

On the Distribution of Integration Error by Randomly-Shifted Lattice Rules

Pierre L'Ecuyer, David Munger and Bruno Tuffin

*DIRO, Université de Montreal,
C.P. 6128, Succ. Centre-Ville,
Montréal (Québec), H3C 3J7, CANADA,
e-mail: lecuyer@iro.umontreal.ca; mungerd@iro.umontreal.ca*

*INRIA Rennes Bretagne Atlantique,
Campus de Beaulieu,
35042 Rennes Cedex, FRANCE,
e-mail: btuffin@irisa.fr*

Abstract: A lattice rule with a randomly-shifted lattice estimates a mathematical expectation, written as an integral over the s -dimensional unit hypercube, by the average of n evaluations of the integrand, at the n points of the shifted lattice that lie inside the unit hypercube. This average provides an unbiased estimator of the integral and, under appropriate smoothness conditions on the integrand, it has been shown to converge faster as a function of n than the average at n independent random points (the standard Monte Carlo estimator). In this paper, we study the behavior of the estimation error as a function of the random shift, as well as its distribution for a random shift, under various settings. While it is well known that the Monte Carlo estimator obeys a central limit theorem when $n \rightarrow \infty$, the randomized lattice rule does not, due to the strong dependence between the function evaluations. We show that for the simple case of one-dimensional integrands, the limiting error distribution is uniform over a bounded interval if the integrand is non-periodic, and has a square root form over a bounded interval if the integrand is periodic. We find that in higher dimensions, there is little hope to precisely characterize the limiting distribution in a useful way for computing confidence intervals in the general case. We nevertheless examine how this error behaves as a function of the random shift from different perspectives and on various examples. We also point out a situation where a classical central-limit theorem holds when the dimension goes to infinity, we provide guidelines on when the error distribution should not be too far from normal, and we examine how far from normal is the error distribution in various instances of a typical real-life problem.

AMS 2000 subject classifications: Primary 62E20; secondary 60F05.

Keywords and phrases: quasi-Monte Carlo, lattice rule, integration error, limit theorem, confidence interval.

1. Introduction

Monte Carlo (MC) and quasi-Monte Carlo (QMC) methods estimate the integral of a function f over the s -dimensional unit hypercube $[0, 1]^s = \{\mathbf{u} = (u_1, \dots, u_s) : 0 \leq u_j < 1 \text{ for all } j\}$, by evaluating f at n points in this hypercube and taking the average. This integral often represents a mathematical

expectation:

$$\mu = \mu(f) = \int_{[0,1]^s} f(\mathbf{u}) \, d\mathbf{u} = \mathbb{E}[f(\mathbf{U})] \quad (1)$$

where $f : [0, 1]^s \rightarrow \mathbb{R}$, $\mathbf{U} = (U_1, \dots, U_s) \sim U(0, 1)^s$, and the latter means that \mathbf{U} is a random vector with the uniform distribution over $(0, 1)^s$. This framework is very general, because the (pseudo)randomness in a simulation is typically obtained from a sequence of independent uniform (pseudo)random numbers over the interval $(0, 1)$. We can take s as an upper bound on the number of calls to the generator and view f as representing all the transformations applied to those uniform random numbers to produce the estimator.

In the MC method, we sample n independent random points $\mathbf{U}_0, \dots, \mathbf{U}_{n-1}$ uniformly over $(0, 1)^s$ and we estimate μ by

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=0}^{n-1} f(\mathbf{U}_i). \quad (2)$$

Under standard assumptions, $\hat{\mu}_n$ obeys the classical central-limit theorem (CLT), so if n is large enough we can compute a confidence interval on μ by assuming that the error follows approximately a normal distribution. The variance can be estimated by the empirical variance, because the observations are independent.

With randomized quasi-Monte Carlo (RQMC), the estimator has the same form:

$$\hat{\mu}_{n,\text{rqmc}} = \frac{1}{n} \sum_{i=0}^{n-1} f(\mathbf{U}_i) \quad (3)$$

and we still have $\mathbf{U}_i \sim U[0, 1]^s$ for each i , but these points are no longer independent. They are constructed in a way that the set $P_n = \{\mathbf{U}_0, \dots, \mathbf{U}_{n-1}\} \subset [0, 1]^s$ covers the unit hypercube more uniformly than typical independent random points. This uniformity is quantified by figures of merit called *discrepancies*, which measure of discrepancy between the empirical distribution of the points and the uniform distribution over $(0, 1)^s$. Two popular ways of constructing RQMC point sets are randomly-shifted lattices and digitally-shifted nets. For general background on quasi-Monte Carlo methods, RQMC, and discrepancies, the reader is referred to [L'Ecuyer \(2009\)](#); [L'Ecuyer and Lemieux \(2002\)](#); [Lemieux \(2009\)](#); [Niederreiter \(1992\)](#); [Owen \(1998\)](#); [Sloan and Joe \(1994\)](#) and the references given there.

The RQMC estimator (3) has expectation μ and variance

$$\text{Var}[\hat{\mu}_{n,\text{rqmc}}] = \mathbb{E}[(\hat{\mu}_{n,\text{rqmc}} - \mu)^2]. \quad (4)$$

This variance is more difficult to estimate than for MC, because of the dependence across observations. The usual way of estimating it and computing a confidence interval on μ is to obtain m independent realizations of $\hat{\mu}_{n,\text{rqmc}}$, say X_1, \dots, X_m , based on m independent randomizations of P_n , and compute their sample mean \bar{X}_m and their sample variance and $S_{x,m}^2$. One has $\mathbb{E}[\bar{X}_m] = \mu$ and $\mathbb{E}[S_{x,m}^2] = m \text{Var}[\bar{X}_m]$ ([L'Ecuyer and Lemieux, 2000](#)). With a variance estimate

in hand, one might be tempted to compute a confidence interval on μ by using a normal approximation for the distribution of \bar{X}_m , as for MC. This is common practice. Even if $\hat{\mu}_{n,\text{rqmc}}$ is not normal, \bar{X}_m will eventually be approximately normal if m is large enough, because the CLT applies when $m \rightarrow \infty$. But in typical RQMC settings, m is small, often no more than 10, and n is taken as large as possible given the available computing budget. Then it is unclear if a normal approximation makes sense.

This motivates the need for a better understanding of the distributional behavior of $\hat{\mu}_{n,\text{rqmc}}$ in the asymptotic regime where m is fixed (and small) and $n \rightarrow \infty$, and also for typical (finite) values of n . This is our aim in the present paper.

For certain RQMC methods that involve a sufficient amount of randomization of the dependent points, a CLT has been proved to hold for $n \rightarrow \infty$. Two such special cases are Latin hypercube sampling (LHS) (Owen, 1992), for which a bound on the total variation convergence to the normal distribution is available, and digital nets with the full nested scrambling of Owen (Loh, 2003). See Loh (2005) for a survey. However, LHS is not one of the most powerful RQMC methods, because it ensures good uniformity only for the one-dimensional projections of the s -dimensional points, and nested scrambling is rarely used because it is very time-consuming.

In this paper, we focus on randomly-shifted lattice rules, a widely-used and effective RQMC technique (L'Ecuyer and Lemieux, 2000; Sloan and Joe, 1994), where the n evaluation points are those of a randomly shifted integration lattice of density n that lie in the unit hypercube, in s dimensions. We assume that the lattice is projection-regular, which means that each of its one-dimensional projections over the interval $[0, 1)$ contains exactly n points. We explore and illustrate how the error behaves as a function of the random shift, and study its asymptotic distribution (for large n), in a variety of settings. We find that when s and m are fixed and $n \rightarrow \infty$, the error does not obey a CLT, in the sense that it does not have a normal asymptotic distribution. We provide some insight on when the normal approximation is likely to be reasonable or not for finite (typical) n . We also have some results for an asymptotic regime where $s \rightarrow \infty$.

Randomly shifting the lattice means that we generate a random shift $\mathbf{U} \sim U(0, 1)^s$ and add it to all the lattice points. This is equivalent to sampling \mathbf{U} in one of the elementary parallelotopes determined by any s vectors that form a basis of the lattice. We show that this is also equivalent to generating it uniformly in $[0, 1)^{s-1} \times [0, 1/n)$, which is much simpler, or more generally in any elementary region of a partition of $[0, 1)^s$ into n pieces of volume $1/n$ if the partition satisfies certain tiling conditions. These properties are useful to understand how the error behaves as a function of the shift. We provide illustrations.

We say that f is c -periodic with respect to coordinate j if its periodic continuation with respect to coordinate j is a continuous function. This means that f

is continuous in $[0, 1]^s$ and that for any $(u_1, \dots, u_{j-1}, u_{j+1}, \dots, u_s) \in [0, 1]^{s-1}$,

$$\lim_{\delta \rightarrow 0^+} f(u_1, \dots, u_{j-1}, 1 - \delta, u_{j+1}, \dots, u_s) = f(u_1, \dots, u_{j-1}, 0, u_{j+1}, \dots, u_s).$$

If it is c -periodic with respect to each coordinate j , f is simply called *c-periodic*. We will see that for continuous c -periodic functions f , the integration error as a function of the shift is a continuous function over the unit hypercube, and that if f is continuous but not c -periodic with respect to coordinate j , the error function has jumps whenever coordinate j of the shift crosses a multiple of $1/n$.

For the simple case of a one-dimensional continuous function ($s = 1$), we develop a generalized Euler-MacLaurin expansion of the integration error, from which we obtain the asymptotic distribution of this error when $n \rightarrow \infty$. As a special case, we find that when f is not c -periodic, the asymptotic distribution of n times the integration error is uniform over a bounded interval centered at 0. If f is c -periodic but its derivative is not, then n^2 times the error converges in distribution to the square of a uniform random variable centered at 0, whose cumulative distribution function (cdf) is a square root function. In general, if f has $k - 2$ c -periodic derivatives but its derivative of order $k - 1$ is not c -periodic, then n^k times the error converges in distribution to a constant times $B_k(U)$ where B_k is the Bernoulli polynomial of degree k and $U \sim U(0, 1)$.

For more general s -dimensional functions, the analysis is more complicated because (among other things) the convergence behavior and even the convergence rate depend on the choice of lattice as a function of n . It is well known from the theory of lattice rules (Kuo and Joe, 2002; Niederreiter, 1992; Sinescu and L'Ecuyer, 2010; Sloan and Reztsov, 2002) that if all the mixed partial derivatives of order one of f are absolutely integrable over $[0, 1]^s$, then for any $\epsilon > 0$ there are rank-1 lattice rules for all n , such that the integration error converges as $O(n^{-1+\epsilon})$ uniformly over the shift \mathbf{U} . These rules are also not difficult to find and they can be constructed explicitly component by component (one coordinate at a time) (Kuo and Joe, 2002; Sloan and Reztsov, 2002). Of course, this convergence rate does not apply to an arbitrary lattice rule. One can construct examples of badly selected rules for a given function f where the integration error is the same for all n . On the other hand, this rate applies to the average over all reasonable rank-1 lattice rules, so it represents the rate for typical rules used in practice. Note that this convergence rate is just an upper bound and does not tell much about the error distribution.

To study this distribution, we will decompose the error as a sum of three terms, $E_1 + E_2 + E_3$, where nE_1 is a linear combination of independent uniform random variables over $(-1/2, 1/2)$, whose cdf is a spline of degree s that does not depend on the choice of lattice, E_2 has a discrete distribution over n^{s-1} values, and E_3 is $O(n^{-2+\epsilon})$ when f is smooth enough, and it typically becomes negligible when $n \rightarrow \infty$. Any of these terms could be zero in certain situations. The term E_1 is the integration error for the best approximation of f by a sum of one-dimensional functions. In situations where this term dominates the total error, the spline is a good approximation of its cdf. For E_2 , we do not have an expression or a characterization of the limiting distribution; it seems that

the cdf of E_2 , and thus the cdf of the error when E_2 dominates, can be almost anything. We also look at the integration error via its Fourier expansion.

Our results for E_1 come from our study of the special case where f is a sum of one-dimensional functions. In practice, it is not unusual that E_1 dominates the error, especially in situations where RQMC works very well, and this is often achieved deliberately by transforming the integrand f via a change of variables [Cafisch, Morokoff and Owen \(1997\)](#); [L'Ecuyer \(2009\)](#); [Sloan and Joe \(1994\)](#). We show that for a sum of s non c -periodic one-dimensional functions, the cdf of n times the integration error converges to a spline (a piecewise-polynomial function) of degree s when $n \rightarrow \infty$, and this limiting distribution has a bounded support (in contrast with the normal distribution). In this sense, the behavior is significantly different than when a classical CLT holds. We also show that when $s \rightarrow \infty$, under certain conditions, a classical CLT holds for a properly standardized version of the error. We give illustrative examples.

Finally, we perform an empirical study of the error distributions for an example inspired from real-life: pricing an Asian option.

The remainder is organized as follows. In [Section 2](#), we recall the basic definitions and results on randomly-shifted lattice rules, and introduce some notation. In [Section 3](#), we show how the random shift can be equivalently generated over a smaller subset of the unit hypercube, and we provide several illustrations. In [Section 4](#), we study the error distribution for one-dimensional integrals. In [Section 5](#), we look at linear combinations of one-dimensional functions and show that a classical CLT holds under certain conditions as the number of functions increases. In [Section 6](#), we study the general s -dimensional case. A real-life example is examined in [Section 7](#).

Preliminary results (mostly for the one-dimensional case) were published in the proceedings of the 2009 Winter Simulation Conference ([L'Ecuyer and Tuffin, 2009](#)).

2. Randomly-Shifted lattice Rules

2.1. Definitions

An *integration lattice* is a discrete vector space of the form

$$L_s = \left\{ \mathbf{v} = \sum_{j=1}^s z_j \mathbf{v}_j \text{ such that each } z_j \in \mathbb{Z} \right\},$$

where the basis vectors $\mathbf{v}_1, \dots, \mathbf{v}_s \in \mathbb{R}^s$ are linearly independent over \mathbb{R} and where L_s contains \mathbb{Z}^s , the set of integer vectors. The *dual lattice* to L_s is

$$L_s^* = \{ \mathbf{h} \in \mathbb{R}^s : \mathbf{h}^t \mathbf{v} \in \mathbb{Z} \text{ for all } \mathbf{v} \in L_s \},$$

where “ t ” means “transposed” and the vectors are assumed to be column vectors whenever it matters (in the text we often write them as row vectors to simplify

the notation). It is easy to see that $\mathbb{Z}^s \subseteq L_s$ if and only if $L_s^* \subseteq \mathbb{Z}^s$, i.e., if all coordinates of all vectors of L_s^* are integers.

Let $P_n = L_s \cap [0, 1]^s = \{\mathbf{u}_0, \dots, \mathbf{u}_{n-1}\}$ be the set of lattice points that belong to the unit hypercube $[0, 1]^s$, where n is the cardinality of this set P_n . One can show that all coordinates of all vectors of L_s must be multiples of $1/n$. A (deterministic) *lattice rule* is a numerical integration method that approximates μ by the average of $f(\mathbf{u}_0), \dots, f(\mathbf{u}_{n-1})$. Further details on integration lattices, lattice rules, and ways of measuring the uniformity of P_n for an integration lattice, can be found in L'Ecuyer (2009); L'Ecuyer and Lemieux (2000); Sloan and Joe (1994); Tuffin (1998), for example. Most integration lattices used in practice are *projection-regular*, which means that each one-dimensional projection of P_n is the set $\{0, 1, \dots, n-1\}$, and we assume here that the lattice has this property. They are also usually *of rank 1*, which means that the basis can be selected so that $\mathbf{v}_j = \mathbf{e}_j$ (the j th unit vector in s -dimensions) for $j = 2, \dots, s$, and we can write

$$P_n = \{\mathbf{v} = i\mathbf{v}_1 \bmod 1, i = 0, \dots, n-1\} = \{(i\mathbf{a}_1 \bmod n)/n, i = 0, \dots, n-1\},$$

where $\mathbf{a}_1 = (a_1, \dots, a_s)$ and $\mathbf{v}_1 = \mathbf{a}_1/n$.

A lattice rule can be turned into a RQMC method by applying a *random shift modulo 1* to P_n , which consists in generating a single point $\mathbf{U} \sim U(0, 1)^s$ and adding it to each point of P_n , modulo 1, coordinate-wise. That is, each point $\mathbf{u}_i = (u_{i,1}, \dots, u_{i,s}) \in L_s \cap [0, 1]^s$ is randomized into $\mathbf{U}_i = (U_{i,1}, \dots, U_{i,s}) = (\mathbf{u}_i + \mathbf{U}) \bmod 1$. This was proposed by Cranley and Patterson (1976) and further studied by L'Ecuyer and Lemieux (2000), among others. This is the same as randomly shifting L_s and then taking the intersection with $[0, 1]^s$. Clearly, $\mathbf{U}_i \sim U[0, 1]^s$ for each i , and the global uniformity of P_n is preserved by the shift in the sense that the points have the same relationship with each other if we forget about the hypercube boundaries.

Figure 1 gives a toy illustration with $s = 2$, $n = 8$, $\mathbf{v}_1 = (1/8, 3/8)$, $\mathbf{v}_2 = (0, 1)$, and $\mathbf{U} = (0.436, 0.233)$.

2.2. Integration Error

The integration error for a lattice rule shifted by \mathbf{u} is given by

$$g_n(\mathbf{u}) = \frac{1}{n} \sum_{i=1}^{n-1} f((\mathbf{u}_i + \mathbf{u}) \bmod 1) - \mu. \quad (5)$$

We are interested in the behavior of $g(\mathbf{u})$ as a function of \mathbf{u} , and the distribution of $g_n(\mathbf{U})$, the integration error for a randomly-shifted lattice rule, for large n .

Suppose that the integrand f has the Fourier expansion

$$f(\mathbf{u}) = \sum_{\mathbf{h} \in \mathbb{Z}^s} \hat{f}(\mathbf{h}) \exp(2\pi i \mathbf{h}^t \mathbf{u}), \quad (6)$$

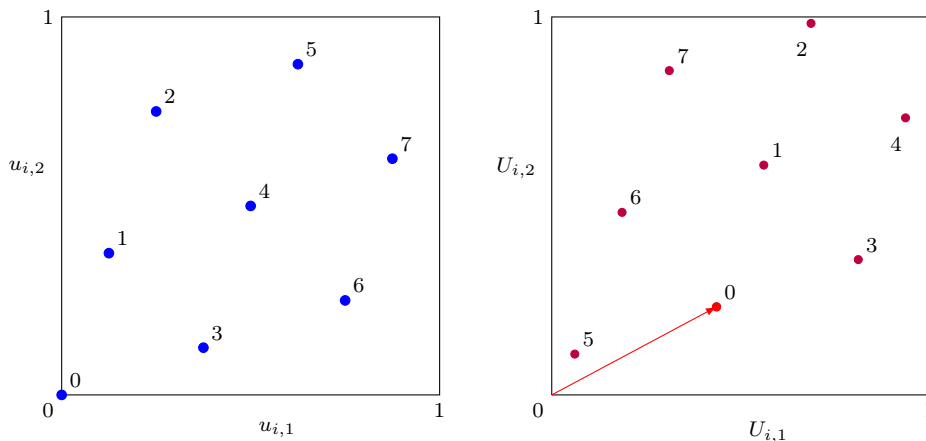


FIG 1. The original eight integration points of a small two-dimensional lattice (left) and the same eight points after a shift modulo 1 by $\mathbf{U} = (0.436, 0.233)$ (right).

with Fourier coefficients

$$\hat{f}(\mathbf{h}) = \int_{[0,1]^s} f(\mathbf{u}) \exp(-2\pi i \mathbf{h}^t \mathbf{u}) \, d\mathbf{u},$$

where $i = \sqrt{-1}$. Then it is known that the Fourier coefficients of g_n satisfy $\hat{g}_n(\mathbf{h}) = \hat{f}(\mathbf{h})$ if $\mathbf{0} \neq \mathbf{h} \in L_s^*$, and $\hat{g}_n(\mathbf{h}) = 0$ otherwise (Niederreiter, 1992; Sloan and Joe, 1994). Therefore, the Fourier expansion of g_n can be written (when it exists) as

$$g_n(\mathbf{u}) = \sum_{\mathbf{0} \neq \mathbf{h} \in L_s^*} \hat{f}(\mathbf{h}) \exp(2\pi i \mathbf{h}^t \mathbf{u}). \quad (7)$$

By the Parseval equality, we also have

$$\text{Var}[\hat{\mu}_{n,\text{rqmc}}] = \text{Var}[g_n(\mathbf{U})] = \mathbb{E}[g_n^2(\mathbf{U})] = \sum_{\mathbf{0} \neq \mathbf{h} \in L_s^*} |\hat{f}(\mathbf{h})|^2$$

whenever f is square integrable (L'Ecuyer and Lemieux, 2000). By making assumptions on how fast the Fourier coefficients converge when the size of \mathbf{h} increases, one can obtain asymptotic bounds on the worst-case error $\sup_{\mathbf{u} \in [0,1]^s} |g_n(\mathbf{u})|$ or on the variance $\text{Var}[g_n(\mathbf{U})]$. For example, let $\alpha > 1/2$, take some non-negative constants $\gamma_1, \dots, \gamma_s$, and consider the class of functions $f : [0,1]^s \rightarrow \mathbb{R}$ such that for all $\mathbf{h} = (h_1, \dots, h_s) \in \mathbb{Z}^s$,

$$|\hat{f}(\mathbf{h})|^2 \leq w(\mathbf{h}) \stackrel{\text{def}}{=} \prod_{\{j:h_j>0\}} \gamma_j |h_j|^{-2\alpha}.$$

It is known that there exists a sequence of lattice rules, indexed by n , such that for any $\delta > 0$, the worst-case error and the variance converge as $O(n^{-\alpha+\delta})$ and

$O(n^{-2\alpha+\delta})$, respectively (Dick *et al.*, 2006; Sloan and Joe, 1994). When α is an integer, the above condition can be written in terms of square integrability of a collection of partial derivatives: f satisfies the condition if for every subset of coordinates, the partial derivative of order α with respect to these coordinates is square integrable, and the partial derivatives of orders 0 to $\alpha - 2$ are c-periodic (Dick *et al.*, 2006; L'Ecuyer, 2009).

Concrete applications where the variance is reduced drastically by randomly-shifted lattice rules can be found in L'Ecuyer (2009); L'Ecuyer and Lemieux (2000); Lemieux (2009), for example.

3. Generating the Shift Over a Smaller Region

3.1. A General Condition on Tiling Patterns

Instead of generating the random shift $\mathbf{U} \sim U(0, 1)^s$, we can equivalently generate it over a smaller subregion R_0 of volume $1/n$. The key observation is that shifting by $\mathbf{u}_i + \mathbf{U}$ (modulo 1) for any $\mathbf{u}_i \in P_n$ is equivalent to shifting by \mathbf{U} , because shifting the lattice L_s by any lattice point just gives the same lattice. The next proposition characterizes the regions R_0 for which this idea works. These regions determine a tiling pattern made of n shifted versions of R_0 that cover the unit hypercube.

Proposition 1. *Let $R_0 \subset [0, 1]^s$ be a region such that the family $\{R_i = (R_0 + \mathbf{u}_i) \bmod 1, i = 0, \dots, n - 1\}$ forms a partition of $[0, 1]^s$ in n regions of volume $1/n$, where R_i is R_0 shifted by \mathbf{u}_i modulo 1. Then, sampling the random shift uniformly in R_i for any fixed i is equivalent to sampling it uniformly in $[0, 1]^s$. Consequently, the error function $g_n(\mathbf{u})$ over any of those regions R_i is the same as over R_0 , in the sense that $g_n(\mathbf{u}) = g_n((\mathbf{u} + \mathbf{u}_i) \bmod 1)$ for all i .*

Proof. Note that if a shifted lattice point \mathbf{u} is in R_k for some k , then for each $i \neq k$, another shifted lattice point $(\mathbf{u}_i - \mathbf{u}_k + \mathbf{u}) \bmod 1$ must be in R_i . This means that each R_i must contain a single shifted lattice point. Consider now a fixed i and suppose (without loss of generality) that \mathbf{u}_i is the lattice point that belongs to R_i before the shift. Let $\tilde{\mathbf{U}}_i$ be the single lattice point that falls in R_i after a shift by $\mathbf{U} \in [0, 1]^s$, and let $\tilde{\mathbf{U}}_k = (\mathbf{u}_k - \mathbf{u}_i + \tilde{\mathbf{U}}_i) \bmod 1$ for $k = 0, \dots, n - 1$. Then a random shift by any of these $\tilde{\mathbf{U}}_k$ is the same as a random shift by $\tilde{\mathbf{U}}_i$, and leads to exactly the same point in R_i . If τ_i denotes the transformation that maps the uniform random shift $\mathbf{U} \in [0, 1]^s$ to $\tilde{\mathbf{U}}_i \in R_i$, then the inverse image by τ_i of any measurable set of volume $\epsilon > 0$ in R_i is the union of n sets of volume ϵ in $[0, 1]^s$, and so has total volume $n\epsilon$. This means that the point $\tilde{\mathbf{U}}_i$ has the uniform density n over R_i , and we can generate it directly from this density. \square

Next, we examine some choices of R_0 that satisfy this proposition.

3.2. Random Shift in a Parallelotope

Consider a set of basis vectors $\mathbf{v}_1, \dots, \mathbf{v}_s$ that define the lattice. These vectors determine a parallelotope (an s -dimensional parallelepiped)

$$P = \left\{ \mathbf{v} = \sum_{j=1}^s c_j \mathbf{v}_j : 0 \leq c_j < 1 \text{ for all } j \right\},$$

of volume $1/n$, with one vertex at the origin, and whose edges connected to that vertex are precisely $\mathbf{v}_1, \dots, \mathbf{v}_s$. With $R_0 = P$, the conditions of Proposition 1 are satisfied (Conway and Sloane, 1999). To sample \mathbf{U} uniformly in P , we can generate $\mathbf{C} = (C_1, \dots, C_s)$ uniformly over $(0, 1)^s$ and put $\mathbf{U} = C_1 \mathbf{v}_1 + \dots + C_s \mathbf{v}_s$. This mapping $\mathbf{C} \rightarrow \mathbf{U}$ is a linear transformation, so its Jacobian is a constant, and therefore it preserves uniformity.

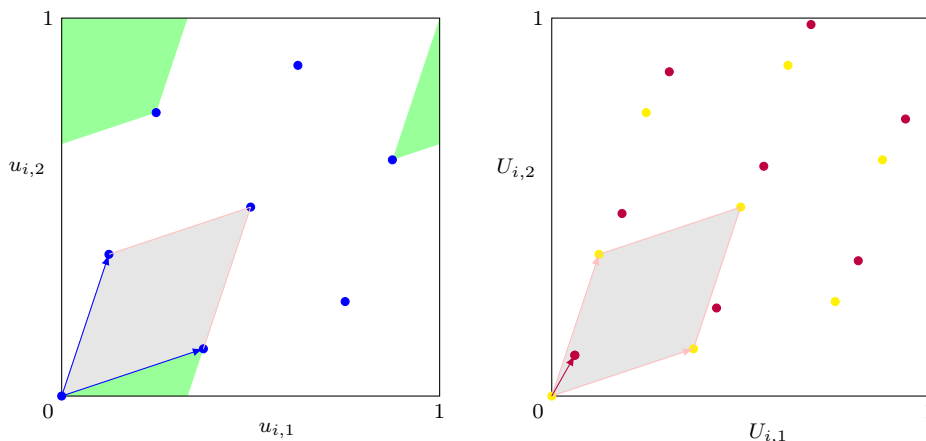


FIG 2. The original eight points and the two vectors \mathbf{v}_1 and \mathbf{v}_2 of a lattice basis (left) and the points shifted by $\tilde{\mathbf{U}} = (0.061, 0.108)$ modulo 1 (right).

Example 1. Figure 2 illustrates this property for the same example as in Figure 1, with the shift $\tilde{\mathbf{U}} \in P$ that corresponds to the shift \mathbf{U} shown in Figure 1. The parallelogram P is shaded in light gray in the figure. The figure also shows the parallelogram $(P + \mathbf{u}_i) \bmod 1$ for $\mathbf{u}_i = (0.875, 0.625)$, which is split into three pieces, shaded in dark green. Sampling uniformly in this parallelogram is easily achieved by sampling uniformly in P and adding \mathbf{u}_i modulo 1. \square

3.3. Random Shift in a Rectangular Slice

A simpler way of partitioning $[0, 1]^s$ is as follows. Fix an arbitrary coordinate $j \in \{1, \dots, s\}$ and, for each $i \in \{0, \dots, n-1\}$, define $R_i = R_{j,i} = \{\mathbf{u} =$

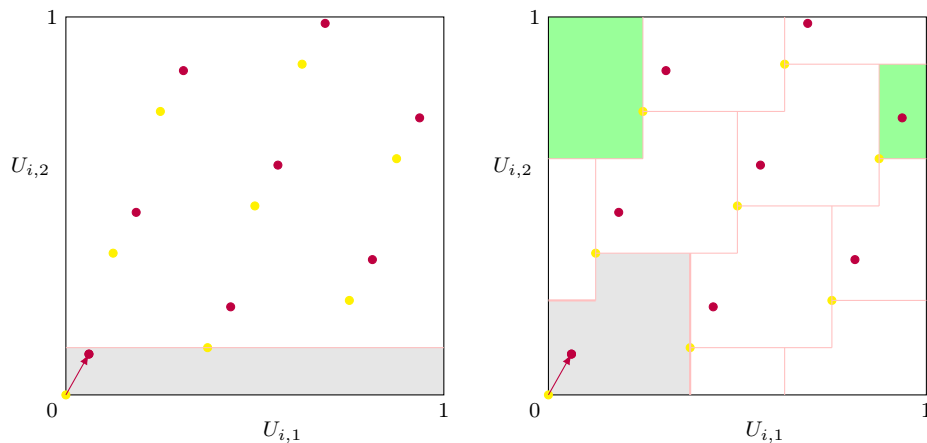


FIG 3. *Left: The points shifted by $\tilde{\mathbf{U}} = (0.061, 0.108)$ modulo 1; this shift was sampled uniformly over the lightly-shaded rectangle. Right: The shift can be equivalently sampled uniformly over the elementary region of a different “tiling” pattern, such as the lightly-shaded region shown here. The same region shifted by $(0.875, 0.625)$ (modulo 1) is shown in green.*

$(u_1, \dots, u_s) \in [0, 1]^s : i/n \leq u_j < (i+1)/n$. This partitions the hypercube into n hyper-rectangles (or slices) R_i of thickness (and volume) $1/n$. The assumption of Proposition 1 holds, because coordinate j of the original points takes each value in $\{0, 1/n, \dots, (n-1)/n\}$ exactly once when we run through all the points, and therefore there is always exactly one point in each R_i .

In one dimension ($s = 1$), this R_i is the interval $[i/n, (i+1)/n)$, so we can generate the random shift over the interval $[0, 1/n)$. Then the error function $g_n(u)$ for $u \in [0, 1)$ is periodic with period $1/n$ (or a divisor of $1/n$).

Example 2. In Example 1, we can generate the random shift \mathbf{U} uniformly over the rectangle $R_0 = [0, 1) \times [0, 1/8)$ instead of over the parallelogram. This is equivalent and obviously simpler. Figure 3 (left) illustrates this situation, with a shift $\mathbf{U} \in R$ equivalent to that of Figure 1. \square

Proposition 1 works for other types of regions than rectangles and paralleleptopes. The shaded region in Figure 3 (right) gives another illustration.

3.4. Discontinuity of g_n for non-periodic functions

Shifting the points by \mathbf{u} modulo 1 is equivalent to shifting the function f by $-\mathbf{u}$ modulo 1. If f is c -periodic, $f(\mathbf{u}_i + \mathbf{u}) \bmod 1$ varies continuously with \mathbf{u} for each i , which means that the error function $g_n(\mathbf{u})$ is also continuous in \mathbf{u} everywhere in $[0, 1)^s$, and is c -periodic.

On the other hand, suppose f is not c -periodic with respect to the j th coordinate and let $\mathbf{u} = (u_1, \dots, u_s)$ denote the shift. When u_j increases while the

other coordinates of \mathbf{u} are fixed, each time u_j crosses a multiple of $1/n$, there is one shifted point whose j th coordinate jumps from 1 to 0, and this causes a jump in the error function $g_n(\mathbf{u})$. Thus, $g_n(\mathbf{u})$ is discontinuous in u_j across the hyperplane $u_j = i/n$ for any integer i , and it is continuous in u_j elsewhere, provided that f is continuous in $(0, 1)^s$. Therefore, in the general non-c-periodic situation, $g_n(\mathbf{u})$ is continuous in each small hypercube of volume n^{-s} determined by these hyperplanes, and is discontinuous across the boundaries of these small hypercubes. In the one-dimensional case, these small hypercubes become the intervals of the form $[i/n, (i + 1)/n)$; the function $g_n(u)$ is the same over each of those intervals and it is discontinuous across them.

If f is a sum of one-dimensional functions, say $f(\mathbf{u}) = \sum_{j=1}^s f_j(u_j)$, then we can analyze the error by examining each f_j separately. The error function for f can be written as

$$g_n(\mathbf{u}) = \sum_{j=1}^s g_{j,n}(u_j)$$

where $g_{j,n}$ is the error function for f_j . This $g_{j,n}$ is periodic with period $1/n$, and it is discontinuous at each multiple of $1/n$ if f_j is not c-periodic. This implies that the function g_n over any of the small hypercubes of the form $\prod_{j=1}^s [i_j/n, (i_j + 1)/n)$, where each $i_j \in \{0, \dots, n - 1\}$, is a replicate of $g_n(\mathbf{u})$ for $\mathbf{u} \in [0, 1/n)^s$. In that case, the random shift \mathbf{U} can be equivalently sampled uniformly in any of the small s -dimensional hypercubes of volume n^{-s} .

The following examples illustrate these properties.

Example 3. Figure 4 shows the error function $g_n(\mathbf{u})$ for $\mathbf{u} \in [0, 1)^2$, for an example where $s = 2$, $f(u_1, u_2) = u_1 + u_2 - 3/2 = (u_1 - 1/2) + 2(u_2 - 1/2)$, $n = 8$, and the lattice points are the same as in Figure 1. This is a sum of two one-dimensional functions, so the error function has exactly the same shape in each small square of the form $[i_1/n, (i_1 + 1)/n) \times [i_2/n, (i_2 + 1)/n)$. It is also discontinuous across the boundaries of those squares, because f is not c-periodic. Here the error varies a lot within each square but it is the same across all squares. \square

Example 4. Figure 5 plots $g_n(\mathbf{u})$ for $\mathbf{u} \in [0, 1)^2$ for $s = 2$, $f(u_1, u_2) = (u_1 - 1/2)(u_2 - 1/2)$, and the same points. In contrast with the previous example, the error varies a lot across the squares and very little within each square. The error function is symmetric with respect to each of the two diagonals. If we partition the unit square into four squares of area $1/4$, the error function is also symmetric with respect to each of the two diagonal in each of those four squares. This high symmetry comes from the symmetry in f . \square

Example 5. Figure 6 plots $g_n(\mathbf{u})$ for $\mathbf{u} \in [0, 1)^2$ for $s = 2$, $f(u_1, u_2) = u_1^2(u_2 - 1) + 1/6$, and the same points. The error has a behavior halfway between that of the previous two examples: It varies substantially both across the squares and within them. Also, it does not exhibit as much symmetry. \square

Example 6. Figure 7 plots $g_n(\mathbf{u})$ for $\mathbf{u} \in [0, 1)^2$ for $s = 2$, $f(u_1, u_2) = u_1 u_2 (u_1 - 1/2)(u_2 - 1/2)$, and the same points. Here, f is c-periodic, so the error




FIG 4. Plot of $ng_n(u_1, u_2)$ for Example 3.

function is continuous and c -periodic. It also has all the symmetries observed in Example 4. \square

4. Error Distribution for One-Dimensional Integrals

4.1. Limit Theorems

One-dimensional integrals are rarely estimated by RQMC methods but studying the distribution of $g_n(U)$ for a randomly-shifted lattice rule in one dimension ($s = 1$ and $U \sim U(0, 1)$) is a first step toward understanding what may happen in general, and the result will be reused as a building block for the general case. It also provides simple counter-examples showing that the CLT does not apply. When $s = 1$, the random shift can be generated uniformly over $(0, 1/n)$, so the randomly-shifted points can be written as $\{U/n, (1 + U)/n, \dots, (n - 1 + U)/n\}$ where $U \sim U(0, 1)$ and the points are shifted by U/n . The corresponding total error is $g_n(U/n)$, and although it is not equal to $g_n(U)$ for the same U , Proposition 1 tells us that it has exactly the same distribution. Thus, we could expand the error over each interval of length $1/n$, and sum up these terms to get a handle on $g_n(U/n)$ and its distribution. This would yield the same result as in the next theorem, which provides an asymptotic expansion of order m for $g_n(U/n)$ when f has at least $m + 1$ integrable derivatives. This theorem appears in a different form in (Abramowitz and Stegun, 1970, Section 23.1.32), without




FIG 5. Plot of $ng_n(u_1, u_2)$ for Example 4.

a proof. For completeness, we provide a proof in Appendix A. We use B_k to denote the k th Bernoulli polynomial, defined by

$$B_k(u) = \sum_{j=0}^k \frac{1}{j+1} \sum_{\ell=0}^j (-1)^\ell \binom{j}{\ell} (u+\ell)^k$$

(Abramowitz and Stegun, 1970). This gives $B_0(u) = 1$, $B_1(u) = u - 1/2$, $B_2(u) = u^2 - u + 1/6$, $B_3(u) = u^3 - (3/2)u^2 + (1/2)u$, $B_4(u) = u^4 - 2u^3 + u^2 - 1/30$, and so on.

Theorem 2 (Generalized Euler–MacLaurin expansion). *If f has $m + 1$ integrable derivatives over $[0, 1]$, then*

$$g_n(U/n) = \sum_{k=1}^m \frac{B_k(U)}{n^k k!} \left[f^{(k-1)}(1) - f^{(k-1)}(0) \right] - \frac{1}{n^{m+1}} \int_0^1 \frac{B_{m+1}((U-u) \bmod 1) - B_{m+1}(U)}{(m+1)!} \left[\frac{1}{n} \sum_{i=0}^{n-1} f^{(m+1)}((i+u)/n) \right] du.$$

The last term is $O(n^{-(m+1)})$. This is the Euler–MacLaurin expansion of $g_n(U/n)$ of order m .




FIG 6. Plot of $ng_n(u_1, u_2)$ for Example 5.

Proposition 3. Suppose that f has $m + 1$ integrable derivatives over $[0, 1]$, $\int_0^1 f^{(m+1)}(u)du < \infty$, and $f^{(k)}(1) = f^{(k)}(0)$ for $k = 0, \dots, m-2$, but $f^{(m-1)}(1) \neq f^{(m-1)}(0)$. Then

$$g_n(U/n) = B_m(U) \frac{[f^{(m-1)}(1) - f^{(m-1)}(0)]}{m!n^m} + O(n^{-(m+1)}),$$

and therefore

$$W_n^{(m)} \stackrel{\text{def}}{=} \frac{m!n^m g_n(U/n)}{f^{(m-1)}(1) - f^{(m-1)}(0)} = B_m(U) + O(1/n)$$

converges in distribution to the random variable $B_m(U)$. If, in addition, f is infinitely differentiable and $f^{(k)}(1) = f^{(k)}(0)$ for all $k \geq m$, then $W_n^{(m)} = B_m(U)$ exactly.

Proof. Under the given assumptions, Theorem 2 gives

$$g_n(U/n) = B_m(U) \frac{f^{(m-1)}(1) - f^{(m-1)}(0)}{m!n^m} + O(n^{-(m+1)}),$$

from which the convergence result follows. For the last part, just take $m \rightarrow \infty$ in Theorem 2. \square




FIG 7. Plot of $ng_n(u_1, u_2)$ for Example 6.

Corollary 4. (i) By taking $m = 1$, we find that if f has an integrable second derivative over $[0, 1]$ and $f(1) \neq f(0)$, then we have $g_n(U/n) = (U - 1/2)(f(1) - f(0))/n + O(n^{-2})$ and

$$W_n^{(1)} \stackrel{\text{def}}{=} \frac{ng_n(U/n)}{f(1) - f(0)} = U - 1/2 + O(1/n),$$

so $W_n^{(1)} + 1/2$ converges in distribution to a $U(0, 1)$ random variable U when $n \rightarrow \infty$.

(ii) For $m = 2$, if f has an integrable third derivative over $[0, 1]$, $f(1) = f(0)$, $f'(1) \neq f'(0)$, we have

$$g_n(U/n) = [(U - 1/2)^2 - 1/12](f'(1) - f'(0))/(2n^2) + O(n^{-3}),$$

and if we define

$$W_n^{(2)} \stackrel{\text{def}}{=} \frac{2n^2 g_n(U/n)}{f'(1) - f'(0)} = (U - 1/2)^2 - 1/12 + O(1/n),$$

then $W_n^{(2)} + 1/12$ converges in distribution to the random variable $(U - 1/2)^2$, which has density $1/\sqrt{x}$ and distribution function $2\sqrt{x}$ for $0 \leq x < 1/4$, and zero density elsewhere. This means that $2[W_n^{(2)} + 1/12]^{1/2}$ converges in distribution to a uniform random variable over $[0, 1]$.

We recall that $g_n(U/n)$ has the same distribution as $\tilde{g}_n(U)$ where $U \sim U(0, 1)$. This corollary says that for non-c-periodic functions, n times the error converges in distribution to a uniform random variable, and for c-periodic functions whose derivative is not c-periodic, n^2 times the error converges in distribution to a random variable with a square root density over a finite interval.

4.2. Symmetric Functions and the Baker's Transformation

An integrand f with $f(0) \neq f(1)$ can be transformed into a c-periodic integrand \tilde{f} having the same integral, by defining $\tilde{f}(1-u) = \tilde{f}(u) = f(2u)$ for $0 \leq u \leq 1/2$. With this, we can switch from case (i) to case (ii) in Corollary 4, i.e., from $O(n^{-1})$ to $O(n^{-2})$ convergence for $g_n(U)$ (Hickernell, 2002). The resulting function \tilde{f} is also symmetric with respect to $u = 1/2$. This transformation is equivalently achieved by keeping f unchanged and transforming the randomized points U_i , for $i = 0, \dots, n-1$, via $\tilde{U}_i = 2U_i$ if $U_i < 1/2$ and $\tilde{U}_i = 2(1-U_i)$ if $U_i \geq 1/2$. This is called the baker's transformation; it stretches the points U_i by a factor of two and then folds back those that exceed 1. After applying this transformation, the lattice points become locally antithetic in each interval of the form $[i/n, (i+2)/n]$ if n and i are even, in the sense that they are at equal distance from the center of the interval, on each side. As a result, they integrate exactly any linear function over this interval. This holds for every such interval, so a piecewise-linear approximation which is linear over each interval is integrated exactly.

4.3. Examples with Smooth Functions

Example 7. Let $f(u) = u^2$. We have $f(1) - f(0) = 1$, so we are in case (i) of Corollary 4 and $W_n^{(1)} + 1/2 = ng_n(U/n) + 1/2$ is approximately $U(0, 1)$, with $O(1/n)$ approximation error. In fact, since $f''(u) = 2$ is a constant, the first two terms in Theorem 2 give the exact error function: $ng_n(U/n) + 1/2 = U + ((U - 1/2)^2 - 1/12)/n$. Thus, the $O(1/n)$ approximation error in the first order approximation is exactly $((U - 1/2)^2 - 1/12)/n$ and its absolute value is bounded by $1/(6n)$. The exact distribution of the error is

$$\mathbb{P}[ng_n(U/n) + 1/2 \leq x] = \frac{n-1}{2} \left[-1 + \sqrt{1 + \frac{4n[x - 1/(6n)]}{(n-1)^2}} \right]$$

when $0 \leq x - 1/(6n) \leq 1$ and is 0 otherwise. Figure 8 (left) shows the exact distribution of $ng_n(U/n) + 1/2$ for one-dimensional lattices with $n = 4$ and $n = 2^{10} = 1024$ points. Recall that these lattices have point sets $\{0, 1, \dots, (n-1)/n\}$. Here the random variable $ng_n(U/n) + 1/2$ has expectation $1/2$ (exactly) but it takes its values in the interval $(1/(6n), 1 + 1/(6n))$. For $n = 4$, the values are in the interval $(1/24, 1 + 1/24)$, which explains the shift observed in the figure. For $n = 1024$, the shift by $1/6144$ is too small to be visible to the eye. \square

Example 8. We apply the baker's transformation to the randomly-shifted points to integrate $f(u) = u^2$, the same function as in Example 7. The transformed function is $\tilde{f}(u) = 4u^2$ for $u \leq 1/2$ and $\tilde{f}(u) = 4(1-u)^2$ for $u > 1/2$. We obtain $\int_0^1 \tilde{f}'(u)du = 0$, but since $f''(u) = 8$ for all u , $\int_0^1 \tilde{f}''(u)du = 8 \neq 0 = \tilde{f}'(1) - \tilde{f}'(0)$, which can be explained by the fact that \tilde{f}' is discontinuous at $u = 1/2$. If n is even, this point of discontinuity lies at the boundary between two intervals, and case (ii) of Corollary 4 still applies with $\tilde{f}'(1) - \tilde{f}'(0)$ replaced by 8. Moreover, $f'''(u) = 0$ for all u . Thus, $n^2 g_n(U/n)/4 + 1/12 = (U - 1/2)^2$, so that $2[n^2 g_n(U/n)/4 + 1/12]^{1/2} \sim U(0, 1)$ (exactly).

If n is odd, the middle term of the summation does not have a continuous derivative, so the theorem does not apply. But if we write $n = 2p + 1$ and define

$$t_p(u) = \begin{cases} \frac{4}{n} \left(\frac{p+u}{n}\right)^2 & \text{if } u \leq 1/2 \\ \frac{4}{n} \left(1 - \frac{p+u}{n}\right)^2 & \text{if } u \geq 1/2, \end{cases}$$

we can compute

$$\begin{aligned} g_n(u/n) &= \sum_{k=0}^{p-1} 4 \left(\frac{k+u}{n}\right)^2 + t_p(u) + \sum_{k=p+1}^n 4 \left(1 - \frac{k+u}{n}\right)^2 - \frac{1}{3} \\ &= \begin{cases} \frac{24u^2 p - 2p + 24u^2 - 1}{3n^3} & \text{if } u \leq 1/2 \\ \frac{24u^2 p + 22p + 24u^2 - 48up + 11 - 24u}{3n^3} & \text{if } u \geq 1/2. \end{cases} \end{aligned}$$

Thus, $n^2 g_n(U/n)/4 + 1/12$ converges to U^2 if $U \leq 1/2$ and to $(1-U)^2$ otherwise, so it has the same limiting distribution as when n is even:

$$\lim_{p \rightarrow \infty} \mathbb{P}[n^2 g_n(U/n)/4 + 1/12 \leq x] = 2\sqrt{x}$$

for $0 \leq x \leq 1/4$. □

Example 9. Take $f(u) = u^3 + u^2 - 2u$. We have $f(0) = f(1) = 0$, $f'(1) - f'(0) = 5$, and f''' does not vanish. Here, $2n^2 g_n(U/n)/5 + 1/12$ has approximately the same distribution as $(U - 1/2)^2$, and $2[2n^2 g_n(U/n)/5 + 1/12]^{1/2}$ is approximately $U(0, 1)$, but this distribution is not exact. Figure 8 (right) shows the empirical distribution \hat{F} of 10^4 independent replications of $2[2n^2 g_n(U)/5 + 1/12]^{1/2}$ for one-dimensional lattices with $n = 64$ and $n = 8192$ points. We see that the approximation is extremely good. □

Example 10. Let $f(u) = B_m(u)$, the Bernoulli polynomial of degree m , for $m \geq 2$. Then $B_m^{(k)}(1) = B_m^{(k)}(0)$ for $k = 0, \dots, m - 2$ and for $k \geq m$, whereas $B_m^{(m-1)}(1) - B_m^{(m-1)}(0) = m!$. It follows that $n^m g_n(U/n) = B_m(U)$. □

4.4. Discontinuous and Unbounded Functions

The integrands f encountered in practice are often discontinuous, and sometimes they are also unbounded and have unbounded derivatives. The next examples give simple illustration of what can happen in these cases, in one dimension.




FIG 8. *Left: Exact distribution function of $ng_n(U) + 1/2$ for one-dimensional lattices with $n = 4$ points (in blue) and $n = 1024$ points (in red), for Example 7. Right: Empirical distribution \hat{F} of 10^4 replicates of $2[2n^2g_n(U)/5 + 1/12]^{1/2}$ for one-dimensional lattices with $n = 64$ points (in blue) and $n = 2^{13} = 8192$ points (in red), for Example 9.*

Example 11. Consider the step function $f(u) = 0$ for $u < a$ and $f(u) = 1$ for $u \geq a$, for some constant $a \in (0, 1)$. For a given n , let $\delta(n) = \lceil na \rceil/n - a$. The integration error with a randomly shifted lattice rule is then $g_n(U/n) = -\delta(n)$ if $U/n < 1/n - \delta(n)$, and $g_n(U/n) = 1/n - \delta(n)$ if $U/n \geq 1/n - \delta(n)$, where U is $U(0, 1)$. Thus, the error is distributed over only two possible values, and the variance is $O(n^{-2})$, unless a is a multiple of $1/n$, in which case there is no error.

To generalize this, suppose that f is piecewise constant with a jump of size d_j (positive or negative) at position a_j , for $j = 1, \dots, k$, for some integer $k > 0$. As for the case of a single jump, changing the position of a jump by any multiple of $1/n$ does not change the distribution of $g_n(U/n)$, so we can assume that all the jumps are in the interval $(0, 1/n)$, sorted in increasing order. Let $a_0 = s_0 = 0$ and $s_j = d_1 + \dots + d_j$, for $j = 1, \dots, k$, and $a_{k+1} = 1/n$. Then the average is $\mu = n \sum_{j=0}^k s_j (a_{j+1} - a_j)$ and the error $g_n(U/n)$ is $(s_j - \mu)/n$ with probability $(a_{j+1} - a_j)n$. Thus, the error is distributed over $k + 1$ possible values only, regardless of n , and the variance is again $O(n^{-2})$. However, these $k + 1$ values depend on n , and the distribution of $ng_n(U/n)$ may vary with n without converging to anything.

Even more generally, suppose f is twice continuously differentiable, except at k points a_j where it has a jump. Then f can be written as $f = f_j + f_c$, where f_j is a piecewise constant function as above and f_c is a smooth function that obeys Proposition 3. The error $g_n(U/n)$ is then the sum of two terms: an $O(1/n)$ error with a discrete distribution over k values as above, and another term obtained by applying Proposition 3 to f_c . \square

Example 12. Consider the unbounded function $f(u) = g(F^{-1}(u))$ where F is the cdf of a random variable with infinite support $[0, \infty)$. Let n_0 be a (fixed) large integer and let $b = 1 - 1/n_0$. Suppose that $n \rightarrow \infty$ while n remains a multiple of

n_0 . The error for the integral over the interval $[0, b]$ has a distribution that obeys a modified version of Proposition 3, with $f^{(k-1)}(1)$ replaced by $f^{(k-1)}(b)$, and n replaced by $n - n/n_0$. Over the last interval $[b, 1)$, the distribution is more complicated, and the error is unbounded. In fact, the error over the interval $[(n-1)/n, 1)$ is distributed as a random variable generated from the tail of the distribution F , so the total error will have a right tail that resembles the right tail of F . It will not be uniform. \square

5. Sums of One-Dimensional Functions

5.1. Setting and Relevance

To study the integration error in the general situation where $s \geq 1$, we will decompose this error in the next section as a sum of three terms, where the first term E_1 is the limiting integration error for a sum of one-dimensional functions. This E_1 is sometimes the dominant term of the error. The present section is devoted to examining the behavior of the error for such a sum. That is, we suppose here that f has the form

$$f(\mathbf{u}) = \sum_{j=1}^s a_j f_j(u_j), \quad (8)$$

where $f_j : [0, 1) \rightarrow \mathbb{R}$ and $a_j \in \mathbb{R}$ for each j . We will obtain the limiting distribution for this case by applying the results of Section 4 to each f_j , and looking at the linear combination.

Our error decomposition is in fact connected with the *functional ANOVA decomposition* of f (Liu and Owen, 2006; Owen, 1998), written as

$$f(\mathbf{u}) = \mu + \sum_{\mathbf{u} \subseteq \{1, \dots, s\}, \mathbf{u} \neq \emptyset} f_{\mathbf{u}}(\mathbf{u}) \quad (9)$$

where the $f_{\mathbf{u}} : (0, 1)^s \rightarrow \mathbb{R}$ depend only on $\{u_i, i \in \mathbf{u}\}$, are defined recursively by $f_{\emptyset} = \mu$ (a constant function) and

$$f_{\mathbf{u}}(\mathbf{u}) = \int_{(0,1)^{s-|\mathbf{u}|}} f(\mathbf{u}) d\mathbf{u}_{\bar{\mathbf{u}}} - \sum_{\mathbf{v} \subset \mathbf{u}} f_{\mathbf{v}}(\mathbf{u})$$

for $\emptyset \neq \mathbf{u} \subseteq \{1, \dots, s\}$, where the first integral is with respect to the coordinates of \mathbf{u} whose indexes are not in \mathbf{u} , denoted by $\mathbf{u}_{\bar{\mathbf{u}}}$. The $f_{\mathbf{u}}$'s integrate to zero and are orthogonal, and the variance $\sigma^2 = \int_{[0,1]^s} (f(\mathbf{u}) - \mu)^2 d\mathbf{u}$ (assumed finite) decomposes as $\sigma^2 = \sum_{\mathbf{u} \subseteq \{1, \dots, s\}} \sigma_{\mathbf{u}}^2$ where $\sigma_{\mathbf{u}}^2 = \text{Var}[f_{\mathbf{u}}(\mathbf{U})]$ for $\mathbf{U} \sim U(0, 1)^s$.

We say that f has *effective dimension d in proportion ρ in the superposition sense* (Owen, 1998) if

$$\sum_{|\mathbf{u}| \leq d} \sigma_{\mathbf{u}}^2 \geq \rho \sigma^2. \quad (10)$$

When (10) holds for $d = 1$ and ρ close to 1, the approximation of $f - \mu$ by the sum of one-dimensional functions $f_{\{j\}}$ for $j = 1, \dots, s$, in the ANOVA decomposition, accounts for most of the variance (and mean square error). Then, the higher-order functions in the decomposition have a small impact and the distribution of the error $g_n(\mathbf{U})$ would be well approximated by the distribution of the sum in (8) with $f_j = f_{\{j\}}$ and $a_j = 1$. Note that low effective dimension can sometimes be achieved by redefining f via a change of variables, without changing the expectation μ , and this often makes RQMC much more effective (Caffisch, Morokoff and Owen, 1997; Glasserman, 2004; L'Ecuyer, 2009).

5.2. Convergence to a spline

Suppose that f has the form (8) and that each f_j satisfies the assumptions of Proposition 3 for some $m = m_j \geq 1$. If $m_j = 1$ for at least one index j , then the error will converge as $O(1/n)$ and the f_j 's for which $m_j > 1$ will have no impact on the asymptotic distribution, because their error contribution converges at a faster rate. So to simplify the exposition, suppose that $m_j = 1$ for all j (the other f_j 's have been removed) and that each f_j has an integrable second derivative. We can then apply Corollary 4(i) to each f_j : It says that if $g_{j,n}(U_j)$ denotes the corresponding integration error, then $n g_{j,n}(U_j/n)$ converges to $[f_j(1) - f_j(0)](U_j - 1/2)$. But the integration error function for f can be written as

$$g_n(\mathbf{u}) = \sum_{j=1}^s a_j g_{j,n}(u_j).$$

This means that the asymptotic distribution of $n g_n(\mathbf{U})$ is the same as the distribution of

$$W = \sum_{j=1}^s a_j [f_j(1) - f_j(0)](U_j - 1/2) = \sum_{j=1}^s b_j (U_j - 1/2), \quad (11)$$

where $b_j = a_j [f_j(1) - f_j(0)]$. This W is a linear combination of independent uniform random variables over the interval $(-1/2, 1/2)$. Theorem 1 of Barrow and Smith (1979), restated in the next proposition, tells us that the (exact) cumulative distribution function of W is a non-decreasing spline of degree s , with $s - 1$ continuous derivatives, and its support is the finite interval $[-b, b]$, where $b = \sum_{j=1}^s |b_j|/2$.

Proposition 5. *The exact distribution function of W is given by*

$$\mathbb{P}[W \leq w] = \frac{1}{s! \prod_{j=1}^s b_j} \sum_{\mathbf{u} \subseteq \{1, \dots, s\}} (-1)^{|\mathbf{u}|} \left[\max \left(0, w + b - \sum_{j \in \mathbf{u}} b_j \right) \right]^s \quad (12)$$

for $-b \leq w \leq b$. Its density is given by a spline of degree $s - 1$:

$$f_w(w) = \frac{1}{(s-1)! \prod_{j=1}^s b_j} \sum_{\mathbf{u} \subseteq \{1, \dots, s\}} (-1)^{|\mathbf{u}|} \left[\max \left(0, w + b - \sum_{j \in \mathbf{u}} b_j \right) \right]^{s-1}$$

where $0^0 = 0$ by convention.

Example 13. As an illustration, let $s = 2$ and $f(u_1, u_2) = 5u_1 + 3u_2 - 4 = 5B_1(u_1) + 3B_1(u_2)$. The constant -4 has no effect on the error. Here, we have $b_j = a_j$, $b = 4$, and the distribution function F_w of $W = ng_n(\mathbf{U}) = ng_n(U_1, U_2)$ (here we have equality) is a spline of degree 2 defined by

$$F_w(w) = \begin{cases} (w+4)^2/30 & \text{for } -4 \leq w \leq -1, \\ (6w+15)/30 & \text{for } -1 \leq w \leq 1, \\ -(w+1)^2/30 + (w-1)/3 + 25/30 & \text{for } 1 < w \leq 4. \end{cases}$$

The corresponding density is a ‘‘table mountain’’ function, defined by

$$f_w(w) = \begin{cases} (4+w)/15 & \text{for } -4 \leq w \leq -1, \\ 1/5 & \text{for } -1 \leq w \leq 1, \\ (4-w)/15 & \text{for } 1 < w \leq 4. \end{cases}$$

□

5.3. Asymptotic Distribution when $s \rightarrow \infty$

If f has an additive form as in (8) and $s \rightarrow \infty$ while n remains fixed, then $g_n(\mathbf{U})$ is a sum of s independent random variables and its properly standardized version may obey a CLT when $s \rightarrow \infty$, under standard conditions on the variances of the terms of the sum. The following version of the Berry–Essen theorem provides simple sufficient conditions for the CLT to hold, together with an explicit bound on the approximation error for the distribution function. It is a direct adaptation of Petrov (1995, Theorem 5.7) reformulated for our setting.

Proposition 6. For $0 < \delta \leq 1$ and $j = 1, \dots, s$, let $X_j = a_j g_{j,n}(U_j)$, $\sigma_j^2 = \text{Var}[X_j] = a_j^2 \mathbb{E}[g_{j,n}^2(U_j)]$, and

$$\nu_j^{2+\delta} = \mathbb{E}[|X_j|^{2+\delta}] = |a_j|^{2+\delta} \mathbb{E}[|g_{j,n}(U_j)|^{2+\delta}].$$

Define $v_s^2 = \sigma_1^2 + \dots + \sigma_s^2$, $w_s^{2+\delta} = \nu_1^{2+\delta} + \dots + \nu_s^{2+\delta}$,

$$Y_s = \frac{X_1 + \dots + X_s}{v_s} = \frac{g_n(\mathbf{U})}{v_s},$$

and let $F_s(x) = \mathbb{P}[Y_s \leq x] = \mathbb{P}[g_n(\mathbf{U})/v_s \leq x]$. Then,

$$\sup_{s \geq 1, x \in \mathbb{R}} |F_s(x) - \Phi(x)| \leq K(\delta)(w_s/v_s)^{2+\delta}. \quad (13)$$

The CLT holds (for fixed n and $s \rightarrow \infty$) if $\lim_{s \rightarrow \infty} w_s/v_s = 0$.

The following condition is less general, but simpler.

Corollary 7. The CLT holds (for fixed n and $s \rightarrow \infty$) if

$$\lim_{s \rightarrow \infty} \frac{1}{v_s} \sup_{j \in \{1, \dots, s\}} |a_j| \sup_{u \in (0,1)} |g_{j,n}(u_j)| = 0.$$

Proof. Because

$$\mathbb{E}[|g_{j,n}(U_j)|^{2+\delta}] \leq \mathbb{E}[g_{j,n}^2(U_j)] \sup_{\mathbf{u} \in (0,1)^s} |g_{j,n}(u_j)|^\delta,$$

we have $v_j^{2+\delta} \leq c_j^\delta \sigma_j^2$, where $c_j = |a_j| \sup_{u \in (0,1)} |g_{j,n}(u_j)|$. Therefore,

$$w_s^{2+\delta} \leq \sum_{j=1}^s c_j^\delta \sigma_j^2 \leq v_s^2 \sup_{j \in \{1, \dots, s\}} c_j^\delta,$$

which we rewrite as

$$\left(\frac{w_s}{v_s}\right)^{1+2/\delta} \leq \frac{1}{v_s} \sup_{j \in \{1, \dots, s\}} c_j.$$

If the right side vanishes when $s \rightarrow \infty$, then w_s/v_s also vanishes. \square

Corollary 8. *If $f_j = f_1$ for all j , the CLT holds (for fixed n and $s \rightarrow \infty$) if*

$$\lim_{s \rightarrow \infty} \frac{\sup_{j \in \{1, \dots, s\}} a_j^2}{\sum_{j=1}^s a_j^2} = 0.$$

Proof. Because the contribution from f_j is the same in every term in the sum, it can be factorized out of the sum, then just ignored because the right-hand side is zero. A similar argument holds for the supremum in the numerator. \square

Example 14. Suppose $f_j(u_j) = B_{m_j}(u_j)$, the Bernoulli polynomial of degree m_j , so that (8) becomes

$$f(\mathbf{u}) = \sum_{j=1}^s a_j B_{m_j}(u_j).$$

Example 10 told us that the error function for the j th term is $g_{j,n}(u_j) = n^{-m_j} B_{m_j}(u_j)$, so the error function for the sum is

$$g_n(\mathbf{u}) = \sum_{j=1}^s a_j n^{-m_j} B_{m_j}(u_j).$$

We also have

$$v_s^2 = \sum_{j=1}^s a_j^2 n^{-2m_j} (m_j!)^2 |b_{2m_j}| / (2m_j)!$$

where b_k is the k -th Bernoulli number, which can be defined as

$$b_k = \sum_{j=0}^k \sum_{\ell=0}^j (-1)^\ell \binom{j}{\ell} \frac{\ell^k}{j+1},$$

and

$$w_s^{2+\delta} = \sum_{j=1}^s |a_j|^{2+\delta} n^{-(2+\delta)m_j} \mathbb{E}[|B_{m_j}(U)|^{2+\delta}].$$

Corollary 7 implies that if

$$\lim_{s \rightarrow \infty} \frac{\sup_{j \in \{1, \dots, s\}} \sup_{u \in (0,1)} a_j^2 n^{-2m_j} B_{m_j}^2(u)}{\sum_{j=1}^s a_j^2 n^{-2m_j} (m_j!)^2 |b_{2m_j}| / (2m_j)!} = 0, \quad (14)$$

then the CLT holds for the standardized version of $g_n(\mathbf{U})$.

It is known (Abramowitz and Stegun, 1970, page 805) that $|b_k| > 2k!/(2\pi)^k$ for $k = 2, 4, \dots$ and that $|B_k(u)| < 2k!/(2\pi)^k(1 - 2^{1-k}) < 4k!/(2\pi)^k$ for $u \in (0, 1)$ and $k \geq 2$. The latter inequality also holds for $k = 1$, since $|B_1(u)| < 1/2 < 4/2\pi$ for $u \in (0, 1)$. Using these inequalities, we can derive an upper bound for the left-hand side of (14), which implies that the CLT holds under the following sufficient condition (in which we have dropped the constant factors):

$$\lim_{s \rightarrow \infty} \frac{\sup_{j \in \{1, \dots, s\}} a_j^2 (2\pi n)^{-2m_j} (m_j!)^2}{\sum_{j=1}^s a_j^2 (2\pi n)^{-2m_j} (m_j!)^2} = 0. \quad (15)$$

For the special case where $m_j = k$ (a positive integer) for all j , (15) is equivalent to Corollary 8. \square

Example 15. As a special case of Example 14, let f be a linear combination of Bernoulli polynomials of degree 1:

$$f(\mathbf{u}) = \sum_{j=1}^s a_j (u_j - 1/2). \quad (16)$$

If $a_j = j^k$ for all j , where $k \geq 0$ is fixed, Corollary 8 applies, because

$$\lim_{s \rightarrow \infty} \frac{s^{2k}}{\sum_{j=1}^s j^{2k}} \leq \lim_{s \rightarrow \infty} \frac{s^{2k}}{\int_0^s u^{2k} du} = \lim_{s \rightarrow \infty} \frac{2k+1}{s} = 0.$$

If $a_j = \rho^j$ for all j , where $0 < \rho < 1$, then

$$(w_s/v_s)^{2+\delta} = c(\delta) \frac{\sum_{j=1}^s \rho^{(2+\delta)j}}{\left(\sum_{j=1}^s \rho^{2j}\right)^{1+\delta/2}} = c(\delta) \frac{1 - \rho^{(2+\delta)(s+1)}}{(1 - \rho^{2(s+1)})^{1+\delta/2}} \rightarrow c(\delta) \neq 0$$

when $s \rightarrow \infty$, where $c(\delta) = 3^{1+\delta/2}/(3 + \delta)$, so Corollary 8 does not apply. One can show that the CLT does not hold in this case; the a_j 's are converging to 0 too fast.

5.4. Normal P-P plots for numerical examples.

We report a few representative illustrations of simulation experiments we made to assess the normality of the error. For these experiments, and all other numerical examples in the rest of the paper (unless otherwise indicated), we use a rank-1 lattice rule in 20 dimensions with $n = 8192$, $\mathbf{v}_1 = \mathbf{a}_1/n$ where

$$\begin{aligned} \mathbf{a}_1 = & (1, 2431, 3739, 1689, 3185, 2609, 4343, 1525, 71, 6083, \\ & 2585, 1583, 5385, 3359, 7759, 3055, 2627, 4535, 4475, 1623), \end{aligned}$$

and $\mathbf{v}_j = \mathbf{e}_j$ (the j th unit vector) for $j = 2, \dots, 20$. When the integrand has $s < 20$ dimensions, we just use the first s coordinates of each point.

In our experiments, we compute m independent replicates of the RQMC estimator $\hat{\mu}_{n,\text{rqmc}}$ in (3) and we subtract the global average (or the exact mean μ when we know it) to each replicate to obtain estimates of the corresponding replicates of $g_n(\mathbf{U})$, say $g_n(\mathbf{U}_1), \dots, g_n(\mathbf{U}_m)$. We then compute their empirical variance $S_{n,m}^2$ and the standardized errors $Z_{n,j} = g_n(\mathbf{U}_j)/S_{n,m}$ for $j = 1, \dots, m$. Finally, we plot the empirical distribution of these standardized errors, defined by

$$\hat{F}_{Z_{n,m}}(z) = \frac{1}{m} \sum_{j=1}^m \mathbb{I}[Z_{n,j} \leq z]$$

against $\Phi(z)$, where Φ is the standard normal distribution function. Unless otherwise specified, we take $m = 10^4$, in order to obtain good accuracy in the estimation of the error distribution.

Example 16. Figure 9 shows normal P-P plots of the error for a function f defined by (16) with $a_j = j^2$ and with $a_j = 2^{-j}$, as in Example 15. As expected, the P-P plots indicate a nearly normal distribution when s is large enough in the first case, for which the CLT holds, but not for the second case. We made a similar plot for $a_j = 1$ (not shown), and the fit was already excellent for $s = 2$ and looked perfect (to the eye) for $s = 10$ and 20. We also performed similar experiments by replacing the first-degree Bernoulli polynomials by second-degree or by third-degree Bernoulli polynomials, and we observed similar behaviors for the same three choices of coefficients a_j . \square




FIG 9. Normal P-P plot of the error for Example 16, with f as in (16) with $a_j = j^2$ (left) and with $a_j = 2^{-j}$ (right). On the left, the curves with $s = 10$ and $s = 20$ almost superpose each other, while on the right, all three curves almost superpose each other.

Example 17. Another interesting special case of Example 14 is when f is a mixed sum of first-degree and of second-degree Bernoulli polynomials with

$a_j = 1$ for all j :

$$f(\mathbf{u}) = \sum_{j=1}^r B_1(u_j) + \sum_{j=r+1}^s B_2(u_j) \quad (17)$$

for some $r < s$. Figure 10 shows normal P-P plots of error for functions f defined as in (17) in $s = 2, 10$ and 20 dimensions, for $r = s - 1$ and for $r = 1$. In both cases, the error $g_n(\mathbf{U})$ is dominated by the errors on the polynomials of degree 1. In the first case ($r = s - 1$), when $s = 10$ or 20 , this dominant part of the error is a sum of several independent terms and the normal approximation provides a very good fit (the P-P plots look perfect) for this reason. For $s = 2$, the normal approximation is not good, as expected. In the second case ($r = 1$), the P-P plots show a bad fit (and almost superpose each other) for all values of s . This can be explained by the fact that the dominant part of the error comes from a single polynomial of degree 1 for all s . By looking at this plot, one may conclude that the CLT does not hold here if $r = 1$ and $s \rightarrow \infty$. But the expression in (15) in that case becomes

$$\lim_{s \rightarrow \infty} \frac{\max[(2\pi n)^{-2}, 4(2\pi n)^{-4}]}{(2\pi n)^{-2} + 4(s-1)(2\pi n)^{-4}} = \frac{1}{1 + (s-1)(\pi n)^{-2}} = 0, \quad (18)$$

so the CLT does apply, but the convergence in s is very slow and does not show up in Figure 10. \square




FIG 10. Normal P-P plot of the error for Example 17, with f as in (17) with $r = s - 1$ (left) and $r = 1$ (right). The behavior of all the curves is somewhat similar to that in Figure 9, except that the non-normality on the right panel is more accentuated.

We say that f has effective dimension d in proportion ρ in the truncation sense (Owen, 1998) if

$$\sum_{\mathbf{u} \subseteq \{1, \dots, d\}} \sigma_{\mathbf{u}}^2 \geq \rho \sigma^2. \quad (19)$$

If (19) holds for small d and ρ close to 1, this means that f depends almost only on the first few variables (or random numbers) U_j . A low effective dimension in the truncation sense is incompatible with meeting (even approximately) the conditions of Proposition 6. Various techniques have been designed to transform the integrand f to reduce its effective dimension in the truncation sense in order to improve the effectiveness of RQMC (L'Ecuyer, 2009). The PCA method in Section 7 is one such technique. One side effect of these techniques is that they tend to reduce significantly the normality of the RQMC estimator. On the other hand, a low effective dimension in the superposition sense is not incompatible with the conditions of Proposition 6. In fact, all the integrands in Examples 16 and 17 have effective dimension 1 in the superposition sense and yet the proposition applies.

6. Error Distribution for Multi-Dimensional Integrals

6.1. A First-Order Error Decomposition

For a general s -dimensional integrand f , we will separate the sum of its one-dimensional ANOVA components from the rest, and apply the results of Section 5 to this sum. If $f_{\{j\}}$ is the j th one-dimensional ANOVA component of f and if each $f_{\{j\}}$ has an integrable second derivative, then the integration error for the sum is a random variable whose distribution can be approximated as in (11) by that of

$$E_1 = E_1(\mathbf{U}) = \sum_{j=1}^s \left(U_j - \frac{1}{2} \right) \frac{f_{\{j\}}(1) - f_{\{j\}}(0)}{n}.$$

This provides a first-order approximation for the situation where the one-dimensional ANOVA components are not all c -periodic. The importance of the contribution of E_1 to the total error $g_n(\mathbf{U})$ will depend on f . On the other hand, it does not depend on the choice of lattice, because all integration lattices have the same one-dimensional projections.

From a different perspective, we have seen that when we randomize the points, we can first choose a small hypercube $[0, 1/n]^s + \mathbf{K}/n$ at random, where \mathbf{K} is a random vector uniformly distributed in $\{0, \dots, n-1\}^s$, and then sample a random shift \mathbf{U}/n uniformly inside that hypercube, where $\mathbf{U} \sim U(0, 1)^s$. Equivalently, we can sample \mathbf{K} uniformly over a restricted subset of n^{s-1} distinct hypercubes that satisfy the conditions of Proposition 1, instead of over all n^s hypercubes. The corresponding error is

$$\tilde{g}_n(\mathbf{K}, \mathbf{U}) = \frac{1}{n} \sum_{i=0}^{n-1} f((\mathbf{u}_i + \mathbf{K}/n) \bmod 1 + \mathbf{U}/n) - \mu \quad (20)$$

and the expected error conditional on \mathbf{K} is

$$E_2 = E_2(\mathbf{K}) = \int_{[0,1]^s} \tilde{g}_n(\mathbf{K}, \mathbf{u}) \, d\mathbf{u}.$$

This E_2 depends only on the choice of \mathbf{K} and represents the error due to the difference of means across the small hypercubes. In contrast, E_1 does not depend on \mathbf{K} and depends only on the random shift \mathbf{U}/n . Thus, E_1 and E_2 are two independent components of the error $g_n(\mathbf{U})$.

An equivalent way of viewing this is to write the first terms of the ANOVA decomposition of $\tilde{g}_n(\mathbf{K}, \mathbf{U})$ after conditioning on \mathbf{K} :

$$\begin{aligned} \tilde{g}_n(\mathbf{K}, \mathbf{U}) &= \tilde{g}_{n,\emptyset}(\mathbf{K}) + \frac{1}{s} \sum_{j=1}^s \tilde{g}_{n,\{j\}}(\mathbf{K}, \mathbf{U}) + \sum_{\substack{\mathbf{u} \subseteq \{1, \dots, s\}, \\ |\mathbf{u}| \geq 2}} \tilde{g}_{n,\mathbf{u}}(\mathbf{K}, \mathbf{U}) \\ &= E_2(\mathbf{K}) + E_1(\mathbf{U}) + O(n^{-2}) + \sum_{\substack{\mathbf{u} \subseteq \{1, \dots, s\}, \\ |\mathbf{u}| \geq 2}} \tilde{g}_{n,\mathbf{u}}(\mathbf{K}, \mathbf{U}), \end{aligned}$$

where we have used the fact (easy to verify) that the one-dimensional ANOVA components conditional on \mathbf{K} , $\tilde{g}_{n,\{j\}}(\mathbf{K}, \mathbf{U})$, do not depend on \mathbf{K} . Let

$$E_3 = E_3(\mathbf{K}, \mathbf{U}) = \tilde{g}_n(\mathbf{K}, \mathbf{U}) - E_1 - E_2$$

be the remainder of the error. We will show in what follows that when f is smooth enough, E_3 becomes negligible compared with E_1 and E_2 when n increases, and E_2 is typically the dominant part of the error when $n \rightarrow \infty$.

6.2. Convergence of the Error Terms

We see from its definition that $|E_1|$ converges as $O(n^{-1})$ (unless it is zero), and that E_1 depends neither on \mathbf{K} nor on the choice of lattice. Its variance is

$$\text{Var}[E_1] = \frac{1}{12n^2} \sum_{j=1}^s [f_{\{j\}}(1) - f_{\{j\}}(0)]^2.$$

In contrast, it is difficult to obtain a simple formula for the variance of E_2 , because it depends on both f and the lattice. We saw earlier that if f is smooth enough, namely if all its mixed partial derivatives of order one are absolutely integrable, then for any $\epsilon > 0$, it is easy to find rank-1 lattice rules for all n , such that for any shift \mathbf{U} , the integration error $g_n(\mathbf{U})$ converges as $O(n^{-1+\epsilon})$. Then, $|E_2|$ and $|E_3|$ will converge at that speed or faster.

Proposition 10 below shows that under a slightly stronger smoothness assumption on f , there are lattice rules for which $|E_3|$ converges almost as $O(n^{-2})$. In fact, this convergence holds on average over rank-1 lattice rules, so it represents the typical behavior as soon as we make some effort to select a good rule. This means that E_3 would typically become negligible compared with $E_1 + E_2$ when $n \rightarrow \infty$. We then obtain $|E_2| = O(n^{-1+\epsilon})$ while $|E_1| = O(n^{-1})$, so E_2 should eventually dominate when $n \rightarrow \infty$. For a typical (moderate) n , however, E_1 may still dominate, and the distribution of $n(E_1 + E_2)$ can be taken as an approximation of the distribution of $n\tilde{g}_n(\mathbf{K}, \mathbf{U})$. We will use the following lemma in the proof.

Lemma 9. *Suppose that f has integrable mixed partial derivatives of second order over $[0, 1]^s$. Then, for any $\epsilon > 0$, there exist rank-1 lattice rules for which*

$$\sup_{\mathbf{K}, \mathbf{U}} \frac{\partial \tilde{g}_n(\mathbf{K}, \mathbf{u})}{\partial u_j} - \frac{f_{\{j\}}(1) - f_{\{j\}}(0)}{n} = O(n^{-2+\epsilon}),$$

and this property actually holds on average over all rank-1 lattice rules.

Proof. Let f_j denote the partial derivative of f with respect to its j th coordinate. The partial derivative of $\tilde{g}_n(\mathbf{K}, \mathbf{u})$ with respect to the j th coordinate of \mathbf{u} can be expanded as

$$\frac{\partial \tilde{g}_n(\mathbf{K}, \mathbf{u})}{\partial u_j} = \frac{1}{n} \sum_{i=0}^{n-1} \frac{f_j((\mathbf{u}_i + \mathbf{K}/n) \bmod 1 + \mathbf{u}/n)}{n}.$$

From the theory of lattice rules, as we have seen earlier, if all the mixed partial derivatives of first order of f_j are absolutely integrable, then for any $\epsilon > 0$, and for all n , on average over all rank-1 lattice rules, the worst-case integration error on f_j by the lattice rule converges as $O(n^{-1+\epsilon})$ uniformly in the shift \mathbf{U} . If we use a rule at least as good as the average, we obtain

$$\begin{aligned} \frac{1}{n} \sum_{i=0}^{n-1} \frac{f_j((\mathbf{u}_i + \mathbf{K}/n) \bmod 1 + \mathbf{u}/n)}{n} &= \frac{1}{n} \int_{[0,1]^s} f_j(\mathbf{u}) d\mathbf{u} + \frac{1}{n} O(n^{-1+\epsilon}) \\ &= \frac{f_{\{j\}}(1) - f_{\{j\}}(0)}{n} + O(n^{-2+\epsilon}). \end{aligned}$$

□

Proposition 10. *If f has integrable mixed partial derivatives of second order over $[0, 1]^s$, then for any $\epsilon > 0$, there exist rank-1 lattice rules for which $E_3(\mathbf{K}, \mathbf{U}) = O(n^{-2+\epsilon})$ uniformly in \mathbf{K} and \mathbf{U} , and this holds on average over all rank-1 lattice rules.*

Proof. Fix the components of \mathbf{u} equal to those of \mathbf{U} , except for the first one, and denote $\mathbf{U}' = (U_2, \dots, U_s)$ and $\mathbf{U}'' = (U_3, \dots, U_s)$. By applying Lemma 12 of the appendix with $m = 1$ to the function $F(u_1) = \tilde{g}_n(\mathbf{K}, u_1, \mathbf{U}')$, given that this function has two integrable derivatives, we obtain that for all $U_1 \in (0, 1)$,

$$\begin{aligned} \int_0^1 \tilde{g}_n(\mathbf{K}, u_1, \mathbf{U}') du_1 &= \tilde{g}_n(\mathbf{K}, U_1, \mathbf{U}') \\ &\quad - (U_1 - 1/2) [\tilde{g}_n(\mathbf{K}, 1, \mathbf{U}') - \tilde{g}_n(\mathbf{K}, 0, \mathbf{U}')] + O(n^{-2}). \end{aligned} \quad (21)$$

In the $O(n^{-2})$ term here and in all the $O(n^{-2})$ terms that follow, the implied constants can be taken independent of \mathbf{K} and \mathbf{U} . The error difference in the square brackets on the right can be approximated by

$$\frac{\partial \tilde{g}_n(\mathbf{K}, \mathbf{U})}{\partial U_1} + O(n^{-2}) = \frac{f_{\{1\}}(1) - f_{\{1\}}(0)}{n} + O(n^{-2+\epsilon}),$$

where we have used, for the left side, the fact that the second-order partial derivatives of \tilde{g}_n are $O(n^{-2})$, and Lemma 9 for the right side. Substituting this approximation into (21), we find

$$\int_0^1 \tilde{g}_n(\mathbf{K}, u_1, \mathbf{U}') du_1 = \tilde{g}_n(\mathbf{K}, U_1, \mathbf{U}') - (U_1 - 1/2) \frac{f_{\{1\}}(1) - f_{\{1\}}(0)}{n} + O(n^{-2+\epsilon}). \quad (22)$$

Next, we let u_2 vary again and integrate with respect to u_2 :

$$\int_0^1 \int_0^1 \tilde{g}_n(\mathbf{K}, u_1, u_2, \mathbf{U}'') du_1 du_2 = \int_0^1 \tilde{g}_n(\mathbf{K}, U_1, u_2, \mathbf{U}'') du_2 - (U_1 - 1/2) \frac{f_{\{1\}}(1) - f_{\{1\}}(0)}{n} + O(n^{-2+\epsilon}).$$

Applying (22) to the integral on the right-hand side of the above equation, but with u_2 taking the place of u_1 , we obtain

$$\int_0^1 \int_0^1 \tilde{g}_n(\mathbf{K}, u_1, u_2, \mathbf{U}'') du_1 du_2 = \tilde{g}_n(\mathbf{K}, U_1, U_2, \mathbf{U}'') - (U_1 - 1/2) \frac{f_{\{1\}}(1) - f_{\{1\}}(0)}{n} - (U_2 - 1/2) \frac{f_{\{2\}}(1) - f_{\{2\}}(0)}{n} + O(n^{-2+\epsilon}).$$

This procedure can be applied repeatedly up to the last coordinate and this yields the result. \square

We just saw that with well-chosen lattice rules, E_2 becomes the dominating part of the error when $n \rightarrow \infty$. However, we do not have an explicit form for the limiting distribution of E_2 , and it seems unlikely that a practically useful form (for computing confidence intervals in general) can be obtained. In fact, it seems that the distribution of E_2 could be almost arbitrary. More specifically, for any pre-specified continuous cdf F with finite support, by taking n large enough and picking any lattice rule with n points, one can construct a smooth function f for which the cdf of $E_2 = E_2(\mathbf{K})$ is arbitrarily close to F . It suffices to specify f so that the average error $E_2(\mathbf{K})$ takes the values $F^{-1}((j + 1/2)n^{-s+1})$ for $j = 0, \dots, n^{s-1}$ when \mathbf{K} runs through its n^{s-1} possibilities, and then smooth the function across the boundaries of the small hypercubes. This means that for large n , the error $\tilde{g}_n(\mathbf{K}, \mathbf{U})$ can have (approximately) essentially any distribution. On the other hand, it is unclear if this can be done for a fixed function f with integrable mixed partial derivatives of order two, and independent of n .

6.3. Two-Dimensional Examples

Figure 11 compares the cdf of nE_2 (in red) and that of $ng_n(\mathbf{U})$ (in blue) for Examples 4 and 5, where $n = 8$. In these figures, the cdf of nE_2 was computed exactly whereas that of $ng_n(\mathbf{U})$ is an empirical version obtained from 10^4 independent replications.

For Example 4, it turns out that $E_1(\mathbf{U}) = 0$ for all \mathbf{U} , so $\tilde{g}_n(K, \mathbf{U}) - E_2(K) = O(n^{-2+\epsilon})$, which explains the good agreement between the two distributions plotted in Figure 11 (top).

For Example 5, we have $E_1(\mathbf{U}) = -(U_1 - 1/2)/(2n) + (U_2 - 1/2)/(3n)$, whose possible values range from $-5/(12n)$ to $5/(12n)$, and whose cdf is a spline of second degree, according to Proposition 5. That is, for any $k \in \{0, \dots, n - 1\}$, $n\tilde{g}_n(k, \mathbf{U})$ can depart from its average in the small square by up to $5/12 \approx 0.417$ (if we neglect the $O(n^{-2+\epsilon})$ term), which is more than twice the largest distance between any two consecutive jumps of the cdf of nE_2 , which is 0.200 (the distance between the first and last jumps of that cdf is approximately 0.571). This explains the fact that the cdf of $ng_n(\mathbf{U})$ here is not well approximated by that of nE_2 . It is much smoother. If we increase n , the steps of nE_2 become smaller and its discrete distribution gets closer to the empirical error distribution, but the gap between these two distributions does not converge to zero very quickly.




FIG 11. Exact cdf of nE_2 (piecewise constant function) and empirical cdf of $ng_n(\mathbf{U})$ for Example 4 (top) and Example 5 (bottom), with $n = 8$.

6.4. Error Distribution via Fourier Expansion

6.4.1. The General Case

We saw earlier that the Fourier expansion of $g_n(\mathbf{u})$, for an arbitrary shift $\mathbf{u} \in [0, 1)^s$, is

$$g_n(\mathbf{u}) = \sum_{\mathbf{0} \neq \mathbf{h} \in L_s^*} \hat{f}(\mathbf{h}) \exp(2\pi i \mathbf{h}^t \mathbf{u}). \quad (23)$$

Let $\mathbf{w}_1, \dots, \mathbf{w}_s$ be a basis of the dual lattice L_s^* and let \mathbf{W} be an $s \times s$ matrix with columns $\mathbf{w}_1, \dots, \mathbf{w}_s$. These vectors have only integer coordinates. Thus, L_s^* is the set of all integer vectors \mathbf{h} that can be written as

$$\mathbf{h} = \mathbf{W}\mathbf{z} = \sum_{j=1}^s z_j \mathbf{w}_j$$

for $\mathbf{z} = (z_1, \dots, z_s)^t \in \mathbb{Z}^s$. By summing over the vectors \mathbf{z} instead of the vectors \mathbf{h} (a change of variable), we can rewrite

$$\begin{aligned} g_n(\mathbf{u}) &= \sum_{\mathbf{0} \neq \mathbf{z} \in \mathbb{Z}^s} \hat{f}(\mathbf{W}\mathbf{z}) \exp(2\pi i \mathbf{u}^t \mathbf{W}\mathbf{z}) \\ &= \sum_{\mathbf{0} \neq \mathbf{z} \in \mathbb{Z}^s} \hat{v}_n(\mathbf{z}) \exp(2\pi i \mathbf{u}^t \mathbf{W}\mathbf{z}) = v_n(\mathbf{u}^t \mathbf{W}) \end{aligned} \quad (24)$$

for some function v_n whose Fourier coefficients are $\hat{v}_n(\mathbf{z}) = \hat{f}(\mathbf{W}\mathbf{z})$. This function v_n depends on n and on the selected lattice, via \mathbf{W} .

If we can recover v_n explicitly from its Fourier coefficients, this would give an explicit formula for $g_n(\mathbf{u})$, the error as a function of the shift. More generally, we may be able to approximate from this expression at least the most important Fourier coefficients in the expansion of g_n , and eventually from this, the distribution of $g_n(\mathbf{U})$.

Note that there is a link between this function v_n and the generation of the shift uniformly in the parallelotope P . Indeed, if \mathbf{U} has the uniform distribution over P , then we can write $\mathbf{U} = \mathbf{V}\tilde{\mathbf{U}}$ where $\tilde{\mathbf{U}} \sim U(0, 1)^s$ and \mathbf{V} is the $s \times s$ matrix whose columns are the basis vectors $\mathbf{v}_1, \dots, \mathbf{v}_s$. Then, $\mathbf{W}^t \mathbf{V} = \mathbf{I}$, and therefore $g_n(\mathbf{U}) = v_n(\mathbf{W}^t \mathbf{U}) = v_n(\mathbf{W}^t \mathbf{V} \tilde{\mathbf{U}}) = v_n(\tilde{\mathbf{U}})$. This implies that if $\mathbf{U} \sim U(0, 1)^s$, then $g_n(\mathbf{U})$ and $v_n(\mathbf{U})$ have the same distribution (but they are not the same function of \mathbf{U}).

6.4.2. A Few Multidimensional Examples

Example 18. In one dimension, the dual lattice contains all multiples of n . Thus, the matrix \mathbf{W} is one by one and contains the single entry n . The Fourier coefficients of v_n are then $\hat{v}_n(z) = \hat{f}(nz)$, for $z \in \mathbb{Z}$, and we have $g_n(u) = v_n(nu)$.

As a special case, if f is the k th Bernoulli polynomial B_k , then its Fourier expansion is

$$B_k(u) = \sum_{0 \neq h \in \mathbb{Z}} \frac{-k!}{(2\pi i h)^k} e^{2\pi i h u}$$

for $u \in (0, 1)$, and we have

$$\hat{v}_n(z) = \frac{-k!}{(2\pi i n z)^k} = \hat{B}_k(z) n^{-k}.$$

We can then conclude that $v_n(u) = n^{-k} B_k(u)$ and $g_n(u) = v_n(nu) = n^{-k} B_k(nu \bmod 1)$ for all $u \in (0, 1)$. But if U is $U(0, 1)$, then $nU \bmod 1$ is also $U(0, 1)$, and therefore $n^k g_n(U)$ has exactly the same distribution as $B_k(U)$. This agrees with Example 10. \square

Example 19. Suppose that $f : [0, 1]^s \rightarrow \mathbb{R}$ is a product of Bernoulli polynomials:

$$f(\mathbf{u}) = f(u_1, \dots, u_s) = \prod_{j=1}^s B_{m_j}(u_j).$$

Its Fourier expansion is then

$$f(\mathbf{u}) = \sum_{0 \neq \mathbf{h} \in \mathbb{Z}} \prod_{j=1}^s \frac{-m_j!}{(2\pi i h_j)^{m_j}} e^{2\pi i h_j u_j} = \frac{(-1)^s m_1! \cdots m_s!}{(2\pi i)^{m_1 + \cdots + m_s}} \sum_{0 \neq \mathbf{h} \in \mathbb{Z}} \frac{1}{h_1^{m_1} \cdots h_s^{m_s}} e^{2\pi i \mathbf{h}^t \mathbf{u}}$$

for $\mathbf{u} = (u_1, \dots, u_s)^t \in (0, 1)^s$. In this case, the Fourier coefficient $\hat{v}_n(\mathbf{z}) = \hat{f}(\mathbf{W}\mathbf{z})$ can be written as

$$\hat{v}_n(\mathbf{z}) = \prod_{j=1}^s \frac{-m_j!}{(2\pi i)^{m_j} (\mathbf{w}^{(j)} \mathbf{z})^{m_j}},$$

where $\mathbf{w}^{(j)}$ is the j -th row of the matrix \mathbf{W} . Here, $(\mathbf{w}^{(j)} \mathbf{z})^{m_j}$ is a multivariate polynomial in \mathbf{z} . Example 20 illustrates the specific case where $m_j = 1$ for all j . \square

Example 20. Figure 12 shows a normal P-P plot of the error for

$$f(\mathbf{u}) = \prod_{j=1}^s (u_j - 1/2), \tag{25}$$

The discrepancy from the normal distribution becomes more obvious for $s = 20$. For this special case, for $s = 2$, the Fourier coefficients of the error function for a rank-1 lattice rule can be written explicitly as

$$\hat{v}_n(z_1, z_2) = \frac{-1}{4\pi^2 n z_1 (-a_1 z_1 + z_2)}.$$

\square




FIG 12. Normal P-P plot of the error for Example 20. The departure from normality is clearly visible only for $s = 20$.

6.5. Sampling the Error in the Dual Space

So far, we saw asymptotic distributions and convergence conditions for the error. Here we point out that expansion formula in (24) can sometimes be used to estimate the error distribution. The idea is to compute a truncated version of this expansion,

$$\tilde{g}_n(\mathbf{u}) = \sum_{0 \neq \mathbf{z} \in \tilde{\mathcal{Z}}} \hat{f}(\mathbf{W}\mathbf{z}) \exp(2\pi i \mathbf{z}^t \mathbf{W}^t \mathbf{u}), \quad (26)$$

where $\tilde{\mathcal{Z}} = \{\mathbf{z} = (z_1, \dots, z_s) : -z_{j,0} \leq z_j \leq z_{j,0}\}$ for some positive integers $z_{j,1}, \dots, z_{j,s}$. If we know the Fourier coefficients $\hat{f}(\mathbf{h})$ of f , then the error distribution can be estimated by the empirical distribution of m independent replicates of $\tilde{g}(\mathbf{U})$ instead of m independent replicates of $\hat{\mu}_{n, \text{rqmc}}$ as explained in Subsection 5.4. This could be advantageous if n is large and $|\hat{f}(\mathbf{W}\mathbf{z})|$ converges quickly to zero as $\|\mathbf{z}\|$ increases. This convergence speed may depend on the choice of \mathbf{W} , i.e., the choice of dual basis.

Example 21. Consider the two five-dimensional functions

$$f_1(\mathbf{u}) = \prod_{j=1}^5 B_1(u_j) \quad \text{and} \quad f_2(\mathbf{u}) = \prod_{j=1}^5 B_2(u_j).$$

We take the same lattice with $n = 8192$ as in Subsection 5.4. For \mathbf{W} , we computed a Minkowski-reduced basis of the dual lattice (Afflerbach and Grothe, 1985); this gives basis vectors which are as short as possible and of almost the same length. Then we approximated the error distribution with $m = 10^5$ independent realizations of $\tilde{g}(\mathbf{U})$ in (26), for both f_1 and f_2 . We did this with $z_{1,0} = z_{2,0} = z_0$ first with $z_0 = 1$ and then with $z_0 = 4$. The Fourier coefficients converge much faster for f_2 than for f_1 , so we expect a better approximation

for f_2 . This is confirmed empirically in Figure 13, which shows a P-P plot of the empirical distribution \hat{F} of the m independent replicates of $n\tilde{g}_n(U)$ against the empirical distribution \tilde{F} of m replicates of $ng_n(U)$ sampled as in Subsection 5.4. Note that here, $\tilde{g}_n(U)$ is much less time-consuming to sample than $g_n(U)$. We observe an excellent fit in the case of f_2 , even with $z_0 = 1$. For f_1 , the approximation with $z_0 = 1$ is poor, while that with $z_0 = 4$ is reasonable but the lack of fit is still visible. \square




FIG 13. P-P plots between \hat{F} and \tilde{F} for Example 21 for f_1 (left) and f_2 (right).

7. A Real-Life Example: Pricing an Asian Option

Here we take an application used as an example in many articles. We compare two ways of applying RQMC, which can be seen as two ways of defining the function f to estimate the same μ . We examine empirically the RQMC error distribution and see how far it is from normal. Based on the asymptotic results of Section 5.3, we expect the error distribution to be generally closer to normal when f depends more or less equally on several coordinates of \mathbf{u} , and farther from normal in the opposite case. Our results agree with this intuition.

Consider the pricing of an Asian option based on a geometric Brownian motion (GBM) $\{S(t), t \geq 0\}$ observed at discrete time points $0 = t_0 < t_1 < \dots < t_s = T$. The GBM process can be written as $S(t) = S(0) \exp[X(t)]$ where $X(t) = (r - \sigma^2/2)t + \sigma B(t)$, $\{B(t), t \geq 0\}$ is a standard Brownian motion, r is the risk-free interest rate, and σ is the volatility parameter. The option's net discounted payoff is $Y = e^{-rT} \max(0, \bar{S} - K)$ where $K > 0$ is a constant called the strike price and $\bar{S} = (S(t_1) + \dots + S(t_s))/d$, and the price of the option is $\mu = \mathbb{E}[Y]$ where the expectation is under the risk neutral measure (Glasserman, 2004; Hull, 2000).

To simulate \bar{S} and Y we need to sample the vector $\mathbf{X} = (X(t_1), \dots, X(t_s))$ from an s -dimensional normal distribution with mean $\mu = (r - \sigma^2/2)(t_1, \dots, t_s)$

and covariance matrix \mathbf{C} with elements $c_{j,k} = \sigma^2 \min(t_j, t_k)$. The standard way of sampling \mathbf{X} with Monte Carlo is to factorize \mathbf{C} as $\mathbf{C} = \mathbf{A}\mathbf{A}^t$, generate $\mathbf{Z} = (Z_1, \dots, Z_s) = (\Phi^{-1}(U_1), \dots, \Phi^{-1}(U_s))$ where $\mathbf{U} = (U_1, \dots, U_s) \sim U(0, 1)^s$, and return $\mathbf{X} = \mu + \mathbf{A}\mathbf{Z}$. This expresses Y as $Y = f(\mathbf{U})$ for some function f , so this example fits our initial framework, and we can compute an RQMC estimator via (3) for some RQMC point set P_n .

The factorization $\mathbf{C} = \mathbf{A}\mathbf{A}^t$ is not unique and its choice has an impact on the performance of RQMC. Perhaps the most standard one is the Cholesky factorization, where \mathbf{A} is lower triangular. Using it turns out to be equivalent to simulating via the recurrence $X(t_j) = X(t_{j-1}) + (r - \sigma^2/2)(t_j - t_{j-1}) + \sigma\sqrt{t_j - t_{j-1}}\Phi^{-1}(U_j)$ for $j = 1, \dots, s$. A second way of factorizing is an eigen-decomposition, for which $\mathbf{A} = \mathbf{P}\mathbf{D}^{1/2}$ where \mathbf{D} is a diagonal matrix that contains the eigenvalues of \mathbf{C} in decreasing order and \mathbf{P} is an orthogonal matrix whose columns are the corresponding unit-length eigenvectors. This is the decomposition used in *principal component analysis* (PCA). It selects \mathbf{A} so that the maximum amount of variance of \mathbf{Y} comes from Z_1 , then the maximum amount of variance conditional on Z_1 comes from Z_2 , and so on. By concentrating the variance in the first coordinates of \mathbf{Z} as much as possible, the method reduces the effective dimension in the truncation sense (if we consider the variance of \mathbf{Y}) and often makes RQMC much more effective for this reason (Glasserman, 2004; L'Ecuyer, 2009). But we expect this to have the side effect of making the error distribution further away from the normal.

For a numerical illustration, we take $S(0) = 100$, $T = 1$, $r = 0.05$, $\sigma = 0.5$, $K = 100$, and $t_j = jT/s$ for $1 \leq j \leq s$. We consider three values of s , namely $s = 2, 10$, and 20 . Figure 14 shows normal P-P plots of the error on the expected payoff $\mu = \mathbb{E}[Y]$, using Cholesky factorization (left) and using PCA decomposition (right). We observe that the standardized error distribution is very close to normal for Cholesky, especially when s increases, but for PCA it is asymmetric and does not get close to normal when we increase s (it seems almost independent of s). Additional insight on why this happens can be obtained by performing an ANOVA decomposition of the resulting function f in each case. We have estimated the variance components by Monte Carlo, using the method explained in Sobol' and Myshetskaya (2007). With Cholesky, for $s = 2, 10$, and 20 , we find that the one-dimensional ANOVA components account for about 87%, 70% and 67% of the variance, respectively. This variance is also spread among all coordinates with comparable magnitude, although slowly decreasing with increasing coordinate index. In other words, the payoff function is well approximated by a sum of univariate functions of independent variables as in (8), with comparable variances, and for this reason the error distribution approaches a normal when s increases. With PCA, in contrast, more than 98% of the variance is concentrated on the first coordinate of \mathbf{u} , which means that f looks more like a single one-dimensional function rather than a sum of several independent one-dimensional functions.




FIG 14. Normal P-P plots of the error for the Asian option payoff, using Cholesky factorization (left) and PCA decomposition (right).

Acknowledgments

This work has been supported by an NSERC-Canada Discovery Grant and a Canada Research Chair to the first author, the EuroNF Network of Excellence to the third author, and INRIA's associated team MOCQUASIN to all authors. This paper is an expanded version of L'Ecuyer and Tuffin (2009) and part of the material was also presented at the Sixth St. Petersburg Workshop on Simulation, in June 2009.

References

- ABRAMOWITZ, M. and STEGUN, I. A. (1970). *Handbook of Mathematical Functions*. Dover, New York.
- AFFLERBACH, L. and GROTHE, H. (1985). Calculation of Minkowski-Reduced Lattice Bases. *Computing* **35** 269–276.
- AVRAMIDIS, A. N. and WILSON, J. R. (1996). Integrated Variance Reduction Strategies for Simulation. *Operations Research* **44** 327–346.
- AVRAMIDIS, A. N. and WILSON, J. R. (1998). Correlation-Induction Techniques for Estimating Quantiles in Simulation Experiments. *Operations Research* **46** 574–591.
- BARROW, D. L. and SMITH, P. W. (1979). Spline Notation Applied to a Volume Problem. *The American Mathematical Monthly* **86** 50–51.
- CAFLISCH, R. E., MOROKOFF, W. and OWEN, A. (1997). Valuation of Mortgage-Backed Securities Using Brownian Bridges to Reduce Effective Dimension. *The Journal of Computational Finance* **1** 27–46.
- CONWAY, J. H. and SLOANE, N. J. A. (1999). *Sphere Packings, Lattices and Groups*, 3rd ed. *Grundlehren der Mathematischen Wissenschaften 290*. Springer-Verlag, New York.

- CRANLEY, R. and PATTERSON, T. N. L. (1976). Randomization of Number Theoretic Methods for Multiple Integration. *SIAM Journal on Numerical Analysis* **13** 904–914.
- DICK, J., SLOAN, I. H., WANG, X. and WOZNAKOWSKI, H. (2006). Good Lattice Rules in Weighted Korobov Spaces with General Weights. *Numerische Mathematik* **103** 63–97.
- ELMAGHRABY, S. (1977). *Activity Networks*. Wiley, New York.
- GLASSERMAN, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- HICKERNELL, F. J. (2002). Obtaining $O(N^{-2+\epsilon})$ Convergence for Lattice Quadrature Rules. In *Monte Carlo and Quasi-Monte Carlo Methods 2000* (K.-T. FANG, F. J. HICKERNELL and H. NIEDERREITER, eds.) 274–289. Springer-Verlag, Berlin.
- HULL, J. C. (2000). *Options, Futures, and Other Derivative Securities*, fourth ed. Prentice-Hall, Englewood-Cliff, N.J.
- KUO, F. Y. and JOE, S. (2002). Component-by-Component Construction of Good Lattice Rules with a Composite Number of Points. *Journal of Complexity* **18** 943–976.
- L'ECUYER, P. (2009). Quasi-Monte Carlo Methods with Applications in Finance. *Finance and Stochastics* **13** 307–349.
- L'ECUYER, P. and LEMIEUX, C. (2000). Variance Reduction via Lattice Rules. *Management Science* **46** 1214–1235.
- L'ECUYER, P. and LEMIEUX, C. (2002). Recent Advances in Randomized Quasi-Monte Carlo Methods. In *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications* (M. Dror, P. L'Ecuyer and F. Szidarovszky, eds.) 419–474. Kluwer Academic, Boston.
- L'ECUYER, P. and TUFFIN, B. (2009). On the Error Distribution for Randomly-Shifted Lattice Rules. In *Proceedings of the 2009 Winter Simulation Conference*. IEEE Press To appear.
- LEMIEUX, C. (2009). *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer-Verlag, New York, NY.
- LIU, R. and OWEN, A. B. (2006). Estimating Mean Dimensionality of Analysis of Variance Decompositions. *Journal of the American Statistical Association* **101** 712–721.
- LOH, W.-L. (2003). On the Asymptotic Distribution of Scramble Nets Quadratures. *Annals of Statistics* **31** 1282–1324.
- LOH, W.-L. (2005). On the Asymptotic Distribution of Some Randomized Quadrature Rules. In *Stein's Method and Applications*, (C. Stein, A. D. Barbour and L. H. Y. Chen, eds.) **5** 209–222. World Scientific.
- NIEDERREITER, H. (1992). *Random Number Generation and Quasi-Monte Carlo Methods*. *SIAM CBMS-NSF Regional Conference Series in Applied Mathematics* **63**. SIAM, Philadelphia, PA.
- OWEN, A. B. (1992). A Central Limit Theorem for Latin Hypercube Sampling. *Journal of the Royal Statistical Society B* **54** 541–551.
- OWEN, A. B. (1998). Latin Supercube Sampling for Very High-Dimensional Simulations. *ACM Transactions on Modeling and Computer Simulation* **8** 71–

102.

- PETROV, V. V. (1995). *Limit Theorems of Probability Theory*. Oxford University Press, Oxford, U.K.
- SINESCU, V. and L'ECUYER, P. (2010). Existence and Construction of Shifted Lattice Rules with an Arbitrary Number of Points and Bounded Weighted Star Discrepancy for General Weights. Submitted for publication.
- SLOAN, I. H. and JOE, S. (1994). *Lattice Methods for Multiple Integration*. Clarendon Press, Oxford.
- SLOAN, I. H. and REZSTOV, A. (2002). Component-by-Component Construction of Good Lattice Rules. *Mathematics of Computation* **71** 262–273.
- SOBOL', I. M. and MYSHETSKAYA, E. E. (2007). Monte Carlo Estimators for Small Sensitivity Indices. *Monte Carlo Methods and Applications* **13** 455–465.
- TUFFIN, B. (1998). Variance Reduction Order Using Good Lattice Points in Monte Carlo Methods. *Computing* **61** 371–378.

Appendix A: Derivation of the generalized Euler–MacLaurin expansion

The following is the derivation of the generalized Euler–MacLaurin expansion given in Theorem 2. Here we use $\{\cdot\}$ to denote the component-wise modulo-1 operator.

Lemma 11. *If F has an integrable $(k + 1)$ -th derivative over $[0, 1]$, then, for $k = 0, 1, \dots$ and any $U \in (0, 1)$,*

$$\int_0^1 \frac{B_k(\{U - u\})}{k!} F^{(k)}(u) \, du = \delta_{k,0} F(U) - \frac{B_{k+1}(U)}{(k+1)!} [F^{(k)}(1) - F^{(k)}(0)] + \int_0^1 \frac{B_{k+1}(\{U - u\})}{(k+1)!} F^{(k+1)}(u) \, du,$$

where $\delta_{k,j} = 1$ if $k = j$ and 0 otherwise.

Proof. We have

$$\int_0^1 \frac{B_k(\{U - u\})}{k!} F^{(k)}(u) \, du = \int_0^U \frac{B_k(U - u)}{k!} F^{(k)}(u) \, du + \int_U^1 \frac{B_k(U - u + 1)}{k!} F^{(k)}(u) \, du.$$

Integration by parts of the first and second terms on the right-hand side yields

$$-\frac{B_{k+1}(0) F^{(k)}(U) - B_{k+1}(U) F^{(k)}(0)}{(k+1)!} + \int_0^U \frac{B_{k+1}(U - u)}{(k+1)!} F^{(k+1)}(u) \, du$$

and

$$-\frac{B_{k+1}(U) F^{(k)}(1) - B_{k+1}(1) F^{(k)}(U)}{(k+1)!} + \int_U^1 \frac{B_{k+1}(U - u + 1)}{(k+1)!} F^{(k+1)}(u) \, du,$$

respectively, where we have used the identity $B'_{k+1}(u) = (k+1)B_k(u)$. Summing the latter two expressions and using $B_1(0) - B_1(1) = 1$, $B_k(0) - B_k(1) = 0$ for $k \geq 2$, Lemma 11 is readily obtained. \square

Lemma 12. *If F has $m + 1$ integrable derivatives over $[0, 1]$, then for any $U \in (0, 1)$,*

$$\int_0^1 F(u) \, du = F(U) - \sum_{k=1}^m \frac{B_k(U)}{k!} [F^{(k-1)}(1) - F^{(k-1)}(0)] + \int_0^1 \frac{B_{m+1}(\{U - u\}) - B_{m+1}(U)}{(m+1)!} F^{(m+1)}(u) \, du.$$

Proof. First note that $\int_0^1 F(u) du = \int_0^1 F^{(0)}(u) B_0(\{U - u\}) du$, then apply Lemma 11 recursively:

$$\int_0^1 F(u) du = F(U) - \sum_{k=1}^{m+1} \frac{B_k(U)}{k!} [F^{(k-1)}(1) - F^{(k-1)}(0)] + \int_0^1 \frac{B_{m+1}(\{U - u\})}{(m+1)!} F^{(m+1)}(u) du.$$

Next, we take the term with $k = m + 1$ out of the sum and into the integral, which yields Lemma 12. \square

Now, to prove Theorem 2, we first write the integral of f as

$$\int_0^1 f(u) du = \sum_{i=0}^{n-1} \int_{i/n}^{(i+1)/n} f(u) du = \int_0^1 \frac{1}{n} \sum_{i=0}^{n-1} f((i+u)/n) du. \quad (27)$$

Next, we set $F(u) = (1/n) \sum_{i=0}^{n-1} f((i+u)/n)$ so that

$$F^{(k)}(u) = (1/n^{k+1}) \sum_{i=0}^{n-1} f^{(k)}((i+u)/n),$$

then we invoke Lemma 12 on the right-hand side of (27):

$$\begin{aligned} \int_0^1 f(u) du &= \frac{1}{n} \sum_{i=0}^{n-1} f((i+U)/n) \\ &\quad - \sum_{k=1}^m \frac{B_k(U)}{n^k k!} \sum_{i=0}^{n-1} [f^{(k-1)}((i+1)/n) - f^{(k-1)}(i/n)] \\ &\quad + \int_0^1 \frac{B_{m+1}(\{U - u\}) - B_{m+1}(U)}{n^{m+2} (m+1)!} \sum_{i=0}^{n-1} f^{(m+1)}((i+u)/n) du. \end{aligned}$$

The sum over the index i in the second term on the right-hand side is telescoping and reduces to $f^{(k-1)}(1) - f^{(k-1)}(0)$. From here, obtaining Theorem 2 is straightforward.

Appendix B: An Additional Example: A Stochastic Activity Network

This additional example was removed from the paper to save space. We put it here just for informational purpose.

We consider the graph of Figure 15, where each arc j has a random length V_j with cdf $F_j(\cdot)$, for $j = 1, \dots, 13$, and these lengths are assumed independent. Let T be the length of the longest path from node 1 to node 9 in this graph. We

are interested in estimating both $\mathbb{E}[T]$ and $q(x) = \mathbb{P}[T > x]$ for a given constant x , by simulation. This type of graph may represent the precedence relationships between the activities of a project, with the arc lengths representing activity durations and T the (random) time needed to complete the project. This particular example is taken from [Elmaghraby \(1977\)](#), where the F_j 's can be found; see also [Avramidis and Wilson \(1996\)](#) and [L'Ecuyer and Lemieux \(2000\)](#).




FIG 15. Example of a stochastic activity network, taken from [Elmaghraby \(1977\)](#)

To estimate $\mathbb{E}[T]$ and $q(x) = \mathbb{P}[T > x]$ by Monte Carlo, it suffices to generate the random variables V_1, \dots, V_{13} and compute T by a standard longest-path algorithm, repeat this n times to obtain n realizations T_1, \dots, T_n of T , and then estimate $\mathbb{E}[T]$ by $\bar{T}_n = (T_1 + \dots + T_n)/n$ and $q(x)$ by $\bar{Y}_n = (1/n) \sum_{i=1}^n \mathbb{I}[T_i > x]$. For RQMC, we generate the n replicates of (V_1, \dots, V_{13}) by inversion from the n points of the lattice rule $Y_j = F_j^{-1}(U_j)$ where U_j is the j th coordinate of the RQMC point, and we repeat this for each point.

Figures 16 and 17 (left) show the normal P-P plots of the error for $m = 10^4$ independent replicates of \bar{T}_n and \bar{Y}_n , respectively, with RQMC. We observe a surprisingly good agreement with the normal distribution for \bar{T}_n . This can be explained (intuitively) by the fact that the length of each path is a sum of independent random variables and some of these random variables already have the normal distribution. On the other hand, the empirical distribution of \bar{Y}_n appears to be concentrated on a small number of values, and that number becomes even smaller (the steps are larger) when $q(x)$ gets closer to 0 or 1. Among the 10^4 independent replicates with RQMC, we have observed exactly 40, 113 and 55 distinct values of \bar{Y}_n (these are the number of jumps in the empirical distribution) for $x = 30, 64, 100$, respectively. Here we have $\mathbb{E}[T] \approx 64$, $q(64) \approx 0.420$ (thus the mean value of T is below the median), $q(30) \approx 0.992$ and $q(100) \approx 0.066$. In the figure, the (horizontal) jumps in $\hat{F}_{Z_n, m}$ corresponds to the numbers (or counts) of observed values. The (vertical) jumps in Φ correspond to spacings between observed values, and since all observed values are equally spaced, the magnitude of a jump in Φ is the probability that a standard normal

variable falls between the two consecutive (standardized) observed values. The biggest jumps in $\hat{F}_{Z_n,m}$ are in the same region (the middle of the plot) as the biggest jumps in Φ . This means that the discrete distribution $\hat{F}_{Z_n,m}$ behaves very much like a normal.




FIG 16. Normal P-P plot of the error on \bar{T}_n for the stochastic activity network example.




FIG 17. Normal P-P plot of the error on \bar{Y}_n for the stochastic activity network example using crude Monte Carlo (left) and CMC (right).

For this example, an estimator with lower variance than \bar{Y}_n can be constructed by *conditional Monte Carlo* (CMC) as follows; see [Avramidis and Wilson \(1998\)](#), Section 4.1, and references therein. Let $\mathcal{L} = \{5, 6, 7, 9, 10\}$ be the set of arcs that separate the nodes $\{1, 2, 3, 4, 5\}$ from the nodes $\{6, 7, 8, 9\}$, and let \mathcal{B} be the eight other arcs. For the CMC estimator, we generate only the arc lengths V_j for $j \in \mathcal{B}$ and we estimate $\mathbb{P}[T > x]$ by the conditional expectation

of $\mathbb{I}[T > x]$ given $\{V_j, j \in \mathcal{B}\}$, that is, by $W = \mathbb{P}[T > x \mid \{V_j, j \in \mathcal{B}\}]$. Let \overline{W}_n be the average of the n replicates of W .

A normal P-P plot for the error on this CMC estimator \overline{W}_n is given in Figure 17 (right). The error distribution is now smooth, and this comes from the fact that W is no longer an indicator function; now it has a density. But it still departs significantly from the normal distribution for $x = 100$; in fact, it is actually not far from a uniform distribution in that case.