}^{t_n,x}(\varphi) \vee \tilde{\lambda}_j^{t_n,x}(\varphi) \wedge \overline{\Gamma}^{t_n,x}(\varphi) \quad (3.32)$$

$$\hat{\Phi}_j^{t_n,x}(\varphi) := \hat{\mathbb{E}}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right] = \underline{\Gamma}^{t_n,x}(\varphi) \vee \hat{\lambda}_j^{t_n,x}(\varphi) \wedge \overline{\Gamma}^{t_n,x}(\varphi) \quad (3.33)$$

where (recalling equation 3.10) $\underline{\Gamma}^{t_n,x}(\varphi)$ and $\overline{\Gamma}^{t_n,x}(\varphi)$ are lower and upper bounds on $\Phi_j^{t_n,x}(\varphi)$:

$$\underline{\Gamma}^{t_n,x}(\varphi) \leq \Phi_j^{t_n,x}(\varphi) \leq \overline{\Gamma}^{t_n,x}(\varphi)$$

Remark 3.5. These definitions are useful to express the dynamic programming equations (3.8), (3.12) and (3.15). Indeed, these equations become:

$$\begin{aligned}
\bar{v}_\Pi(T, x, i) &= g(T, x, i) \\
\bar{v}_\Pi(t_n, x, i) &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n, x, j) - k(t_n, i, j) + \Phi_j^{t_n,x}(\bar{v}_\Pi)\right\},\ n = N-1, \dots, 0
\end{aligned}$$

$$\begin{aligned}
\tilde{v}_\Pi(T, x, i) &= g(T, x, i) \\
\tilde{v}_\Pi(t_n, x, i) &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n, x, j) - k(t_n, i, j) + \tilde{\Phi}_j^{t_n,x}(\tilde{v}_\Pi)\right\},\ n = N-1, \dots, 0
\end{aligned}$$

$$\begin{aligned}
\hat{v}_\Pi(T, x, i) &= g(T, x, i) \\
\hat{v}_\Pi(t_n, x, i) &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n, x, j) - k(t_n, i, j) + \hat{\Phi}_j^{t_n,x}(\hat{v}_\Pi)\right\},\ n = N-1, \dots, 0
\end{aligned}$$

Remark 3.6. For $\varphi = \bar{v}_\Pi$, we can easily explicit bounding functions $\underline{\Gamma}^{t_n,x}(\bar{v}_\Pi)$ and $\overline{\Gamma}^{t_n,x}(\bar{v}_\Pi)$ of $\Phi_j^{t_n,x}(\bar{v}_\Pi)$. Indeed, using the growth conditions on f and g, the nonnegativity of k and the definition of $R(T,\varepsilon)$ (see Paragraph 3.1.3), there exists $C>0$ such that $\forall (t_n,x,j) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$:

$$\begin{aligned} |\bar{v}_\Pi(t_n,x,j)| &\leq Ce^{-\rho t_n}(1+R(T,\varepsilon)) && (3.34)\\ \left|\Phi_j^{t_n,x}(\bar{v}_\Pi)\right| &\leq \Gamma^{t_n}(\bar{v}_\Pi) := Ce^{-\rho t_n}(1+R(T,\varepsilon)) && (3.35)\end{aligned}$$

Moreover, the same is true for $\varphi = \tilde{v}_\Pi$: there exists $C>0$ such that $\forall (t_n,x,j) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$:

$$\begin{aligned} |\tilde{v}_\Pi(t_n,x,j)| &\leq Ce^{-\rho t_n}(1+R(T,\varepsilon)) && (3.36)\\ \left|\tilde{\Phi}_j^{t_n,x}(\tilde{v}_\Pi)\right| &\leq \Gamma^{t_n}(\tilde{v}_\Pi) := Ce^{-\rho t_n}(1+R(T,\varepsilon)) && (3.37)\end{aligned}$$

Finally, we impose the same bound for the definition of $\hat{v}_\Pi$, i.e. $\Gamma^{t_n}(\hat{v}_\Pi) := \Gamma^{t_n}(\bar{v}_\Pi)$.

Now we can start the assessment of the regression error.

Lemma 3.3. *Consider a measurable function $\varphi : \Pi \times \mathbb{R}^d \times \mathbb{I}_q \to \mathbb{R}$. Suppose that, for a fixed $t_{n+1} \in \Pi$, it is Lipschitz with constant C_{n+1}, uniformly in j: $\forall (x_1,x_2,j) \in \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{I}_q$*

$$|\varphi(t_{n+1},x_1,j) - \varphi(t_{n+1},x_2,j)| \leq C_{n+1}|x_1-x_2| \tag{3.38}$$

Then $\Phi_j^{t_n,x}(\varphi)$ is Lipschitz with constant $C_{n+1}(1+Lh)$, uniformly in j, where $L := C_b + \frac{C_\sigma^2}{2} > 0$.

Proof. Choose $(t_n,j,x_1,x_2) \in \Pi \times \mathbb{I}_q \times \mathbb{R}^d \times \mathbb{R}^d$. Then:

$$\begin{aligned} \left|\Phi_j^{t_n,x_1}(\varphi) - \Phi_j^{t_n,x_2}(\varphi)\right| &= \left|\mathbb{E}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_1},j\right) - \varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_2},j\right)\right]\right| \\ &\leq \left\|\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_1},j\right) - \varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_2},j\right)\right\|_{L_1} \\ &\leq \left\|\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_1},j\right) - \varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_2},j\right)\right\|_{L_2}\end{aligned}$$

Now, using equations (3.38) and (3.5), and G denoting a d-dimensional standard Gaussian random variable, we have

$$\begin{aligned} &\mathbb{E}\left[\left(\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_1},j\right) - \varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x_2},j\right)\right)^2\right] \\ \leq\ & C_{n+1}^2\mathbb{E}\left[\left(\bar{X}_{t_{n+1}}^{t_n,x_1} - \bar{X}_{t_{n+1}}^{t_n,x_2}\right)^2\right] \\ \leq\ & C_{n+1}^2\mathbb{E}\left[\left(x_1-x_2+h\left(b(t_n,x_1)-b(t_n,x_2)\right)+\sqrt{h}\left(\sigma(t_n,x_1)-\sigma(t_n,x_2)\right)G\right)^2\right] \\ =\ & C_{n+1}^2\left\{\left(x_1-x_2+h\left(b(t_n,x_1)-b(t_n,x_2)\right)\right)^2 + h\mathbb{E}\left[\left(\left(\sigma(t_n,x_1)-\sigma(t_n,x_2)\right)G\right)^2\right]\right\} \\ \leq\ & C_{n+1}^2(x_1-x_2)^2\left\{1+\left(2C_b+C_\sigma^2\right)h + C_b^2h^2\right\} \\ \leq\ & C_{n+1}^2(x_1-x_2)^2\left(C_b+\frac{C_\sigma^2}{2}\right)^2\left(\frac{1}{C_b+\frac{C_\sigma^2}{2}}+h\right)^2.\end{aligned}$$

Thus:

$$\left|\Phi_j^{t_n,x_1}(\varphi) - \Phi_j^{t_n,x_2}(\varphi)\right| \leq C_{n+1}\left(1+\left(C_b+\frac{C_\sigma^2}{2}\right)h\right)|x_1-x_2|$$

□

Lemma 3.4. *Consider again a function $\varphi : \Pi \times \mathbb{R}^d \times \mathbb{I}_q \to \mathbb{R}$ such that (3.38) holds for a given $t_{n+1} \in \Pi$. Then, $\forall (x,j) \in \mathcal{D} \times \mathbb{I}_q$:*

$$\left|\tilde{\lambda}_j^{t_n,x}(\varphi) - \Phi_j^{t_n,x}(\varphi)\right| \leq C_{n+1}\delta(1+Lh).$$

In particular:

$$\left|\tilde{\Phi}_j^{t_n,x}(\varphi) - \Phi_j^{t_n,x}(\varphi)\right| \leq C_{n+1}\delta(1+Lh) \tag{3.39}$$

Proof. Recalling the definitions of $B_{t_n}(x)$, of $\tilde{\lambda}_j^{t_n,x}(\varphi)$ (equation (3.29)) and of $\Phi_j^{t_n,x}(\varphi)$ (equation (3.31)), simply remark that:

$$\begin{aligned}\min_{\tilde{x}\in B_{t_n}(x)}\Phi_j^{t_n,\tilde{x}}(\varphi) &\leq \Phi_j^{t_n,x}(\varphi) \leq \max_{\tilde{x}\in B_{t_n}(x)}\Phi_j^{t_n,\tilde{x}}(\varphi)\\ \min_{\tilde{x}\in B_{t_n}(x)}\Phi_j^{t_n,\tilde{x}}(\varphi) &\leq \tilde{\lambda}_j^{t_n,x}(\varphi) \leq \max_{\tilde{x}\in B_{t_n}(x)}\Phi_j^{t_n,\tilde{x}}(\varphi)\ .\end{aligned}$$

Now, using Lemma 3.3:

$$\begin{aligned}\left|\tilde{\lambda}_j^{t_n,x}(\varphi)-\Phi_j^{t_n,x}(\varphi)\right| &\leq \max_{\tilde{x}\in B_{t_n}(x)}\Phi_j^{t_n,\tilde{x}}(\varphi)-\min_{\tilde{x}\in B_{t_n}(x)}\Phi_j^{t_n,\tilde{x}}(\varphi)\\ &\leq C_{n+1}(1+Lh)\max_{(x_1,x_2)\in B_{t_n}(x)^2}|x_1-x_2|\\ &\leq C_{n+1}(1+Lh)\,\delta\end{aligned}$$

□

Lemma 3.5. *$\forall (t_n,x_1,x_2,i)\in\Pi\times\left(\mathbb{R}^d\right)^2\times\mathbb{I}_q$:*

$$|\bar{v}_\Pi(t_n,x_1,i)-\bar{v}_\Pi(t_n,x_2,i)|\leq C_n|x_1-x_2| \tag{3.40}$$

where:

$$\begin{aligned}C_N &= e^{-\rho t_N}C_g\\ C_n &= hC_f e^{-\rho t_n}+C_{n+1}(1+Lh)\ ,\ n=N-1,\dots,0\end{aligned} \tag{3.41}$$

In particular, $\exists C>0$ such that $\forall n=0,1,\dots,N$:

$$C_n\leq Ce^{-\rho t_n}e^{L(T-t_n)} \tag{3.42}$$

Proof. Recall Remark 3.5. We prove the lemma by induction. First, remark that, using hypothesis (3.21), it holds for $n=N$. Now, suppose that it holds for some $(n+1)\in[1,\dots,N]$. Then, using Lemma 3.3:

$$\begin{aligned}&\bar{v}_\Pi(t_n,x_1,i)\\ &=\max_{j\in\mathbb{I}_q}\left\{hf(t_n,x_1,j)-k(t_n,i,j)+\Phi_j^{t_n,x_1}(\bar{v}_\Pi)\right\}\\ &=\max_{j\in\mathbb{I}_q}\left\{hf(t_n,x_2,j)-k(t_n,i,j)+\Phi_j^{t_n,x_2}(\bar{v}_\Pi)+h(f(t_n,x_1,j)-f(t_n,x_2,j))+\left(\Phi_j^{t_n,x_1}(\bar{v}_\Pi)-\Phi_j^{t_n,x_2}(\bar{v}_\Pi)\right)\right\}\\ &\leq\max_{j\in\mathbb{I}_q}\left\{hf(t_n,x_2,j)-k(t_n,i,j)+\Phi_j^{t_n,x_2}(\bar{v}_\Pi)+he^{-\rho t_n}C_f|x_1-x_2|+C_{n+1}(1+Lh)|x_1-x_2|\right\}\\ &=\bar{v}_\Pi(t_n,x_2,i)+\left(he^{-\rho t_n}C_f+C_{n+1}(1+Lh)\right)|x_1-x_2|\end{aligned}$$

Symmetrically, the same inequality holds for $\bar{v}_\Pi(t_n,x_2,i)-\bar{v}_\Pi(t_n,x_1,i)$, yielding equations (3.40) and (3.41). Finally, use the discrete version of Gronwall's inequality to obtain equation (3.42) □

Proposition 3.5. *$\exists C>0$ s.t. $\forall(t,x,i)\in\Pi\times\mathbb{R}^d\times\mathbb{I}_q$:*

$$|\bar{v}_\Pi(t,x,i)-\tilde{v}_\Pi(t,x,i)|\leq C\frac{\delta}{h}e^{-\rho t}\,.$$

Proof. For each $t_n\in\Pi$, we look for an upper bound E_n, independent of x and i, of the quantity $|\bar{v}_\Pi(t_n,x,i)-\tilde{v}_\Pi(t_n,x,i)|$. First:

$$|\bar{v}_\Pi(T,x,i)-\tilde{v}_\Pi(T,x,i)|=|g(T,x,i)-g(T,x,i)|=0$$

Hence $E_N=0$. Fix now $n\in[0,N-1]$. Using Remark 3.5:

$$\begin{aligned}\tilde{v}_\Pi(t_n,x,i) &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j)-k(t_n,i,j)+\tilde{\Phi}_j^{t_n,x}(\tilde{v}_\Pi)\right\}\\ &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j)-k(t_n,i,j)+\Phi_j^{t_n,x}(\bar{v}_\Pi)\right.\\ &\qquad+\tilde{\Phi}_j^{t_n,x}(\bar{v}_\Pi)-\Phi_j^{t_n,x}(\bar{v}_\Pi)\\ &\qquad\left.+\tilde{\Phi}_j^{t_n,x}(\tilde{v}_\Pi)-\tilde{\Phi}_j^{t_n,x}(\bar{v}_\Pi)\right\}\end{aligned}$$

Using Lemmas 3.4 and 3.5, $\tilde{\Phi}_j^{t_n,x}(\bar{v}_\Pi) - \Phi_j^{t_n,x}(\bar{v}_\Pi) \leq C_{n+1}\delta(1+Lh)$ where C_{n+1} is the Lipschitz constant of $\bar{v}_\Pi$ at time t_{n+1} (see Lemma 3.5). Moreover,

$$\begin{aligned}\tilde{\Phi}_j^{t_n,x}(\tilde{v}_\Pi) - \tilde{\Phi}_j^{t_n,x}(\bar{v}_\Pi) &\leq \mathbb{E}\left[\tilde{v}_\Pi\left(t_{n+1},\bar{X}_{t_{n+1}},j\right) - \bar{v}_\Pi\left(t_{n+1},\bar{X}_{t_{n+1}},j\right)\middle| X_{t_n}\in B_{t_n}(x)\right]\\ &\leq E_{n+1}\,.\end{aligned}$$

Hence:

$$\tilde{v}_\Pi(t_n,x,i) \leq \bar{v}_\Pi(t_n,x,i) + C_{n+1}\delta(1+Lh) + E_{n+1}$$

Symmetrically, the same inequality holds for $\bar{v}_\Pi(T,x,i) - \tilde{v}_\Pi(t_n,x,i)$, leading to:

$$|\bar{v}_\Pi(t_n,x,i) - \tilde{v}_\Pi(t_n,x,i)| \leq E_n$$

where:

$$\begin{aligned}E_N &= 0\\ E_n &= C_{n+1}\delta(1+Lh) + E_{n+1}\,.\end{aligned}$$

Consequently, using equation (3.42):

$$E_n = \delta(1+Lh)\sum_{k=n+1}^{N} C_k \leq C\frac{\delta}{h}e^{-\rho t_n}$$

where $C>0$ does not depend on t_n nor T. □

The following lemma measures the regression error. It is an extension of Lemma 3.8 in [33] (itself inspired by Theorem 5.1 in [6]).

Lemma 3.6. *Consider a measurable function $\varphi:\Pi\times\mathbb{R}^d\times\mathbb{I}_q\to\mathbb{R}$. For any $p\geq 1$, there exists $C_p\geq 0$ such that $\forall(t_n,l,j)\in\Pi\times[1,M]\times\mathbb{I}_q$:*

$$\left\|\hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right\|_{L_p} \leq \frac{C_p}{\sqrt{M}}\frac{\Gamma^{t_n}(\varphi)+\bar{\varphi}^{t_n}}{\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^{1-\frac{1}{p\vee 2}}} + \frac{C_p}{M}\frac{\bar{\varphi}^{t_n}}{\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)} \tag{3.43}$$

where $\bar{\varphi}^{t_n}\in\mathbb{R}_+$ is such that $\left|\varphi\left(t_{n+1},\bar{X}_{t_{n+1}},j\right)\right|\leq\bar{\varphi}^{t_n}$ a.s. .

Proof. Define the following centered random variables:

$$\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi) := \frac{1}{M}\sum_{m=1}^{M}\varphi\Big(t_{n+1},\bar{X}^m_{t_{n+1}},j\Big)\mathbf{1}\{\bar{X}^m_{t_n}\in B_{t_n}(\bar{X}^l_{t_n})\} - \mathbb{E}\Big[\varphi\Big(t_{n+1},\bar{X}^m_{t_{n+1}},j\Big)\mathbf{1}\{\bar{X}^m_{t_n}\in B_{t_n}(\bar{X}^l_{t_n})\}\Big]$$

$$\varepsilon^{t_n,\bar{X}^l_{t_n}}(1) := \frac{1}{M}\sum_{m=1}^{M}\mathbf{1}\left\{\bar{X}^m_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\} - \mathbb{P}\left(\bar{X}^m_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)$$

Then:

$$\left|\hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right| = \left|\hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right| \wedge 2\Gamma^{t_n}(\varphi)$$

$$\leq \left|\hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right|\mathbf{1}\left\{\frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)}\leq\frac{1}{2}\right\} + 2\Gamma^{t_n}(\varphi)\mathbf{1}\left\{\frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)}>\frac{1}{2}\right\}$$

and:

$$\left|\hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right|\mathbf{1}\left\{\frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)}\leq\frac{1}{2}\right\}$$

$$= \left| \hat{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi) \frac{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\}} - \right.$$

$$\left. \tilde{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi) \frac{\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\}} \right| \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \leq \frac{1}{2} \right\}$$

$$\leq \left\{ \frac{\left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right|}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\}} \wedge 3\Gamma^{t_n}(\varphi) + \right.$$

$$\left. \left|\tilde{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right| \frac{\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\}} \right\} \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \leq \frac{1}{2} \right\}$$

$$\leq \frac{2}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \left\{ \left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) + \left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right| \Gamma^{t_n}(\varphi) \right\} \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \leq \frac{1}{2} \right\}$$

Now, for any $p \geq 1$:

$$\left|\hat{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right|^p$$
$$\leq \frac{2^{3p-2}}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)^p} \left\{ \left\{ \left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) \right\}^p + \left\{ \left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right| \Gamma^{t_n}(\varphi) \right\}^p \right\} \times$$
$$\mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \leq \frac{1}{2} \right\} + 2^{2p-1} \left(\Gamma^{t_n}(\varphi)\right)^p \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} > \frac{1}{2} \right\}$$

and:

$$\mathbb{E}\left[\left|\hat{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right|^p \right]$$
$$\leq \frac{2^{3p-2}}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)^p} \left\{ \mathbb{E}\left[\left\{ \left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) \right\}^p \right] + \left(\Gamma^{t_n}(\varphi)\right)^p \mathbb{E}\left[\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|^p \right] \right\}$$
$$+ 2^{2p-1} \left(\Gamma^{t_n}(\varphi)\right)^p \mathbb{P}\left(\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|^p > \frac{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)^p}{2^p} \right)$$
$$\leq \frac{8^p}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)^p} \left\{ \mathbb{E}\left[\left\{ \left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) \right\}^p \right] + \left\{\Gamma^{t_n}(\varphi)\right\}^p \mathbb{E}\left[\left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(1)\right|^p \right] \right\} \quad (3.44)$$

using Markov's inequality. We then obtain upper bounds for $\mathbb{E}\left[\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|^p\right]$ and $\mathbb{E}\left[\left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right|^p\right]$ using Lemma A.1 in Appendix A. Suppose that $\exists \bar{\varphi}^{t_n} \in \mathbb{R}_+$ s.t. $\left|\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)\right| \leq \bar{\varphi}^{t_n}$ a.s. . Then, using Lemma A.1, $\exists C_p > 0$ such that:

$$\mathbb{E}\left[\left|\varepsilon^{t_n,\bar{X}_{t_n}^l}(1)\right|^p \right] \leq \frac{C_p}{M^{\frac{p}{2}}} \mathbb{E}\left[\left| \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} - \mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right) \right|^{p \vee 2} \right]^{\frac{p}{p \vee 2}} \quad (3.45)$$

$$\mathbb{E}\left[\left|\varepsilon_j^{t_n,\bar{X}_{t_n}^l}(\varphi)\right|^p \right] \leq C_p \left\{ \frac{\left(\bar{\varphi}^{t_n}\right)^p}{M^p} + \frac{1}{M^{\frac{p}{2}}} \mathbb{E}\left[\left| \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} \right.\right.\right.$$
$$\left.\left.\left. - \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} \right] \right|^{p \vee 2} \right]^{\frac{p}{p \vee 2}} \right\} \quad (3.46)$$

where, for the second inequality, the term $m = l$ in the sum was treated separately. Then:

$$\mathbb{E}\left[\left| \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} - \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} \right] \right|^{p \vee 2} \right]^{\frac{p}{p \vee 2}}$$
$$\leq \left(2^{p \vee 2 - 1} \mathbb{E}\left[\left(\bar{\varphi}^{t_n}\right)^{p \vee 2} \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} + \mathbb{E}\left[\left(\bar{\varphi}^{t_n}\right)^{p \vee 2} \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} \right] \right] \right)^{\frac{p}{p \vee 2}}$$
$$\leq 2^p \left(\bar{\varphi}^{t_n}\right)^p \mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)^{\frac{p}{p \vee 2}} \quad (3.47)$$

In a similar manner:

$$\mathbb{E}\left[\left|\mathbf{1}\left\{\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\}-\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)\right|^{p\vee 2}\right]^{\frac{p}{p\vee 2}}\leq 2^p\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^{\frac{p}{p\vee 2}} \tag{3.48}$$

Finally, the combination of inequalities (3.44), (3.45), (3.46), (3.47) and (3.48) proves equation (3.43). □

We now apply Lemma 3.6 to $\bar{v}_\Pi$ in the following Corollary:

Corollary 3.1. *For every $p\geq 1$, there exists $C_p\geq 0$ s.t. $\forall\,(t_n,l,j)\in\Pi\times[1,M]\times\mathbb{I}_q$:*

$$\left\|\hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}\left(\tilde{v}_\Pi\right)-\tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}\left(\tilde{v}_\Pi\right)\right\|_{L_p}\leq C_p e^{-\rho t_n}\frac{1+R\left(T,\varepsilon\right)}{\sqrt{M}p\left(T,\delta,\varepsilon\right)^{1-\frac{1}{p\vee 2}}}\left(1+\frac{1}{\sqrt{M}p\left(T,\delta,\varepsilon\right)^{\frac{1}{p\vee 2}}}\right)$$

Proof. First, recall from equation (3.36) and (3.37) that there exists $C>0$ such that for every $(t_n,j)\in\Pi\times\mathbb{I}_q$:

$$\begin{aligned}\Gamma_j^{t_n}\left(\tilde{v}_\Pi\right) &= Ce^{-\rho t_n}\left(1+R\left(T,\varepsilon\right)\right)\\ \left|\tilde{v}_\Pi\left(t_{n+1},\bar{X}_{t_{n+1}},j\right)\right| &\leq Ce^{-\rho t_n}\left(1+R\left(T,\varepsilon\right)\right)\end{aligned}$$

Hence one can apply Lemma 3.6 to $\tilde{v}_\Pi$ with these upper bounds. The final step is to recall that the minimum probability $p\left(T,\delta,\varepsilon\right)$ defined in equation (3.16) is a lower bound on $\mathbb{P}\left(\bar{X}_{t_n}\in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)$ for any $(t_n,l)\in\Pi\times[1,M]$. □

Using this result, we can now assess the error between $\tilde{v}_\Pi$ and $\hat{v}_\Pi$.

Proposition 3.6. *$\forall p\geq 1$, $\exists C_p>0$ s.t. $\forall\,(t_n,l)\in\Pi\times[1,M]$:*

$$\left\|\sup_{i\in\mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(t,\bar{X}^l_{t_n},i\right)-\hat{v}_\Pi\left(t,\bar{X}^l_{t_n},i\right)\right|\right\|_{L_p}\leq C_p e^{-\rho t_n}\frac{1+R\left(T,\varepsilon\right)}{h\sqrt{M}p\left(T,\delta,\varepsilon\right)^{1-\frac{1}{p\vee 2}}}\left(1+\frac{1}{\sqrt{M}p\left(T,\delta,\varepsilon\right)^{\frac{1}{p\vee 2}}}\right)$$

where $\mathbb{I}_q^N$ is the set of $\mathcal{F}_{t_N}$-measurable random variables taking values in $\mathbb{I}_q$.

Proof. For each $t_n\in\Pi$, we look for an upper bound E_n, independent of l, such that:

$$\left\|\sup_{i\in\mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(t,\bar{X}^l_{t_n},i\right)-\hat{v}_\Pi\left(t,\bar{X}^l_{t_n},i\right)\right|\right\|_{L_p}\leq E_n\ .$$

First:

$$\left\|\sup_{i\in\mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(T,\bar{X}^l_T,i\right)-\hat{v}_\Pi\left(T,\bar{X}^l_T,i\right)\right|\right\|_{L_p}=\left\|\sup_{i\in\mathbb{I}_q^N}\left|g\left(T,\bar{X}^l_T,i\right)-g\left(T,\bar{X}^l_T,i\right)\right|\right\|_{L_p}=0$$

Hence $E_N=0$. Fix now $n\in[0,N-1]$. Recall the dynamic programming equations from Remark 3.5, and, for every $(i,l)\in\mathbb{I}_q^N\times[1,M]$, introduce $\tilde{j}^*$ (resp. $\hat{j}^*$) the arg max for $\tilde{v}_\Pi$ (resp. $\hat{v}_\Pi$) at point $\bar{X}^l_{t_n}$, i.e.:

$$\begin{aligned}\tilde{v}_\Pi\left(t_n,\bar{X}^l_{t_n},i\right) &= hf\left(t_n,\bar{X}^l_{t_n},\tilde{j}^*\right)-k\left(t_n,i,\tilde{j}^*\right)+\tilde{\Phi}_{\tilde{j}^*}^{t_n,\bar{X}^l_{t_n}}\left(\tilde{v}_\Pi\right)\\ \hat{v}_\Pi\left(t_n,\bar{X}^l_{t_n},i\right) &= hf\left(t_n,\bar{X}^l_{t_n},\hat{j}^*\right)-k\left(t_n,i,\hat{j}^*\right)+\hat{\Phi}_{\hat{j}^*}^{t_n,\bar{X}^l_{t_n}}\left(\hat{v}_\Pi\right)\end{aligned}$$

Now:

$$\begin{aligned}
\hat{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right) &= hf\left(t_n, \bar{X}^l_{t_n}, \hat{j}^*\right) - k\left(t_n, i, \hat{j}^*\right) + \hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_{\hat{j}^*}\left(\hat{v}_\Pi\right) \\
&= \left\{hf\left(t_n, \bar{X}^l_{t_n}, \hat{j}^*\right) - k\left(t_n, i, \hat{j}^*\right) + \tilde{\Phi}^{t_n,\bar{X}^l_{t_n}}_{\hat{j}^*}\left(\tilde{v}_\Pi\right)\right\} \\
&\quad + \left\{\hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_{\hat{j}^*}\left(\tilde{v}_\Pi\right) - \tilde{\Phi}^{t_n,\bar{X}^l_{t_n}}_{\hat{j}^*}\left(\tilde{v}_\Pi\right)\right\} + \left\{\hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_{\hat{j}^*}\left(\hat{v}_\Pi\right) - \hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_{\hat{j}^*}\left(\tilde{v}_\Pi\right)\right\} \\
&\leq \tilde{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right) + \sum_{j\in\mathbb{I}_q}\left|\hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right) - \tilde{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right)\right| \\
&\quad + \sup_{j\in\mathbb{I}^N_q}\left|\hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right) - \hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right)\right|
\end{aligned}$$

Symmetrically:

$$\begin{aligned}
\tilde{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right) &\leq \hat{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right) + \sum_{j\in\mathbb{I}_q}\left|\tilde{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right) - \hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right)\right| \\
&\quad + \sup_{j\in\mathbb{I}^N_q}\left|\hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right) - \hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\hat{v}_\Pi\right)\right|
\end{aligned}$$

Combining these two inequalities:

$$\begin{aligned}
\sup_{i\in\mathbb{I}^N_q}\left|\tilde{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right) - \hat{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right)\right| &\leq \sum_{j\in\mathbb{I}_q}\left|\hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right) - \tilde{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right)\right| \\
&\quad + \sup_{j\in\mathbb{I}^N_q}\left|\hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\hat{v}_\Pi\right) - \hat{\Phi}^{t_n,\bar{X}^l_{t_n}}_j\left(\tilde{v}_\Pi\right)\right|
\end{aligned}$$

Hence, using the triangular inequality, Corollary 3.1, equation (3.30), and the induction hypothesis:

$$\begin{aligned}
\left\|\sup_{i\in\mathbb{I}^N_q}\left|\tilde{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right) - \hat{v}_\Pi\left(t_n, \bar{X}^l_{t_n}, i\right)\right|\right\|_{L_p} \leq E_n := \; & C_p e^{-\rho t_n}\frac{1+R\left(T,\varepsilon\right)}{\sqrt{M}p\left(T,\delta,\varepsilon\right)^{1-\frac{1}{p\vee 2}}} \\
& + C_p e^{-\rho t_n}\frac{1+R\left(T,\varepsilon\right)}{Mp\left(T,\delta,\varepsilon\right)} + E_{n+1}
\end{aligned}$$

for some constant $C_p > 0$ which depends only on p. Consequently:

$$E_n \leq C_p e^{-\rho t_n}\frac{1+R\left(T,\varepsilon\right)}{h\sqrt{M}p\left(T,\delta,\varepsilon\right)^{1-\frac{1}{p\vee 2}}}\left(1+\frac{1}{\sqrt{M}p\left(T,\delta,\varepsilon\right)^{\frac{1}{p\vee 2}}}\right)$$

where $C_p > 0$ depends only on p. □

Finally, the combination of Propositions 3.1 3.2, 3.3, 3.5 and 3.6 at time $t = t_0$ proves Theorem 3.1.

4 Complexity analysis and memory reduction

4.1 Complexity

4.1.1 Computational complexity

The number of operations required by the algorithm described below is in $\mathcal{O}\left(q^2 \cdot N \cdot M\right)$, where we recall that q is the number of possible switches, N is the number of time steps and M is the number of Monte Carlo trajectories.

- The q^2 term stems from the fact that for every $i \in \mathbb{I}_q$, one has to compute a maximum on $j \in \mathbb{I}_q$ (see equation (3.15)). However, this q^2 can be reduced to q as soon as the two following conditions are satisfied:

1. (Irreversibility) The controlled variable can only be increased (or, symmetrically, can only be decreased)

2. (Cost Separability) There exists two functions k_1 and k_2 such that $\forall\,(t,i,j) \in \mathbb{R}_+ \times \mathbb{I}_q \times \mathbb{I}_q$, $k\,(t,i,j) = k_1\,(t,i) + k_2\,(t,j)$. For instance, this is true of affine costs.

Indeed, under those two conditions, equation (3.15) becomes:

$$\hat{v}_\Pi\,(t_n,x,i)+k_1\,(t_n,i) = \max_{j\in\mathbb{I}_q,\, j\geq i} \left\{ hf\,(t_n,x,j) - k_2\,(t_n,j) + \hat{\mathbb{E}}\left[\hat{v}_\Pi\left(t_{n+1}, \bar{X}^{t_n,x}_{t_{n+1}}, j\right)\right]\right\}\,,\ n = N-1,\dots,0$$

These maxima can be computed in $\mathcal{O}(q)$ instead of $\mathcal{O}\left(q^2\right)$ by starting from the biggest element $i = i_q$ down to the smallest element $i = i_1$ (in lexicographical order) and keeping track of the partial maxima.

Note that these two conditions hold for the numerical application from Section 5, providing the improved complexity $\mathcal{O}(q \cdot N \cdot M)$.

- The N term comes from the backward time induction.

- The M term corresponds to the cost of a regression, which is in $\mathcal{O}\,(M)$ (by using either the Cholesky decomposition or the more stable Thin SVD decomposition) .

4.1.2 Memory complexity

The memory size required for solving optimal switching problems (as well as the simpler American option problems and the more complex BSDE problems) by Monte Carlo methods is often said to be in $\mathcal{O}(N \cdot M)$, because, as the Euler scheme is a forward scheme and the dynamic programming principle is a backward scheme, the storage of the Monte Carlo trajectories seems inescapable. This fact is the major limitation of such methods, as acknowledged in [10] for instance.

Since such a complexity would be unbearable in high dimension, we describe below a general memory reduction method to obtain a much more amenable $\mathcal{O}(N + M)$ complexity (or, more precisely, of $\mathcal{O}(m \cdot N + q \cdot M)$ with $m \ll M$). This improvement really opens the door to the use of Monte Carlo methods for American options, optimal switching and BSDEs on high-dimensional practical applications. Note that this tool can be combined with all the existing Monte Carlo backward methods which (seem to) require the storage of all the trajectories.

A drawback of this tool is that it is limited to Markovian processes. However, one can usually circumvent this restriction by increasing the dimension of the state variable.

4.2 General memory reduction method

4.2.1 Description

The memory reduction method for Monte Carlo pricing of American options was pioneered by [12] for the geometric Brownian motion, and was subsequently extended to multi-dimensional geometric Brownian motions ([13]) as well as exponential Lévy processes ([14]). These papers take advantage of the additivity property of the processes considered. However, as briefly hinted in [37], the memory reduction trick can be extended to more general processes. In particular, it can be combined with any discretization scheme, for instance the Euler scheme or Milstein scheme, as long as the value of the stochastic process at one time step can be expressed via its value at the subsequent time step.

From a practical point of view, the production of "random" sequences usually involves wisely chosen deterministic sequences, with statistical properties as close as possible to true randomness (cf.

[25] for instance for an overview). These sequences are usually set using a *seed*, i.e. a (possibly multidimensional) fixed value aimed at initializing the algorithm which produces the sequence:

$$\begin{array}{ccccccccc} \{\text{set seed } s\} & \to & s_1 & \to & s_2 & \to & \dots & \to & s_n \\ & \text{rand()} & & \text{rand()} & & \text{rand()} & & \text{rand()} & \\ & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\ & \varepsilon_1 & & \varepsilon_2 & & \varepsilon_3 & & \varepsilon_n & \end{array} \tag{4.1}$$

The rand() produces a new random value ε and changes the internal seed value s. The internal value of the seed can be read (getseed()) and changed (setseed()). Now two useful aspects can be stressed. The first is that one can usually recover the current seed at any stage of the sequence. The second is that, if the seed is set later to, say, once again the seed s from equation (4.1), then the following elements of the sequence will be once again ε_1, ε_2, ... In other words, one can recover any previously produced subsequence of the sequence $(\varepsilon_n)_{n\geq 1}$, provided one stored beforehand the seed at the beginning of the subsequence. This feature is at the core of the memory reduction method, which we are going to discuss below in a general setting.

Consider a Markovian stochastic process $(X_t)_{t\geq 0}$, for instance the solution of the stochastic differential equation (2.2), recalled below:

$$\begin{aligned} X_0 &= x_0 \in \mathbb{R}^d \\ dX_s &= b(s, X_s)\, ds + \sigma(s, X_s)\, dW_s \end{aligned}$$

The application of the Euler scheme to this equation can be denoted as follows:

$$x^j_{t_{i+1}} = f\left(x^j_{t_i}, \varepsilon^j_i\right) \tag{4.2}$$

$$f(x, \varepsilon) := x + b(t_i, x)\, h + \sigma(t_i, x)\, \varepsilon\sqrt{h} \tag{4.3}$$

where $\forall i \in [0, N-1]$ and $\forall j \in [1, M]$, $\varepsilon^j_i \in \mathbb{R}^d$ is drawn from a d-dimensional Gaussian random variable. Suppose that for any $\varepsilon \in \mathbb{R}^d$, the function $x \mapsto f(x, \varepsilon)$ is invertible (call f_{inv} its inverse). Then, starting from the final value $x^j_{t_N}$ of the sequence (4.2), one can recover the whole trajectory of X:

$$x^j_{t_i} = f_{\text{inv}}\left(x^j_{t_{i+1}}, \varepsilon^j_i\right) \tag{4.4}$$

as long as one can recover the previous draws ε^j_{N-1}, ..., ε^j_0. The following pseudo-code describes an easy way to do it.

Algorithm 1 Euler Scheme — Inverse Euler Scheme

```
% Initialization
for j from 1 to M
        X[j] <- xj
end for

% LOOP 1: Euler scheme
for i from 0 to N-1
   S[i] <- getseed()
   for j from 1 to M
      E <- rand(d)
      X[j] <- f(X[j],E)
   end for
end for
S[N] <- getseed()
```

```
% LOOP 2: Inverse Euler scheme
for i from N-1 down to 0
   setseed(S[i])
   for j from 1 to M
      E <- rand(d)
      X[j] <- finv(X[j],E)
   end for
end for
setseed(S[N])
```

The first column of Algorithm 1 corresponds to the Euler scheme, with the addition of the storage of the seeds. At the end of the first colum, the vector X contains the last values X^j_T, $j = 1, \dots, M$. From this point, one can recover the previous values $X^j_{t_i}$, $i = N-1, \dots, 0$, $j = 1, \dots, M$ as done in the second column.

Inside this last loop, one can perform the estimation of the conditional expectations required by the resolution algorithm of our stochastic control problem (equation (2.10)). Compared to the standard storage of the full trajectories $X_{t_i}^j$, $i = 0, \dots, N$, $j = 1, \dots, M$, the pros and cons are the following:

- The number of calls to the rand() function is doubled.
- The memory needed is brought down from $\mathcal{O}(M \times N)$ to $\mathcal{O}(M + N)$ (storage of the vector space and the seeds).

In other words, at the price of doubling the computation time, one can bring down the required memory storage by the factor $\min(M, N)$, which is a very significant saving. Moreover, the theoretical additional computation time can be insignificant in practice, as the availability of much more physical memory makes the resort to the slower virtual memory much less likely.

Remark 4.1. Even though the storage of the seeds does take $\mathcal{O}(N)$ in memory size, the constant may be much greater than 1. For instance, on Matlab®, a seed from the Combined Multiple Recursive algorithm (refer for instance to [25] for a description of several random generators) is made of 12 uint32 (32-bit unsigned integer), a seed from the Multiplicative Lagged Fibonacci algorithm is made of 130 uint64, and a seed from the popular Mersenne Twister algorithm is made of 625 uint32.

In order to relieve the storage of the seeds, we now provide a finer memory reduction algorithm (Algorithm 2). Although Algorithm 2 requires three main loops, it enables to perform the last loop without fiddling the seed of the random generator, and without any vector of seeds locked in memory, which will thus be fully dedicated to the regressions and other resolution operations. Moreover, the first two main loops can be performed beforehand once and for all, storing only the last values of the vector X as well as the first seed S[0]. Finally, if the random generator is able to leapfrop a given number of steps, the first loop can be drastically reduced.

Algorithm 2 General Memory Reduction Method

```
% LOOP 1: Seeds storage
for i from 0 to N-1
   S[i] <- getseed()
   for j from 1 to M
      E <- rand(d)
   end for
end for

% Initialization
for j from 1 to M
        X[j] <- xj
end for
%
%
%
%
%
```

```
% LOOP 2: Euler scheme
for i from 0 to N-1
   setseed(S[N-i-1])
   for j from 1 to M
      E <- rand(d)
      X[j] <- f(X[j],E)
   end for
end for
setseed(S[0]) ; free(S)

% LOOP 3: Inverse Euler scheme
for i from N-1 down to 0
   for j from 1 to M
      E <- rand(d)
      X[j] <- finv(X[j],E)
   end for
end for
```

4.2.2 Numerical stability

Theoretically, the trajectories produced by the Euler scheme (4.2) and the inverse Euler scheme (4.4) are exactly the same. In practice however, a discrepancy may appear, the cause of which is discussed below.

On a computer, not all real numbers can be reproduced. Indeed, they must be stored on a finite number of bits, using a predefined format (usually the IEEE Standard for Floating-Point Arithmetic (IEEE 754)). In particular, there exists an incompressible distance $\varepsilon > 0$ between two different numbers stored. This causes rounding errors when performing operations on real numbers.

For instance, consider $x \in \mathbb{R}$ and an invertible function $f : \mathbb{R} \mapsto \mathbb{R}$. Compute $y = f(x)$ and then compute $\hat{x} = f_{\text{inv}}(y)$. One would expect that $\hat{x} = x$, but in practice, because of rounding effects, one may get $\hat{x} = x + \epsilon z$ for a small $\epsilon > 0$, where z is a discrete variable, which can be deemed random, taking values around zero. This phenomenon is illustrated on Figure 4.1, which displays a histogram of $\hat{x} - x$ for $n = 10^7$ different values of $x \in [0, 1]$ and for the simple linear function $f(x) = 2x + 3$.

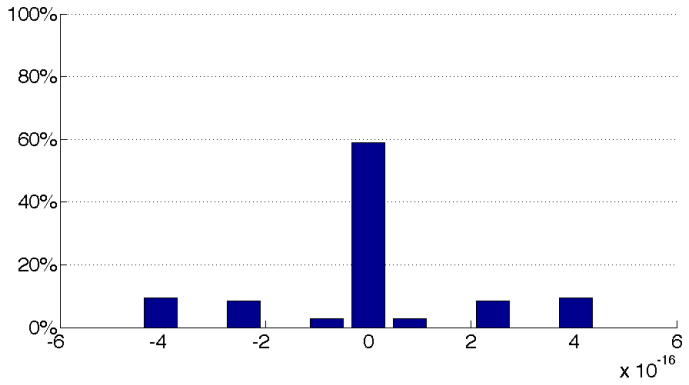

Figure 4.1: Histogram of rounding errors

We now describe how this affects our memory reduction method. Recall equation 4.2:

$$x^j_{t_{i+1}} = f\left(x^j_{t_i}, \varepsilon^j_i\right)$$

Now, instead of equation (4.4), the inverse Euler scheme will provide something like:

$$\begin{aligned} y^j_{t_N} &= x^j_{t_N} \\ y^j_{t_i} &= f_{\text{inv}}\left(y^j_{t_{i+1}}, \varepsilon^j_i\right) + \epsilon z^j_i \end{aligned} \tag{4.5}$$

for a small $\epsilon > 0$, where z^j_i, $i = 0, \ldots, N$, $j = 1, \ldots, M$, can be deemed realizations of a discrete random variable Z, independent of W. The distribution of Z is unknown, but data suggests it may be innocuously assumed centered, symmetric, and with finite moments.

We are now interested in studying the compound rounding error $y_{t_i} - x_{t_i}$ as a function of ϵ. Of course, its behaviour depends on the choice of f (equation (4.3)). Below, we explicit this error on two simple examples: an arithmetic Brownian motion and an Ornstein-Uhlenbeck process. These two examples illustrate how the compound rounding error can vary dramatically w.r.t. f.

First example: arithmetic Brownian motion Consider first the case of an arithmetic Brownian motion with drift parameter μ and volatility parameter σ. Here f and its inverse are given by:

$$\begin{aligned} f(x, \varepsilon) &= x + \mu h + \sigma\sqrt{h}\varepsilon \\ f_{\text{inv}}(x, \varepsilon) &= x - \mu h - \sigma\sqrt{h}\varepsilon \end{aligned}$$

Hence, using equation (4.5), for every $j = 1, \ldots, M$:

$$y^j_{t_i} - x^j_{t_i} = \epsilon \sum_{k=i}^{N-1} z^j_k$$

In other words, the compound rounding error behaves as a random walk, multiplied by the small parameter ϵ. Hence, as long as $\epsilon \ll h$ (which is always the case as real numbers smaller than ϵ cannot be handled properly on a computer), this numerical error is harmless.

Remark that a similar numerical error arises from the algorithms proposed in [12] , [13] and [14], but, fortunately, as discussed above, this error is utterly negligible.

Second example: Ornstein-Uhlenbeck process Now, consider the case of an Ornstein-Uhlenbeck process with mean reversion $\alpha > 0$, long-term mean μ and volatility σ. Here:

$$\begin{aligned} f(x,\varepsilon) &= x + \alpha(\mu - x)h + \sigma\sqrt{h}\varepsilon \\ f_{\text{inv}}(x,\varepsilon) &= \frac{1}{1-\alpha h}\left(x - \alpha\mu h - \sigma\sqrt{h}\varepsilon\right) \end{aligned}$$

Using equation (4.5), for every $j = 1, \dots, M$ the compound error is given by:

$$y_{t_i}^j - x_{t_i}^j = \epsilon \sum_{k=i}^{N-1} \frac{1}{(1-\alpha h)^{k-i}} z_k^j$$

As $(1-\alpha h)^{-N} \sim \exp(\alpha T)$ when $h \to 0$, one can see that, as soon as $T > -\frac{\ln(\epsilon)}{\alpha}$, this error may become overwhelming. This phenomenon is illustrated on Figure 4.2a on a sample of 100 trajectories.

In order to mitigate this effect, we propose to modify the Algorithm 2 as follows: in its second loop (usual Euler scheme), instead of saving only the last values x_T^j, one may define a small subset $\tilde{\Pi} \subset \Pi$ and save the intermediate values $x_{t_i}^j$, $t_i \in \tilde{\Pi}$. Then, in the last loop (inverse Euler scheme), every time that $t_i \in \tilde{\Pi}$, the current value of the set $x_{t_i}^j$ may be recovered from this previous storage.

Figure 4.2b illustrates the new behaviour of the compound rounding error with this mended algorithm, on an example with $T = 10$ years and 4 intermediate saves (in addition to the final values).

The drawback of this modification, of course, is that it multiplies the required storage space by the factor $\#\tilde{\Pi}$. However, this remains much smaller than the $\mathcal{O}(M \times N)$ required by the naive full storage algorithm.

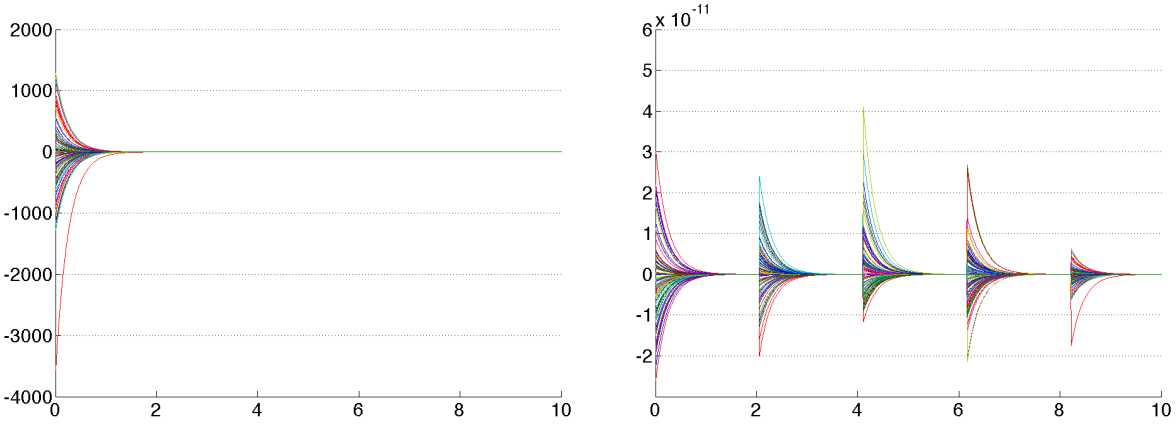

(a) Without intermediate saves

(b) With intermediate saves

Figure 4.2: Compound rounding error for the Ornstein-Uhlenbeck process

5 Application to investment in electricity generation

This section is devoted to an application of the resolution method studied in Section 3 to an investment problem in electricity generation.

Since our intention here is to show that the algorithm described in Section 3 can handle high-dimensional problems, our modeling of the electric system focuses on the various fundamental drivers of the electricity spot price formation mechanism that are electricity demand, available capacities and above all fuel prices.

Thus, were neglected some strategic aspects of investment, like construction delays and network constraints. We did not consider dynamic constraints of production either, which are known to

increase spot price during peak hours and to decrease them during off peak hours (see [26]), as we consider these effects to be negligible compared to the effect induced by a lack or an excess of capacity.

We based our model on [3, 1] where the electricity spot price is defined as a combination of fuel prices adjusted by a scarcity factor. This model exhibits the main feature wanted here, which is that the spot price, being determined both by the fuel prices and the residual capacity, is directly affected by the evolution of the installed capacity. When the residual capacity tends to decrease, spot prices will tend to increase, making investment valuable. Thus, in this model, investments are undertaken not on the specific purpose of satisfying the demand but as soon as they are profitable. Energy non-served and loss of load probability may still be adjusted through the price cap on the spot market.

In this section, we first detail the chosen modeling and objective function (which will be shown to be encompassed in the general optimal multiple switching problem (2.1)), and then solve it numerically using the general algorithm developed in the previous sections.

5.1 Modeling

The key variable in order to describe our electricity generation investment problem is the price of electricity. More precisely, the key quantities are the spreads between the prices of electricity and other energies. To model these spreads accurately, it may be worth considering a structural model for electricity (cf. the survey [8]). Here we choose such a model, mainly inspired by those introduced in [3] and [1], albeit amended and customized for a long-term time horizon. All the variables involved are detailed below.

5.1.1 Electricity demand

The electricity demand, or electricity load, at time t on the given geographical zone considered is modelled by an exogenous stochastic process $(D_t)_{t\geq 0}$:

$$D_t = f_0(t) + Z_t^0 \tag{5.1}$$

where Z^0 is an Ornstein-Uhlenbeck (henceforth O.U.) process:

$$dZ_t^0 = -\alpha_0 Z_t^0 dt + \beta_0 dW_t^D$$

where α_0 and β_0 are constants, and f_0 is a deterministic function that takes into account demand seasonalities.

5.1.2 Production capacities

Let d' be the number of different production technologies. Denote as $I_t = \left(I_t^1, \dots, I_t^{d'}\right)$ the installed production capacities at time t. They represent the maximum amount of electricity that is physically possible to produce. These fleets can be modified: at a given time τ_n, one can decide to build (or dismantle) an amount ζ_n of capacities:

$$I_{\tau_n} = I_{\tau_n^-} + \zeta_n\,,\ n \geq 0 \tag{5.2}$$

Denote as $\alpha = (\tau_n, \zeta_n)_{n\geq 1}$ the corresponding impulse control strategy, where $(\tau_n)_{n\geq 0}$ is an increasing sequence of stopping times with $\tau_n \nearrow \infty$ when $n \to \infty$, and $(\zeta_n)_{n\geq 0}$ is a sequence of vectors corresponding to the increases (or decreases) in capacities. Apart from these variations, I_t will be deemed constant, i.e.:

$$I_t = I_{0-} + \sum_{n,\,\tau_n \leq t} \zeta_n\,. \tag{5.3}$$

Now, denote as $C_t = \left(C_t^1, \dots, C_t^{d'}\right)$ the available production capacities. Because of spinning reserves, maintenance and random outages, these quantities are lower than the installed capacities I_t, which represent their physical maximum. In other terms, C_t is a fraction of I_t:

$$C_t^i = I_t^i \times A_t^i \tag{5.4}$$

for every $1 \leq i \leq d'$, where A_t^i corresponds to the rate of availability of the i$^{\text{th}}$ production technology. Therefore one must choose a model for the process A_t that ensures that it stays within the interval $[0,1]$. One could use the bounded Jacobi process (cf. for instance [36] and references therein), but here we choose a simpler modeling. Adapting the (bounded) wind power infeed efficiency model from [38], we model $\left(A_t^i\right)_{t\geq 0}^{1\leq i\leq d'}$ as follows:

$$A_t^i := \mathcal{T}\left(f_i\left(t\right) + Z_t^i\right) \tag{5.5}$$

where Z, f and $\mathcal{T}$ are chosen as follows:

- Z^i is an O.U. process :
$$dZ_t^i = -\alpha_i Z_t^i dt + \beta_i dW_t^{Z^i}$$
where $\alpha_i > 0$, $\beta_i > 0$ and $\left(W_t^{Z^i}\right)_{t\geq 0}$ is a Brownian motion.

- The deterministic function f_i accounts for the seasonality in the availability of production capacities, which stems from the maintenance plannings, which usually mimic the long term seasonality of demand (which in turn originates in the seasonality of temperature).

- The mapping $\mathcal{T} : \mathbb{R} \to [0,1]$ is here to ensure that $\forall t \geq 0$, $A_t \in [0,1]^{d'}$.

5.1.3 Fuels and CO_2 prices

For each technology i, denote as S_t^i the price of the fuel i to produce electricity at time t. In the particular case of renewable energies, which, *per se,* do not involve traded fuels, the corresponding S_t^i can be chosen to be zero. Moreover, define S_t^0 as the price of CO_2. Denote as S_t the full vector $\left(S_t^0, S_t^1, \ldots, S_t^{d'}\right)$.

Now, we introduce the multiplicative constants needed to convert theses quantities into €/MWh. For each technology $i = 1, \ldots, d'$, let h_i denote its heat rate, and h_i^0 denote its CO_2 emission rate. Hence, the quantity

$$\tilde{S}_t^i := h_i^0 S_t^0 + h_i S_t^i \tag{5.6}$$

expressed in €/MWh, corresponds to the price in € to pay in order to produce 1MWh of electricity using the ith technology. We note $h^0 = \left(h_1^0, \ldots, h_{d'}^0\right) \in \mathbb{R}^{d'}$ and $h = \left(h_1, \ldots, h_{d'}\right) \in \mathbb{R}^{d'}$.

Remark 5.1. One can choose to add a fixed cost into the definition of $\tilde{S}_t^i$. This is all the more so relevant for technologies whose fixed costs outweigh the cost of fuel (e.g. nuclear).

Over long time horizons, it is crucial to take into account the existence of long-term relationships between energy prices (c.f.[30] for instance). Thus, extending the model of cointegrated Brownian motions from [4], we model S_t as cointegrated geometric Brownian motions:

$$dS_t = \Xi S_t dt + \text{diag}\left(S_t\right) \Sigma dW_t^S$$

where Ξ is the $(d'+1) \times (d'+1)$ cointegration matrix (which models the long term relations), Σ is the $(d'+1) \times (d'+1)$ covariance matrix (which models the short term behaviour), and $\left(W_t^S\right)_{t\geq 0}$ is a $(d'+1)$-dimensional Brownian motion. We assume that $1 \leq \text{rank}\left(\Xi\right) < d'$ (so as to produce "true" cointegration, see [4]), and that for every $i \neq j$, $\Xi_{i,j} \geq 0$ (so as to ensure that the process S stays positive, see Appendix B).

5.1.4 Electricity price

We model the price of electricity using a long-term structural model. We model it as the sum of two building blocks: the marginal cost of producing electricity (cf. [3] for more details) plus a power law *scarcity premium* (along the lines of [1]), this sum being capped at a fixed upper bound [1].

[1] Indeed, in the French, German and Austrian markets for instance, power prices cannot be set outside the $[-3000, 3000]$€/MWh range, see http://www.epexspot.com/en/product-info/auction..

For any time $t \geq 0$, define the permutation $(1), \ldots (d')$ of the numbers $1, \ldots, d'$, such that $\widetilde{S}_t^{(1)} \leq \ldots \leq \widetilde{S}_t^{(d')}$. Then, define $\overline{C}_t^{(i)}$ as the total capacity available at time t from the i first technologies, i.e. $\overline{C}_t^{(i)} := \sum_{j \leq i} C_t^{(j)}$.

Now, from two points (x_1, y_1) and (x_2, y_2) in $\mathbb{R}^2$, one can always find three positive constants $a := a(x_1, x_2, y_1, y_2)$, $b := b(x_1, x_2, y_1, y_2)$ and $c := c(x_1, x_2, y_1, y_2)$ such that the function:

$$p(x) := p(x; x_1, x_2, y_1, y_2) = \frac{x}{a - bx} + c \tag{5.7}$$

satisfies $p(x_1) = y_1$ and $p(x_2) = y_2$.

Using this notation, we model the price P_t of electricity as follows:

$$\begin{aligned} P_t \quad := \quad & \widetilde{S}_t^{(1)} \mathbf{1}\{D_t < 0\} + \left\{ \widetilde{S}_t^{(1)} + p\left(D_t; 0, \overline{C}_t^{(1)}, \widetilde{S}_t^{(1)}, \widetilde{S}_t^{(2)}\right) \right\} \mathbf{1}\left\{ 0 \leq D_t < \overline{C}_t^{(1)} \right\} \\ & \sum_{i=2}^{d'-1} \left\{ \widetilde{S}_t^{(i)} + p\left(D_t; \overline{C}_t^{(i-1)}, \overline{C}_t^{(i)}, \widetilde{S}_t^{(i)}, \widetilde{S}_t^{(i+1)}\right) \right\} \mathbf{1}\left\{ \overline{C}_t^{(i-1)} \leq D_t < \overline{C}_t^{(i)} \right\} \\ & + \left\{ \widetilde{S}_t^{(d')} + p\left(D_t; \overline{C}_t^{(d'-1)}, \overline{C}_t^{(d')}, \widetilde{S}_t^{(d')}, P_{\max}\right) \right\} \mathbf{1}\left\{ \overline{C}_t^{(d'-1)} \leq D_t \right\} \end{aligned} \tag{5.8}$$

where $P_{\max} > 0$ is a fixed upper bound on the price of electricity. In particular, the last term, the one involving $P_{\max}$, enables price spikes to occur (when the residual capacity is small). Remark that the price of CO_2 emissions is explicitly included in the marginal cost (through equation (5.6)). Finally, remark that thanks to the knitting function (5.7), the electricity price P is a Lipschitz continuous function of the structural variables D, C and S.

5.1.5 Objective function

We now explicit the objective function of the investor in electricity generation. Suppose that, at time t, the level of installed capacity of type $j \in [1, d']$ is changed from I_{t-}^j to $I_s^j = I_{t-}^j + \zeta^j$, $s \geq t$. It generates the cost:

$$k\left(\zeta^j\right) := \begin{cases} \kappa_j^{f+} + \zeta^j \kappa_j^{p+} & , \zeta^j > 0 \\ 0 & , \zeta^j = 0 \\ \kappa_j^{f-} - \zeta^j \kappa_j^{p-} & , \zeta^j < 0 \end{cases}$$

where κ_j^{f+} and κ_j^{p+} are the fixed and proportional costs of building new plants of type j, and κ_j^{f-} and κ_j^{p-} are the fixed and proportional costs of dismantling old plants of type j.

Summing up the gains of the whole fleet of power plants on a given geographical zone, discounted to time 0 using a constant interest rate $\rho > 0$, and maximising its expectation along the potential new plants yield the following value function (cf.[2] for more details):

$$v(t, x, i) = \sup_{\alpha \in \mathcal{A}_{t,i}} \mathbb{E}\left[\sum_{j=1}^{d'} \int_t^\infty e^{-\rho s} \left(\min\left\{ C_s^j, D_s - \overline{C}_s^{(j-1)} \right\} \times \left(P_s - \widetilde{S}_s^j \right)^+ - \kappa_i \right) ds - \sum_{\tau_n \geq t} e^{-\rho \tau_n} k\left(\zeta^j\right) \right] \tag{5.9}$$

where the strategies α affect the installed capacities (equations (5.3)), hence also the available capacities (equation (5.4)) as well as the power price (equation (5.8)).

Replacing P in (5.9) by its definition (5.8), it is patent that this objective function fits into the mould studied thoroughly in Section 3. In Subsection 5.2 below, we apply our algorithm to this specific objective function.

5.2 Numerical results

Finally, we solve the control problem described in Subsection 5.1 on a numerical example, using the algorithm detailed in Subsection 3 combined with the general memory reduction method described in Subsection 4.2.

Our purpose here is not to perform a full study of investments in electricity markets, but a more modest attempt at illustrating the practical feasibility of our approach, with some possible outputs that the algorithm can provide.

We consider a numerical example including two cointegrated fuels (in addition to the price of CO_2): one "base fuel" and one "peak fuel", starting respectively from 40€/MWh and 80€/MWh. Hence, using the notations from Subsection 5.1, $d' = 2$ (two technologies) and $d = 6$ (electricity demand, CO_2 price, two fuel prices and two availability rates). The main choices of parameters for this application (initial fuel prices and volatilities, initial fleet and proportional costs of new power plants) are summed up in Table 5.1. Moreover, the demand process starts from $D_0 = 70$GW and does not integrate any linear trend.

i	S_0^i	σ^i	I_0^i	κ_i^{p+}
1	40€/MWh	5%	67GW	0.24 10^9€/GW
2	80€/MWh	15%	33GW	2.00 10^9€/GW

Table 5.1: Model parameters

In order to take into account the minimum size of one power plant we restrict the values of the installed capacity process(5.3) to a (bi-dimensional) fixed grid $\Lambda^{d'}$, with a mesh of 1GW. We make the simplifying assumptions that investments are irreversible, and that no dismantling can occur (recall from Subsection 4.1 the computational gain provided by this assumption).

Remark 5.2. If such a grid is indeed manageable in dimension $d' = 2$, it may less be the case if additional technologies were considered. However, as discussed in [33] equation (3.2), instead of performing one regression for each $i \in \Lambda^{d'}$, one can solve equation (3.15) at time t_i by only one $(d+d')$-dimensional regression, by choosing an a priori law for the randomized control ζ_{t_i}. The error analysis from Section 2 can be generalized to such regressions in higher dimension.

Finally, we consider the following numerical parameters. We choose a time horizon $T = 40$ years and a time step $h = \frac{1}{730}$ (i.e. two time steps per day, allowing for some intraday pattern in the demand process) but allow for only one investment decision per year. For the regression, we consider a basis of $b = 2^d = 64$ adaptative local functions, chosen piecewise linear on each hypercube (which is a bit more refined than the piecewise constant basis studied in Section 3) on a sample of $M = 5000$ trajectories.

With these numerical parameters, we obtain a non-parametric confidence interval of $[3.731, 3.752] \times 10^8$ for the value function $v(0, x_0, i_0)$ at time 0 (cf. Appendix C on how these bounds are computed), i.e. a relative error smaller than 1%, which is sufficiently small for the numerical results obtained, displayed on Figures 5.1 and 5.2, to be considered relevant.

First, Figure 5.1 deals with the optimal strategies. Figure 5.1a displays the time evolution of the average as well as the variability of the optimal fleet (only the new plants are shown). One can distinguish a first short phase characterised by the construction of several GW of peak load assets, followed by a much slower second phase involving the construction of both base load and peak load assets. Moreover, the variability of the optimal fleet increases over time. The detailed histogram of the optimal strategy at time $T = 40$ years is displayed on Figure 5.1b, where it is combined with the price of fuel. One can see that the more the peak fuel is expensive (and hence both fuels are expensive on average, as they are cointegrated), the more constructions of base load plants occur.

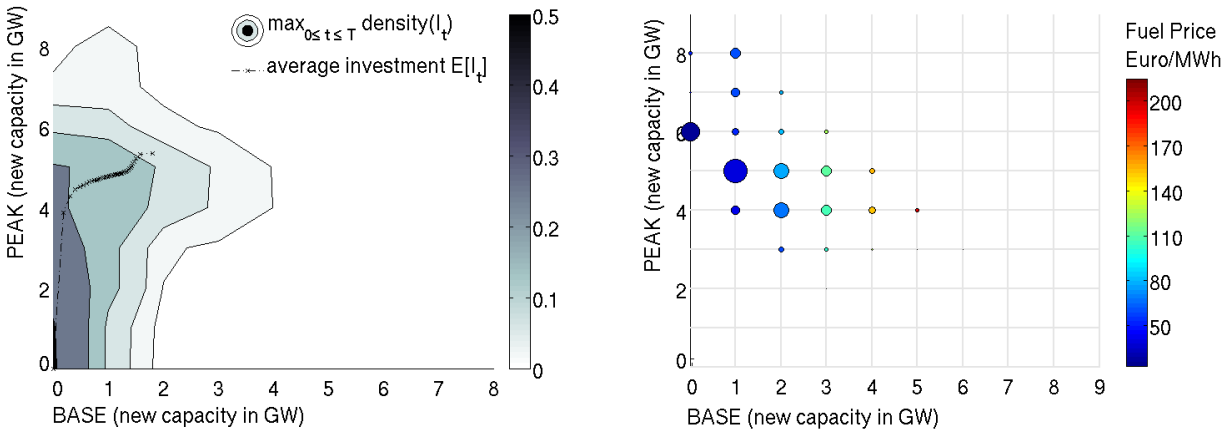

(a) Time evolution of new capacities

(b) Final fleet distribution

Figure 5.1: Optimal strategies

The fact that the average fleet seem to converge is related to the fact that this numerical example does not consider any growth trend in the electricity demand. Otherwise, more investments would occur, indeed, to keep the pace with consumption.

Then Figure 5.2 provides information on the price of electricity. Figure 5.2a displays the time evolution of the electricity spot price density. For better readability, each density covers one whole year. One can see how the density moves away from the initial bimodal density (with prices clustering around the initial prices of the two fuels) towards a more diffuse density. Moreover, the downward effect of investments on prices can be noticed. This downward effect is even more visible on Figure 5.2b. It compares the effect on electricity prices of three different strategies: the optimal strategy, the optimal deterministic strategy (computed as the average of the optimal strategy), and the do-nothing strategy. For each strategy, the joint time-evolution of the yearly median price and the yearly interquartile range are drawn. As expected, prices tend to be higher and more scattered without any new plant. Nevertheless, on this specific example, the price distribution under the optimal deterministic strategy is close to that under the optimal strategy (only slightly more scattered).

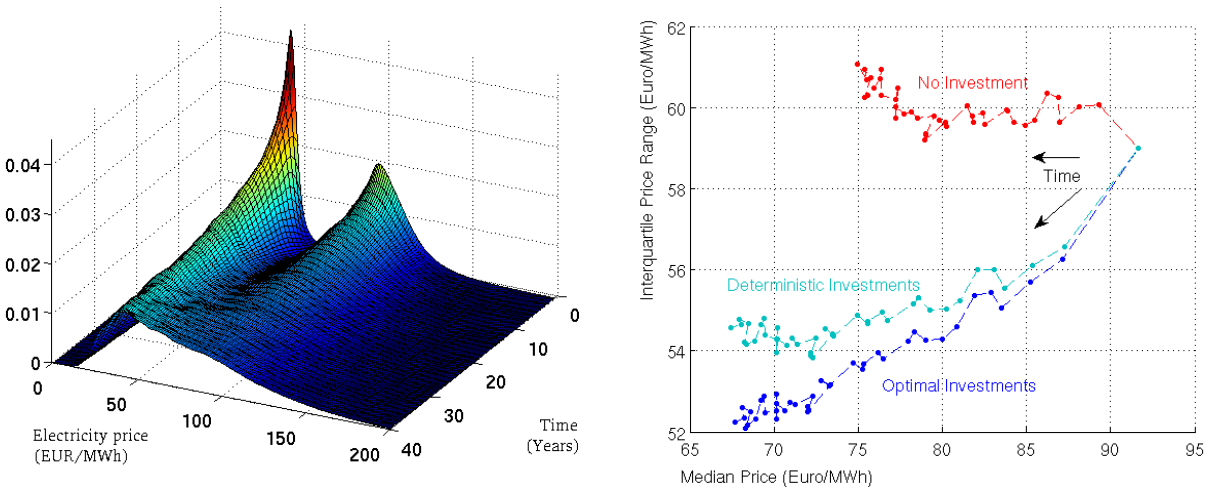

(a) Time evolution of electricity spot price density

(b) Comparison between investment strategies

Figure 5.2: Electricity spot price

These few pictures illustrate the kind on information that can be be extracted from the resolution of

this control problem. Of course, as a by-product of the resolution, much more can be extracted and analyzed (distribution of income, CO_2 emissions, optimal exercise frontiers, etc) if needed.

6 Conclusion

In this paper, we presented a probabilistic method to solve optimal multiple switching problems. We showed on a realistic investment model for electricity generation that it can efficiently provide insight into the distribution of future generation mixes and electricity spot prices. We intend to develop this work in several directions in the future. First, we wish to take into account more generation technologies, most notably wind farms, nuclear production, as well as solar distributed production. These additions would raise the dimension of the problem from eight to fifteen. Yet another range of innovations in numerical methods will be necessary to overcome this increase in dimension. Second, we wish to take time-to-build into account. And last but not least, we wish to adapt the problem to a continuous-time multiplayer game and contribute to the quest for an efficient algorithm to solve it.

References

[1] R. Aïd, L. Campi, and N. Langrené. A structural risk-neutral model for pricing and hedging power derivatives. *Mathematical Finance*, 23(3):387–438, July 2013. (Cited on pages 1, 26, and 27)

[2] R. Aïd, L. Campi, N. Langrené, and H. Pham. A probabilistic numerical method for optimal multiple switching problem in high dimension. Preprint, 2012. (Cited on page 28)

[3] R. Aïd, L. Campi, A. Nguyen Huu, and N. Touzi. A structural risk-neutral model of electricity prices. *International Journal of Theoretical and Applied Finance*, 12(7):925–947, 2009. (Cited on pages 1, 26, and 27)

[4] G. Benmenzer, E. Gobet, and C. Jérusalem. Arbitrage free cointegrated models in gas and oil future markets. Technical report, GDF SUEZ and Laboratoire Jean Kuntzmann, Grenoble, 2007. (Cited on page 27)

[5] A. Botterud, M. Ilic, and I. Wangensteen. Optimal investments in power generation under centralized and decentralized decision making. *IEEE Transactions on Power Systems*, 20(1):254–263, 2005. (Cited on page 2)

[6] B. Bouchard and N. Touzi. Discrete-time approximation and Monte-Carlo simulation of backward stochastic differential equations. *Stochastic Processes and their Applications*, 111(2):175–206, 2004. (Cited on pages 1, 8, 12, and 17)

[7] B. Bouchard and X. Warin. Monte-Carlo valorisation of American options: facts and new algorithms to improve existing methods. In R. Carmona, P. Del Moral, P. Hu, and N. Oudjane, editors, *Numerical Methods in Finance*, volume 12 of *Springer Proceedings in Mathematics*, 2012. (Cited on pages 1, 5, 8, 10, and 35)

[8] R. Carmona and M. Coulon. A survey of commodity markets and structural models for electricity prices. In F. Benth, V. Kholodnyi, and P. Laurence, editors, *Quantitative Energy Finance: Modeling, Pricing and Hedging in Energy and Commodity Markets*, 2013. (Cited on page 26)

[9] R. Carmona, P. Del Moral, N. Oudjane, and P. Hu. *Numerical Methods in Finance*. Springer, 2012. (Cited on page 1)

[10] R. Carmona and M. Ludkovski. Pricing asset scheduling flexibility using optimal switching. *Applied Mathematical Finance*, 15(5):405–447, 2008. (Cited on pages 1, 5, 12, and 21)

[11] J. Carriere. Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics*, 19(1):19–30, 1996. (Cited on page 5)

[12] R. Chan, Y. Chen, and K.-M. Yeung. A memory reduction method in pricing American options. *Journal of Statistical Computation and Simulation*, 74(7):501–511, 2004. (Cited on pages 1, 21, and 24)

[13] R. Chan, C.-Y. Wong, and K.-M. Yeung. Pricing multi-asset American-style options by memory reduction Monte Carlo methods. *Applied Mathematics and Computation*, 179(2):535–544, 2006. (Cited on pages 1, 21, and 24)

[14] R. Chan and T. Wu. Memory-reduction method for pricing American-style options under exponential Lévy processes. *East Asian Journal on Applied Mathematics*, 1(1):20–34, 2011. (Cited on pages 1, 21, and 24)

[15] J. Chassagneux, R. Elie, and I. Kharroubi. *Annals of Applied Probability*, (3):971–1007. (Cited on page 12)

[16] B. El Asri. Optimal multi-modes switching problem in infinite horizon. *Stochastics and Dynamics*, 10(2):231–261, 2010. (Cited on page 5)

[17] M. Fischer and G. Nappo. On the moments of the modulus of continuity of Itô processes. *Stochastic Analysis and Applications*, 28(1):103–122, 2009. (Cited on page 12)

[18] P. Gassiat, I. Kharroubi, and H. Pham. Time discretisation and quantization methods for optimal multiple switching problem. *Stochastic Processes and their Applications*, 122(5):2019–2052, 2012. (Cited on pages 1 and 12)

[19] E. Gobet, J.-P. Lemor, and X. Warin. A regression-based Monte Carlo method to solve Backward Stochastic Differential Equations. *The Annals of Applied Probability*, 15(3):2172–2202, 2005. (Cited on page 8)

[20] S. Hamadène. Optimal switching systems of reflected bsdes and systems of variational inequalities with interconnected obstacles. In J. Blath, P. Imkeller, and S. Roelly, editors, *Surveys in Stochastic Processes*, Series of Congress Reports, 2011. (Cited on page 5)

[21] S. Hamadène, J.-P. Lepeltier, and Z. Wu. Infinite horizon reflected bsdes and applications in mixed control and game problems. *Probability and Mathematical Statistics*, 19(2):211–234, 1999. (Cited on page 5)

[22] Y. Hu and S. Tang. Multi-dimensional BSDE with oblique reflection and optimal switching. *Probability Theory and Related Fields*, 147:89–121, 2010. (Cited on page 12)

[23] P. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Stochastic Modelling and Applied Probability*. Springer, 3rd edition, 1999. (Cited on pages 3 and 4)

[24] M. Kohler. A review on regression based Monte Carlo methods for pricing American options. In L. Devroye, B. Karasözen, M. Kohler, and R. Korn, editors, *Recent Developments in Applied Probability and Statistics*, Physica-Verlag, pages 39–61, 2010. (Cited on page 5)

[25] D. Kroese, T. Taimre, and Z. Botev. *Handbook of Monte Carlo methods*, volume 706 of *Wiley series in probability and statistics*. Wiley, 2011. (Cited on pages 22 and 23)

[26] N. Langrené, W. van Ackooij, and F. Bréant. Dynamic constraints for aggregated units: Formulation and application. *IEEE Transactions on Power Systems*, 26(3):1349–1356, August 2011. (Cited on page 26)

[27] J.-P. Lemor, E. Gobet, and X. Warin. Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations. *Bernoulli*, 12(5):889–916, 2006. (Cited on page 10)

[28] F. Longstaff and E. Schwartz. Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies*, 14(1):113–147, 2001. (Cited on pages 1 and 5)

[29] B. Mo, J. Hegge, and I. Wangensteen. Stochastic generation expansion planning by means of stochastic dynamic programming. *IEEE Transactions on Power Systems*, 6(2):662–668, 1991. (Cited on page 2)

[30] J. Obermayer. An analysis of the fundamental price drivers of EU ETS carbon credits. Master's thesis, KTH Royal Institute of Technology, Stockholm, 2009. (Cited on page 27)

[31] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*. Springer, 3rd edition, 1999. (Cited on page 34)

[32] R. Seydel. Existence and uniqueness of viscosity solutions for QVI associated with impulse control of jump-diffusions. *Stochastic Processes and their Applications*, 119(10):3719–3748, 2009. (Cited on page 5)

[33] X. Tan. A splitting method for fully nonlinear degenerate parabolic PDEs. *Electronic Journal of Probability*, 18(15):1–24, 2013. (Cited on pages 1, 8, 17, and 29)

[34] N. Todorović. *Bewertung Amerikanischer Optionen mit Hilfe von regressionbasierten Monte-Carlo-Verfahren*. PhD thesis, University of Saarland, 2007. (Cited on page 5)

[35] J. Tsitsiklis and B. Van Roy. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703, 2001. (Cited on pages 1 and 5)

[36] A. Veraart and L. Veraart. Stochastic volatility and stochastic leverage. *Annals of Finance*, 8(2-3):205–233, 2012. (Cited on page 27)

[37] V. Volpe. *The Electricity price modelling and derivatives pricing in the Nord Pool market*. PhD thesis, Università della Svizzera italiana, 2009. (Cited on page 21)

[38] A. Wagner. Residual demand modeling and application to electricity pricing. Technical Report 213, Fraunhofer ITWM, Department for Financial Mathematics, 2012. (Cited on page 27)

A L_p convergence speed of empirical mean

Lemma A.1. *For every $p \geq 1$, there exists $C_p > 0$ such that for any i.i.d. sample $X_1, \ldots, X_M$ of $\mathbb{R}$-valued random variables such that $\mathbb{E}[X_1] = 0$ and $\mathbb{E}\left[|X_1|^{p \vee 2}\right] < \infty$, the following holds:*

$$\left\| \frac{1}{M} \sum_{m=1}^{M} X_m \right\|_{L_p} \leq \frac{C_p}{\sqrt{M}} \|X_1\|_{L_{p \vee 2}} \tag{A.1}$$

Proof. Using Marcinkiewicz-Zygmund's inequality, there exists $C_p > 0$ such that:

$$\mathbb{E}\left[\left| \sum_{m=1}^{M} X_m \right|^p \right] \leq C_p \mathbb{E}\left[\left(\sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \right]$$

Multiplying both sides by $\frac{1}{M^p}$:

$$\mathbb{E}\left[\left| \frac{1}{M} \sum_{m=1}^{M} X_m \right|^p \right] \leq \frac{C_p}{M^{\frac{p}{2}}} \mathbb{E}\left[\left(\frac{1}{M} \sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \right] \tag{A.2}$$

If $p \geq 2$, then $\frac{p}{2} \geq 1$ and, using Jensen's inequality:

$$\left(\frac{1}{M} \sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \leq \frac{1}{M} \sum_{m=1}^{M} \left(|X_m|^2 \right)^{\frac{p}{2}} = \frac{1}{M} \sum_{m=1}^{M} |X_m|^p$$

Taking expectations on both sides:

$$\mathbb{E}\left[\left(\frac{1}{M}\sum_{m=1}^{M}|X_m|^2\right)^{\frac{p}{2}}\right] \leq \mathbb{E}\left[|X_1|^p\right] \tag{A.3}$$

Now, if $p < 2$, then $\frac{p}{2} < 1$ and, using Jensen's inequality:

$$\mathbb{E}\left[\left(\frac{1}{M}\sum_{m=1}^{M}|X_m|^2\right)^{\frac{p}{2}}\right] \leq \mathbb{E}\left[\left(\frac{1}{M}\sum_{m=1}^{M}|X_m|^2\right)\right]^{\frac{p}{2}} = \mathbb{E}\left[|X_1|^2\right]^{\frac{p}{2}} \tag{A.4}$$

Then combine inequalities (A.2), (A.3) and (A.4) and take the power $\frac{1}{p}$ to obtain inequality (A.1). □

B Positivity of cointegrated geometric Brownian motions

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Consider the following d-dimensional process:

$$\begin{aligned} dS_t &= \Xi S_t dt + \text{diag}\left(S_t\right)\Sigma dW_t \\ S_0 &> 0 \end{aligned}$$

where W is a $\mathcal{F}$-adapted, d-dimensional Brownian motion, Ξ is a $d \times d$ cointegration matrix, and Σ is a $d \times d$ covariance matrix.

Proposition B.1. *$S > 0$ a.s. if and only if $\forall i \neq j$, $\Xi_{i,j} \geq 0$*

Proof. First, suppose that $\forall i, j = 1, \ldots, d$, $i \neq j$, $\Xi_{i,j} \geq 0$. Consider the following stopping time:

$$\tau = \inf\left\{t \geq 0\,;\, \exists j \in [1,d] \text{ } s.t. \, S_\tau^j = 0\right\}$$

i.e. τ is the first time when one component of S reaches 0. In particular, $S_t \geq 0$ a.s. $\forall t \in [0, \tau]$.

Now, suppose that $\tau < \infty$. There exists at least one component i such that $S_\tau^i = 0$. Recall the dynamics of S^i:

$$dS_t^i = \left(\sum_{j=1}^{d}\Xi_{i,j}S_t^j\right)dt + S_t^i\left(\sum_{j=1}^{d}\Sigma_{i,j}dW_t^j\right)$$

By Girsanov's theorem, there exists a probability measure $\mathbb{Q}^i$, equivalent to $\mathbb{P}$, such that

$$dS_t^i = \left(\sum_{1\leq j\leq d; j\neq i}\Xi_{i,j}S_t^j\right)dt + S_t^i\left(\sum_{j=1}^{d}\Sigma_{i,j}d\tilde{W}_t^j\right)$$

where $\tilde{W}$ is a d-dimensional $\mathbb{Q}^i$-Brownian motion. Then, using Proposition (2.3) from [31] (Chapter IX):

$$S_t^i = \mathcal{E}\left(\tilde{X}^i\right)_t\left\{S_0^i + \int_0^t \mathcal{E}\left(\tilde{X}^i\right)_s^{-1}\left(\sum_{1\leq j\leq d; j\neq i}\Xi_{i,j}S_s^j\right)ds\right\} \tag{B.1}$$

where $\tilde{X}_t^i := \sum_{j=1}^{d}\Sigma_{i,j}d\tilde{W}_t^j$, and $\mathcal{E}\left(\tilde{X}^i\right)$ denotes the exponential martingale of $\tilde{X}^i$. At time τ, it yields:

$$0 = S_\tau^i = \mathcal{E}\left(\tilde{X}^i\right)_\tau\left\{S_0^i + \int_0^\tau \mathcal{E}\left(\tilde{X}^i\right)_s^{-1}\left(\sum_{1\leq j\leq d; j\neq i}\Xi_{i,j}S_s^j\right)ds\right\} > 0$$

using the positivity of S_0^i and of the exponential martingale, as well as the non-negativity of $\Xi_{i,j}$, $i \neq j$, and of S before τ. This contradiction means that $S_\tau^i > 0$. As the same reasoning can be applied for every $i \in [1,d]$, this means that $\tau = \infty$, i.e. that $S > 0$ a.s. .

Next, suppose that $S > 0$ a.s. . Choose $i \in [1,d]$. Using equation (B.1) and the positivity of S, we obtain:

$$S_0^i + \sum_{1\leq j\leq d; j\neq i}\Xi_{i,j}\int_0^t \mathcal{E}\left(\tilde{X}^i\right)_s^{-1}S_s^j ds > 0 \text{ a.s.}$$

As $S_0^i > 0$ and the coefficients $\int_0^t \mathcal{E}\left(\tilde{X}^i\right)_s^{-1}S_s^j ds$ are a.s. positive with support $\mathbb{R}_+$, the only possibility for the above inequality to hold a.s. is that $\Xi_{i,j} \geq 0$ for all $i \neq j$. □

C Empirical confidence intervals

This Appendix describes how to obtain an empirical confidence interval for $v(0, x_0, i_0)$. Here we adapt arguments from [7] to our optimal switching problem.

We assume that the parameters T (time localisation) and h (discretization) are chosen such that the error between v and $\bar{v}_\Pi$ is negligible (the space localization being redundant in practice), and focus on the error between $\bar{v}_\Pi$ and $\hat{v}_\Pi$.

First, from equation (3.8), the dynamic programming principle for the process $\bar{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)$ reads:

$$\begin{aligned}
\bar{v}_\Pi\left(T, \bar{X}_T, i\right) &= g\left(T, \bar{X}_T, i\right) \\
\bar{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right) &= \sup_{j \in \mathbb{I}_q^n} \left\{ hf\left(t_n, \bar{X}_{t_n}, j\right) - k\left(t_n, i, j\right) + \mathbb{E}\left[\bar{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) | \mathcal{F}_{t_n}\right] \right\}, n = N-1, \ldots, 0
\end{aligned} \tag{C.1}$$

where $\mathbb{I}_q^n$ is the set of $\mathcal{F}_{t_n}$-measurable random variables taking values in $\mathbb{I}_q$. Suppose that the approximated conditional expectation $\hat{\mathbb{E}}\left[. | \mathcal{F}_{t_n}\right]$ is unbiased, i.e. that

$$\mathbb{E}\left[\hat{\mathbb{E}}\left[. | \mathcal{F}_{t_n}\right]\right] = \mathbb{E}\left[. | \mathcal{F}_{t_n}\right]$$

Then, using equation (3.15) and Jensen's inequality, the following holds:

$$\mathbb{E}\left[\hat{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)\right] \geq \sup_{j \in \mathbb{I}_q^n} \left\{ hf\left(t_n, \bar{X}_{t_n}, j\right) - k\left(t_n, i, j\right) + \mathbb{E}\left[\hat{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) | \mathcal{F}_{t_n}\right] \right\}, n = N-1, \ldots, 0 \tag{C.2}$$

Combining equations (C.1) and (C.2), an induction argument yields:

$$\mathbb{E}\left[\hat{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)\right] \geq \mathbb{E}\left[\bar{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)\right]$$

In particular, $\mathbb{E}\left[\hat{v}_\Pi\left(0, x_0, i\right)\right] \geq \bar{v}_\Pi\left(0, x_0, i\right)$. This reasoning means that $\hat{v}_\Pi\left(0, x_0, i\right)$ can be used approximatively as an asymptotic upper bound for $\bar{v}_\Pi\left(0, x_0, i\right)$.

For the lower bound, simply use the estimated optimal control $\hat{\alpha}$, which is a side-product of the computation of $\hat{v}_\Pi$, and compute equation (3.6) by replacing the supremum over every control α by this specific $\hat{\alpha}$. By definition of the supremum, this yields a lower bound for $\bar{v}_\Pi\left(0, x_0, i\right)$.