

CENTRAL LIMIT THEOREM FOR A RANDOMIZED QUASI-MONTE CARLO ESTIMATOR OF A SMOOTH FUNCTION OF MEANS

Marvin K. Nakayama¹, Bruno Tuffin², and Pierre L'Ecuyer³

¹Dept. of Computer Science, New Jersey Institute of Technology, Newark, NJ, USA

²Inria, Univ Rennes, CNRS, IRISA Campus de Beaulieu, 35042 Rennes, FRANCE

³DIRO, Université de Montréal, Montréal, QC, CANADA

ABSTRACT

Consider estimating a known smooth function (such as a ratio) of unknown means. Our paper accomplishes this by first estimating each mean via randomized quasi-Monte Carlo and then evaluating the function at the estimated means. We prove that the resulting plug-in estimator obeys a central limit theorem by first establishing a joint central limit theorem for a triangular array of estimators of the vector of means and then employing the delta method.

1 INTRODUCTION

Scientists and engineers frequently need to compute a performance measure α expressed as a known smooth function of unknown means. Examples include the standard deviation of a random variable, which is a function of the first and second moments (that is, two means), or the correlation of two random variables, expressed in terms of the first two moments and mixed moment of the two random variables. Other examples are the expected hitting time of a set in a regenerative context or a conditional expectation, both expressed as a ratio of two means, or the derivative of a ratio, which can be written as a function of four expectations (Nakayama and Tuffin 2023; Glynn et al. 1991). Sections 3.1–3.4 of Serfling (1980) and Section III.3 of Asmussen and Glynn (2007) provide additional settings that fit this framework.

When the means are analytically intractable, as is typically the case for situations arising in practice, a commonly used computational approach estimates them via *Monte Carlo* (MC) simulation and evaluates the known function at the resulting estimators to obtain a plug-in estimator of α . This estimator of α often obeys a *central limit theorem* (CLT), established by first showing a joint CLT for the MC estimator of the vector of means and then applying the delta method (see Serfling 1980, p. 122, or Asmussen and Glynn 2007, p. 75). We then can exploit the CLT to obtain a *confidence interval* (CI) for α to provide a measure of the error of the plug-in estimator. The delta method has been employed extensively in the MC literature, as, e.g., in the references mentioned in the previous paragraph. Also, application of the delta method and constructing a CI rely on the existence and continuity of a derivative, which Rhee and Glynn (2023) study in the context of general state space Markov chains.

As an alternative to MC to estimate the means, *quasi-Monte Carlo* (QMC) methods constitute a class of deterministic (quadrature) algorithms to numerically compute a multidimensional integral over a unit hypercube, so the integral can correspond to a mean. QMC averages the integrand at carefully placed deterministic points that are more evenly spaced over the integration domain than a typical MC sample, enabling faster convergence to the expected value. However, QMC suffers from a serious drawback: there is no practical way to provide a computable measure of error of the QMC estimator.

Randomized QMC (RQMC) methods randomize a QMC point set without losing its good distribution property, and repeating the process $r \geq 2$ *independent and identically distributed* (i.i.d.) times gives a simple way to construct an approximate CI for a single mean through a sample variance (L'Ecuyer 2018). Nakayama and Tuffin (2024) provide sufficient conditions on the joint growth of the number m of QMC

points and the number r of randomizations to ensure that the RQMC estimator for a single mean obeys a CLT and to obtain an asymptotically valid CI. The analysis requires a triangular-array formulation to accommodate that the resulting RQMC estimator averages the r randomizations of estimators based on a (possibly) growing number m of QMC points, and the fact that the distribution of these r estimators changes when m increases.

The goal of this paper is to provide sufficient conditions for a CLT when estimating α , a known smooth function of unknown means by RQMC, not just a single mean as in Nakayama and Tuffin (2024). To achieve this, we first establish reasonable conditions for a specific multivariate triangular-array CLT for the RQMC estimator of the vector of means. This multivariate CLT is stated in Theorem 1, whose proof is a key part of our contribution. Applying the delta method to the function of asymptotically Gaussian mean estimators then produces the desired CLT for the RQMC plug-in estimator, given by Theorem 2.

The rest of the paper unfolds as follows. Section 2 defines the mathematical problem of estimating a function of means, the main notations, and the existing results for an evaluation by MC simulation. Section 3 recalls the basics of QMC and RQMC methods, and Section 4 presents the main results, establishing CLTs for the RQMC estimators of the mean vector and for α . Finally, Section 5 briefly concludes and provides suggestions for extensions of the work.

2 MATHEMATICAL SETUP

Our goal is to estimate

$$\alpha = g(\boldsymbol{\mu}) \tag{1}$$

for a known smooth function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ of an unknown d -vector $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_d)$, so we need to estimate $\boldsymbol{\mu}$ via some computational method. (All vectors are of column type, although we depict them as row vectors to save space. If a column vector needs to be converted into a row vector, we use a superscript \top for transpose, which also applies to matrices.) For each $t = 1, 2, \dots, d$, we assume that

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

with $h_t : [0, 1]^s \rightarrow \mathbb{R}$ a given function (integrand) for some fixed dimension $s \geq 1$, random vector $\mathbf{U} \sim \mathcal{U}[0, 1]^s$ with $\mathcal{U}[0, 1]^s$ the uniform distribution on the s -dimensional unit hypercube $[0, 1]^s$, \sim means “is distributed as”, and \mathbb{E} denotes the expectation operator. Integrating over $[0, 1]^s$ is the standard (R)QMC setting, and it is often possible (e.g., through a change of variables) to express the mean of many stochastic models in this way. We can regard each integrand h_t as a complicated simulation program that converts s independent univariate uniform random numbers into observations from specified input distributions (possibly with dependencies and different marginals), which are used to produce an output of the stochastic model, where the output has mean μ_t .

We assume that the hypercube $[0, 1]^s$ has the same dimension s for each integrand h_t . When each h_t has domain $[0, 1]^{s_t}$ with dimension s_t dependent on t , one can simply take $s = \max_{t=1,2,\dots,d} s_t$, and for h_t with $s_t < s$, just ignore coordinates $s_t + 1, s_t + 2, \dots, s$. This is similar to what is done in Section 10 of L'Ecuyer and Lemieux (2000) to apply QMC and RQMC to problems of infinite dimension, e.g., a random time horizon with no finite upper bound.

We now review how to estimate $\boldsymbol{\mu}$ via Monte Carlo. MC employs an i.i.d. sequence $(\mathbf{U}_i)_{i \geq 1}$ of uniform random vectors in $[0, 1]^s$ to construct an estimator of μ_t for each $t = 1, \dots, d$. One possibility is to use independent sequences $(\mathbf{U}_i)_{i \geq 1}$ for the different values of t . Another one, the *common random numbers* (CRN) strategy (Asmussen and Glynn 2007, Section V.6), employs the same sequence for all t , so the d estimators become dependent. CRN may either decrease or increase the variance of the estimator of α , depending on how the random numbers are used in the estimators. For instance, if $d = 2$ and h_1 and h_2 are increasing functions, then for $g(\mu_1, \mu_2) = \mu_2 - \mu_1$ or $g(\mu_1, \mu_2) = \mu_1/\mu_2$, CRN decreases the variance,

while for $g(\mu_1, \mu_2) = \mu_2 + \mu_1$, it increases the variance. See Section 6.4.4 of L'Ecuyer (2024) for more on this. In general, it is also possible to use CRN for only certain subsets of the values of t , which may be appropriate in some cases. In the rest of the paper, we assume that CRN is employed, i.e., we have a single sequence of uniform random points for all t . This is easy to generalize, but it would require more complicated notation. The MC estimator of μ_t is then

$$\hat{\mu}_{n,t}^{\text{MC}} = \frac{1}{n} \sum_{i=1}^n h_t(\mathbf{U}_i) \quad (2)$$

and the plug-in MC estimator of α is

$$\hat{\alpha}_n^{\text{MC}} = g(\hat{\mu}_{n,1}^{\text{MC}}, \hat{\mu}_{n,2}^{\text{MC}}, \dots, \hat{\mu}_{n,d}^{\text{MC}}).$$

Nakayama and Tuffin (2023) consider more general possibilities for the dependencies among $\hat{\mu}_{n,t}^{\text{MC}}$ ($t = 1, \dots, d$), e.g., allowing for *measure-specific importance sampling* (Goyal et al. 1992), where importance sampling is applied to estimate some means, and independently using naive MC for other means, which is useful when some means relate to rare events but others do not.

Let $\Sigma = (\Sigma_{t,t'} : t, t' = 1, 2, \dots, d)$ be the covariance matrix of the random vector $(h_1(\mathbf{U}), \dots, h_d(\mathbf{U}))$, assumed finite and positive definite. The vector $\hat{\boldsymbol{\mu}}_n^{\text{MC}} = (\hat{\mu}_{n,1}^{\text{MC}}, \dots, \hat{\mu}_{n,d}^{\text{MC}})$ then obeys a multivariate CLT (Billingsley 1995, Theorem 29.5)

$$\sqrt{n} [\hat{\boldsymbol{\mu}}_n^{\text{MC}} - \boldsymbol{\mu}] \Rightarrow \Sigma^{1/2} \mathcal{N}_d \quad \text{as } n \rightarrow \infty, \quad (3)$$

where $\Sigma^{1/2}$ is a $d \times d$ matrix satisfying $\Sigma^{1/2}(\Sigma^{1/2})^\top = \Sigma$ and \mathcal{N}_d is a d -dimensional standard normal random vector (with mean vector $\mathbf{0} = (0, 0, \dots, 0)$ and covariance matrix equal to the $d \times d$ identity matrix). Let $\nabla g(\cdot) = (g_1(\cdot), g_2(\cdot), \dots, g_d(\cdot))$ be the gradient of $g(\cdot)$, where $g_t(\cdot)$ is the partial derivative of $g(\cdot)$ with respect to its t -th argument, for $t = 1, \dots, d$. Under the assumption that g has a nonzero differential at $\boldsymbol{\mu}$, the delta method (Serfling 1980, p. 124) leads to a CLT for $\hat{\alpha}_n^{\text{MC}}$:

$$\sqrt{n} [\hat{\alpha}_n^{\text{MC}} - \alpha] \Rightarrow \tau_{\text{MC}} \mathcal{N}_1 \quad \text{as } n \rightarrow \infty, \quad (4)$$

with asymptotic variance

$$\tau_{\text{MC}}^2 = \nabla g(\boldsymbol{\mu})^\top \Sigma \nabla g(\boldsymbol{\mu}) = \sum_{t=1}^d \sum_{t'=1}^d g_t(\boldsymbol{\mu}) g_{t'}(\boldsymbol{\mu}) \Sigma_{t,t'}. \quad (5)$$

A sufficient condition for the differential of $g(\cdot)$ (also known as a total derivative) to be nonzero at $\boldsymbol{\mu}$ is that the gradient $\nabla g(\cdot)$ exists in a neighborhood of $\boldsymbol{\mu}$ and is continuous at $\boldsymbol{\mu}$, with $\nabla g(\boldsymbol{\mu}) \neq \mathbf{0}$ (Apostol 1974, Theorem 12.11).

3 RANDOMIZED QUASI-MONTE CARLO

QMC methods replace the sequence of independent uniformly distributed vectors $(\mathbf{U}_i)_{i \geq 1}$ in the estimators (2) by a *deterministic* sequence $\boldsymbol{\Xi} = (\boldsymbol{\xi}_i)_{i \geq 1}$ with $\boldsymbol{\xi}_i = (\xi_{i,1}, \dots, \xi_{i,d}) \in [0, 1]^d$. The sequence $\boldsymbol{\Xi}$ is carefully designed to cover $[0, 1]^d$ more evenly than a typical sequence of independent uniform random vectors. The unevenness of the spread of the first n points of $\boldsymbol{\Xi}$ can be measured by, e.g., the star discrepancy $D_n^*(\boldsymbol{\Xi})$ (Niederreiter 1992; Lemieux 2009), and $\boldsymbol{\Xi}$ is called a *low-discrepancy sequence* when $D_n^*(\boldsymbol{\Xi}) = O(n^{-1}(\ln n)^s)$ as $n \rightarrow \infty$, where the notation $f_1(n) = O(f_2(n))$ as $n \rightarrow \infty$ for functions f_1 and f_2 means that there exist constants $c_0 > 0$ and n_0 such that $|f_1(n)| \leq c_0 |f_2(n)|$ for all $n \geq n_0$. We also write $f_1(n) = \Theta(f_2(n))$ as $n \rightarrow \infty$ to denote that both $f_1(n) = O(f_2(n))$ and $f_2(n) = O(f_1(n))$ as $n \rightarrow \infty$. Several QMC error bounds

exist in the literature, the best-known being the Koksma-Hlawka bound. To describe it for a single integrand $h_t(\cdot)$, let $V(h_t)$ be its variation in the sense of Hardy and Krause. If $V(h_t) < \infty$ (a property known as “bounded variation in the sense of Hardy and Krause” (BVHK)), then the estimator $\widehat{\mu}_{n,t}^Q = \frac{1}{n} \sum_{i=1}^n h_t(\xi_i)$ satisfies

$$\left| \widehat{\mu}_{n,t}^Q - \mu_t \right| \leq V(h_t) D_n^*(\Xi), \quad (6)$$

leading to, when Ξ is a low-discrepancy sequence, a convergence rate of $O(n^{-1}(\ln n)^s)$, faster than MC’s rate of $\Theta(n^{-1/2})$ for its root-mean-square error from the CLT. Niederreiter (1992) and Lemieux (2009) provide more details on QMC and low-discrepancy sequences. While there are several families of low-discrepancy sequences (digital nets, etc.), we just provide one as an illustration. A *lattice rule* of rank-1 for a number n of points selects a generating vector $\mathbf{a} = (a_1, \dots, a_d)$ and uses $\xi_i = ((i-1)\mathbf{a} \bmod 1)/n$ for $1 \leq i \leq n$, where the modulo is applied coordinate-wise. The software tool LatNet Builder (L’Ecuyer et al. 2022) provides good ways to select \mathbf{a} .

Despite its appealing fast convergence rate, QMC suffers the drawback that it is usually not possible to compute an actual estimate of the integration error. While, e.g., the Koksma-Hlawka inequality (6) provides a theoretically attractive bound on the QMC error, computing the exact values of $V(h_t)$ and $D_n^*(\Xi)$ is very difficult, and even if we could, the resulting bound typically overspecifies the actual error by orders of magnitude for reasonable values of n .

This motivated the development of RQMC to try to obtain computable error bounds via a CLT through $r \geq 2$ i.i.d. randomizations of a QMC point set. Specifically, in a single randomization, RQMC randomizes a low-discrepancy sequence Ξ of m points without losing its good repartition property (L’Ecuyer 2018). Let $(\mathbf{U}'_i)_{i \geq 1}$ be the randomized low-discrepancy sequence built from Ξ , such that each \mathbf{U}'_i is uniformly distributed over $[0, 1]^s$. One approach to do this (mostly used for lattice point sets) is through a *random shift*: generate a single uniform $\mathbf{U} \sim \mathcal{U}[0, 1]^s$ and add it to each point of Ξ , resulting in the randomized points $\mathbf{U}'_i = (\xi_i + \mathbf{U}) \bmod 1$, for $i = 1, 2, \dots, m$. Typically, the points \mathbf{U}'_i are correlated (because each \mathbf{U}'_i uses the same uniform \mathbf{U}) but the low-discrepancy property is preserved, e.g., when the random shift is applied to Ξ from a lattice rule. We can preserve the finer structure of other types of low-discrepancy sequences by applying alternative randomization methods (L’Ecuyer 2018).

RQMC repeats this $r \geq 1$ times, independently, computing an estimator from each randomization. Specifically, let $\mathbf{U}'_{i,j} \in [0, 1]^s$ be the i -th point of the j -th i.i.d. randomization ($i = 1, 2, \dots, m$, and $j = 1, 2, \dots, r$). The RQMC estimator of α is

$$\widehat{\alpha}_{m,r}^{\text{RQ}} = g(\widehat{\mu}_{m,r,1}^{\text{RQ}}, \widehat{\mu}_{m,r,2}^{\text{RQ}}, \dots, \widehat{\mu}_{m,r,d}^{\text{RQ}}),$$

where for each $t = 1, 2, \dots, d$, the RQMC estimator of μ_t is

$$\widehat{\mu}_{m,r,t}^{\text{RQ}} = \frac{1}{r} \sum_{j=1}^r X_{j,t}, \quad \text{where} \quad X_{j,t} = \frac{1}{m} \sum_{i=1}^m h_t(\mathbf{U}'_{i,j}), \quad (7)$$

with $X_{j,t}$ as the estimator of μ_t from randomization $j = 1, 2, \dots, r$, of m points. Since the randomizations are independent across the r replicates, the random vectors $(X_{j,1}, X_{j,2}, \dots, X_{j,d})$, $j = 1, 2, \dots, r$, are always i.i.d. However, their components are generally not independent. We will assume that the same randomized points $\mathbf{U}'_{i,j}$, $i = 1, 2, \dots, m$, are used across all values of t , which corresponds to CRN. Then, $X_{j,1}, X_{j,2}, \dots, X_{j,d}$ are dependent. We could make them independent by using independent randomizations for $X_{j,1}, X_{j,2}, \dots, X_{j,d}$ instead. In that case, the analysis in Section 4 simplifies significantly but we may not benefit from the application of CRN described in Section 2. It is hoped that as m or r (or both) grows large, the overall RQMC estimator $\widehat{\alpha}_{m,r}^{\text{RQ}}$ obeys a Gaussian CLT.

To analyze the asymptotic behavior of our RQMC estimator as the number of integrand evaluations grows large, we define a “budget parameter” n corresponding to using at most nd integrand evaluations

across all d integrands h_t , $t = 1, 2, \dots, d$; i.e., we must have $mr \leq n$. We now define the RQMC estimator in (7) with $(m, r) = (m_n, r_n)$, where $r_n \geq 1$ is the number of randomizations and $m_n \geq 1$ is the number of points used from each randomized sequence, so the total number of evaluations of the integrand h_t is $m_n r_n$, and the total number of evaluations across all d integrands is at most $m_n r_n d$. Note that the RQMC point set may change completely when m_n increases. We assume the following:

Assumption 1 $m_n r_n \leq n$ for each $n \geq 1$, with $m_n r_n / n \rightarrow 1$ and $r_n \rightarrow \infty$ as $n \rightarrow \infty$.

Under Assumption 1, the RQMC estimator of μ_t , $t = 1, 2, \dots, d$, in (7) becomes

$$\hat{\mu}_{m_n, r_n}^{\text{RQ}, t} = \frac{1}{r_n} \sum_{j=1}^{r_n} X_{n, j, t}, \quad \text{where} \quad X_{n, j, t} = \frac{1}{m_n} \sum_{i=1}^{m_n} h_t(\mathbf{U}'_{i, j}), \quad (8)$$

so $X_{n, j, t}$ is the estimator from randomization $j = 1, 2, \dots, r_n$, of m_n points, where $m_n \leq n$. We then get an estimator

$$\hat{\alpha}_{m_n, r_n}^{\text{RQ}} = g(\hat{\mu}_{m_n, r_n, 1}^{\text{RQ}}, \hat{\mu}_{m_n, r_n, 2}^{\text{RQ}}, \dots, \hat{\mu}_{m_n, r_n, d}^{\text{RQ}}). \quad (9)$$

The question we aim to answer in the next section is: how should m_n and r_n increase with n so that the resulting estimator $\hat{\alpha}_{m_n, r_n}^{\text{RQ}}$ satisfies a weak convergence result with a Gaussian limit? Assumption 1 requires $r_n \rightarrow \infty$ as $n \rightarrow \infty$ because otherwise, a Gaussian CLT may not hold. For example, L'Ecuyer et al. (2010) show that when applying RQMC using a lattice rule and the random shift, the resulting estimator can obey a limit theorem with non-Gaussian limit as $m_n \rightarrow \infty$ for fixed $r_n \geq 1$. The only existing Gaussian CLT for r_n fixed and $m_n \rightarrow \infty$ is established for a particular (more costly) type of RQMC, nested digital scrambling (Loh 2003; Basu and Mukherjee 2017; He and Zhu 2017). Our goal here is to express conditions on (m_n, r_n) under which a Gaussian CLT can be proved for all randomization methods. Nakayama and Tuffin (2024) perform this type of analysis for a single mean $\hat{\alpha}_{m_n, r_n}^{\text{RQ}} = \hat{\mu}_{m_n, r_n, t}^{\text{RQ}}$ (i.e., $t = d = 1$ and g the identity function in (1)). While Assumption 1 specifies that $r_n \rightarrow \infty$ as $n \rightarrow \infty$, it does not require that $m_n \rightarrow \infty$ as $n \rightarrow \infty$; we may take $m_n = m_0$ for some fixed $m_0 \geq 1$, where $m_0 = 1$ corresponds to MC. Our analysis in the next section covers all these special cases.

4 CENTRAL LIMIT THEOREMS

To establish a CLT for $\hat{\alpha}_{m_n, r_n}^{\text{RQ}}$ in (9), our approach will be to first obtain a multivariate CLT for

$$\hat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}} = (\hat{\mu}_{m_n, r_n, 1}^{\text{RQ}}, \hat{\mu}_{m_n, r_n, 2}^{\text{RQ}}, \dots, \hat{\mu}_{m_n, r_n, d}^{\text{RQ}})$$

as $n \rightarrow \infty$, and then apply the delta method (Serfling 1980, Theorem 3.3A). The multivariate CLT requires a triangular-array formulation, as in Theorem 5 on p. 19 of Gikhman and Skorokhod (1996), which is originally stated for random functions but we specialize to random vectors. We next set up the necessary framework to do this.

Define the random vector $\mathbf{X}_{n, j} = (X_{n, j, 1}, X_{n, j, 2}, \dots, X_{n, j, d})$, where each $X_{n, j, t}$ is from (8). This vector has a multivariate distribution that depends only on $m_n = m$ but not on n nor r_n , conditionally on m_n . Component t has mean $\mu_t = \mathbb{E}[X_{n, j, t}]$ and variance

$$\sigma_{m_n, t}^2 = \text{Var}[X_{n, j, t}], \quad (10)$$

where $\text{Var}[\cdot]$ denotes the variance operator. To avoid uninteresting situations, we assume the strict positiveness and finiteness of the variances when the budget parameter n is large enough:

Assumption 2 For each $t = 1, 2, \dots, d$,

$$\sigma_{m_n, t}^2 \in (0, \infty) \quad \text{for all } n \text{ sufficiently large.} \quad (11)$$

In the following, we assume throughout that n is sufficiently large in the sense of (11). Both μ_t and $\sigma_{m_n,t}^2$ do not depend on j because as noted earlier,

$$\begin{aligned} & \mathbf{X}_{n,1}, \mathbf{X}_{n,2}, \dots, \mathbf{X}_{n,r_n} \text{ are i.i.d. random vectors,} \\ & \text{where each } \mathbf{X}_{n,j} = (X_{n,j,1}, X_{n,j,2}, \dots, X_{n,j,d}) \text{ has the same joint distribution } F_n. \end{aligned} \quad (12)$$

This triangular-array setup allows for the distribution F_n to change with n , as is the case in (8). In the special case where $m_n = m_0$ is fixed and only r_n changes with n , then F_n does not depend on n and we readily have a CLT similar to (4), with n there replaced by r_n , $\Sigma_{t,t'}$ replaced by $\text{Cov}[X_{n,j,t}, X_{n,j,t'}]$, where $\text{Cov}[\cdot, \cdot]$ denotes the covariance operator, and $\widehat{\alpha}_n^{\text{MC}}$ replaced by $\widehat{\alpha}_{m_0, r_n}^{\text{RQ}}$.

The general case when m_n varies with n requires a more complicated analysis that involves the covariance structure. We could write the covariance matrix of $\mathbf{X}_{n,j}$ with elements $\text{Cov}[X_{n,j,t}, X_{n,j,t'}]$ for $1 \leq t, t' \leq d$, but for the proof of Theorem 1, it will be more convenient to use the *correlation* matrix $\Sigma_n = \text{Corr}(\mathbf{X}_{n,j})$ instead. Its (t, t') entry is

$$\Sigma_{n,t,t'} = \text{Corr}[X_{n,1,t}, X_{n,1,t'}] = \frac{\text{Cov}[X_{n,1,t}, X_{n,1,t'}]}{\sigma_{m_n,t} \sigma_{m_n,t'}}, \quad (13)$$

which is 1 when $t = t'$ (i.e., on the diagonal). The covariance matrix is then

$$\text{Cov}(\mathbf{X}_{n,j}) = \mathbf{\Gamma}_n \Sigma_n \mathbf{\Gamma}_n, \quad \text{where } \mathbf{\Gamma}_n = \text{diag}(\sigma_{m_n,1}, \sigma_{m_n,2}, \dots, \sigma_{m_n,d}), \quad (14)$$

i.e., $\mathbf{\Gamma}_n$ is a $d \times d$ diagonal matrix that contains the standard deviations from (10).

Assumption 3 There exists a positive-definite $d \times d$ correlation matrix Σ_0 such that $\Sigma_n \rightarrow \Sigma_0$ as $n \rightarrow \infty$.

We now discuss the conditions of this assumption. When m_n is fixed (independent of n), Σ_n is the same for all n , and we just need to assume that Σ_n is positive definite. Another special case where things simplify is if we use independent randomizations across the different values of t . In that case, the correlation matrix is just the identity for all n . For the more general case, the correlations outside the diagonal can be anywhere in $[-1, 1]$. On the right side of (13), the denominator typically converges to 0 when $m_n \rightarrow \infty$, so the numerator must shrink to 0 as well, but it is conceivable that the ratio might not converge. Assumption 3 rules out this possibility. Moreover, requiring the limit Σ_0 to be positive definite simplifies the development by preventing degenerate cases from occurring.

We also assume that each coordinate $X_{n,1,t}$ of $\mathbf{X}_{n,1}$, $t = 1, 2, \dots, d$, obeys a marginal Lindeberg condition (Billingsley 1995, p. 359) specialized to our setting (12), where $I(\cdot)$ denotes the indicator function:

Assumption 4 For each $t = 1, 2, \dots, d$,

$$\frac{1}{\sigma_{m_n,t}^2} \mathbb{E} \left[(X_{n,1,t} - \mu_t)^2 I(|X_{n,1,t} - \mu_t| > w \sqrt{r_n} \sigma_{m_n,t}) \right] \rightarrow 0, \quad \text{as } n \rightarrow \infty, \quad \forall w > 0. \quad (15)$$

A sufficient condition to secure (15) is the Lyapunov condition (Billingsley 1995, Theorem 27.3): there exists $\varepsilon > 0$ such that

$$\frac{1}{r_n^{\varepsilon/2} \sigma_{m_n,t}^{2+\varepsilon}} \mathbb{E} \left[|X_{n,1,t} - \mu_t|^{2+\varepsilon} \right] \rightarrow 0, \quad \text{as } n \rightarrow \infty. \quad (16)$$

Nakayama and Tuffin (2024) also consider conditions (15) and (16) in the univariate case of a single integrand. These conditions impose restrictions on the allocation (m_n, r_n) , on each integrand h_t , and on the RQMC method through $X_{n,1,t}$ in (8) and $\sigma_{m_n,t}^2$ in (10), which Nakayama and Tuffin (2024) further explore.

Theorem 1 Suppose that Assumptions 1, 2, 3 and 4 hold. Then $\widehat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}}$ obeys a d -dimensional CLT:

$$r_n^{1/2} \boldsymbol{\Gamma}_n^{-1} \left(\widehat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}} - \boldsymbol{\mu} \right) \Rightarrow \mathcal{N}_d(\mathbf{0}, \boldsymbol{\Sigma}_0), \quad \text{as } n \rightarrow \infty, \quad (17)$$

where $\boldsymbol{\Gamma}_n$ is from (14), and $\mathcal{N}_d(\mathbf{0}, \boldsymbol{\Sigma}_0)$ represents a d -dimensional random vector with mean vector $\mathbf{0}$ and covariance matrix $\boldsymbol{\Sigma}_0$.

Proof. We will apply Theorem 5 on p. 19 of Gikhman and Skorokhod (1996), which has three conditions (denoted here as a, b, c) that we will show hold. Rather than utilizing that result on the triangular array for $\mathbf{X}_{n,j}$ having n th row from (12), we will consider instead another triangular array with random vectors $\mathbf{Y}_{n,j} = (Y_{n,j,1}, Y_{n,j,2}, \dots, Y_{n,j,d})$, with scaled and centered components defined by

$$Y_{n,j,t} = \frac{X_{n,j,t} - \mu_t}{\sqrt{r_n} \sigma_{m_n,t}}, \quad (18)$$

where $\sigma_{m_n,t}$ is from (10). Specifically, for any given n ,

$$\mathbf{Y}_{n,1}, \mathbf{Y}_{n,2}, \dots, \mathbf{Y}_{n,r_n} \text{ are i.i.d. random vectors, where each } \mathbf{Y}_{n,j} \text{ has the same joint distribution } G_n, \quad (19)$$

which defines the n th row of a triangular array.

The squared denominator in (18) satisfies

$$r_n \sigma_{m_n,t}^2 = \text{Var} \left[\sum_{j=1}^{r_n} X_{n,j,t} \right] \quad (20)$$

which by (10) is the variance of the sum of the t -th components across the r_n i.i.d. random vectors in (12). Thus,

$$\mathbb{E}[Y_{n,j,t}] = 0, \quad \text{Var}[Y_{n,j,t}] = \frac{1}{r_n}, \quad \text{and} \quad s_{n,t}^2 \equiv \text{Var} \left[\sum_{j=1}^{r_n} Y_{n,j,t} \right] = 1, \quad (21)$$

where each does not depend on j . Also, note that

$$\begin{aligned} \text{Cov} \left[\sum_{j=1}^{r_n} Y_{n,j,t}, \sum_{j=1}^{r_n} Y_{n,j,t'} \right] &= \sum_{j=1}^{r_n} \sum_{\ell=1}^{r_n} \text{Cov}[Y_{n,j,t}, Y_{n,\ell,t'}] = \sum_{j=1}^{r_n} \text{Cov}[Y_{n,j,t}, Y_{n,j,t'}] \\ &= \sum_{j=1}^{r_n} \frac{\text{Cov}[X_{n,j,t}, X_{n,j,t'}]}{r_n \sigma_{m_n,t} \sigma_{m_n,t'}} = \frac{\text{Cov}[X_{n,1,t}, X_{n,1,t'}]}{\sigma_{m_n,t} \sigma_{m_n,t'}} = \Sigma_{n,t,t'} \end{aligned}$$

for $\Sigma_{n,t,t'}$ in (13), where the second equality stems for the independence of the $Y_{n,j,t}$ across the values of j , and the last one from the identical distribution for each j ($j = 1, \dots, r_n$). For each $t = 1, 2, \dots, d$, (21) implies that $\max_{j=1,2,\dots,r_n} \text{Var}[Y_{n,j,t}] = 1/r_n \rightarrow 0$ as $n \rightarrow \infty$ since $r_n \rightarrow \infty$ by Assumption 1. Thus, condition a holds by Assumption 2, (21), and (19).

Our Assumption 3 is slightly stronger than condition b, as the latter does not assume that $\boldsymbol{\Sigma}_0$ is positive-definite.

Next we turn to condition c. Recall that $s_{n,t}^2 = \text{Var}[\sum_{j=1}^{r_n} Y_{n,j,t}]$ in (21). Condition c requires a marginal Lindeberg condition for each component t of the random vectors in (19): i.e., we need that for each $w > 0$,

$$q_{n,t}(w) \equiv \frac{1}{s_{n,t}} \sum_{j=1}^{r_n} \mathbb{E} [Y_{n,j,t}^2 I(|Y_{n,j,t}| > w s_{n,t})] \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

But since $s_{n,t}^2 = 1$ by (21) and because $Y_{n,j,t}$, $j = 1, 2, \dots, r_n$, are i.i.d. by (19), we can use (18) to write

$$\begin{aligned} q_{n,t}(w) &= r_n \mathbb{E} \left[Y_{n,1,t}^2 I(|Y_{n,1,t}| > w) \right] \\ &= r_n \mathbb{E} \left[\left(\frac{X_{n,1,t} - \mu_t}{\sqrt{r_n} \sigma_{m_n,t}} \right)^2 I \left(\left| \frac{X_{n,1,t} - \mu_t}{\sqrt{r_n} \sigma_{m_n,t}} \right| > w \right) \right] \\ &= \frac{1}{\sigma_{m_n,t}^2} \mathbb{E} \left[(X_{n,1,t} - \mu_t)^2 I(|X_{n,1,t} - \mu_t| > w \sqrt{r_n} \sigma_{m_n,t}) \right], \end{aligned}$$

where $\sigma_{m_n,t}^2$ is defined in (10). Thus, Assumption 4 ensures that condition c holds.

Since all of the conditions of Theorem 5 on p. 19 of Gikhman and Skorokhod (1996) hold, we obtain

$$\sum_{j=1}^{r_n} \mathbf{Y}_{n,j} \Rightarrow \mathcal{N}_d(\mathbf{0}, \mathbf{\Sigma}_0), \quad \text{as } n \rightarrow \infty, \quad (22)$$

for $\mathbf{\Sigma}_0$ from Assumption 3. Note that (22) does not need any scaling by $\sqrt{r_n}$ or $1/\sqrt{r_n}$ because each $Y_{n,j,t}$ has standard deviation $1/\sqrt{r_n}$ by (18) and (20), so each of the r_n i.i.d. summands in (22) already includes the appropriate scaling for a CLT.

Now we want to convert (22) into a multivariate CLT for $\hat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}} = (1/r_n) \sum_{j=1}^{r_n} \mathbf{X}_{n,j}$, where each $\mathbf{X}_{n,j}$ is as in (12) and (8). Doing this requires some care because each component $Y_{n,j,t}$ in (18) of each random vector $\mathbf{Y}_{n,j}$ in (22) is centered by a different μ_t and scaled by a different $\sqrt{r_n} \sigma_{m_n,t}^2$. By (18), we first rewrite (22) as

$$\left(\sum_{j=1}^{r_n} \frac{X_{n,j,t} - \mu_t}{\sqrt{r_n} \sigma_{m_n,t}} : t = 1, 2, \dots, d \right) \Rightarrow \mathcal{N}_d(\mathbf{0}, \mathbf{\Sigma}_0), \quad \text{as } n \rightarrow \infty, \quad (23)$$

which can then alternatively be expressed in terms of $\hat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}}$ as in (17), where the $r_n^{1/2}$ on the left side of (17) comes from $r_n \hat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}} = \sum_{j=1}^{r_n} \mathbf{X}_{n,j}$ and the $\sqrt{r_n}$ in the denominator on the left side of (23). This completes the proof. \square

We can bring the $\mathbf{\Sigma}_0$ to the left side of (17) by premultiplying each side of (17) by $\mathbf{\Sigma}_0^{-1/2}$, whose existence is guaranteed by Assumption 3, leading to $r_n^{1/2} \mathbf{\Sigma}_0^{-1/2} \mathbf{\Gamma}_n^{-1} (\hat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}} - \boldsymbol{\mu}) \Rightarrow \mathcal{N}_d(\mathbf{0}, \mathbf{I}_d)$ as $n \rightarrow \infty$, where \mathbf{I}_d is the $d \times d$ identity matrix. Also, we can further replace $\mathbf{\Sigma}_0^{-1/2}$ on the left side with $\mathbf{\Sigma}_n^{-1/2}$ because Assumption 3 ensures for n sufficiently large the existence of $\mathbf{\Sigma}_n^{-1/2}$ and its convergence to $\mathbf{\Sigma}_0^{-1/2}$ as $n \rightarrow \infty$, so applying Slutsky's theorem (Serfling 1980, Theorem 1.5.4) yields

$$r_n^{1/2} \mathbf{\Sigma}_n^{-1/2} \mathbf{\Gamma}_n^{-1} (\hat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}} - \boldsymbol{\mu}) \Rightarrow \mathcal{N}_d(\mathbf{0}, \mathbf{I}_d), \quad \text{as } n \rightarrow \infty. \quad (24)$$

In the CLTs in (23), (17), and (24), each $\sigma_{m_n,t} \rightarrow 0$ as $n \rightarrow \infty$ when $m_n \rightarrow \infty$, where $\sigma_{m_n,t}$ appears as a diagonal entry in $\mathbf{\Gamma}_n$ from (14) in (17) and (24). This is needed to ensure a Gaussian limit by counteracting that $\hat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}} - \boldsymbol{\mu}$ can shrink faster than $1/r_n^{1/2}$ because typically RQMC converges faster than MC when $m_n \rightarrow \infty$. For example, if h_t has BVHK, then we often have that $\sigma_{m_n,t} = O((\ln m_n)^s / m_n)$ as $m_n \rightarrow \infty$ (e.g., see eqs. (12) and (37) of Nakayama and Tuffin 2024), but recall that this $O(\cdot)$ notation only gives an upper bound; the actual convergence rate for $\sigma_{m_n,t}$ can be faster than that. Faster rates can also be proved under various conditions (L'Ecuyer 2018).

Finally we use the delta method to get a CLT for $\hat{\boldsymbol{\alpha}}_{m_n, r_n}^{\text{RQ}} = g(\hat{\boldsymbol{\mu}}_{m_n, r_n, 1}^{\text{RQ}}, \hat{\boldsymbol{\mu}}_{m_n, r_n, 2}^{\text{RQ}}, \dots, \hat{\boldsymbol{\mu}}_{m_n, r_n, d}^{\text{RQ}})$ in (9).

Theorem 2 Suppose that Assumptions 1, 2, 3 and 4 hold. Also suppose that the function g in (1) has non-zero differential at $\boldsymbol{\mu}$. Then $\widehat{\alpha}_{m_n, r_n}^{\text{RQ}}$ in (9) obeys a CLT

$$\frac{\sqrt{r_n}}{\tau_n} [\widehat{\alpha}_{m_n, r_n}^{\text{RQ}} - \alpha] \Rightarrow \mathcal{N}_1(0, 1), \quad \text{as } n \rightarrow \infty, \quad (25)$$

where

$$\tau_n^2 = \nabla g(\boldsymbol{\mu})^\top \boldsymbol{\Gamma}_n \boldsymbol{\Sigma}_n \boldsymbol{\Gamma}_n \nabla g(\boldsymbol{\mu}) = \sum_{t=1}^d \sum_{t'=1}^d g_t(\boldsymbol{\mu}) \sigma_{m_n, t} g_{t'}(\boldsymbol{\mu}) \sigma_{m_n, t'} \Sigma_{n, t, t'}. \quad (26)$$

For example, if each h_t has BVHK, then $\tau_n/\sqrt{r_n}$ is typically $O((\ln m_n)^s/(m_n \sqrt{r_n}))$, which shrinks faster (as $n \rightarrow \infty$) than $\tau_{\text{MC}}/\sqrt{n} = \Theta(1/\sqrt{m_n r_n})$ from (4) when using MC with an equal budget parameter $n = m_n r_n$. It shows the gain that can be obtained with RQMC. As with MC, using CRN across the different values of t improves the variance compared to using independent randomizations if and only if τ_n^2 as given in (26) is smaller than the one we get if we replace $\boldsymbol{\Sigma}_n$ by the identity in that equation.

We now give a heuristic justification of Theorem 2. By (8), we have for each t that $\sqrt{r_n}[\widehat{\mu}_{m_n, r_n, t}^{\text{RQ}} - \mu_t] = \sum_{j=1}^{r_n} (X_{n, j, t} - \mu_t)/\sqrt{r_n}$ is roughly equal in distribution to $\sigma_{m_n, t} Z_t$ for $Z_t \sim \mathcal{N}_1(0, 1)$ for large n by the CLT (23) since $\Sigma_{n, t, t} = \Sigma_{0, t, t} = 1$ from (21). Also, $(Z_1, Z_2, \dots, Z_d) \sim \mathcal{N}_d(\mathbf{0}, \boldsymbol{\Sigma}_n) \approx \mathcal{N}_d(\mathbf{0}, \boldsymbol{\Sigma}_0)$ for large n by Assumption 3. Then a first-order Taylor approximation for $\widehat{\alpha}_{m_n, r_n}^{\text{RQ}} = g(\widehat{\boldsymbol{\mu}}_{m_n, r_n, 1}^{\text{RQ}}, \widehat{\boldsymbol{\mu}}_{m_n, r_n, 2}^{\text{RQ}}, \dots, \widehat{\boldsymbol{\mu}}_{m_n, r_n, d}^{\text{RQ}})$ suggests that for large n ,

$$\sqrt{r_n} [\widehat{\alpha}_{m_n, r_n}^{\text{RQ}} - \alpha] \approx \sum_{t=1}^d g_t(\boldsymbol{\mu}) \sqrt{r_n} [\widehat{\mu}_{m_n, r_n, t}^{\text{RQ}} - \mu_t] \approx \sum_{t=1}^d g_t(\boldsymbol{\mu}) \sigma_{m_n, t} Z_t,$$

whose variance is given by (26). Also, in (26) note that the factors $\sigma_{m_n, t} \rightarrow 0$ and $\sigma_{m_n, t'} \rightarrow 0$ as $n \rightarrow \infty$ when $m_n \rightarrow \infty$, leading to $\tau_n \rightarrow 0$ as $n \rightarrow \infty$. So dividing by τ_n in (25) is needed to ensure a Gaussian limit to counteract that $\widehat{\alpha}_{m_n, r_n}^{\text{RQ}} - \alpha$ can shrink faster than $1/\sqrt{r_n}$ because typically RQMC converges faster than MC when $m_n \rightarrow \infty$.

5 CONCLUDING REMARKS

We considered an RQMC plug-in estimator of the estimand α in (1), which is a smooth function of a vector $\boldsymbol{\mu}$ of unknown means. We applied RQMC to estimate $\boldsymbol{\mu}$ using r_n i.i.d. randomizations of m_n points from a low-discrepancy sequence, and we established a CLT (Theorem 2) for the RQMC plug-in estimator as $n \approx m_n r_n \rightarrow \infty$. Our argument first showed a multivariate CLT for the RQMC estimator of $\boldsymbol{\mu}$ (Theorem 1), and then applied the delta method to obtain Theorem 2, which gives a CLT for the estimator of α .

There are several directions for future work. For the case of the RQMC estimator of a single mean, Nakayama and Tuffin (2024) provide CLT refinements based on the allocation (m_n, r_n) in terms of properties of a single integrand (e.g., BVHK or simply bounded) and the randomized low-discrepancy sequence. Also, that paper establishes the asymptotic validity of confidence intervals. In our Theorem 2 we may similarly want in practice to replace τ_n in (25) by an estimator. A natural solution is to estimate τ_n^2 in (26) by say

$$\widehat{\tau}_n^2 = \nabla g(\widehat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}})^\top \widehat{\boldsymbol{C}}_n \nabla g(\widehat{\boldsymbol{\mu}}_{m_n, r_n}^{\text{RQ}}) \quad (27)$$

where $\widehat{\boldsymbol{C}}_n$ is the matrix of empirical covariances, used to estimate the covariance matrix $\boldsymbol{\Gamma}_n \boldsymbol{\Sigma}_n \boldsymbol{\Gamma}_n$. Showing that the CLT still holds and that we obtain an *asymptotically valid confidence interval* (AVCI) is of interest. Our triangular-array setting in (12) presents a theoretical complication in establishing an AVCI: it is *not* sufficient to merely show that $\widehat{\tau}_n^2$ consistently estimates τ_n^2 in the sense that $\widehat{\tau}_n^2 - \tau_n^2 \Rightarrow 0$ as $n \rightarrow \infty$. Because $\tau_n^2 \rightarrow 0$ when $m_n \rightarrow \infty$ (see the paragraph before Section 5), it could be the case that $\widehat{\tau}_n^2$ shrinks at a different

rate than τ_n^2 does, which may lead to $\tau_n/\widehat{\tau}_n \not\Rightarrow 1$ and (25) not holding when $\widehat{\tau}_n$ replaces τ_n . Instead, what needs to be shown is that $\widehat{\tau}_n^2/\tau_n^2 \Rightarrow 1$, which is more complicated than simply proving consistency. We plan to also carry out analyses for the estimation of α similar to the principles developed in Nakayama and Tuffin (2024) when estimating a single mean by RQMC. In particular, we want to specialize Assumption 1 to consider allocations $(m_n, r_n) = (n^c, n^{1-c})$ for some $c \in [0, 1]$, and determine the largest values of c (to gain the benefits of the fast convergence rate of RQMC) for which a CLT and AVCI hold.

ACKNOWLEDGMENTS

This work has been supported in part by the National Science Foundation under Grant No. CMMI-2345330, and by the Natural Sciences and Engineering Research Council of Canada (NSERC) Grant RGPIN-2024-05864. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or NSERC.

REFERENCES

- Apostol, T. M. 1974. *Mathematical Analysis*. 2nd ed. Reading, MA: Addison-Wesley.
- Asmussen, S., and P. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*. New York: Springer.
- Basu, K., and R. Mukherjee. 2017. "Asymptotic Normality of Scrambled Geometric Net Quadrature". *Annals of Statistics* 45(4):1759–1788.
- Billingsley, P. 1995. *Probability and Measure*. 3rd ed. New York: John Wiley and Sons.
- Gikhman, I. I., and A. V. Skorokhod. 1996. *Introduction to the Theory of Random Processes*. Dover books on mathematics. Mineola, NY: Dover Publications.
- Glynn, P. W., P. L'Ecuyer, and M. Adès. 1991. "Gradient Estimation for Ratios". In *1991 Winter Simulation Conference (WSC)*, 986–993 <https://doi.org/10.1109/WSC.1991.185714>.
- Goyal, A., P. Shahabuddin, P. Heidelberger, V. Nicola, and P. W. Glynn. 1992. "A Unified Framework for Simulating Markovian Models of Highly Dependable Systems". *IEEE Transactions on Computers* C-41(1):36–51.
- He, Z., and L. Zhu. 2017. "Asymptotic Normality of Extensible Grid Sampling". *Statistics and Computing* 29(1):53–65.
- L'Ecuyer, P. 2018. "Randomized Quasi-Monte Carlo: An Introduction for Practitioners". In *Monte Carlo and Quasi Monte Carlo Methods: MCQMC 2016*, edited by P. W. Glynn and A. B. Owen, 29–52. Berlin: Springer-Verlag.
- L'Ecuyer, P. 2024. "Stochastic Simulation and Monte Carlo Methods". Draft Textbook, <https://www-labs.iro.umontreal.ca/~lecuyer/ift6561/book.pdf>, accessed 10th April.
- L'Ecuyer, P., and C. Lemieux. 2000. "Variance Reduction via Lattice Rules". *Management Science* 46(9):1214–1235.
- L'Ecuyer, P., P. Marion, M. Godin, and F. Puchhammer. 2022. "A Tool for Custom Construction of QMC and RQMC Point Sets". In *Monte Carlo and Quasi-Monte Carlo Methods: MCQMC 2020*, edited by A. Keller, 51–70. Berlin: Springer.
- L'Ecuyer, P., D. Munger, and B. Tuffin. 2010. "On the Distribution of Integration Error by Randomly-Shifted Lattice Rules". *Electronic Journal of Statistics* 4:950–993.
- Lemieux, C. 2009. *Monte Carlo and Quasi-Monte Carlo Sampling*. Series in Statistics. New York: Springer.
- Loh, W.-L. 2003. "On the Asymptotic Distribution of Scrambled Net Quadrature". *Annals of Statistics* 31(4):1282–1324.
- Nakayama, M. K., and B. Tuffin. 2023. "Efficiency of Estimating Functions of Means in Rare-Event Contexts". In *2023 Winter Simulation Conference (WSC)*, 351–362 <https://doi.org/10.1109/WSC60868.2023.10407217>.

- Nakayama, M. K., and B. Tuffin. 2024. "Sufficient Conditions for Central Limit Theorems and Confidence Intervals for Randomized Quasi-Monte Carlo Methods". *ACM Transactions on Modeling and Computer Simulation* 34(3):1–38.
- Niederreiter, H. 1992. *Random Number Generation and Quasi-Monte Carlo Methods*, Volume 63 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Philadelphia: SIAM.
- Rhee, C.-H., and P. W. Glynn. 2023. "Lyapunov Conditions for Differentiability of Markov Chain Expectations". *Mathematics of Operations Research* 48(4):2019–2042.
- Serfling, R. J. 1980. *Approximation Theorems of Mathematical Statistics*. New York: John Wiley and Sons.

AUTHOR BIOGRAPHIES

MARVIN K. NAKAYAMA is a professor in the Department of Computer Science at the New Jersey Institute of Technology. He received an M.S. and a Ph.D. in operations research from Stanford University. His research interests include simulation, modeling, statistics, and risk analysis. His email: marvin@njit.edu.

BRUNO TUFFIN received his PhD degree in applied mathematics from the University of Rennes 1, France, in 1997. Since then, he has been with Inria in Rennes. His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for performance evaluation. His email: bruno.tuffin@inria.fr.

PIERRE L'ECUYER is a Professor in the Département d'Informatique et de Recherche Opérationnelle, Université de Montréal, Canada. He has published over 300 scientific articles and has developed software libraries and systems for random number generation and stochastic simulation. His email address is lecuyer@iro.umontreal.ca and his website is <https://www.iro.umontreal.ca/~lecuyer>