

Optimal Quantization for Finance: From Random Vectors to Stochastic Processes

Gilles Pagès

*Laboratoire de Probabilités et Modèles aléatoires, UMR 7599,
Université Paris 6, case 188, 4, pl. Jussieu, F-75252 Paris Cedex 5, France.
E-mail address: gpa@ccr.jussieu.fr*

Jacques Printems

*LAMA, Université Paris 12,
E-mail address: printems@univ-paris12.fr.*

Abstract

In this chapter, we present an overview of the recent developments of vector quantization and functional quantization and their applications as a numerical method in finance, with an emphasis on the quadratic case. Quantization is a way to approximate a random vector or a stochastic process, viewed as a Hilbert-valued random variable, using a nearest neighbor projection on a finite codebook. We make a review of cubature formulas to approximate expectation, an conditional expectation, including the introduction of a quantization-based Richardson–Romberg extrapolation method. The optimal quadratic quantization of the Brownian motion is presented in full detail. A special emphasis is made on the computational aspects and the numerical applications, in particular, the pricing of different kinds of options in various fields (swing options on gas and options in a Heston stochastic volatility model).

1. Introduction

Quantization is a way to discretize the path space of a random phenomenon: a random vector in finite dimension and a stochastic process in infinite dimension. Optimal vector quantization theory (finite dimensional) of random vectors finds its origin in the early 1950s in order to discretize some emitted signal (see GERSHO and GRAY [1992] or GRAF and LUSCHGY [2000]). It was further developed by specialists in signal processing and in information theory.

The infinite-dimensional case started to be extensively investigated in the early 2000s by several authors (see PAGÈS [2000], LUSCHGY and PAGÈS [2002, 2004, 2006], DEREICH and SCHEUTZOW [2003, 2006], WILBERTZ [2005], GRAF, LUSCHGY and PAGÈS [2007]).

Let us consider a Hilbertian setting. One considers a random vector X defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ taking its values in a separable Hilbert space $(H, (\cdot|\cdot)_H)$ (equipped with its natural Borel σ -algebra) and satisfying $\mathbb{E}|X|^2 < +\infty$. When H is an Euclidean space (\mathbb{R}^d) , one speaks about vector quantization. When H is an infinite-dimensional space like $L^2_T := L^2([0, T], dt)$ (endowed with the usual Hilbertian norm $|f|_{L^2_T} := (\int_0^T f^2(t)dt)^{\frac{1}{2}}$), one speaks of functional quantization (denoted by L^2_T from now on). A (bimeasurable) stochastic process $(X_t)_{t \in [0, T]}$ defined on $(\Omega, \mathcal{A}, \mathbb{P})$ satisfying $|X(\omega)|_{L^2_T} < +\infty$ $\mathbb{P}(d\omega)$ -a.s. can always be seen, once possibly modified on a \mathbb{P} -negligible set, as an L^2_T -valued random variable. Although we will focus on the Hilbertian framework, other choices are possible for H , in particular, some more general Banach settings like $L^p([0, T], dt)$ or $\mathcal{C}([0, T], \mathbb{R})$ spaces.

This chapter is organized as follows: in Section 2 and in its subsections we introduce quadratic quantization in a Hilbertian setting. In Section 3, we focus on optimal quantization, including some extensions to nonquadratic quantization. Section 4 is devoted to some quantized cubature formulae. Section 5.1 provides some classical background on the quantization rate in finite dimension. Section 6 deals with functional quantizations of Gaussian processes, like the Brownian motion, with a special emphasis on the numerical aspects. We present here what is, to our guess, the first large-scale numerical optimization of the quadratic quantization of the Brownian motion. We compare it to the optimal product quantization, formerly investigated in a study by PAGÈS and PRINTEMS [2005]. In Section 7, we propose a constructive approach to the functional quantization of scalar or multidimensional diffusions (in the Stratanovich sense). In Section 8, we show how to use functional quantization to price path-dependent options like Asian options (in a Heston stochastic volatility model). We conclude by some recent results showing how to derive universal (often optimal) functional quantization rate from time regularity of a process described in Section 9 and by a few examples in Section 10 about the specific methods that produce some lower bounds (this important subject as many others like the connections with small deviation theory is not treated in this numerically oriented overview). As concerns statistical applications of functional quantization, we refer to the studies by TARPEY and KINATEDER [2003], TARPEY, PETKOVA, and OGDEN [2003].




FIG. 1.1 A two-dimensional 10-quantizer $\Gamma = \{x_1, \dots, x_{10}\}$ and its Voronoi diagram.

NOTATIONS.

- $a_n \approx b_n$ means $a_n = O(b_n)$ and $b_n = O(a_n)$; $a_n \sim b_n$ means $a_n = b_n + o(a_n)$.
- If $X : (\Omega, \mathcal{A}, \mathbb{P}) \rightarrow (H, |\cdot|_H)$ (Hilbert space), then $\|X\|_2 = (\mathbb{E}|X|_H^2)^{\frac{1}{2}}$.
- $[x]$ denotes the integral part of the real x .

2. What is quadratic quantization?

Let $(H, (\cdot|\cdot)_H)$ denote a separable Hilbert space. Let $X \in L^2_H(\mathbb{P})$, that is, a random vector $X : (\Omega, \mathcal{A}, \mathbb{P}) \mapsto H$ (H is endowed with its Borel σ -algebra) such that $\mathbb{E}|X|_H^2 < +\infty$. An N -quantizer (or N -codebook) is defined as a subset

$$\Gamma := \{x_1, \dots, x_N\} \subset H$$

with $\text{card } \Gamma = N$. In numerical applications, Γ is also called *grid*. Then, one can quantize (or simply discretize) X by $q(X)$, where $q : H \mapsto \Gamma$ is a Borel function. It is straightforward that

$$\forall \omega \in \Omega, \quad |X(\omega) - q(X(\omega))|_H \geq d(X(\omega), \Gamma) = \min_{1 \leq i \leq N} |X(\omega) - x_i|_H$$

so that the best pointwise approximation of X is provided by considering for q a nearest neighbor projection on Γ , denoted by Proj_Γ . Such a projection is in one-to-one correspondence with the Voronoi partitions (or diagrams) of H induced by Γ , that is, the Borel partitions of H satisfying

$$C_i(\Gamma) \subset \left\{ \xi \in H : |\xi - x_i|_H = \min_{1 \leq j \leq N} |\xi - x_j|_H \right\} = \bar{C}_i(\Gamma), \quad i = 1, \dots, N,$$

where $\overline{C}_i(\Gamma)$ denotes the closure of $C_i(\Gamma)$ in H (this heavily uses the Hilbert structure). Then,

$$\text{Proj}_\Gamma(\xi) := \sum_{i=1}^N x_i \mathbf{1}_{C_i(\Gamma)}(\xi)$$

is a nearest neighbor projection on Γ . These projections only differ on the boundaries of the Voronoi cells $C_i(\Gamma)$, $i = 1, \dots, N$. All Voronoi partitions have the same boundary contained in the union of the median hyperplanes defined by the pairs (x_i, x_j) , $i \neq j$. Fig. 1.1 represents the Voronoi diagram defined by a (random) 10-tuple in \mathbb{R}^2 . Then, one defines a Voronoi N -quantization of X by setting for every $\omega \in \Omega$,

$$\hat{X}^\Gamma(\omega) := \text{Proj}_\Gamma(X(\omega)) = \sum_{i=1}^N x_i \mathbf{1}_{C_i(\Gamma)}(X(\omega)).$$

One clearly has, still for every $\omega \in \Omega$, that

$$|X(\omega) - \hat{X}^\Gamma(\omega)|_H = \text{dist}_H(X(\omega), \Gamma) = \min_{1 \leq i \leq N} |X(\omega) - x_i|_H.$$

The mean (quadratic) quantization error is then defined by

$$e(\Gamma, X, H) = \|X - \hat{X}^\Gamma\|_2 = \sqrt{\mathbb{E} \left(\min_{1 \leq i \leq N} |X - x_i|_H^2 \right)}. \tag{2.1}$$

The distribution of \hat{X}^Γ as a random vector is given by the N -tuple $(\mathbb{P}(X \in C_i(\Gamma)))_{1 \leq i \leq N}$ of the Voronoi cells. This distribution clearly depends on the choice of the Voronoi




FIG. 2.1 Two N -quantizers (and their Voronoi diagram) related to bi-variate normal distribution $\mathcal{N}(0; I_2)$ ($N = 500$); which one is the best?

partition as emphasized by the following elementary situation: if $H = \mathbb{R}$, the distribution of X is given by $\mathbb{P}_X = \frac{1}{3}(\delta_0 + \delta_{1/2} + \delta_1)$, $N = 2$ and $\Gamma = \{0, 1\}$ since $1/2 \in \partial C_0(\Gamma) \cap \partial C_1(\Gamma)$. However, if \mathbb{P}_X weights no hyperplane, the distribution of \hat{X}^Γ depends only on Γ .

As concerns terminology, vector quantization is concerned with the finite-dimensional case, when $\dim H < +\infty$, and is a rather old story, going back to the early 1950s when it was designed in the field of signal processing and then mainly developed in the community of information theory. The term *functional quantization*, probably introduced by LUSCHGY and PAGÈS [2002], PAGÈS [2000], deals with the infinite-dimensional case including the more general Banach-valued setting. The term *functional* comes from the fact that a typical infinite-dimensional Hilbert space is the function space $H = L^2_T$. Then, any (bimeasurable) process $X: ([0, T] \times \Omega, \text{Bor}([0, T]) \otimes \mathcal{A}) \rightarrow (\mathbb{R}, \text{Bor}(\mathbb{R}))$ can be seen as a random vector taking values in the set of Borel functions on $[0, T]$. Furthermore, $((t, \omega) \mapsto X_t(\omega)) \in L^2(dt \otimes d\mathbb{P})$ if and only if $(\omega \mapsto X \cdot(\omega)) \in L^2_H(\mathbb{P})$ since

$$\int_{[0, T] \times \Omega} X_t^2(\omega) dt \mathbb{P}(d\omega) = \int_{\Omega} \mathbb{P}(d\omega) \int_0^T X_t^2(\omega) dt = \mathbb{E} \|X \cdot\|_{L^2_T}^2.$$

3. Optimal (quadratic) quantization

At this stage, we are lead to wonder whether it is possible to design some optimally fitted grids to a given distribution \mathbb{P}_X , that is, which induce the lowest possible mean quantization error among all grids of size at most N (see **e.g.** Fig. 2.1). This amounts to the following optimization problem

$$e_N(X, H) := \inf_{\Gamma \subset H, \text{card}(\Gamma) \leq N} e(\Gamma, X, H). \tag{3.1}$$

It is convenient at this stage to make a correspondence between quantizers of size at most N and N -tuples of H^N : to any N -tuple $x := (x_1, \dots, x_N)$ corresponds a quantizer $\Gamma := \Gamma(x) = \{x_i, i = 1, \dots, N\}$ (of size at most N). One introduces the quadratic distortion, denoted by D_N^X , defined by H^N as a (symmetric) function by

$$\begin{aligned} D_N^X &: H^N \longrightarrow \mathbb{R}_+ \\ (x_1, \dots, x_N) &\longmapsto \mathbb{E} \left(\min_{1 \leq i \leq N} |X - x_i|_H^2 \right). \end{aligned}$$

Note that combining (2.1) and the definition of the distortion show that

$$D_N^X(x_1, \dots, x_N) = \mathbb{E} \left(\min_{1 \leq i \leq N} |X - x_i|_H^2 \right) = \mathbb{E} \left(d(X, \Gamma(x))^2 \right) = \|X - \hat{X}^{\Gamma(x)}\|_2^2$$

so that

$$e_N(X, H) = \inf_{(x_1, \dots, x_N) \in H^N} \sqrt{D_N^X(x_1, \dots, x_N)}.$$

The following proposition shows the existence of an optimal N -tuple $x^{(N,*)} \in H^N$ such that $e_N(X, H) = \sqrt{D_N^X(x^{(N,*)})}$. The corresponding optimal quantizer at level N is denoted by $\Gamma^{(N,*)} := \Gamma(x^{(N,*)})$. In finite dimensions, we refer to POLLARD [1982] and in infinite-dimensional settings to CUESTA-ALBERTOS and MATRÁN [1988] and PÄRNA [1990]; one may also refer to PAGÈS [1993], GRAF and LUSCHGY [2000], and LUSCHGY and PAGÈS [2002]. For recent developments on existence and pathwise regularity of optimal quantizer, see GRAF ET AL. [2007].

PROPOSITION 3.1.

- (a) The function D_N^X is lower semicontinuous for the product weak topology on H^N .
- (b) The function D_N^X reaches a minimum at a N -tuple $x^{(N,*)}$ (so that $\Gamma^{(N,*)}$ is an optimal quantizer at level N).
 - If $\text{card}(\text{supp}(\mathbb{P}_X)) \geq N$, the quantizer has full size N (i.e., $\text{card}(\Gamma^{(N,*)}) = N$) and $e_N(X, H) < e_{N-1}(X, H)$.
 - If $\text{card}(\text{supp}(\mathbb{P}_X)) \leq N$, $e_N(X, H) = 0$.

Furthermore, $\lim_N e_N(X, H) = 0$.

- (c) Any optimal (Voronoi) quantization at level N , $\hat{X}^{\Gamma^{(N,*)}}$ satisfies

$$\hat{X}^{\Gamma^{(N,*)}} = \mathbb{E}(X \mid \sigma(\hat{X}^{\Gamma^{(N,*)}})), \tag{3.2}$$

where $\sigma(\hat{X}^{\Gamma^{(N,*)}})$ denotes the σ -algebra generated by $\hat{X}^{\Gamma^{(N,*)}}$.

- (d) Any optimal (quadratic) quantization at level N is a best least square (i.e., $L^2(\mathbb{P})$) approximation of X among all H -valued random variables taking at most N values:

$$e_N(X, H) = \|X - \hat{X}^{\Gamma^{(N,*)}}\|_2 = \min\{\|X - Y\|_2, Y : (\Omega, \mathcal{A}) \rightarrow H, \text{card}(Y(\Omega)) \leq N\}.$$

PROOF. (sketch of):

- (a) The claim follows from the l.s.c. of $\xi \mapsto |\xi|_H$ for the weak topology and Fatou's lemma.
- (b) One proceeds by induction on N . If $N = 1$, the optimal one-quantizer is $x^{(N,*)} = \{\mathbb{E} X\}$ and $e_2(X, H) = \|X - \mathbb{E} X\|_2$.

□

Assume now that an optimal quantizer $x^{(N,*)} = (x_1^{(N,*)}, \dots, x_N^{(N,*)})$ does exist at level N .

- If $\text{card}(\text{supp}(\mathbb{P})) \leq N$, then the $N + 1$ -tuple $(x^{(N,*)}, x_N^{(N,*)})$ (among other possibilities) is also optimal at level $N + 1$ and $e_{N+1}(X, H) = e_N(X, H) = 0$.
- Otherwise, $\text{card}(\text{supp}(\mathbb{P})) \geq N + 1$, hence $x^{(N,*)}$ has pairwise distinct components and there exists $\xi_{N+1} \in \text{supp}(\mathbb{P}_X) \setminus \{x_i^{(N,*)}, i = 1, \dots, N\} \neq \emptyset$.

Then, with obvious notations,

$$D_{N+1}^X((x^{(N,*)}, \xi_{N+1})) < D_N^X(x^{(N,*)}).$$

Then, the set $F_{N+1} := \{x \in H^{N+1} \mid D_{N+1}^X(x) \leq D_{N+1}^X((x^{(N,*)}, \xi_{N+1}))\}$ is nonempty, weakly closed since D_{N+1}^X is l.s.c.. Furthermore, it is bounded in H^{N+1} . Otherwise, there would exist a sequence $x_{(m)} \in H^{N+1}$ such that $|x_{(m),i_m}|_H = \max_i |x_{(m),i}|_H \rightarrow +\infty$ as $m \rightarrow \infty$. Then, by Fatou's lemma, one checks that

$$\liminf_{m \rightarrow \infty} D_{N+1}^X(x_{(m)}) \geq D_N^X(x^{(N,*)}) > D_{N+1}^X((x^{(N,*)}, \xi_{N+1})).$$

Consequently, F_{N+1} is weakly compact and the minimum of D_{N+1}^X on F_{N+1} is clearly its minimum over the whole space H^{N+1} . In particular,

$$e_{N+1}(X, H) \leq D_{N+1}^X((x^{(N,*)}, \xi_{N+1})) < e_N(X, H).$$

If $\text{card}(\text{supp}(\mathbb{P})) = N + 1$, set $x^{(N+1,*)} = \text{supp}(\mathbb{P})$ (as sets) so that $t X = \hat{X}^{\Gamma(N+1,*)}$, which implies $e_{N+1}(X, H) = 0$.

To establish that $e_N(X, H)$ goes to 0, one considers an everywhere dense sequence $(z_k)_{k \geq 1}$ in the separable space H . Then, $d(\{z_1, \dots, z_N\}, X(\omega))$ goes to 0 as $N \rightarrow \infty$ for every $\omega \in \Omega$. Furthermore, $d(\{z_1, \dots, z_N\}, X(\omega))^2 \leq |X(\omega) - z_1|_H^2 \in L^1(\mathbb{P})$. One concludes by the Lebesgue dominated convergence theorem that $D_N^X(z_1, \dots, z_N)$ goes to 0 as $N \rightarrow \infty$.

(c) and (d) temporarily set $\hat{X}^* := \hat{X}^{\Gamma(N,*)}$ for convenience. Let $Y : (\Omega, \mathcal{A}) \rightarrow H$ be a random vector taking at most N values. Set $\Gamma := Y(\Omega)$. Since \hat{X}^Γ is a Voronoi quantization of X induced by Γ ,

$$|X - \hat{X}^\Gamma|_H = d(X, \Gamma) \leq |X - Y|_H$$

so that

$$\|X - \hat{X}^\Gamma\|_2 \leq \|X - Y\|_2.$$

On the other hand, the optimality of $\Gamma^{(N,*)}$ implies

$$\|X - \hat{X}^*\|_2 \leq \|X - \hat{X}^\Gamma\|_2.$$

Consequently,

$$\|X - \hat{X}^*\|_2 \leq \min \{ \|X - Y\|_2, Y : (\Omega, \mathcal{A}) \rightarrow H, \text{card}(Y(\Omega)) \leq N \}.$$

The inequality holds as an equality since \hat{X}^* takes at most N values. Furthermore, considering random vectors of the form $Y = g(\hat{X})$ (which take at most as many values as the size of $\Gamma^{(N,*)}$) shows, going back to the very definition of conditional expectation, that $\hat{X}^* = \mathbb{E}(X \mid \hat{X}^*)$ \mathbb{P} -a.s. \diamond

Item (c) introduces a very important notion in (quadratic) quantization.

DEFINITION 3.1. A quantizer $\Gamma \subset H$ is *stationary* (or *self-consistent*) if (there is a nearest-neighbor projection such that $\hat{X}^\Gamma = \text{Proj}_\Gamma(X)$ satisfying)

$$\hat{X}^\Gamma = \mathbb{E} \left(X \mid \hat{X}^\Gamma \right). \tag{3.3}$$

Note, in particular, that any stationary quantization satisfies $\mathbb{E}X = \mathbb{E}\hat{X}^\Gamma$.

As shown by Proposition 3.1(c) any quadratic optimal quantizer at level N is stationary. Usually, at least when $d \geq 2$, there are other stationary quantizers: indeed, the distortion function D_N^X is $|\cdot|_H$ -differentiable at N -quantizers $x \in H^N$ with pairwise distinct components and

$$\nabla D_N^X(x) = 2 \left(\int_{C_i(x)} (x_i - \xi) \mathbb{P}_X(d\xi) \right)_{1 \leq i \leq N} = 2 \left(\mathbb{E}(\hat{X}^{\Gamma(x)} - X) \mathbf{1}_{\{\hat{X}^{\Gamma(x)}=x_i\}} \right)_{1 \leq i \leq N}.$$

Hence, any critical point of D_N^X is a stationary quantizer.

REMARKS and COMMENTS.

- In fact (see GRAF and LUSCHGY [2000], theorem 4.2, pp 38), the Voronoi partitions of $\Gamma^{(N,*)}$ always have a \mathbb{P}_X -negligible boundary so that (3.3) holds for any Voronoi diagram induced by Γ .
- The problem of the uniqueness of optimal quantizer (viewed as a set) is not mentioned in the above proposition. In higher dimension, this essentially never occurs. In one dimension, uniqueness of the optimal N -quantizer was first established by FLEISCHER [1964] with strictly log-concave density function. This was successively extended by KIEFFER [1983] and TRUSHKIN [1982] and lead to the following criterion (for more general “loss” functions than the square function):

If the distribution of X is absolutely continuous with a log-concave density function, then, for every $N \geq 1$, there exists only one stationary quantizer of size N , which turns out to be the optimal quantizer at level N .

More recently, a more geometric approach to uniqueness based on the Mountain Pass lemma first developed by LAMBERTON and PAGÈS [1996] and then generalized by COHORT [1998] provided a slight extension of the above criterion (in terms of loss functions).

This log-concavity assumption is satisfied by many families of probability distributions like the uniform distribution on compact intervals, the normal distributions, and the gamma distributions. There are examples of distributions with a non-log-concave density function having a unique optimal quantizer for every $N \geq 1$ (see the Pareto distribution in FORT and PAGÈS [2004]). On the other hand, simple examples of scalar distributions having multiple optimal quantizers at a given level can be found in the study by GRAF and LUSCHGY [2000].

- A stationary quantizer can be suboptimal. This will be emphasized in Section 6 for the Brownian motion (but it is also true for finite-dimensional Gaussian random

vectors), where some families of suboptimal quantizers—the product quantizers designed from the Karhunen–Loève (K-L) basis—are stationary quantizers.

- For the uniform distribution over an interval $[a, b]$, there is a closed form for the optimal quantizer at level N given by $\Gamma^{(N,*)} = \{a + (2k - 1)\frac{b-a}{N}, k = 1, \dots, N\}$. This N -quantizer is optimal not only in the quadratic case but also for any L^r -quantization (see a definition further on). In general, there is no such closed form, either in one or in higher dimension. However, a study FORT and PAGÈS [2004] obtained some semiclosed forms for several families of (scalar) distributions including the exponential and the Pareto distributions: all the optimal quantizers can be expressed using a single underlying sequence $(a_k)_{k \geq 1}$ defined by an induction $a_{k+1} = F(a_k)$.
- In one dimension, as soon as the optimal quantizer at level N is unique (as a set or as an N -tuple with increasing components), it is generally possible to compute it as the solution of the stationarity Eq. (3.2) either by a zero search (Newton–Raphson gradient descent) or by a fixed point (like the specific Lloyd I procedure, see KIEFFER [1982]) procedure.
- In higher dimension, deterministic optimization methods become intractable, and one uses stochastic procedures to compute optimal quantizers. We decided to postpone the short overview on these aspects to Section 6, devoted to the optimal functional quantization of the Brownian motion, where the case of Gaussian vectors with (diagonal) covariance matrix is considered. All stochastic optimization approaches rely on some repeated nearest-neighbor searches: our procedures include some fast (exact) algorithms for that purpose (like K - d -tree, see FRIEDMAN, BENTLEY and FINKEL [1977]). So far, the most efficient methods are also based on the so-called splitting method, which increases progressively the quantization level N (this must be understood when looking for a systematic quantization of a distribution). This method is directly inspired by the induction developed in the proof of claim (b) of Proposition 3.1 since one designs the starting value of the optimization procedure at size $N + 1$ by “merging” the optimized N -quantizer obtained at level N with one further point of \mathbb{R}^d , usually randomly sampled with respect to an appropriate distribution (see PAGÈS and PRINTEMS [2003] for a discussion). For normal distributions $\mathcal{N}(0; I_d)$, alternative starting values living on a sphere with an appropriate radius seem to yield the same accuracy for a given size N without splitting (see PAGÈS and SAGNA [2007]).
- As concerns functional quantization, for example, $H = L^2_T$, there is a close connection between the regularity of optimal (or even stationary) quantizers and that of $t \mapsto X_t$ from $[0, T]$ into $L^2(\mathbb{P})$. Furthermore, as concerns optimal quantizers of Gaussian processes, one shows (see LUSCHGY and PAGÈS [2002]) that they belong to the reproducing space of their covariance operator, for example, to the Cameron–Martin space $H^1 = \{\int_0^\cdot \dot{h}_s ds, \dot{h} \in L^2_T\}$ when $X = W$. Other properties of optimal quantization of Gaussian processes are established by LUSCHGY and PAGÈS [2002].

Extensions to the $L^r(\mathbb{P})$ -quantization of random variables. In this chapter, we focus on the purely quadratic framework (L^2_T and $L^2(\mathbb{P})$ -norms), essentially because it is a natural (and somewhat easier) framework for the computation of optimized grids for

the Brownian motion and for some first applications (like the pricing of path-dependent options, see Section 8). But a more general and natural framework is to consider the functional quantization of random vectors taking values in a separable Banach space $(E, |\cdot|_E)$. Let $X : (\Omega, \mathcal{A}, \mathbb{P}) \rightarrow (E, |\cdot|_E)$ such that $\mathbb{E} |X|_E^r < +\infty$ for some $r \geq 1$ (the case $0 < r < 1$ can also be taken into consideration).

The N -level $(L^r(\mathbb{P}), |\cdot|_E)$ -quantization problem for $X \in L^r_E(\mathbb{P})$ reads

$$e_{N,r}(X, E) := \inf \{ \|X - \hat{X}^\Gamma\|_r, \Gamma \subset E, \text{card}(\Gamma) \leq N \}.$$

The main examples for $(E, |\cdot|_E)$ are the non-Euclidean norms on \mathbb{R}^d , the functional spaces $L^p_T(\mu) := L^p([0, T], \mu(dt))$, $1 \leq p \leq \infty$, equipped with its usual norm, $(E, |\cdot|_E) = (\mathcal{C}([0, T]), \|\cdot\|_{\text{sup}})$, etc. As concerns the existence of an optimal quantizer, it holds true for reflexive Banach spaces (see PÄRÄ [1990]) and $E = L^1_T$, but otherwise it may fail even when $N = 1$ (see GRAF, LUSCHGY and PAGÈS [2007]). In finite dimension, the Euclidean feature is not crucial (see GRAF and LUSCHGY [2000]). In the functional setting, many results originally obtained in a Hilbert setting have been extended to the Banach setting either for existence or for regularity results (see GRAF, LUSCHGY and PAGÈS [2007]) or for rates (see DEREICH [2005a], DEREICH and SCHEUTZOW [2006], LUSCHGY and PAGÈS [2004], LUSCHGY and PAGÈS [2007]).

4. Cubature formulae: conditional expectation and numerical integration

Let $F : H \rightarrow \mathbb{R}$ be a continuous functional (with respect to the norm $|\cdot|_H$) and let $\Gamma \subset H$ be an N -quantizer. It is natural to approximate $\mathbb{E}(F(X))$ by $\mathbb{E}(F(\hat{X}^\Gamma))$. This quantity $\mathbb{E}(F(\hat{X}^\Gamma))$ is simply the finite-weighted sum

$$\mathbb{E}(F(\hat{X}^\Gamma)) = \sum_{i=1}^N F(x_i) \mathbb{P}(\hat{X}^\Gamma = x_i). \tag{4.1}$$

Numerical computation of $\mathbb{E}(F(\hat{X}^\Gamma))$ is possible as soon as $F(\xi)$ can be computed at any $\xi \in H$ and the distribution $(\mathbb{P}(\hat{X} = x_i))_{1 \leq i \leq N}$ of \hat{X}^Γ is known. The induced quantization error $\|X - \hat{X}^\Gamma\|_2$ is used to control the error (see below). These quantities related to the quantizer Γ are also called *companion parameters*.

Likewise, one can consider a priori the $\sigma(\hat{X}^\Gamma)$ -measurable random variable $F(\hat{X}^\Gamma)$ as a good approximation of the conditional expectation $\mathbb{E}(F(X) | \hat{X}^\Gamma)$.

4.1. Lipschitz functionals

Assume that the functional F is Lipschitz continuous on H . Then,

$$\left| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) \right| \leq [F]_{\text{Lip}} \mathbb{E}(|X - \hat{X}^\Gamma| | \hat{X}^\Gamma)$$

so that, for every real exponent $r \geq 1$,

$$\|\mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma)\|_r \leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_r,$$

(where we applied conditional Jensen inequality to the convex function $u \mapsto u^r$). In particular, using $\mathbb{E} F(X) = \mathbb{E}(\mathbb{E}(F(X) | \hat{X}^\Gamma))$, one derives (with $r = 1$) that

$$\begin{aligned} \left| \mathbb{E} F(X) - \mathbb{E} F(\hat{X}^\Gamma) \right| &\leq \|\mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma)\|_1 \\ &\leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_1. \end{aligned}$$

Finally, using the monotony of the $L^r(\mathbb{P})$ -norms as a function of r yields

$$\left| \mathbb{E} F(X) - \mathbb{E} F(\hat{X}^\Gamma) \right| \leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_1 \leq [F]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_2. \quad (4.2)$$

In fact, considering the Lipschitz functional $F(\xi) := d(\xi, \Gamma)$ shows that

$$\|X - \hat{X}^\Gamma\|_1 = \sup_{[F]_{\text{Lip}} \leq 1} \left| \mathbb{E} F(X) - \mathbb{E} F(\hat{X}^\Gamma) \right|. \quad (4.3)$$

By the Lipschitz functionals making up a characterizing family for the weak convergence of probability measures on H , one derives that, for any sequence of N -quantizers Γ^N satisfying $\|X - \hat{X}^{\Gamma^N}\|_1 \rightarrow 0$ as $N \rightarrow \infty$,

$$\sum_{1 \leq i \leq N} \mathbb{P}(\hat{X}^{\Gamma^N} = x_i^N) \delta_{x_i^N} \xrightarrow{(H)} \mathbb{P}_X,$$

where $\xrightarrow{(H)}$ denotes the weak convergence of probability measures on $(H, |\cdot|_H)$.

4.2. Convex functionals

If $F : H \rightarrow \mathbb{R}$ is a convex functional and \hat{X} is a stationary quantization of X , a straightforward application of Jensen inequality yields

$$\mathbb{E} \left(F(X) | \hat{X} \right) \geq F(\hat{X})$$

so that $\mathbb{E} \left(F(\hat{X}) \right) \leq \mathbb{E} (F(X))$.

4.3. Differentiable functionals with Lipschitz differentials

Assume now that F is differentiable on H , with a Lipschitz continuous differential DF , and that the quantizer Γ is stationary (see Eq. (3.3)). A Taylor expansion yields

$$\left| F(X) - F(\hat{X}^\Gamma) - DF(\hat{X}^\Gamma).(X - \hat{X}^\Gamma) \right| \leq [DF]_{\text{Lip}} |X - \hat{X}^\Gamma|^2.$$

Taking conditional expectation, given \hat{X}^Γ , yields

$$\left| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) - \mathbb{E} \left(DF(\hat{X}^\Gamma).(X - \hat{X}^\Gamma) | \hat{X}^\Gamma \right) \right| \leq [DF]_{\text{Lip}} \mathbb{E}(|X - \hat{X}^\Gamma|^2 | \hat{X}^\Gamma).$$

Now, using that the random variable $DF(\hat{X}^\Gamma)$ is $\sigma(\hat{X}^\Gamma)$ -measurable, one has

$$\mathbb{E}(DF(\hat{X}^\Gamma).(X - \hat{X}^\Gamma)) = \mathbb{E}(DF(\hat{X}^\Gamma).\mathbb{E}(X - \hat{X}^\Gamma | \hat{X}^\Gamma)) = 0$$

so that

$$\left| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) \right| \leq [DF]_{\text{Lip}} \mathbb{E} \left(|X - \hat{X}^\Gamma|^2 | \hat{X}^\Gamma \right).$$

Then, for every real exponent $r \geq 1$,

$$\left\| \mathbb{E}(F(X) | \hat{X}^\Gamma) - F(\hat{X}^\Gamma) \right\|_r \leq [DF]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_{2r}^2. \tag{4.4}$$

In particular, when $r = 1$, one derives like in the former setting

$$\left| \mathbb{E}F(X) - \mathbb{E}F(\hat{X}^\Gamma) \right| \leq [DF]_{\text{Lip}} \|X - \hat{X}^\Gamma\|_2^2. \tag{4.5}$$

In fact, the above inequality holds provided F is C^1 with Lipschitz differential on every Voronoi cell $C_i(\Gamma)$. A characterization similar to (4.3) based on these functionals could be established.

Some variant of these cubature formulae can be found in PAGÈS and PRINTEMS [2003] or GRAF ET AL. [2006] for functions or functionals F having only some local Lipschitz regularity.

4.4. Quantized approximation of $\mathbb{E}(F(X) | Y)$

Let X and Y be two H -valued random vectors defined on the same probability space $(\Omega, \mathcal{A}, \mathbb{P})$ and $F : H \rightarrow \mathbb{R}$ be a Borel functional. The natural idea is to approximate $\mathbb{E}(F(X) | Y)$ by the quantized conditional expectation $\mathbb{E}(F(\hat{X}) | \hat{Y})$, where \hat{X} and \hat{Y} are quantizations of X and Y , respectively.

Let $\varphi_F : H \rightarrow \mathbb{R}$ be a (Borel) version of the conditional expectation, that is, satisfying

$$\mathbb{E}(F(X) | Y) = \varphi_F(Y).$$

Usually, no closed form is available for the function φ_F but some regularity property can be established, especially in a (Feller) Markovian framework. Thus, assume that both F and φ_F are Lipschitz continuous with Lipschitz coefficients $[F]_{\text{Lip}}$ and $[\varphi_F]_{\text{Lip}}$. Then,

$$\mathbb{E}(F(X) | Y) - \mathbb{E}(F(\hat{X}) | \hat{Y}) = \mathbb{E}(F(X) | Y) - \mathbb{E}(F(X) | \hat{Y}) + \mathbb{E}(F(X) - F(\hat{X}) | \hat{Y}).$$

Hence, assuming that \hat{Y} is $\sigma(Y)$ -measurable and that conditional expectation is an L^2 -contraction,

$$\begin{aligned} \left\| \mathbb{E}(F(X) | Y) - \mathbb{E}(F(X) | \hat{Y}) \right\|_2 &= \left\| \mathbb{E}(F(X) | Y) - \mathbb{E}(\mathbb{E}(F(\hat{X}) | \hat{Y}) | Y) \right\|_2 \\ &\leq \left\| \varphi_F(Y) - \mathbb{E}(F(X) | \hat{Y}) \right\|_2 \\ &= \left\| \varphi_F(Y) - \mathbb{E}(\varphi_F(Y) | \hat{Y}) \right\|_2 \\ &\leq \left\| \varphi_F(Y) - \varphi_F(\hat{Y}) \right\|_2. \end{aligned}$$

The last inequality follows from the definition of conditional expectation, given \hat{Y} as the best quadratic approximation among $\sigma(\hat{Y})$ -measurable random variables. On the other hand, still assuming that $\mathbb{E}(\cdot | \sigma(\hat{Y}))$ is an L^2 -contraction and this time that F is Lipschitz continuous yields

$$\|\mathbb{E}(F(X) - F(\hat{X}) | \hat{Y})\|_2 \leq \|F(X) - F(\hat{X})\|_2 \leq [F]_{\text{Lip}} \|X - \hat{X}\|_2.$$

Finally,

$$\|\mathbb{E}(F(X) | Y) - \mathbb{E}(F(\hat{X}) | \hat{Y})\|_2 \leq [F]_{\text{Lip}} \|X - \hat{X}\|_2 + [\varphi_F]_{\text{Lip}} \|Y - \hat{Y}\|_2.$$

In the nonquadratic case, the above inequality remains valid provided $[\varphi_F]_{\text{Lip}}$ is replaced by $2[\varphi_F]_{\text{Lip}}$.

5. Vector quantization

5.1. Vector quantization rate ($H = \mathbb{R}^d$)

The fact that $e_N(X, \mathbb{R}^d)$ is a nonincreasing sequence that goes to 0 as N goes to ∞ is a rather simple result established in Proposition 3.1. Its (sharp) rate of convergence to 0 is a much more challenging problem. An answer is provided by the so-called Zador theorem stated below.

This theorem was first stated and established for distributions with compact supports by Zador (see ZADOR [1963, 1982]). Then, a first extension to general probability distributions on \mathbb{R}^d is developed by BUCKLEW and WISE [1982]. The first mathematically rigorous proof can be found in a study by GRAF and LUSCHGY [2000], and relies on a random quantization argument (called upon in a step of the proof sometimes called Pierce lemma). We also provide a nonasymptotic error bound that can be seen as simple reformulation of this Pierce lemma. It turns out to be very useful for applications.

THEOREM 5.1 (a) SHARP RATE (see GRAF and LUSCHGY [2000]). *Let $r > 0$ and $X \in L^{r+\eta}(\mathbb{P})$ for some $\eta > 0$. Let $\mathbb{P}_X(d\xi) = \varphi(\xi) d\xi + \nu(d\xi)$ be the canonical decomposition of the distribution of X (ν and the Lebesgue measure are singular). Then (if $\varphi \not\equiv 0$),*

$$e_{N,r}(X, \mathbb{R}^d) \sim \tilde{J}_{r,d} \times \left(\int_{\mathbb{R}^d} \varphi^{\frac{d}{d+r}}(u) du \right)^{\frac{1}{d} + \frac{1}{r}} \times N^{-\frac{1}{d}} \quad \text{as } N \rightarrow +\infty, \quad (5.1)$$

where $\tilde{J}_{r,d} \in (0, \infty)$.

(b) NONASYMPTOTIC UPPER BOUND (see LUSCHGY and PAGÈS [2007]). *Let $d \geq 1$. There exists $C_{d,r,\eta} \in (0, \infty)$ such that, for every \mathbb{R}^d -valued random vector X ,*

$$\forall N \geq 1, \quad e_{N,r}(X, \mathbb{R}^d) \leq C_{d,r,\eta} \|X\|_{r+\eta} N^{-\frac{1}{d}}.$$

REMARKS.

- The real constant $\tilde{J}_{r,d}$ clearly corresponds to the case of the uniform distribution over the unit hypercube $[0, 1]^d$ for which the slightly more precise statement holds

$$\lim_N N^{\frac{1}{d}} e_{N,r}(X, \mathbb{R}^d) = \inf_N N^{\frac{1}{d}} e_{N,r}(X, \mathbb{R}^d) = \tilde{J}_{r,d}.$$

The proof is based on a self-similarity argument. The value of $\tilde{J}_{r,d}$ depends on the reference norm on \mathbb{R}^d . When $d = 1$, elementary computations show that $\tilde{J}_{r,1} = (r + 1)^{-\frac{1}{r}}/2$. When $d = 2$, with the canonical Euclidean norm, one shows (see NEWMAN [1982] for a proof (see also GRAF and LUSCHGY [2000]) that $\tilde{J}_{2,d} = \sqrt{\frac{5}{18\sqrt{3}}}$. Its exact value is unknown for $d \geq 3$ but, still for the canonical Euclidean norm, one has (see GRAF and LUSCHGY [2000]) using some random quantization arguments,

$$\tilde{J}_{2,d} \sim \sqrt{\frac{d}{2\pi e}} \approx \sqrt{\frac{d}{17,08}} \quad \text{as } d \rightarrow +\infty.$$

- When $\varphi \equiv 0$, the distribution of X is purely singular. The rate (5.1) still holds in the sense that $\lim_N N^{\frac{1}{d}} e_{r,N}(X, \mathbb{R}^d) = 0$. Consequently, this is not the right asymptotics. The quantization problem for singular measures (like uniform distribution on fractal compact sets) has been extensively investigated by several authors, leading to the definition of a quantization dimension in connection with the rate of convergence of the quantization error on these sets. For more details, we refer to GRAF and LUSCHGY [2000], GRAF and LUSCHGY [2005] and the references therein.
- A more naive way to quantize the uniform distribution on the unit hypercube is to proceed by product quantization, that is, by quantizing the marginals of the uniform distribution. If $N = m^d$, $m \geq 1$, one easily proves that the best quadratic product quantizer (for the canonical Euclidean norm on \mathbb{R}^d) is the “midpoint square grid”

$$\Gamma^{sq,N} = \left(\frac{2i_1 - 1}{2m}, \dots, \frac{2i_d - 1}{2m} \right)_{1 \leq i_1, \dots, i_d \leq m},$$

which induces a quadratic quantization error equal to

$$\sqrt{\frac{d}{12}} \times N^{-\frac{1}{d}}.$$

Consequently, product quantizers are still rate optimal in every dimension d . Moreover, note that the ratio of these two rates remains bounded as $d \uparrow \infty$.

- For a brief discussion and comparison with quasi-Monte Carlo methods, we refer to PAGÈS [2007] and the references therein. Let us simply recall that sequences (or sets) with low discrepancy are uniformly distributed sequence over the unit d -dimensional hypercube $[0, 1]^d$. When used instead of (pseudo-)random numbers to integrate a function f with bounded variations, the rate of convergence is

theoretically “almost dimension free”: it is the product of the variation of f by the discrepancy, which behaves like $O\left(\frac{\log(N)^d}{N}\right)$ (for sequences). However, such functions become less and less “standard” in higher dimension. When implemented with Lipschitz continuous functions, the quasi-Monte Carlo (QMC) method does face the curse of dimensionality with theoretical performances which seem to be worse than optimal quantizers of the uniform distribution over $[0, 1]^d$, namely $O\left(\frac{\log(N)}{N^{\frac{1}{d}}}\right)$ (still for sequences) owing to Proinov’s theorem (PROINOV [1988]).

- The nonasymptotic Zador theorem stated and established by LUSCHGY and PAGÈS [2007] is essentially a variant of the so-called Pierce lemma (see GRAF and LUSCHGY [2000]). Many developments and heuristics about the rate of convergence of some quantization-based algorithms for American option pricing, stochastic control, or nonlinear filtering (see PAGÈS, PHAM and PRINTEMS [2003]) can be significantly simplified or established rigorously by calling upon this result. This is emphasized by the example below devoted to swing options.

5.2. Examples of application of optimal vector quantization

5.2.1. Numerical integration (II): Richardson–Romberg extrapolation versus curse of dimensionality

Combining the above cubature formula (4.1) and the rate of convergence of the (optimal) quantization error, the theoretical critical dimension to use quantization-based cubature formulae seems to be $d = 4$ when compared to Monte Carlo simulation (at least for continuously differentiable functions). Several numerical tests have been carried out and reported by PAGÈS, PHAM and PRINTEMS [2003] and PAGÈS and PRINTEMS [2003] to evaluate more precisely the effect of the so-called curse of dimensionality. The benchmark was made of several European payoffs on a geometric index made of d independent assets in a Black–Scholes model: *vanilla put and put spread options and their smoothed versions*. No control variate was used. The absence of correlation is not a realistic assumption in finance but is clearly more challenging as a benchmark for numerical integration. Once the dimension d and the quantizer size N have been chosen, we compared the resulting integration error to a symmetric confidence interval with total length equal to two standard deviations of a Monte Carlo (MC) estimator based on N simulated data $\left(\frac{\sigma_{\text{payoff}}}{\sqrt{N}}\right)$. Furthermore, σ_{payoff} has been computed by a Monte Carlo simulation on 10^4 simulated data of the payoff.

The results turned out to be more favorable to quantization than predicted by theoretical bounds, mainly because we carried out our tests with rather small values of N , whereas curse of dimensionality is an asymptotic bound. Until the dimension 4, the larger N is, the more quantization outperforms MC simulation. When the dimension $d \geq 5$, quantization always outperforms MC (in the above sense) until a critical size $N_c(d)$, which decreases as d increases.

RICHARDSON–ROMBERG (R-R) EXTRAPOLATION. In this section, we provide a method to push ahead these critical sizes, at least for *smooth enough functionals*. Let $F : \mathbb{R}^d \rightarrow \mathbb{R}$ be a twice differentiable functional with Lipschitz–Hessian $D^2 F$. Let $(\hat{X}^{(N)})_{N \geq 1}$ be a

sequence of optimal quadratic quantizations. Then,

$$\mathbb{E}(F(X)) = \mathbb{E}(F(\hat{X}^{(N)})) + \frac{1}{2}\mathbb{E}\left(D^2F(\hat{X}^{(N)}).(X - \hat{X}^{(N)})^{\otimes 2}\right) + O\left(\mathbb{E}|X - \hat{X}|^3\right) \tag{5.2}$$

Under some assumptions that are satisfied by most usual distributions (including the normal one), it is proved by GRAF, LUSCHGY and PAGÈS [2006] as a special case of a general theorem about the asymptotic behavior of L^s of sequences of optimal L^r -quantizers for $s \in (r, r + d)$ that

$$\mathbb{E}|X - \hat{X}|^3 = O(N^{-\frac{3}{d}}) \quad \text{if } d \geq 2,$$

or $\mathbb{E}|X - \hat{X}|^3 = O(N^{-\frac{3-\varepsilon}{d}})$, $\varepsilon > 0$, if $d = 2$. Furthermore, if we make the conjecture that

$$\mathbb{E}(D^2F(\hat{X}^{(N)}).(X - \hat{X}^{(N)})^{\otimes 2}) = c_{F,X}N^{-\frac{2}{d}} + O(N^{-\frac{3}{d}}), \tag{5.3}$$

it becomes possible to implement an R-R extrapolation to compute $\mathbb{E}(F(X))$. Namely, one considers two sizes N_1 and N_2 (in practice, one often sets $N_1 = N/2$ and $N_2 = N$). Then, combining (5.2) with N_1 and N_2 ,

$$\mathbb{E}(F(X)) = \frac{N_2^{\frac{2}{d}}\mathbb{E}(F(\hat{X}^{(N_2)})) - N_1^{\frac{2}{d}}\mathbb{E}(F(\hat{X}^{(N_1)}))}{N_2^{\frac{2}{d}} - N_1^{\frac{2}{d}}} + O\left(\frac{1}{(N_1 \wedge N_2)^{\frac{1}{d}}(N_2^{\frac{2}{d}} - N_1^{\frac{2}{d}})}\right).$$

In Section 8.1, a similar procedure is tested in an infinite-dimensional setting: \mathbb{R}^d is replaced by the Hilbert space $H = L^2([0, T], dt)$ viewed as a state of paths for a stochastic process X (namely, the Brownian motion).

NUMERICAL ILLUSTRATION: In order to evaluate the effect of the R-R technique described above, numerical computations have been carried out in the case of the regularized versions of some put spread options on geometric indices in dimension $d = 4, 6, 8, 10$. By “regularized,” we mean that the payoff at maturity $T = 1$ has been replaced by its price function at time $T' < T$. Numerical integration was performed using the Gaussian optimal grids of size $N = 2^k$, $k = 2, \dots, 12$ (available at the Web site www.quantize.maths-fi.com).

We consider again one of the test functions implemented by PAGÈS and PRINTEMS [2003] (pp 152). These test functions were borrowed from classical option pricing in mathematical finance: one considers d independent traded assets S^1, \dots, S^d following a d -dimensional Black–Scholes dynamics (under its risk-neutral probability)

$$S_t^i = s_0^i \exp\left(\left(r - \frac{\sigma^2}{2}\right)t + \sigma\sqrt{t}Z^{i,t}\right), \quad i = 1, \dots, d,$$

where $Z^{i,t} = W_t^i/\sqrt{t}$ and $W = (W^1, \dots, W^d)$ is a d -dimensional standard Brownian motion. We also assume that $S_0^i = s_0 > 0$, $i = 1, \dots, d$ and that the d assets share the

same volatility $\sigma^i = \sigma > 0$. One considers the geometric index $I_t = (S_t^1 \dots S_t^d)^{\frac{1}{d}}$. One shows that $e^{-\frac{\sigma^2}{2}(\frac{1}{d}-1)t} I_t$ has itself a risk-neutral Black–Scholes dynamics. We want to test the regularized put spread option on this geometric index with strikes $K_1 < K_2$ (at time $T/2$). Let $\psi(s_0, K_1, K_2, r, \sigma, T)$ the premium at time 0 of a put spread on any of the assets S^i .

$$\begin{aligned} \psi(x, K_1, K_2, r, \sigma, T) &= \pi(x, K_2, r, \sigma, T) - \pi(x, K_1, r, \sigma, T) \\ \pi(x, K, r, \sigma, T) &= K e^{-rT} \operatorname{erf}(-d_2) - x \operatorname{erf}(-d_1), \\ d_1 &= \frac{\log(x/K) + (r + \frac{\sigma^2}{2d})T}{\sigma\sqrt{T/d}}, \quad d_2 = d_1 - \sigma\sqrt{T/d}. \end{aligned}$$

Using the martingale property of the discounted value of the premium of a European option yields that the premium $e^{-rT} \mathbb{E}((K_1 - I_T)_+ - (K_2 - I_T)_+)$ of the put spread option on I satisfies, on the one hand,

$$e^{-rT} \mathbb{E}((K_1 - I_T)_+ - (K_2 - I_T)_+) = \psi(s_0 e^{\frac{\sigma^2}{2}(\frac{1}{d}-1)T}, K_1, K_2, r, \sigma/\sqrt{d}, T)$$

and, on the other hand,

$$e^{-rT} \mathbb{E}((K_1 - I_T)_+ - (K_2 - I_T)_+) = \mathbb{E} g(Z),$$

where

$$g(Z) = e^{-rT/2} \psi(s_0 e^{\frac{\sigma^2}{2}(\frac{1}{d}-1)\frac{T}{2}} I_{\frac{T}{2}}, K_1, K_2, r, \sigma, T/2)$$

and $Z = (Z^1, \frac{T}{2}, \dots, Z^d, \frac{T}{2}) \stackrel{d}{=} \mathcal{N}(0; I_d)$. The numerical specifications of the function g are as follows:

$$s_0 = 100, \quad K_1 = 98, \quad K_2 = 102, \quad r = 5\%, \quad \sigma = 20\%, \quad T = 2.$$

The results are shown below (see Fig. 5.1) in a log-log scale for the dimensions $d = 4, 6, 8, 10$.

First, we recover the theoretical rates (namely, $-2/d$) of convergence for the error bounds. Indeed, some slopes $\beta(d)$ can be derived (using a regression) for the quantization errors and we found $\beta(4) = -0.48$, $\beta(6) = -0.33$, $\beta(8) = -0.25$, and $\beta(10) = -0.23$ for $d = 10$ (see Fig. 5.1). These rates plead for the implementation of R-R extrapolation. Also note that, as already reported by PAGÈS and PRINTEMS [2003], when $d \geq 5$, quantization still outperforms MC simulations (in the above sense) up to a critical number $N_c(d)$ of points ($N_c(6) \sim 5000$, $N_c(7) \sim 1000$, $N_c(8) \sim 500$, etc).

As concerns the R-R extrapolation method itself, note first that it always gives better results than crude quantization. As regards, the comparison with Monte Carlo simulation, no critical number of points $N_{\text{Romb}}(d)$ comes out beyond which MC simulation outperforms R-R extrapolation. This means that $N_{\text{Romb}}(d)$ is greater than the range of use of quantization-based cubature formulas in our benchmark, namely 5000.

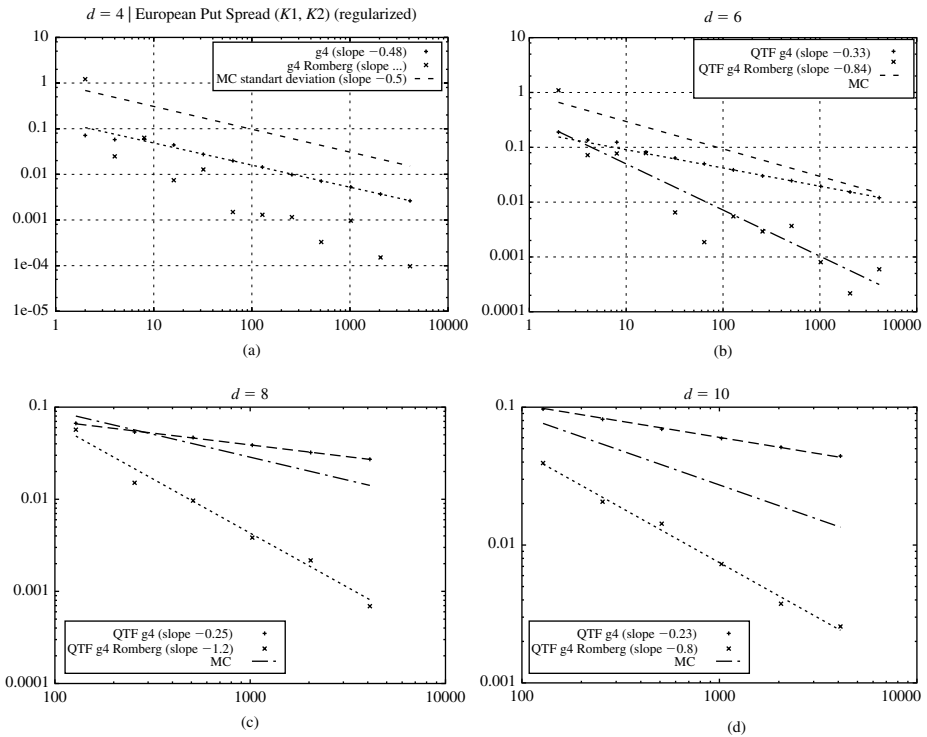


FIG. 5.1 Errors and standard deviations as functions of the number of points N in a log-log scale. The quantization error is shown by the symbol $+$ and the R - R extrapolation error by the symbol \times . The dashed line without crosses denotes the standard deviation of the Monte Carlo estimator. (a) $d = 4$, (b) $d = 6$, (c) $d = 8$, and (d) $d = 10$.

The R - R extrapolation techniques are commonly known to be unstable, and indeed, it has not been always possible to estimate satisfactorily its rate of convergence on our benchmark. But when a significant slope (in a log-log scale) can be estimated from the R - R errors (like for $d = 8$ and $d = 10$ in Fig. 5.1 (c), (d)), its absolute value is larger than $1/2$, and so, these extrapolations always outperform the MC method even for large values of N . As a by-product, our results plead in favor of the conjecture (5.3) and lead to think that R - R extrapolation is a powerful tool to accelerate numerical integration by optimal quantization, even in higher dimension.

5.3. An application to the pricing of swing options

Optimal-quantization-based algorithms have been already devised to solve several multi-dimensional nonlinear problems, from multiasset American style options (BALLY, PAGÈS and PRINTEMS [2001, 2003, 2005], BALLY and PAGÈS [2003a,b] to nonlinear filtering and portfolio management (see PAGÈS, PHAM and PRINTEMS [2003]). Here, we present a new application developed by Gaz de France (French gas company) to price swing options contracts. For a detailed version of this section, we refer to the original works

by BARDOU, BOUTHEMY and PAGÈS [2007a,b]. The holder of such a contract daily purchases a quantity of gas, say q_{t_k} at time $t_k, k = 0, \dots, n - 1$, at a price K_{t_k} , which can be deterministic or random (e.g., an oil-based index). These quantities are subject to two kinds of constraints, some local daily constraints $q_{\min} \leq q_{t_k} \leq q_{\max}$ and some global constraint about the total amount of purchased gas $Q_{\min} \leq q_0 + \dots + q_{t_{n-1}} \leq Q_{\max}$ (with $n q_{\min} \leq Q_{\min} \leq Q_{\max} \leq n q_{\max}$). The spot price S_{t_k} at time t_k of gas is usually not quoted on gas markets but is usually approximated by the day ahead price. The dynamics of the gas price itself is usually multifactorial, which makes it non-Markovian. However, it depends on a multidimensional underlying Markov structure process, which can be Gaussian or not. For the sake of simplicity, we assume here that (S_{t_k}) has a Markov dynamics and that the exercise prices K_{t_k} are deterministic (and there is no interest rate). Then, given $q_{\min}, q_{\max}, Q_{\min}, Q_{\max}$ and if $\bar{q}_{t_k} := q_0 + q_{t_1} + \dots + q_{t_{k-1}}$ denotes the purchased quantity prior to time t_k , the price of this contract at time t_k is given by

$$P(t_k, \bar{q}_{t_k}, S_{t_k}) := \inf \left\{ \mathbb{E} \left(\sum_{\ell=k}^{n-1} q_{t_\ell} (S_{t_\ell} - K_{t_\ell}) \mid S_{t_k} \right), q_{t_\ell} \in \mathcal{F}_{t_{\ell-1}}^S, \right. \\ \left. q_{\min} \leq q_{t_\ell} \leq q_{\max}, Q_{\min} \leq \bar{q}_{t_\ell} \leq Q_{\max} \right\}$$

where $\mathcal{F}_{t_\ell}^S = \sigma(S_0, S_{t_1}, \dots, S_{t_\ell})$ and \in stands for measurability with respect to a σ -field. This formula shows that this pricing problem, is a stochastic control problem, where the purchased quantity process appears as the control variable. This price satisfies the following dynamic programming principle

$$P(t_k, \bar{q}, S_{t_k}) = \sup_{q \in [q_{\min}, q_{\max}]} \left(q(S_{t_k} - K_{t_k}) + \mathbb{E} \left(P(t_{k+1}, \bar{q} + q, S_{t_{k+1}}) \mid S_{t_k} \right) \right).$$

It is shown by BARDOU, BOUTHEMY and PAGÈS [2007b] that the optimal control does exist under mild integrability assumptions but is not always *bang-bang* in general (owing to prediction errors). However, if $\frac{Q_{\max} - n q_{\min}}{q_{\max} - q_{\min}}$ and $\frac{Q_{\min} - n q_{\min}}{q_{\max} - q_{\min}}$ are integers, then the optimal control $q_{t_k}^*$ is always $\{q_{\min}, q_{\max}\}$ -valued. Then, one defines the quantized dynamic programming formula by setting

$$\hat{P}(t_k, Q, \hat{S}_{t_k}) := \max_{q=q_{\min}, q_{\max}} \left(q(S_{t_k} - K_{t_k}) + \mathbb{E} \left(\hat{P}(t_{k+1}, Q + q, \hat{S}_{t_{k+1}}) \mid \hat{S}_{t_k} \right) \right),$$

where \hat{S}_{t_k} is an N_k -quantization of S_{t_k} obtained by a nearest-neighbor projection on an optimal (quadratic) N_k -quantizer $\Gamma_k = \{x_1^k, \dots, x_{N_k}^k\}$. This quantization approach amounts to approximating the transitions $\mathcal{L}(S_{t_{k+1}} \mid S_{t_k})$ by the finitely valued transition $\mathcal{L}(\hat{S}_{t_{k+1}} \mid \hat{S}_{t_k})$. All these grids and transition weights make up a so-called quantization tree. In a Gaussian framework, some grids can be obtained from precomputed normalized optimal grids (available at the Web site www.quantize.maths-fi.com). Otherwise, this quantization optimization step can be performed by some stochastic optimization procedures (like randomized Lloyd's I and *CLVQ* procedures, see Section 6.3 for an example in a Gaussian framework). The next step is the computation of the quantized

transitions $\hat{\pi}_{ij}^k := \mathbb{P}(\hat{S}_{t_{k+1}} = x_j^{k+1} \mid \hat{S}_{t_k} = x_i^k)$ by a Monte Carlo simulation. Both steps are mainly based on repeated nearest-neighbor searches. They can be carried out off-line since they do not depend on the payoff characteristics (the K_{t_k} s). However, using some fast nearest-neighbor procedures, typically the K - d -tree algorithm introduced by FRIEDMAN, BENTLEY and FINKEL [1977] or some improved versions like (Principal Axis Tree (see McNAMES [2001])) drastically reduces the complexity of this phase (hence its duration) in higher dimension. Now with modern computing devices it becomes possible in many applications to include both phases in the online computations, which makes it as flexible as Monte Carlo-based methods like regression methods.

Assume that $S_0 = s_0 \in (0, \infty)$ and that $\hat{S}_0 = s_0$. It is proved BARDOU, BOUTHEMY and PAGÈS [2007b] that

$$|\hat{P}(0, 0, s_0) - P(0, 0, s_0)| \leq C \sum_{k=0}^{n-1} \|S_{t_k} - \hat{S}_{t_k}\|_2.$$

If all the quantizations \hat{S}_{t_k} of S_{t_k} are optimal (with size $N_k = \bar{N}$), this provides a $O(\frac{n}{\bar{N}^d})$ rate of convergence. In fact, numerical evidences show that the observed rate is usually

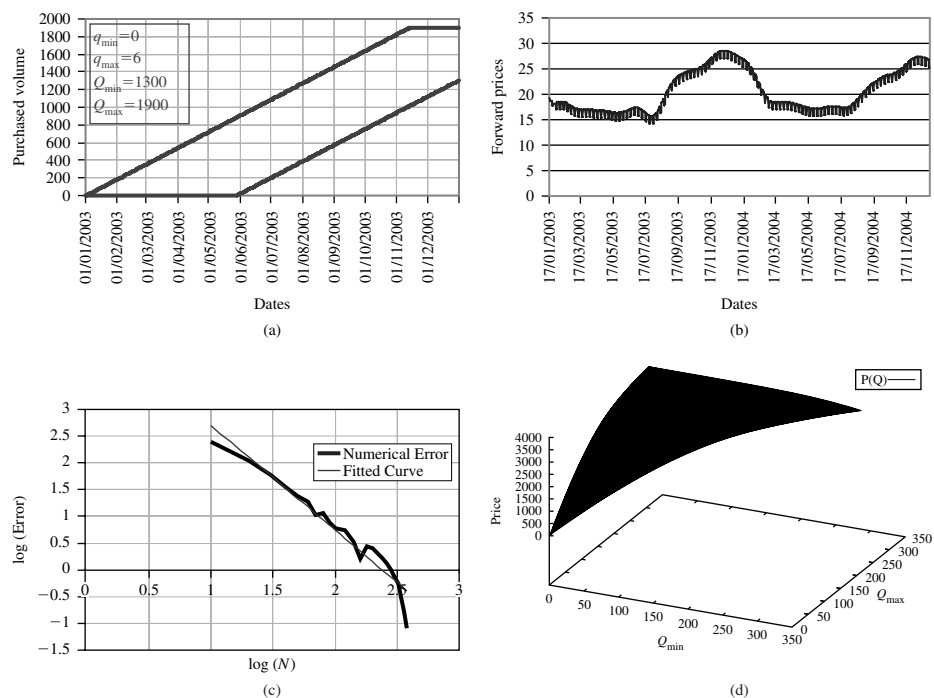


FIG. 5.2 The parameters are those given in the numerical illustration. (a) Constraint set, (b) daily forward curve, (c) numerical convergence as a function of the optimal grid size (log-log scale) and (d) the graph of the price as a function of the global constraints.

$O(\frac{n}{N^{\frac{1}{d}}})$, that is, somewhat similar to the rates obtained in the cubature formula for differentiable functions. The choice of the N_k s can be refined like for American options following the lines of BALLY and PAGÈS [2003a].

A comparison carried out by BARDOU, BOUTHEMY and PAGÈS [2007b] suggests that the quantization approach (including the transition computation) is significantly faster than the least squares regression (LSR) methods à la Longstaff–Schwartz.

NUMERICAL ILLUSTRATION: We consider the one-factor Toy model given by

$$S_t = F_{0,t} \exp\left(\sigma \int_0^t e^{-\alpha(t-s)} dW_s - \frac{1}{2} \frac{\sigma^2}{2\alpha} (1 - e^{-2\alpha t})\right),$$

where $\sigma = 70\%$, $\alpha = 4$, and $t_k = k/n$. The future prices are real data (January 17, 2003) corresponding to the first part of the curve in Fig. 5.2(b). The contract parameters are $q_{\min} = 0$, $q_{\max} = 6$, $Q_{\min} = 1300$, $Q_{\max} = 1900$, $K_{t_k} = K$, and $n = 365$ (1 year).

Note the slope in Fig. 5.2(d) is $1.96 \approx 2$. The main asset of quantization is to directly approximate the underlying Markov dynamics (in particular, when dealing with multifactor models): it is a model-driven method, which has its ability to “capture” automatically the correlation structure of the asset, which becomes quickly impossible with multinomial trees as the number of factors increases.

6. Optimal quadratic functional quantization of Gaussian processes

Optimal quadratic functional quantization of Gaussian processes is closely related to their so-called K-L expansion, which can be seen in some sense as some infinite-dimensional principal component analysis of a (Gaussian) process. Before stating a general result for Gaussian processes, we start by the standard Brownian motion: it is the most important example in view of (numerical) applications and for this process, everything can be made explicit.

6.1. Brownian motion

One considers the Hilbert space $H = L^2_T := L^2([0, T], dt)$, $(f|g)_2 = \int_0^T f(t)g(t)dt$, $\|f\|_{L^2_T} = \sqrt{(f|f)_2}$. The covariance operator C_W of the Brownian motion $W = (W_t)_{t \in [0, T]}$ is defined on L^2_T by

$$C_W(f) := \mathbb{E}((f, W)_2 W) = \left(t \mapsto \int_0^T (s \wedge t) f(s) ds\right).$$

It is a symmetric positive trace class operator, which can be diagonalized in the so-called K-L orthonormal basis $(e_n^W)_{n \geq 1}$ of L^2_T , with eigenvalues $(\lambda_n)_{n \geq 1}$, given by

$$e_n^W(t) = \sqrt{\frac{2}{T}} \sin\left(\pi\left(n - \frac{1}{2}\right)\frac{t}{T}\right), \quad \lambda_n = \left(\frac{T}{\pi\left(n - \frac{1}{2}\right)}\right)^2, \quad n \geq 1.$$

This classical result can be established as a simple exercise by solving the functional equation $C_W(f) = \lambda f$. In particular, one can expand W itself on this basis so that

$$W \stackrel{L_T^2}{=} \sum_{n \geq 1} (W|e_n^W)_2 e_n^W.$$

Now, the orthonormality of the (K-L) basis implies, using Fubini's theorem,

$$\mathbb{E}((W|e_k^W)_2 (W|e_\ell^W)_2) = (e_k^W | C_W(e_\ell^W))_2 = \lambda_\ell \delta_{k\ell},$$

where $\delta_{k\ell}$ denotes the Kronecker symbol. Hence, the Gaussian sequence $((W|e_n^W)_2)_{n \geq 1}$ is pairwise noncorrelated, which implies that these random variables are independent. The above identity also implies that $\text{Var}((W|e_n^W)_2) = \lambda_n$. Finally, this shows that

$$W \stackrel{L_T^2}{=} \sum_{n \geq 1} \sqrt{\lambda_n} \xi_n e_n^W, \tag{6.1}$$

where $\xi_n := (W|e_n^W)_2 / \sqrt{\lambda_n}$, $n \geq 1$, is an i.i.d. sequence of $\mathcal{N}(0; 1)$ -distributed random variables. Furthermore, this K-L expansion converges in a much stronger sense since $\sup_{t \in [0, T]} |W_t - \sum_{k=1}^n \sqrt{\lambda_k} \xi_k e_k^W(t)| \rightarrow 0$ \mathbb{P} -a.s. and

$$\| \sup_{[0, T]} |W_t - \sum_{1 \leq k \leq n} \sqrt{\lambda_k} \xi_k e_k^W(t)| \|_2 = O(\sqrt{\log n/n})$$

(see LUSCHGY and PAGÈS [2007]). Similar results (with various rates) hold true for a wide class of Gaussian processes expanded on “admissible” basis (see LUSCHGY and PAGÈS [2007]).

THEOREM 6.1 (LUSCHGY and PAGÈS [2002], LUSCHGY and PAGÈS [2004], and LUSCHGY, PAGÈS and WILBERTZ [2007]). *Let Γ^N , $N \geq 1$, be a sequence of optimal N -quantizers for W .*

(a) *For every $N \geq 1$, $\text{span}(\Gamma^N) = \text{span}\{e_1^W, \dots, e_{d(N)}^W\}$ with $d(N) \gtrsim \frac{1}{2} \log N$.*

(b) *$e_N(W, L_T^2) = \|W - \hat{W}^{\Gamma^N}\|_2 \sim \frac{T\sqrt{2}}{\pi} \frac{1}{\sqrt{\log N}}$ as $N \rightarrow \infty$.*

REMARKS.

- The fact, confirmed by numerical experiments (see Section 6.3, Fig. 6.4), that $d(N) \sim \log N$ holds as a conjecture.

- Denoting Π_d the orthogonal projection on $\text{span}\{e_1^W, \dots, e_d^W\}$, one derives from (a) that $\hat{W}^{\Gamma^N} = \widehat{\Pi_{d(N)}(W)}^{\Gamma^N}$ (optimal quantization at level N) and

$$\begin{aligned} \|W - \hat{W}^{\Gamma^N}\|_2^2 &= \|\Pi_{d(N)}(W) - \widehat{\Pi_{d(N)}(W)}^{\Gamma^N}\|_2^2 + \|W - \Pi_{d(N)}(W)\|_2^2 \\ &= e_N \left(Z_{d(N)}, \mathbb{R}^{d(N)} \right)^2 + \sum_{n \geq d(N)+1} \lambda_n, \end{aligned}$$

where $Z_{d(N)} \stackrel{\mathcal{L}}{\sim} \Pi_{d(N)}(W) \sim \bigotimes_{k=1}^{d(N)} \mathcal{N}(0; \lambda_k)$.

6.2. Centered Gaussian processes

Theorem 6.1 devoted to the standard Brownian motion is a particular case of a more general theorem, which holds for a wide class of Gaussian processes.

THEOREM 6.2 (LUSCHGY and PAGÈS [2002], LUSCHGY and PAGÈS [2004]). *Let $X = (X_t)_{t \in [0, T]}$ be a Gaussian process with K - L eigensystem $(\lambda_n^X, e_n^X)_{n \geq 1}$ (with $\lambda_1 \geq \lambda_2 \geq \dots$ is nonincreasing). Let $\Gamma^N, N \geq 1$, be a sequence of quadratic optimal N -quantizers for X . Assume*

$$\lambda_n^X \sim \frac{\kappa}{n^b} \quad \text{as } n \rightarrow \infty \quad (b > 1).$$

- (a) $\text{span}(\Gamma^N) = \text{span}\{e_1^X, \dots, e_{d^X(N)}^X\}$ and $d^X(N) \gtrsim \frac{1}{b^{1/(b-1)}} \frac{2}{b} \log N$.
- (b) $e_N(X, L_T^2) = \|X - \hat{X}^{\Gamma^N}\|_2 \sim \sqrt{\kappa} \sqrt{b^b(b-1)^{-1}} (2 \log N)^{-\frac{b-1}{2}}$.

REMARKS.

- The above result admits an extension to the case $\lambda_n^X \sim \varphi(n)$ as $n \rightarrow \infty$ with φ regularly varying, index $-b \leq -1$ (see LUSCHGY and PAGÈS [2004]). In LUSCHGY and PAGÈS [2002], upper or lower bounds are also established when

$$(\lambda_n^X \leq \varphi(n), \quad n \geq 1) \quad \text{or} \quad (\lambda_n^X \geq \varphi(n), \quad n \geq 1).$$

- The sharp asymptotics $d^X(N) \sim \frac{2}{b} \log N$ holds as a conjecture.

Applications to classical (centered) Gaussian processes.

- Brownian bridge: $X_t := W_t - \frac{t}{T} W_T, t \in [0, T]$ and $e_n^X(t) = \sqrt{2/T} \sin(\pi n \frac{t}{T}), \lambda_n = (\frac{T}{\pi n})^2$ so that $e_N(X, L_T^2) \sim T \frac{\sqrt{2}}{\pi} (\log N)^{-\frac{1}{2}}$.
- Fractional Brownian motion with Hurst constant $H \in (0, 1)$

$$e_N(W^H, L_T^2) \sim T^{H+\frac{1}{2}} c(H) (\log N)^{-H},$$

where $c(H) = \left(\frac{\Gamma(2H)\sin(\pi H)(1+2H)}{\pi}\right)^{\frac{1}{2}}\left(\frac{1+2H}{2\pi}\right)^H$ and $\Gamma(t)$ denotes the Gamma function at $t > 0$.

- Some further explicit sharp rates can be derived from the above theorem for other classes of Gaussian stochastic processes (see LUSCHGY and PAGÈS [2004]) like the fractional Ornstein–Uhlenbeck processes, the Gaussian diffusions, a wide class Gaussian stationary processes (the quantization rate is derived from the high-frequency asymptotics of its spectral density, assumed to be square integrable on the real line) for the m -folded integrated Brownian motion, the fractional Brownian sheet, etc.
- Of course, some upper bounds can be derived for some even wider classes of processes, based on the first remark (see LUSCHGY and PAGÈS [2002]).

Extensions to $r, p \neq 2$ When the processes have some self-similarity properties, it is possible to obtain some sharp rates in the nonpurely quadratic case: this has been done for fractional Brownian motion DEREICH and SCHEUTZOW [2006] using some quite different techniques in which self-similarity properties play crucial role. It leads to the following sharp rates, for $p \in [1, +\infty]$ and $r \in (0, \infty)$

$$e_{N,r}(W^H, L_T^p) \sim T^{H+\frac{1}{2}}c(r, H)(\log N)^{-H}, \quad c(r, H) \in (0, +\infty).$$

6.3. Numerical optimization of quadratic functional quantization

Thanks to the scaling property of Brownian motion, one may focus on the normalized case $T = 1$. The numerical approach to optimal quantization of the Brownian motion is essentially based on Theorem 6.1 and the remark that follows: indeed, these results show that quadratic optimal functional quantization of a centered Gaussian process reduces to a finite-dimensional optimal quantization problem for a Gaussian distribution with a diagonal covariance structure. Namely, the optimization problem at level N reads

$$(\mathcal{O}_N) \equiv \left\{ \begin{array}{l} e_N(W, L_T^2)^2 := e_N(Z_{\mathbf{d}(N)}, \mathbb{R}^{\mathbf{d}(N)})^2 + \sum_{k \geq \mathbf{d}(N)+1} \lambda_k \\ \text{where } Z_{\mathbf{d}(N)} \stackrel{\mathcal{L}}{\sim} \bigotimes_{k=1}^{\mathbf{d}(N)} \mathcal{N}(0, \lambda_k). \end{array} \right.$$

Moreover, if $\beta^N := \{\beta_1^N, \dots, \beta_N^N\}$ denotes an optimal N -quantizer of $Z_{\mathbf{d}(N)}$, then the optimal N -quantizer Γ^N of W reads $\Gamma^N = \{x_1^N, \dots, x_N^N\}$ with

$$x_i^N(t) = \sum_{1 \leq \ell \leq \mathbf{d}(N)} (\beta_i^N)^\ell e_\ell^W(t), \quad i = 1, \dots, N. \tag{6.2}$$

The good news is that (\mathcal{O}_N) is in fact a finite-dimensional quantization optimization problem for each $N \geq 1$. The bad news is that the problem is somewhat ill-conditioned since the decrease of the eigenvalues of W is very steep for small values of n : $\lambda_1 =$

$0.40528 \dots, \lambda_2 = 0.04503 \dots \approx \lambda_1/10$. This is probably one reason for which former attempts to produce good quantizations of the Brownian motion first focused on other kinds of quantizers like scalar product quantizers (see PAGÈS and PRINTEMS [2005] and Section 6.4) or d -dimensional block product quantizations (see WILBERTZ [2005], LUSCHGY, PAGÈS and WILBERTZ [2007]).

Optimization of the (quadratic) quantization of \mathbb{R}^d -valued random vectors has been extensively investigated since the early 1950s, first in one-dimension, then in higher dimension when the cost of numerical Monte Carlo simulation was drastically cut down (see GERSHO and GRAY [1992]). Recent application of optimal vector quantization to numerics turned out to be much more demanding in terms of accuracy. In that direction, one may cite PAGÈS and PRINTEMS [2003], MRAD and BEN HAMIDA [2006] (mainly focused on numerical optimization of the quadratic quantization of normal distributions). To apply these methods, it is more convenient to rewrite our optimization problem with respect to the standard d -dimensional distribution $\mathcal{N}(0; I_d)$ by simply considering the Euclidean norm derived from the covariance matrix $\text{Diag}(\lambda_1, \dots, \lambda_{d(N)})$, that is,

$$(\mathcal{O}_N) \Leftrightarrow \begin{cases} N\text{-optimal quantization of } \bigotimes_{k=1}^{d(N)} \mathcal{N}(0, 1) \\ \text{for the covariance norm } |(z_1, \dots, z_{d(N)})|^2 = \sum_{k=1}^{d(N)} \lambda_k z_k^2. \end{cases}$$

The main point is, of course, that the dimension $d(N)$ is unknown. However (see Fig. 6.4), one clearly verifies on small values of N that the conjecture $d(N) \sim \log N$ is most likely true. Then, for higher values of N , one relies on it to shift from one dimension to another following the rule $d(N) = d, N \in \{e^d, \dots, e^{d+1} - 1\}$.

6.3.1. A toolbox for quantization optimization: a short overview

Here is a short overview of stochastic optimization methods to compute optimal or at least locally optimal quantizers in finite dimension. For more detail, we refer to PAGÈS and PRINTEMS [2003] and the references therein. Let $Z \stackrel{\mathcal{L}}{\sim} \mathcal{N}(0; I_d)$.

Competitive learning vector quantization (CLVQ). This procedure is a recursive stochastic approximation gradient descent based on the integral representation of the gradient $\nabla D_N^Z(x), x \in H^n$ (temporarily coming back to N -tuple notation), of the distortion as the expectation of a local gradient, that is,

$$\forall x^N \in H^N, \quad \nabla D_N^Z(x^N) = \mathbb{E}(\nabla D_N^Z(x^N, \zeta)), \quad \zeta_k \text{ i.i.d., } \zeta_1 \stackrel{\mathcal{L}}{\sim} \mathcal{N}(0, I_d)$$

so that, starting from $x^N(0) \in (\mathbb{R}^d)^N$, one sets

$$\forall k \geq 0, \quad x^N(k+1) = x^N(k) - \frac{c}{k+1} \nabla D_N^Z(x^N(k), \zeta_{k+1}),$$

where $c \in (0, 1]$ is a real constant to be tuned. As set, this looks quite formal but the operating CLVQ procedure consists of two, phases at each iteration.

- (i) Competitive Phase: Search of the nearest-neighbor $x^{N(k)}_{i^*(k+1)}$ of ζ_{k+1} among the components of $x^N(k)_i, i = 1, \dots, N$ (using a “winning convention” in case of conflict on the boundary of the Voronoi cells).
- (ii) Cooperative Phase: One moves the winning component toward ζ_{k+1} using a dilatation, that is, $x^{N(k+1)}_{i^*(k+1)} = \text{dilatation}_{\zeta_{k+1}, 1 - \frac{c}{k+1}}(x^N(k)_{i^*(k+1)})$.

This procedure is useful for small or medium values of N . For an extensive study of this procedure, which turns out to be singular in the world of recursive stochastic approximation algorithms, we refer to PAGÈS [1998]. For general background on stochastic approximation, we refer to BENVENISTE, MÉTIVIER and PRIOURET [1990], KUSHNER and YIN [2003].

The randomized “Lloyd I procedure.” This is the randomization of the stationarity-based fixed-point procedure since any optimal quantizer satisfies (3.3):

$$\hat{Z}^{x^{N(k+1)}} = \mathbb{E}(Z \mid \hat{Z}^{x^{N(k)}}), \quad x^N(0) \subset \mathbb{R}^d.$$

At every iteration, the conditional expectation $\mathbb{E}(Z \mid \hat{Z}^{x^{N(k)}})$ is computed using a Monte Carlo simulation. For more details about practical aspects of Lloyd I procedure, we refer to PAGÈS and PRINTEMS [2003]. In MRAD and BEN HAMIDA [2006], an approach based on genetic evolutionary algorithms is developed.

For both procedures, one may substitute a sequence of quasi-random numbers to the usual pseudorandom sequence. This often speeds up the rate of convergence of the method, although this can only be proved (see BENVENISTE, MÉTIVIER and PRIOURET [1990]) for a very specific class of stochastic algorithms (to which *CLVQ* does not belong).

The most important step to preserve the accuracy of the quantization as N (and $d(N)$) increase is to use the so-called splitting method, which finds its origin in the proof of the existence of an optimal N -quantizer: once the optimization of a quantization grid of size N is achieved, one specifies the starting grid for the size $N + 1$ or more generally $N + \nu, \nu \geq 1$, by merging the optimized grid of size N resulting from the former procedure with ν points sampled independently from the normal distribution with probability density proportional to $\varphi^{\frac{d}{d+2}}$, where φ denotes the p.d.f. of $\mathcal{N}(0; I_d)$. This rather unexpected choice is motivated by the fact that this distribution provides the lowest in average random quantization error (see COHORT [1998]).

As a result, to be downloaded on the Web site PAGÈS and PRINTEMS [2005]: www.quantize.maths-fi.com

- Optimized stationary codebooks for W : in practice, the N -quantizers β^N of the distribution $\otimes_{k=1}^{d(N)} \mathcal{N}(0; \lambda_k), N = 1$ up to 10,000 ($d(N)$ runs from 1 to 9).
- Companion parameters:
 - distribution of $\hat{W}^{\Gamma^N}: \mathbb{P}(\hat{W}^{\Gamma^N} = x_i^N) = \mathbb{P}(\hat{Z}_{d(N)}^{\beta^N} = \beta_i^N) (\leftarrow \text{in } \mathbb{R}^{d(N)})$.
 - The quadratic quantization error: $\|W - \hat{W}^{\Gamma^N}\|_2$.

See Figs. 6.1, 6.2 and 6.3 for some examples of optimal quantizers Γ^N (and their counter parts β^N in $\mathbb{R}^{d(N)}$).




FIG. 6.1 Optimized functional quantization of the Brownian motion W for $N = 10, 15$ ($d(N) = 2$). Top: β^N depicted in \mathbb{R}^2 . Bottom: the optimized N -quantizer Γ^N .




FIG. 6.2 Optimized functional quantization of the Brownian motion W . The N -quantizers Γ^N . Left: $N = 48$ ($d(N) = 3$). Right: $N = 96, d(96) = 4$.

REMARKS.

- Both stochastic optimization procedures that we described above can, of course, be implemented to produce optimal (or optimized) grids of any multidimensional probability distribution on \mathbb{R}^d , having in mind that as the dimension d increases the second one becomes the most efficient.
- These procedures are based on a nearest-neighbor search among N points. A naive implementation of such a procedure has a linear complexity in N and becomes very demanding in higher dimension. So, to drastically reduce this optimization phase as well as that devoted to the weight estimation of the resulting optimal quantizer, one can call upon some fast nearest-neighbor procedure like that originally developed




FIG. 6.3 Optimized N -quantizer Γ^N of the Brownian motion W with $N = 400$. The grey level of the paths codes their weights.

by Bentley and analyzed in a seminal paper FRIEDMAN, BENTLEY and FINKEL [1977], which is based on the notion of k - d -tree introduced for that purpose by the authors. It reduces the complexity of the search down to $O(\log N)$. The (relative) efficiency of the method increases as the dimension of the state space increases.

6.4. An alternative: product functional quantization

Scalar product functional quantization is a quantization method that produces rate optimal suboptimal quantizers. They were used by LUSCHGY and PAGÈS [2002] to provide exact rate (although not sharp) for a very large class of processes. The first attempt to use functional quantization for numerical computation with the Brownian motion was achieved with these quantizers (see PAGÈS and PRINTEMS [2005]). We will see further on their assets. What follows is presented for the Brownian motion but would work for a large class of centered Gaussian processes.

Let us consider again the expansion of W in its K - L basis:

$$W \stackrel{L^2}{=} \sum_{n \geq 1} \sqrt{\lambda_n} \xi_n e_n^W,$$




FIG. 6.4 Optimal functional quantization of the Brownian motion. $N \mapsto \log N (e_N(W, L_T^2))^2$, $N \in \{6, \dots, 160\}$. Vertical dashed lines: critical dimensions for $d(N)$, $e^2 \approx 7$, $e^3 \approx 20$, $e^4 \approx 55$, $e^5 \approx 148$.

where $(\xi_n)_{n \geq 1}$ is an i.i.d. sequence $\mathcal{N}(0; 1)$ -distributed random variables (keep in mind this convergence also holds a.s. uniformly in $t \in [0, T]$). The idea is simply to quantize these (normalized) random coordinates ξ_n : for every $n \geq 1$, one considers an optimal N_n -quantization of ξ_n , denoted by $\hat{\xi}_n^{(N_n)}$ ($N_n \geq 1$). For $n > m$, set $N_n = 1$ and $\hat{\xi}_n^{(N_n)} = 0$ (which is the optimal one-quantization). The integer m is called the *length* of the product quantization. Then, one sets

$$\hat{W}_t^{(N_1, \dots, N_m, \text{prod})} := \sum_{n \geq 1} \sqrt{\lambda_n} \hat{\xi}_n^{(N_n)} e_n^W(t) = \sum_{n=1}^m \sqrt{\lambda_n} \hat{\xi}_n^{(N_n)} e_n^W(t).$$

Such a quantizer takes $\prod_{n=1}^m N_n \leq N$ values.

If one denotes by $\alpha^M = \{\alpha_1^M, \dots, \alpha_M^M\}$ the (unique) optimal quadratic M -quantizer of the $\mathcal{N}(0; 1)$ -distribution, the underlying quantizer of the above quantization $\hat{W}^{(N_1, \dots, N_m, \text{prod})}$ can be expressed as follows (if one introduces the appropriate multiindexation): for every multiindex $\underline{i} := (i_1, \dots, i_m) \in \prod_{n=1}^m \{1, \dots, N_n\}$, set

$$x_{\underline{i}}^{(N)}(t) := \sum_{n=1}^m \sqrt{\lambda_n} \alpha_{i_n}^{(N_n)} e_n^W(t) \text{ and } \Gamma^{N_1, \dots, N_m, \text{prod}} := \left\{ x_{\underline{i}}^{(N)}, \underline{i} \in \prod_{n=1}^m \{1, \dots, N_n\} \right\}.$$

Then, the product quantization $\hat{W}^{(N_1, \dots, N_m, \text{prod})}$ can be rewritten as

$$\hat{W}_t^{(N_1, \dots, N_m, \text{prod})} = \sum_{\underline{i}} \mathbf{1}_{\{W \in C_{\underline{i}}(\Gamma^{N_1, \dots, N_m, \text{prod}})\}} x_{\underline{i}}^{(N)}(t),$$

where the Voronoi cell of $x_{\underline{i}}^{(N)}$ is given by

$$C_{\underline{i}}(\Gamma^{N_1, \dots, N_m, \text{prod}}) = \prod_{n=1}^m (\alpha_{i_n - \frac{1}{2}}^{(N_n)}, \alpha_{i_n + \frac{1}{2}}^{(N_n)}), \quad \alpha_{i \pm \frac{1}{2}}^{(M)} := \frac{\alpha_i^{(M)} + \alpha_{i \pm 1}^{(M)}}{2},$$

$$\alpha_0 = -\infty, \alpha_{M+1} = +\infty.$$

6.4.1. Quantization rate by product quantizers

It is clear that the optimal product quantizer is the solution to the optimal integral bit allocation

$$\min \left\{ \|W - \hat{W}^{(N_1, \dots, N_m, \text{prod})}\|_2, N_1, \dots, N_m \geq 1, N_1 \times \dots \times N_m \leq N, m \geq 1 \right\}. \quad (6.3)$$

Expanding $\|W - \hat{W}^{(N_1, \dots, N_m, \text{prod})}\|_2^2 = \| |W - \hat{W}^{(N_1, \dots, N_m, \text{prod})}|_{L_T^2} \|_2^2$ yields

$$\|W - \hat{W}^{(N_1, \dots, N_m, \text{prod})}\|_2^2 = \sum_{n \geq 1} \lambda_n \|\hat{\xi}_n^{(N_n)} - \xi_n\|_2^2 \quad (6.4)$$

$$= \sum_{n=1}^m \lambda_n (e^2_{N_n}(\mathcal{N}(0; 1), \mathbb{R}) - 1) + \frac{T^2}{2} \quad (6.5)$$

since

$$\sum_{n \geq 1} \lambda_n = \mathbb{E} \sum_{n \geq 1} (W |e_n^W|_2)^2 = \mathbb{E} \int_0^T W_t^2 dt = \int_0^T t dt = \frac{T^2}{2}.$$

THEOREM 6.3 (see LUSCHGY and PAGÈS [2002]). *For every $N \geq 1$, there exists an optimal scalar product quantizer of size at most N (or at level N), denoted by $\hat{W}^{(N, \text{prod})}$, of the Brownian motion defined as the solution to the minimization problem (6.3). Furthermore, these optimal product quantizers make up a rate optimal sequence: there exists a real constant $c_W > 0$ such that*

$$\|W - \hat{W}^{(N, \text{prod})}\|_2 \leq \frac{c_W T}{(\log N)^{\frac{1}{2}}}.$$

Proof (sketch of). By scaling, one may assume without loss of generality that $T = 1$. Combining (6.4) and Zador's theorem shows

$$\begin{aligned} \|W - \hat{W}^{(N_1, \dots, N_m, \text{prod})}\|_2^2 &\leq C \left(\sum_{n=1}^m \frac{1}{n^2 N_n^2} \right) + \sum_{n \geq m+1} \lambda_n \\ &\leq C' \left(\sum_{n=1}^m \frac{1}{n^2 N_n^2} + \frac{1}{m} \right), \end{aligned}$$

with $\prod_n N_n \leq N$. Setting $m := m(N) = \lfloor \log N \rfloor$ and $N_k = \lfloor \frac{(m!N)^{\frac{1}{m}}}{k} \rfloor \geq 1$, $k = 1, \dots, m$, yields the announced upper bound. \diamond

REMARKS.

- One can show that the length $m(N)$ of the optimal quadratic product quantizer satisfies

$$m(N) \sim \log N \quad \text{as} \quad N \rightarrow +\infty.$$

- The most striking fact is that very few ingredients are necessary to make the proof work as far as the quantization rate is concerned. We only need the basis of L_T^2 on which W is expanded to be orthonormal or the random coordinates to be orthogonal in $L^2(\mathbb{P})$. This robustness of the proof has been used to obtain some upper bounds for very wide classes of Gaussian processes by considering alternative orthonormal basis of L_T^2 like the Haar basis for processes having self-similarity properties (see LUSCHGY and PAGÈS [2002]), or trigonometric basis for stationary processes (see LUSCHGY and PAGÈS [2002]). More recently, combined with the nonasymptotic Zador's theorem, it was used to provide some connections between mean regularity of stochastic processes and quantization rate (see Section 9 and LUSCHGY and PAGÈS [2007]).
- Block quantizers combined with large deviation estimates were used to provide the sharp rate obtained in Theorem 6.1 LUSCHGY and PAGÈS [2004].
- d -dimensional block quantization is also possible, possibly with varying block size, providing a constructive approach to sharp rate see WILBERTZ [2005] and LUSCHGY, PAGÈS and WILBERTZ [2007].
- A similar approach can also provide some $L^r(\mathbb{P})$ -rates for product quantization with respect to the sup-norm over $[0, T]$, see LUSCHGY and PAGÈS [2007].

6.4.2. How to use product quantizers for numerical computations?

For numerics, one can assume by a scaling argument that $T = 1$. To use product quantizers for numerics, we need to have access to the quantizers (or grid) at a given level N , their weights (and the quantization error). All these quantities are available with product quantizers. In fact, the first attempts to use functional quantization for numerics (path-dependent option pricing) were carried out with product quantizers (see PAGÈS and PRINTEMS [2005]).

TABLE 6.1
Optimal product quantization of the Brownian motion: optimal allocations for $N = 10^k, k = 0, \dots, 5$.

N	N_{rec}	Quantization error	Optimal allocation
1	1	0.7071	1
10	10	0.3138	5-2
100	96	0.2264	12-4-2
1 000	966	0.1881	23-7-3-2
10 000	9 984	0.1626	26-8-4-3-2-2
100 000	97 920	0.1461	34 - 10 - 6 - 4 - 3 - 2 - 2

- The optimal product quantizers (denoted by $\Gamma^{(N, \text{prod})}$) at level N are explicit, given the optimal quantizers of the scalar normal distribution $\mathcal{N}(0; 1)$. In fact, the optimal allocation of the size N_i of each marginal has been already achieved up to very high values of N . Some typical optimal allocation (and the resulting quadratic quantization error) is reported in Table 6.1 below. The integer N_{rec} denotes the effective size of the optimal product quantizer.
- The weights $\mathbb{P}(\hat{W}^{(N, \text{prod})} = x_i)$ are explicit too: the normalized coordinates ξ_n of W in its K - L basis are independent, consequently,

$$\begin{aligned} \mathbb{P}(\hat{W}^{(N, \text{prod})} = x_i) &= \mathbb{P}(\hat{\xi}_n^{(N_n)} = \alpha_{i_n}^{(N_n)}, n = 1, \dots, m(N)) \\ &= \prod_{n=1}^{m(N)} \underbrace{\mathbb{P}(\hat{\xi}_n^{(N_n)} = \alpha_{i_n}^{(N_n)})}_{1D \text{ (tabulated) weights}}. \end{aligned}$$

- Eq. (6.5) shows that the (squared) quantization error of a product quantizer can be straightforwardly computed as soon as one knows the eigenvalues and the (squared) quantization error of the normal distributions for the N_i s.

The optimal allocations up to $N = 12\,000$ can be downloaded on the Web site (PAGÈS and PRINTEMS [2005]) as well as the necessary one-dimensional optimal quantizers (including the weights and the quantization error) of the scalar normal distribution (up to a size of 500, which is enough for this purpose). Some examples of optimal product quantizers are displayed in Figs. 6.5 and 6.6.

For numerical purpose, we are also interested in the stationarity property since such quantizers produce lower (weak) errors in cubature formulas.

PROBLEM 6.1 (see PAGÈS and PRINTEMS [2005]). The product quantizers obtained from the K - L basis are stationary quantizers (although suboptimal).

PROOF. First, note that

$$\hat{W}^{N, \text{prod}} = \sum_{n \geq 1} \sqrt{\lambda_n} \hat{\xi}_n^{(N_n)} e_n(t)$$

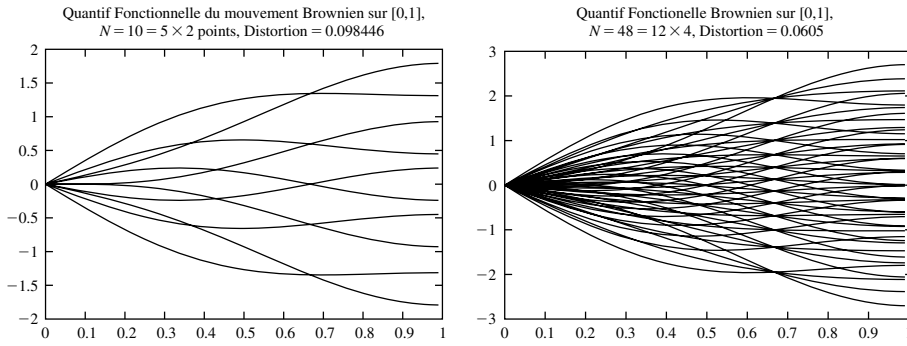


FIG. 6.5 Product quantization of the Brownian motion: the N_{rec} -quantizer $\Gamma^{(N, \text{prod})}$. $N = 10$: $N_{\text{rec}} = 10$ and $N = 50$: $N_{\text{rec}} = 12 \times 4 = 48$.

so that $\sigma(\hat{W}^{N, \text{prod}}) = \sigma(\hat{\xi}_k^{(N_k)}, k \geq 1)$. Consequently,

$$\begin{aligned} \mathbb{E}(W \mid \hat{W}^{N, \text{prod}}) &= \mathbb{E}(W \mid \sigma(\hat{\xi}_k^{(N_k)}, k \geq 1)) \\ \mathbb{E}(W \mid \hat{W}^{N, \text{prod}}) &= \sum_{n \geq 1} \sqrt{\lambda_n} \mathbb{E}(\xi_n \mid \sigma(\hat{\xi}_k^{(N_k)}, k \geq 1)) e_n^W \\ &\stackrel{i.i.d.}{=} \sum_{n \geq 1} \sqrt{\lambda_n} \mathbb{E}(\xi_n \mid \hat{\xi}_n^{(N_n)}) e_n^W \\ &= \sum_{n \geq 1} \sqrt{\lambda_n} \hat{\xi}_n^{(N_n)} e_n^W = \hat{W}. \end{aligned}$$

□

REMARKS.

- This result is no longer true for product quantizers based on other orthonormal basis.
- This shows the existence of nonoptimal stationary quantizers.

6.5. Optimal versus product quadratic functional quantization ($T = 1$)

(NUMERICAL) OPTIMIZED QUANTIZATION: By scaling, we can assume without loss of generality that $T = 1$. We carried out a huge optimization task in order to produce some optimized quantization grids for the Brownian motion by solving numerically (\mathcal{O}_N) for $N = 1$ up to $N = 10\,000$.

$$e_N(W, L_T^2)^2 \approx \frac{0.2195}{\log N}, \quad N = 1, \dots, 10\,000.$$

