

EFFICIENT COMPUTATION FOR STRATIFIED SPLITTING

Peter W. Glynn

Department of Management Science
and Engineering
Stanford University
Stanford, CA 94305, USA

Zeyu Zheng

Department of Industrial Engineering
and Operations Research
University of California, Berkeley
Berkeley, CA 94720, USA

ABSTRACT

Many applications in the areas of finance, environments, service systems, and statistical learning require the computation for a general function of the expected performances associated with one or many random objects. When complicated stochastic systems are involved, such computation needs to be done by stochastic simulation and the computational cost can be expensive. We design simulation algorithms that exploit the common random structure shared by all random objects, referred to as stratified splitting. We discuss the optimal simulation budget allocation problem for the proposed algorithms. A brief numerical experiment is conducted to illustrate the performance of the proposed algorithm with various budget allocation rules.

1 INTRODUCTION

Computing a function of the expectations of many random objects finds various applications in the area of operations research. In financial risk management, a dealer needs to evaluate the summation of the expected risk exposures by all clients under certain market conditions. In disaster risk management, it is needed to evaluate the overall expected losses of all households that suffer from a regional wildfire. In inventory management with a large number of warehouses across the nation, an interest is to compute the expected aggregated stock-out risks for a critical product from all warehouses under some market conditions, such as the outbreak of a pandemic.

Quantitatively, this task is formulated as computing the following quantity given a number of random objects $X(1), X(2), \dots, X(d)$,

$$\alpha = g(\mu_1, \mu_2, \dots, \mu_d), \quad (1)$$

where g is a deterministic function that maps \mathbf{R}^d to \mathbf{R} and $\mu_i \triangleq \mathbb{E}X(i)$ for $1 \leq i \leq d$. The dimension parameter d is a positive integer that represents the number of random objects in consideration. Even the special case where $g(\mu_1, \mu_2, \dots, \mu_d) = \sum_{i=1}^d \mu_i$ already observes a number of applications. In these applications, the random objects often involve complicated stochastic systems and the computation of $\mu_i = \mathbb{E}X(i)$ needs to be done by Monte Carlo simulation, so as the computation of $\alpha = g(\mu_1, \mu_2, \dots, \mu_d)$. In presence of stochastic errors, to obtain an accurate enough estimate of α may need a large number of simulation replications for all random objects. The computational task can therefore be computationally expensive and it is beneficial to explore approaches that reduce the computational cost needed to achieve a certain accuracy of computing α . For many applications, including the ones listed above, the random objects in consideration share some common random structure Z .

In financial risk management, all clients are affected by the same complicated market condition. The market condition itself is an outcome of a complicated stochastic process. In disaster risk management, all households are affected by the path and severity of the regional wildfire, where the regional wildfire observes complicated stochastic uncertainties. When computing the function of the expectations of the risk of all households using simulation, the shared common random structure enables us to couple the

simulations of all households by first simulating the common random structure (wildfire) and then simulating the conditional expectations for each household.

Motivated by these applications, we consider settings where the random object $X(i)$'s are generated conditional on some common random quantity Z . As described in Bratley et al. (1987), this computational structure is also called ‘‘stratified splitting.’’ The simulation of Z , as a part of the simulation task of each $X(i)$, is also expensive, though shared by all random objects. In this work, we develop a simulation algorithm to compute $\alpha = g(\mu_1, \mu_2, \dots, \mu_d)$ by first simulating Z and then simulating the conditional expectations $\mathbb{E}(X(i)|Z)$'s. The proposed algorithm naturally gives rise to a simulation budget allocation problem, in terms of how to allocate computation resources to the simulation of Z and the simulation of each conditional expectation. Our contribution in this work is to discuss the optimal budget allocation for the one-dimensional case ($d = 1$) and the general-dimensional case (arbitrary d).

This computational task is described by Bratley, Fox, and Schrage (1987) only for the case $d = 1$. However, even for $d = 1$, no optimal budget allocation appears to have been documented in the literature. In the literature of nested simulation (for example, Gordy and Juneja (2010), Broadie et al. (2011), Dang et al. (2019)), a computational task is also divided into two parts, for which they call an inner simulation and outer simulation. Their computation target is different from ours and one of their main focus is on the bias variance trade-off. Our work has a different computation target and there is no biased estimator involved. Our work is also broadly connected to the concept of conditional Monte Carlo methods, where the general idea is to strategically identify conditions so that part of the computation can be done in closed form; see Asmussen and Glynn (2007) and Fu and Hu (2012) for reference.

The rest of this paper is organized as follows. Section 2 discusses the simulation algorithm and the asymptotic efficiency. Section 3 discusses the optimal budget allocation that achieves the optimal asymptotic efficiency when $d = 1$. Section 4 extends Section 3 to the general-dimensional case and propose a heuristic solution for the optimal budget allocation problem. Section 5 provides a brief numerical example and Section 6 concludes.

2 ALGORITHM AND ASYMPTOTIC EFFICIENCY

In this section we describe the simulation algorithm that is used to compute

$$\alpha = g(\mu_1, \mu_2, \dots, \mu_d), \quad (2)$$

by exploiting the common random structure Z and we then show the asymptotic efficiency of the proposed algorithm. The algorithm is given as follows.

1. Simulate a copy of Z .
2. Conditional on Z , for each $i = 1, 2, \dots, d$, simulate m_i independent copies of $X(i)$ from the distribution $P(X(i) \in \cdot | Z)$, denoted as $X_1(Z, i), X_2(Z, i), \dots, X_{m_i}(Z, i)$.
3. Compute the sample average

$$\bar{X}_{m_i}(Z, i) \triangleq \frac{1}{m_i} \sum_{j=1}^{m_i} X_j(Z, i)$$

for $i = 1, 2, \dots, d$.

4. Repeat 1-3 for n times independently, collect n independent copies of Z as Z_1, Z_2, \dots, Z_n , and then compute

$$\left(\frac{1}{n} \sum_{k=1}^n \bar{X}_{m_1}(Z_k, 1), \frac{1}{n} \sum_{k=1}^n \bar{X}_{m_2}(Z_k, 2), \dots, \frac{1}{n} \sum_{k=1}^n \bar{X}_{m_d}(Z_k, d) \right).$$

5. Compute the estimator $\hat{\alpha}_n$ for α as

$$\hat{\alpha}_n = g \left(\frac{1}{n} \sum_{k=1}^n \bar{X}_{m_1}(Z_k, 1), \frac{1}{n} \sum_{k=1}^n \bar{X}_{m_2}(Z_k, 2), \dots, \frac{1}{n} \sum_{k=1}^n \bar{X}_{m_d}(Z_k, d) \right).$$

With a given sample size n , this simulation algorithm generates an estimator $\hat{\alpha}_n$. To study the efficiency of the estimator (and the algorithm), we need to specify the computational cost of the estimator and the variance of the estimator. We introduce the following notation.

1. Let a denote the expected time to generate a copy of Z , which is assumed to be positive and finite.
2. Let b_i denote the expected time to generate a copy of $X_1(Z, i)$ for $i = 1, 2, \dots, d$, which is assumed to be positive and finite.
3. Let σ^2 denote the variance of $\nabla g(\mu) \cdot (\mathbb{E}(X(1)|Z), \mathbb{E}(X(2)|Z), \dots, \mathbb{E}(X(d)|Z))^\top$, which is assumed to be positive and finite.
4. Let σ_i^2 denote $\mathbb{E}[\text{Var}(\frac{\partial g(\mu)}{\partial x_i} X_1(Z, i)|Z)]$ for $i = 1, 2, \dots, d$, which is assumed to be positive and finite.

Now, let the computation time to generate the estimator $\hat{\alpha}_n$ be T_n . Note that T_n can be a random variable itself, is correlated with $\hat{\alpha}_n$, and is increasing in n . Given a computer time budget $c > 0$, let $N(c)$ be

$$N(c) = \max\{n \geq 0 : T_n \leq c\}.$$

The estimator for α that is available with computational budget c is then $\hat{\alpha}_{N(c)}$. Using the result in (Glynn and Whitt 1992), a central limit theorem, based on the computer time budget c , is available as follows.

Proposition 1 Assume that $a, b_i, \sigma^2, \sigma_i^2, m_i$ are all positive and finite for $i = 1, 2, \dots, n$. When $c \rightarrow \infty$,

$$c^{\frac{1}{2}}(\hat{\alpha}_{N(c)} - \alpha) \Rightarrow (\lambda^{-1}\eta^2)^{\frac{1}{2}}N(0, 1)$$

where

$$\lambda^{-1} = a + \sum_{i=1}^d m_i b_i$$

$$\eta^2 = \sigma^2 + \sum_{i=1}^d \frac{\sigma_i^2}{m_i}.$$

With this central limit theorem in hand, we are able to construct asymptotically valid confidence intervals for α . In particular, we denote \tilde{a} as a constant such that $P(-\tilde{a} \leq N(0, 1) \leq \tilde{a}) = 0.9\%$ (say), then the interval

$$\left[\hat{\alpha}_{N(c)} - \tilde{a}(\lambda^{-1}\eta^2)^{\frac{1}{2}} \frac{1}{\sqrt{c}}, \hat{\alpha}_{N(c)} + \tilde{a}(\lambda^{-1}\eta^2)^{\frac{1}{2}} \frac{1}{\sqrt{c}} \right]$$

is an asymptotically 90% confidence interval for α .

Moreover, the asymptotic efficiency of this estimator $\hat{\alpha}_{N(c)}$ is determined by the asymptotic variance $\lambda^{-1}\eta^2$, which is controlled by the choices of integers m_1, m_2, \dots, m_d given the model primitives. The asymptotic efficiency is higher if the asymptotic variance is smaller. This criterion can guide the design of the simulation algorithm and especially the budget allocation.

3 OPTIMAL BUDGET ALLOCATION FOR $d = 1$

In Section 2, the proposed simulation algorithm adopts the flexibility in the choices of integers m_1, m_2, \dots, m_d . They reflect the number of replications implemented on each random object $X(i)$ conditional on one copy of Z . In this section, we consider the special case of $d = 1$ and provide an optimal choice of $m_1^* \in \mathbf{Z}_+$ that minimizes the asymptotic variance of the associated estimator constructed by the algorithm.

The optimization problem is given as follows.

$$\min_{m_1 \in \mathbf{Z}_+} f(m_1) \triangleq (a + bm_1) \left(\sigma^2 + \frac{\sigma_1^2}{m_1} \right).$$

Note that this optimization problem gives the best possible allocation m_1^* in terms of asymptotic variance when all model parameters $a, b, \sigma^2, \sigma_1^2$ are known. In practice when these parameters are not known, it is standard to develop two-stage procedures to first estimate the parameters and then commit the allocation based on the estimated parameters. It is also of interest to develop fully adaptive procedure where the estimate of the the parameters is adaptively updated throughout the algorithm as more simulation replications are revealed. Such two-stage or fully adaptive procedures, of course, bring a gap compared to the optimal procedure when all parameters are known a priori. The practical performance of two-stage and fully adaptive procedures are deferred to a future work.

We denote m_1^* as $\operatorname{argmin}_{m_1 \in \mathbf{Z}_+} f(m_1)$, which is an arbitrary optimizer of $f(\cdot)$ in all positive integers. We show the following result.

Proposition 2 The optimizing integer m_1^* as $\operatorname{argmin}_{m_1 \in \mathbf{Z}_+} f(m_1)$ satisfies

$$-\frac{1}{2} + \sqrt{\frac{\sigma_1^2 a}{\sigma^2 b} + \frac{1}{4}} \leq m_1^* \leq \frac{1}{2} + \sqrt{\frac{\sigma_1^2 a}{\sigma^2 b} + \frac{1}{4}}.$$

Proof. In this one-dimensional case, it is evident that the function $f(x) = (a + bx)(\sigma^2 + \sigma_1^2/x)$, if extended to the domain of $x \in (0, +\infty)$, is a convex function of x . Therefore in order to explicitly identify m_1^* , it suffices for m_1^* to satisfy that

$$\begin{aligned} f(m_1^*) &\leq f(m_1^* + 1) \\ f(m_1^*) &\leq f(m_1^* - 1), \end{aligned}$$

assuming that $m_1^* \geq 1$. This set of inequalities reduce to

$$\begin{aligned} (a + bm_1^*)(\sigma^2 + \frac{\sigma_1^2}{m_1^*}) &\leq (a + b(m_1^* + 1))(\sigma^2 + \frac{\sigma_1^2}{m_1^* + 1}) \\ (a + bm_1^*)(\sigma^2 + \frac{\sigma_1^2}{m_1^*}) &\leq (a + b(m_1^* - 1))(\sigma^2 + \frac{\sigma_1^2}{m_1^* - 1}), \end{aligned}$$

and then to

$$\begin{aligned} 0 &\leq b\sigma^2 m_1^*(m_1^* + 1) - \sigma_1^2 a \\ 0 &\leq -b\sigma^2 m_1^*(m_1^* - 1) + \sigma_1^2 a. \end{aligned}$$

Note that

$$0 \leq b\sigma^2 m_1^*(m_1^* + 1) - \sigma_1^2 a$$

implies that

$$0 \leq b\sigma^2((m_1^*)^2 + m_1^* + \frac{1}{4}) - \sigma_1^2 a - \frac{1}{4}b\sigma^2$$

and therefore

$$\sqrt{\frac{\sigma_1^2 a}{\sigma^2 b} + \frac{1}{4}} - \frac{1}{2} \leq m_1^*.$$

Similarly, the inequality $0 \leq -b\sigma^2 m_1^*(m_1^* - 1) + \sigma_1^2 a$ implies that

$$m_1^* \leq \frac{1}{2} + \sqrt{\frac{\sigma_1^2 a}{\sigma^2 b} + \frac{1}{4}}.$$

Hence, m_1^* is an integer that satisfies

$$-\frac{1}{2} + \sqrt{\frac{\sigma_1^2 a}{\sigma^2 b} + \frac{1}{4}} \leq m_1^* \leq \frac{1}{2} + \sqrt{\frac{\sigma_1^2 a}{\sigma^2 b} + \frac{1}{4}},$$

the range of which covers exactly a unique integer if $\sqrt{\frac{\sigma_1^2 a}{\sigma^2 b} + \frac{1}{4}}$ is not in $\{\mathbf{Z}_+ + 1/2\}$. \square

4 HEURISTIC SOLUTION OF OPTIMAL BUDGET ALLOCATION FOR ARBITRARY d

In Section 3, we provided a closed-form optimal budget allocation for the case $d = 1$. In this section, we discuss the optimal budget allocation for arbitrary $d \in \mathbf{Z}_+$. The optimization problem to identify the optimal budget allocation $(m_1^*, m_2^*, \dots, m_d^*)$ is given by

$$\begin{aligned} \min f(m_1, m_2, \dots, m_d) &\triangleq (a + \sum_{i=1}^d b_i m_i) (\sigma^2 + \sum_{j=1}^d \frac{\sigma_j^2}{m_j}) \\ \text{s.t. } (m_1, m_2, \dots, m_d) &\in \mathbf{Z}_+^d. \end{aligned} \quad (3)$$

The optimal solution among all the integer points is not straightforward to derive for the general-dimension case. We first present a heuristic solution. This heuristic solution is given by

$$(\tilde{m}_1, \tilde{m}_2, \dots, \tilde{m}_d) \triangleq (\lfloor x_1^* \rfloor, \lfloor x_2^* \rfloor, \dots, \lfloor x_d^* \rfloor)$$

where $\lfloor \cdot \rfloor$ represents the operation of taking the nearest integer, and x_i^* is given by

$$x_i^* = \sqrt{\frac{\sigma_i^2 a}{\sigma^2 b_i}}$$

for $i = 1, 2, \dots, d$. This heuristic solution is motivated by the following result.

Proposition 3 Define a function $g : \mathbb{R}_+^d \rightarrow \mathbb{R}$ as

$$g(x) \triangleq (a + \sum_{i=1}^d b_i x_i) (\sigma^2 + \sum_{j=1}^d \frac{\sigma_j^2}{x_j})$$

where $x = (x_1, x_2, \dots, x_d) \in \mathbb{R}_+^d$. The minimum value of $g(x)$ over $x \in \mathbb{R}_+^d$ is achieved at the point

$$x^* = (x_1^*, x_2^*, \dots, x_d^*) = \left(\sqrt{\frac{\sigma_1^2 a}{\sigma^2 b_1}}, \sqrt{\frac{\sigma_2^2 a}{\sigma^2 b_2}}, \dots, \sqrt{\frac{\sigma_d^2 a}{\sigma^2 b_d}} \right).$$

Proof. This result can be proved by applying an extended version of Cauchy–Schwarz inequality. Specifically,

$$(a + \sum_{i=1}^d b_i x_i) (\sigma^2 + \sum_{j=1}^d \frac{\sigma_j^2}{x_j}) \geq (\sqrt{a\sigma^2} + \sum_{i=1}^d \sqrt{b_i x_i \frac{\sigma_i^2}{x_i}})^2 = (\sqrt{a\sigma^2} + \sum_{i=1}^d \sqrt{b_i \sigma_i^2})^2.$$

The right-hand side does not depend on x_i 's and the inequality binds only if

$$\frac{a}{\sigma^2} = \frac{b_i x_i^2}{\sigma_i^2}$$

for all $i = 1, 2, \dots, d$. Therefore,

$$x_i^* = \sqrt{\frac{\sigma_i^2 a}{\sigma^2 b_i}}$$

for $i = 1, 2, \dots, d$. □

We note that the heuristic solution $(\tilde{m}_1, \tilde{m}_2, \dots, \tilde{m}_d) \triangleq (\lfloor x_1^* \rfloor, \lfloor x_2^* \rfloor, \dots, \lfloor x_d^* \rfloor)$ may not be the optimal solution of (3) and it is challenging to identify the optimal budget allocation $(m_1^*, m_2^*, \dots, m_d^*)$. One key obstacle is that the function $h(x)$ is not guaranteed to be convex in x in general.

5 NUMERICAL EXPERIMENT

In this section, we conduct a brief numerical experiment to illustrate the performance of the proposed algorithm with different choices of the budget allocation (m_1, m_2, \dots, m_d) . Consider the task of computing the aggregated number of expected arrivals from d different sources that share a common baseline factor. That is,

$$\alpha = g(\mu_1, \mu_2, \dots, \mu_d) = \sum_{i=1}^d \mathbb{E}(X(i)),$$

where for $i = 1, 2, \dots, d$

$$X(i) = N_i(2^{-i}Z),$$

in which $(N_i(t) : t \geq 0)$ is a unit-rate Poisson process that is independent of all other randomness and Z is a random variable that is shared by the definition of all $X(i)$'s. Specifically, the random variable Z is the terminal state Y_t of an auto-regressive stochastic process $Y = (Y_i : i = 0, 1, 2, \dots, t)$ where t is a positive integer. In this experiment, we consider the stochastic process Y defined by

$$\begin{aligned} Y_0 &= 1000, \\ Y_{i+1} &= \frac{1}{2}Y_i + \frac{1}{2}\varepsilon_{i+1}, \quad i = 1, 2, \dots, t-1, \end{aligned}$$

where $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_t$ is a sequence of i.i.d. random variables with probability distribution $N(1000, 10)$. We make such stylized distribution assumptions so that the probability distribution of Y_t has a closed-form representation, but we have in mind general settings so that the generation of Y_t needs to be done sequentially period by period. For simplicity, the computational cost is taken as the number of random variables that are needed to simulate in the generation of the associated random quantity. In this task, the computational cost of generating a copy of Z is then given by

$$a = t,$$

and the computational cost of generating a copy of $X(i)$ conditional on Z is given by

$$b_i = 1$$

for $i = 1, 2, \dots, d$. In this experiment, we consider the parameters

$$d = 10, t = 10.$$

We implement three algorithms. The first algorithm uses budget allocation

$$(m_1, m_2, \dots, m_d) = (1, 1, \dots, 1),$$

which is crude Monte carlo. The second algorithm uses budget allocation

$$(m_1, m_2, \dots, m_d) = (22, 22, \dots, 22).$$

The third algorithm uses the heuristic budget allocation proposed in Section 4, integrated with a two-stage procedure. Specifically, given a budget c , the first \sqrt{c} budget is used to construct estimators of σ^2 and σ_i^2 for $i = 1, 2, \dots, d$. The purpose for selecting \sqrt{c} is such that it tends to infinity when c tends to infinity, while the proportion among the entire budget tends to zero. (We note that other choices of budget that is allocated to construct estimators may yield better performances, and we do not target on an optimal allocation in this work.) Then, the estimated value of σ^2 and σ_i^2 's are used to compute an approximation of the heuristic solution $(\tilde{m}_1, \tilde{m}_2, \dots, \tilde{m}_d)$ proposed in Section 4. For the rest of the $c - \sqrt{c}$ budget, the third algorithm implements the estimator with the estimated heuristic solution for the budget allocation. Note that the true heuristic solution $(\tilde{m}_1, \tilde{m}_2, \dots, \tilde{m}_d)$ in this setting (which is not known a priori by the third algorithm) is given by

$$(\tilde{m}_1, \tilde{m}_2, \dots, \tilde{m}_d) = (22, 16, 11, 8, 6, 4, 3, 2, 1, 1).$$

For each algorithm, given a budget c , the estimator $\hat{\alpha}_{N(c)}$ is independently constructed 1000 times to construct a 90% confidence interval. Figure 1 presents the results for four different choices of budget $c = 1000, 2000, 4000, 8000$. For each given budget, the three vertical lines are respectively the 90% confidence intervals for the estimator $\hat{\alpha}_{N(c)}$ constructed by Algorithm 1 (the first algorithm), Algorithm 2 (the second algorithm), Algorithm 3 (the third algorithm), from left to right. It is clear that Algorithm 3 achieves a variance reduction compared to both Algorithm 1 and Algorithm 2 for all given choices of the budget.

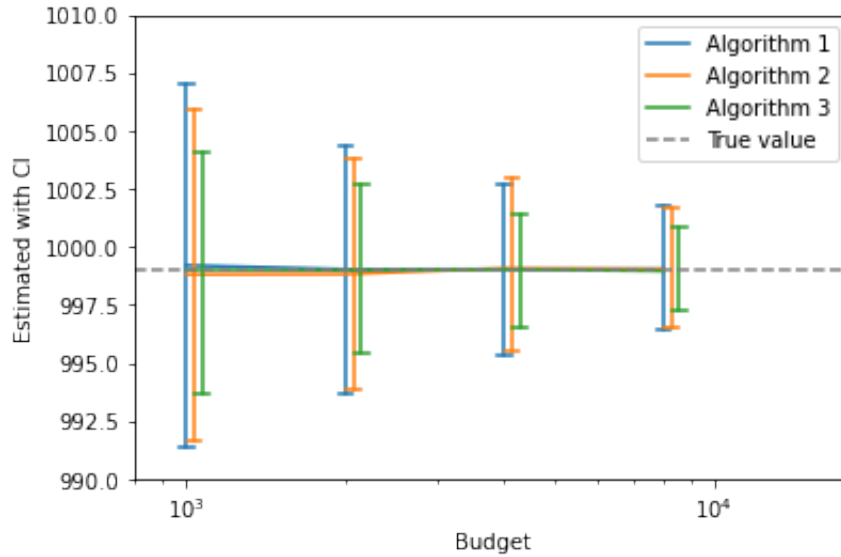


Figure 1: Estimated value with confidence intervals

6 CONCLUSION

Many applications in the areas of finance, environments, service systems and dependability systems require the computation for a general function of the expected performances associated with one or many random objects. When complicated stochastic systems are involved, such computation needs to be done by stochastic simulation and the computational cost can be expensive. We design simulation algorithms that exploit the common random structure shared by all random objects. We discuss the optimal simulation budget allocation

for the proposed algorithms. A brief numerical experiment is conducted to illustrate the performance of the proposed algorithm with various budget allocation rules. Future work includes designing fully adaptive simulation algorithms and analyzing the optimality gap.

ACKNOWLEDGMENTS

The authors are grateful to the two anonymous reviewers for their very helpful and constructive comments and suggestions. One reviewer suggests the analysis of the choice of n and its impact, which we greatly appreciate and will leave it to future work. The authors are also grateful to the track chairs and program committee for their help and support.

REFERENCES

- Asmussen, S., and P. W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*, Volume 57. Springer Science & Business Media.
- Bratley, P., B. L. Fox, and L. E. Schrage. 1987. *A Guide to Simulation*. Springer Science & Business Media.
- Broadie, M., Y. Du, and C. C. Moallemi. 2011. "Efficient risk estimation via nested sequential simulation". *Management Science* 57(6):1172–1194.
- Dang, J. O., B. M. Feng, and M. R. Hardy. 2019. "Efficient nested simulation of tail risk measures". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 938–949. Piscataway NJ: IEEE.
- Fu, M. C., and J.-Q. Hu. 2012. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*, Volume 392. Springer Science & Business Media.
- Glynn, P. W., and W. Whitt. 1992. "The asymptotic efficiency of simulation estimators". *Operations Research* 40(3):505–520.
- Gordy, M. B., and S. Juneja. 2010. "Nested simulation in portfolio risk measurement". *Management Science* 56(10):1833–1848.

AUTHOR BIOGRAPHIES

PETER W. GLYNN is the Thomas Ford Professor in the Department of Management Science and Engineering (MS&E) at Stanford University. He is a Fellow of INFORMS and of the Institute of Mathematical Statistics, has been co-winner of Best Publication Awards from the INFORMS Simulation Society in 1993, 2008, and 2016, and was the co-winner of the John von Neumann Theory Prize from INFORMS in 2010. In 2012, he was elected to the National Academy of Engineering. His research interests lie in stochastic simulation, queueing theory, and statistical inference for stochastic processes. His email address is glynn@stanford.edu.

ZEYU ZHENG is an assistant professor in the Department of Industrial Engineering & Operations Research at the University of California Berkeley. He has done research in simulation and non-stationary stochastic modeling and decision making. His email address is zyzheng@berkeley.edu.