

UNBIASED ESTIMATION OF THE RECIPROCAL MEAN FOR NON-NEGATIVE RANDOM VARIABLES

Sarat Babu Moka
Dirk P. Kroese

Sandeep Juneja

School of Mathematics and Physics
The University of Queensland
Brisbane, 4072, AUSTRALIA

School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, 400005, INDIA

ABSTRACT

In recent years, Monte Carlo estimators have been proposed that can estimate the ratio of two expectations without bias. We investigate the theoretical properties of a Taylor-expansion based estimator of the reciprocal mean of a non-negative random variable. We establish explicit expressions for the computational efficiency of this estimator and obtain optimal choices for its parameters. We also derive corresponding practical confidence intervals and show that they are asymptotically equivalent to the maximum likelihood (biased) ratio estimator as the simulation budget increases.

1 INTRODUCTION

Over the past few years, unbiased Monte Carlo estimation methods have received significant attention, due to both theoretical interest and practical applications; see, for example, (Rhee and Glynn 2015; Blanchet et al. 2015; Blanchet and Glynn 2015; Jacob and Thiery 2015; McLeish 2011). Efficient unbiased estimation of non-linear functions of expectations of random variables is challenging and has several applications; see, for example, (Blanchet et al. 2015; Jacob and Thiery 2015). An important “canonical” case is the unbiased estimation of $1/\mathbb{E}Z$ for a non-negative random variable Z . Applications include regenerative simulation, estimating parameters involving densities with unknown normalizing constants, and Bayesian inference.

Motivated by these applications, we study properties of an unbiased estimator of $\beta = 1/\mathbb{E}Z$ proposed by Blanchet et al. (2015) (which is in turn based on the ideas proposed by Rhee and Glynn (2015) in the context of stochastic differential equations). The estimator is obtained as follows. Write $\beta = \frac{1}{\mathbb{E}Z} = w \sum_{n=0}^{\infty} (1 - w\mathbb{E}Z)^n$ for $w < 2\beta$; here the condition $w < 2\beta$ guarantees the convergence of the geometric series $\sum_{n=0}^{\infty} (1 - w\mathbb{E}Z)^n$. Further, let $\{Z_i, i \geq 0\}$ be a sequence of iid copies of Z , and let N be a non-negative integer-valued random variable with $q_n := \mathbb{P}(N = n) > 0$, for all $n \geq 0$. Then

$$\frac{1}{\mathbb{E}Z} = w \sum_{n=0}^{\infty} q_n \frac{(1 - w\mathbb{E}Z)^n}{q_n} = w \sum_{n=0}^{\infty} q_n \frac{\mathbb{E} \prod_{i=1}^n (1 - wZ_i)}{q_n} = w \mathbb{E} \left[\frac{1}{q_N} \prod_{i=1}^N (1 - wZ_i) \right].$$

Define,

$$\widehat{\beta}(w) := \frac{w}{q_N} \prod_{i=1}^N (1 - wZ_i). \quad (1)$$

Clearly, $\mathbb{E}\widehat{\beta}(w) = \beta$ and thus $\widehat{\beta}(w)$ is an unbiased estimator of β . Note that if $Z \leq b$ almost surely for a constant b , then with the choice $w < 1/b$, $\widehat{\beta}(w)$ becomes non-negative. In this paper, the goal is to study optimal choices for w and $\{q_n, n \geq 0\}$ that make $\widehat{\beta}(w)$ *efficient*. In particular, a brief description of our contributions is as follows:

- When $\{q_n, n \geq 0\}$ is the variance-minimizing distribution for a fixed w , we show that as $w \searrow 0$, the expected cost to construct $\hat{\beta}(w)$ increases to ∞ , while both the variance and the expected time variance product of $\hat{\beta}(w)$ decrease.
- As a consequence, we argue that for any w , instead of approximating β with a sample mean of iid copies of $\hat{\beta}(w)$, it is optimal to approximate it by just one outcome of $\hat{\beta}(w^*)$, where w^* is such that $w^* < w$ and the expected cost of constructing $\hat{\beta}(w^*)$ is the same as that of the sample mean.
- We study the asymptotic distribution of $\hat{\beta}(w)$ as $w \searrow 0$ (i.e., as the expected computational cost for the estimator goes to ∞). We establish a central limit theorem type convergence result that is useful for finding asymptotically valid confidence intervals.
- We compare the asymptotic performance of the unbiased estimator $\hat{\beta}(w)$ with that of the maximum likelihood (biased) ratio estimator, where β is estimated using the reciprocal of a sample mean of iid copies of Z .
- The above results are studied under the assumption that N has the variance-minimizing distribution. Generating samples from this distribution is impossible as it involves unknown parameters. Since $\hat{\beta}(w)$ is unbiased even for a different distribution of N , we develop a method to implement the estimator by proposing a distribution for N (using samples of Z) that closely resembles to the variance-minimizing distribution .

Background: Several applications of Monte Carlo simulation involve the estimation of $\beta = 1/\mathbb{E}Z$ for a non-negative random variable Z . In some applications it is a desirable property to have an unbiased estimator of β when the magnitudes of the available biased estimators are unknown a priori. Examples include the estimation of a steady-state parameter $\alpha = \mathbb{E}R/\mathbb{E}\tau$ for a *regenerative* stochastic process, where τ denotes the length of a regenerative cycle and R denotes the *cumulative reward* obtained over the regenerative cycle; see, e.g. (Glynn 2006; Asmussen and Glynn 2007; Moka and Juneja 2015). It is evident that we have an unbiased estimator of α if we have an unbiased estimator of $1/\mathbb{E}\tau$. A similar case is where parameters can be expressed as $\alpha = \mathbb{E}[h(X)f(X)]/\mathbb{E}f(X)$ for some real-valued function h and probability density f , where f is known up to the normalizing constant $\mathbb{E}f(X)$. Such densities occur, for example, in Gibbs point processes (Møller and Waagepetersen 2003); and a standard method to estimate such parameters is by using Markov Chain Monte Carlo (MCMC) methods, see Asmussen and Glynn (2007), Rubinstein and Kroese (2017). However, in many situations it is difficult to bound the bias of the MCMC estimator, as the mixing time of the Markov chain is unknown. An alternative approach is to use a ratio estimator, where α is approximated by ratio of the sample means of the numerator and the denominator. However, this still returns a biased estimator and the bias decreases at a rate that is inversely proportional to the sample size; see Remark 2 and also (Asmussen and Glynn 2007). Therefore, it is desirable to have an unbiased estimator for $1/\mathbb{E}f(X)$ (and equally for $1/\mathbb{E}\tau$) that has the same order of complexity as that of the ratio estimator.

Most importantly, in some applications, having an unbiased estimator of β is essential. For example, in the study of doubly intractable models in Bayesian inference, it is assumed that the observations follow a distribution with a density of the form $f(y|\theta) = \frac{g(y,\theta)}{\int g(y,\theta)dy}$, where $g(y,\theta)$ can be evaluated point-wise up to the normalizing constant $\int g(y,\theta)dy$; see, for example, Lyne et al. (2015), Walker (2011), Jacob and Thiery (2015). Standard Metropolis–Hastings algorithms to obtain posterior estimates are not applicable due to the intractability of the normalizing constant. However, an exact inference method called *pseudo-marginal Metropolis–Hastings* proposed by Andrieu and Roberts (2009) can be implemented if a non-negative unbiased estimator of $1/\int g(y,\theta)dy$ is available; also see Beaumont (2003), Jacob and Thiery (2015), Walker (2011). In particular, Jacob and Thiery (2015) highlight the importance of the estimators of the form (1).

A standard method called *Russian roulette truncation* can be used for unbiased estimation of β . This method is first proposed in the physics literature (Carter and Cashwell 1975; Lux and Koblinger 1991) and

further studied by McLeish (2011), Glynn and Rhee (2014), Lyne et al. (2015), Wei and Murray (2016). The key drawback of these estimators is that they can take negative values with positive probability even when Z is bounded.

In Section 2, we study the properties of the estimator as a function of w , under the assumption that N has the variance minimizing estimator. An implementable method is proposed in Section 3. A conclusion of the paper is given in Section 4. All the results are proved in Appendix A.

2 PROPERTIES OF THE ESTIMATOR

Without loss of generality, assume that Z is non-degenerate. As the estimator $\hat{\beta}(w)$ in (1) is unbiased, a sample mean of independent copies of $\hat{\beta}(w)$ is an unbiased estimator of β as well. It is well known that the sample mean has square-root convergence rate if $\text{Var}\hat{\beta}(w) < \infty$; see, e.g., (Asmussen and Glynn 2007). Thus, a simple strategy is to seek values of w and $\{q_n, n \geq 0\}$ that minimize $\text{Var}\hat{\beta}(w)$. Using the Cauchy–Schwarz inequality, (Blanchet et al. 2015) show that for any $w < 2\mathbb{E}Z/\mathbb{E}Z^2$, $\text{Var}\hat{\beta}(w)$ is finite and is minimal if N has a geometric distribution on the non-negative integers with success probability

$$p_w = 1 - \sqrt{\mathbb{E}(1 - wZ)^2} = 1 - \sqrt{1 - 2w\mathbb{E}Z + w^2\mathbb{E}Z^2};$$

that is, if

$$q_n = (1 - p_w)^n p_w, \quad n \geq 0, \tag{2}$$

where the assumption $w < 2\mathbb{E}Z/\mathbb{E}Z^2$ guarantees that $p_w > 0$. Unfortunately, p_w depends on $\mathbb{E}Z$ and $\mathbb{E}Z^2$, which are unknown. However, $\hat{\beta}(w)$ is unbiased even when N has a different distribution. Therefore, in the implementation of $\hat{\beta}(w)$, we can replace p_w with an estimate of p_w ; see Section 3. In this section, we study the properties of $\hat{\beta}(w)$ under the assumption that N has the distribution (2), because it offers an understanding of what is the best that can be expected from the estimator.

Note that the variance of $\hat{\beta}(w)$ is given by,

$$\text{Var}\hat{\beta}(w) = w^2 \sum_{n=0}^{\infty} \frac{(\mathbb{E}(1 - wZ)^2)^n}{q_n} - \beta^2 = \frac{w^2}{p_w^2} - \beta^2, \tag{3}$$

for all $0 < w < 2\mathbb{E}Z/\mathbb{E}Z^2$. Further, observe that $\mathbb{E}N = 1/p_w - 1$. Now we can ask what is the value of $w \in (0, 2\mathbb{E}Z/\mathbb{E}Z^2)$ that minimizes (3). This question is not addressed in the existing literature. In addition to the variance, it is often important to include the running time to construct the estimator to determine its efficiency; see (Glynn and Whitt 1992). In that case, we need to select w for which the *expected time variance product*, $\mathbb{E}T \text{Var}\hat{\beta}(w)$, is minimal, where T is the time required to construct $\hat{\beta}(w)$. From (1) (since Z_1, Z_2, \dots are iid), it is reasonable to assume that T is proportional to the number of Z_i 's used for constructing $\hat{\beta}(w)$. Since $N = N(w)$ samples of Z are used in the construction of $\hat{\beta}(w)$, we assume that the expected time variance product is $\mathbb{E}N \text{Var}\hat{\beta}(w)$.

Theorem 1 Suppose that N has the geometric distribution given in (2). Then the following hold true.

- (i) The success probability p_w is a strictly concave function of $w \in (0, 2\mathbb{E}Z/\mathbb{E}Z^2)$ with a maximum value of $1 - \sqrt{1 - (\mathbb{E}Z)^2/\mathbb{E}Z^2}$ attained at $w = \mathbb{E}Z/\mathbb{E}Z^2$, and

$$\lim_{w \searrow 0} p_w = \lim_{w \nearrow 2\mathbb{E}Z/\mathbb{E}Z^2} p_w = 0.$$

- (ii) The variance $\text{Var}\hat{\beta}(w)$ is a strictly increasing convex function of $w \in (0, 2\mathbb{E}Z/\mathbb{E}Z^2)$, with

$$\lim_{w \searrow 0} \text{Var}\hat{\beta}(w) = 0 \quad \text{and} \quad \lim_{w \nearrow 2\mathbb{E}Z/\mathbb{E}Z^2} \text{Var}\hat{\beta}(w) = \infty.$$

(iii) The expected time variance product $\mathbb{E}N(w) \text{Var} \hat{\beta}(w)$ is a strictly increasing function of $w \in (0, 2\mathbb{E}Z/\mathbb{E}Z^2)$, with

$$\lim_{w \searrow 0} \mathbb{E}N(w) \text{Var} \hat{\beta}(w) = \frac{\text{Var} Z}{(\mathbb{E}Z)^4}, \quad \text{and} \quad \lim_{w \nearrow 2\mathbb{E}Z/\mathbb{E}Z^2} \mathbb{E}N(w) \text{Var} \hat{\beta}(w) = \infty.$$

Remark 1 To understand the implications of Theorem 1, suppose we select a $w \in (0, 2\mathbb{E}Z/\mathbb{E}Z^2)$, giving an expected computational cost of $\mathbb{E}N(w)$ to obtain $\hat{\beta}(w)$. Further, let $\bar{\beta}_k(w)$ be the sample mean of k iid copies of $\hat{\beta}(w)$. Then the expected time variance product for $\bar{\beta}_k(w)$ is

$$k \mathbb{E}N(w) \text{Var} \bar{\beta}_k(w) = \mathbb{E}N(w) \text{Var} \hat{\beta}(w).$$

Now suppose that $w^* < w$ is selected so that the average cost to generate one outcome of $\hat{\beta}(w^*)$ is equal to the average cost to construct $\bar{\beta}_k(w)$; that is, $k \mathbb{E}N(w)$. Then, from Theorem 1 (iii), for the same computational effort, $\hat{\beta}(w^*)$ has a smaller variance than $\bar{\beta}_k(w)$ for any feasible w selected as above, and is therefore a better estimator.




Figure 1: An example to illustrate the dependency of the performance of the unbiased estimator on parameter w . Panels (a) and (b) show, respectively, the relative variance $\frac{\text{Var} \hat{\beta}(w)}{\beta^2}$ and the expected time relative variance product $\frac{\mathbb{E}N(w) \text{Var} \hat{\beta}(w)}{\beta^2}$, as functions of w .

To illustrate the results of Theorem 1, consider an example where $Z = \mathbb{I}(A)$ for an event A with probability $\mathbb{P}(A) = \mathbb{E}Z = 0.001$. Since $\mathbb{E}Z^2 = 0.001$, the relative variance $\frac{\text{Var} Z}{(\mathbb{E}Z)^2} = 999$. By substituting the values of $\mathbb{E}Z$ and $\mathbb{E}Z^2$, we can calculate p_w , $\mathbb{E}N$ and $\text{Var} \hat{\beta}(w)$ for any $w < 2\mathbb{E}Z/\mathbb{E}Z^2 = 2$. Figure 1 illustrate the effect of w on the efficacy of the estimator $\hat{\beta}(w)$. As expected, both $\frac{\text{Var} \hat{\beta}(w)}{\beta^2}$ and $\frac{\mathbb{E}N(w) \text{Var} \hat{\beta}(w)}{\beta^2}$ are decreasing as $w \searrow 0$ with the limits 0 and $\frac{\text{Var} Z}{(\mathbb{E}Z)^2} = 999$, respectively. \square

Remark 1 motivates us to study the asymptotic distributional properties of $\hat{\beta}(w)$ as $w \searrow 0$, when N has the geometric distribution given in (2). Theorem 2 is crucial for establishing confidence intervals that are asymptotically valid as $w \searrow 0$.

Theorem 2 Suppose that N has the distribution given by (2) and Z is bounded. Then, as $w \searrow 0$,

$$(i) \quad \hat{\beta}(w) \xrightarrow{d} \beta, \quad \text{and} \quad (ii) \quad \frac{(\hat{\beta}(w) - \beta)}{\sqrt{w \mathbb{E}Z}} \xrightarrow{d} \left(\sqrt{\frac{\text{Var} Z}{(\mathbb{E}Z)^4}} \mathcal{E}(1) \right) \mathcal{N}(0, 1),$$

where \xrightarrow{d} denotes convergence in distribution, and $\mathcal{E}(1)$ and $\mathcal{N}(0, 1)$ are independent random variables from respectively the standard (mean 1) exponential and standard normal distributions.

We show later in Section A.1 that $\frac{p_w}{w\mathbb{E}Z} \rightarrow 1$ as $w \searrow 0$ (see (8)). Since $\mathbb{E}N = 1/p_w - 1$ and $\lim_{w \rightarrow 0} p_w = 0$, $w\mathbb{E}Z\mathbb{E}N \rightarrow 1$ as $w \searrow 0$. That means, an alternative expression for Theorem 2 (ii) is

$$\sqrt{\mathbb{E}N} \left(\widehat{\beta}(w) - \beta \right) \xrightarrow{d} \sigma \sqrt{\mathcal{E}(1)} \mathcal{N}(0, 1), \quad \text{as } w \searrow 0,$$

where $\sigma = \sqrt{\frac{\text{Var}Z}{(\mathbb{E}Z)^4}}$. The above expression has more resemblance to the standard central limit theorem, since $\mathbb{E}N$ is the computational cost of the estimator. It is not difficult to show that $\sqrt{\mathcal{E}(1)}\mathcal{N}(0, 1)$ is a random variable with density $f(x) = \frac{1}{\sqrt{2}} \exp(-\sqrt{2}|x|)$, which is the density of a mean zero *Laplace* (or double exponential) distribution with scale $1/\sqrt{2}$. These observations are useful for constructing asymptotically valid confidence intervals as follows. For any given $\alpha \in (0, 1)$, by solving $\int_0^t f(x) dx = (1 - \alpha)/2$ for t , we get $t = -\log(\alpha)/\sqrt{2}$. Then using Theorem 2, we can say that the interval

$$\left(\widehat{\beta}(w) + \frac{\log(\alpha)}{\sqrt{2}} \sigma \sqrt{w\mathbb{E}Z}, \quad \widehat{\beta}(w) - \frac{\log(\alpha)}{\sqrt{2}} \sigma \sqrt{w\mathbb{E}Z} \right)$$

is an asymptotic $1 - \alpha$ confidence interval for β .

Remark 2 (Comparison with the ratio estimator) A standard (biased) estimator of β is $1/\bar{Z}_n$, where \bar{Z}_n is the sample mean of n iid copies of Z . Using Taylor's theorem for the function $1/x$ about $\mathbb{E}Z$, we can easily show that the bias of $1/\bar{Z}_n$ is approximately $\frac{1}{n} \frac{\text{Var}Z}{(\mathbb{E}Z)^3}$ for large n , while, on the other hand, $\widehat{\beta}(w)$ has zero bias. Furthermore, using the same Taylor's theorem, we can show that the asymptotic time variance product of $1/\bar{Z}_n$ is $\text{Var}Z/(\mathbb{E}Z)^4$ as $n \rightarrow \infty$. From Theorem 1 (iii), the unbiased estimator $\widehat{\beta}(w)$ has the same asymptotic expected time variance product. However, unbiasedness of $\widehat{\beta}(w)$ comes at cost. As an application of the delta method, we can show that the ratio estimator satisfies the central limit theorem: $\sqrt{n}(1/\bar{Z}_n - \beta) \xrightarrow{d} \sqrt{\frac{\text{Var}Z}{(\mathbb{E}Z)^4}} \mathcal{N}(0, 1)$. That is, the ratio estimator is asymptotically normal. On the other hand, the asymptotic distribution of the unbiased estimator $\widehat{\beta}(w)$ is Laplace, which has more slowly decaying tails than a normal distribution. In conclusion, the ratio estimator can have narrower confidence intervals than the unbiased estimator.

Remark 3 (Importance sampling) Just like in the case of the ratio estimator, from Theorems 1 and 2, the relative variance of Z is the key factor influencing the asymptotic properties of the unbiased estimator. The smaller the value is of the relative variance of Z , the better is the reliability of the unbiased estimator. One of the most effective technique of variance reduction is *importance sampling*. We can improve the performance of the estimator if we can implement an importance sampling technique for the random variable Z .

Remark 4 (The time variance product minimizing distribution for $N(w)$) We have assumed that for a given w the random variable $N(w)$ has the variance minimizing distribution given by (2). However, when the criteria for the optimality of $\widehat{\beta}(w)$ is the minimization of the expected time variance product, we need to seek a distribution $\{q_n, n \geq 0\}$ that minimizes $\mathbb{E}N(w)\text{Var}\widehat{\beta}(w)$. It is shown in (Blanchet et al. 2015) that the distribution that minimizes the expected time variance is given by

$$\tilde{q}_n = \frac{w(1-p_w)^n}{\sqrt{\beta^2 + d_w n}}, n \geq 0, \tag{4}$$

where d_w is the unique (positive) number satisfying $\sum_{n=0}^{\infty} \frac{w(1-p_w)^n}{\sqrt{\beta^2 + d_w n}} = 1$. When compared to the distribution (2), drawing samples from (4) has an extra difficulty of finding d_w by solving an equation that contains

an infinite sum. Even if we overcome this difficulty, the reduction in the expected time variance product is typically insignificant, because for small values of w ,

$$\mathbb{E}\tilde{N}(w) \mathbb{V}\text{ar}\tilde{\beta}(w) = \mathbb{E}N(w) \mathbb{V}\text{ar}\hat{\beta}(w)(1 + O(w)), \tag{5}$$

when $N(w)$ has the variance minimizing distribution (2); see Section A.3 for a proof of (5). □

3 AN IMPLEMENTATION

Recall that the success probability p_w of N is a function of unknown quantities $\mathbb{E}Z$ and $\mathbb{E}Z^2$. However, fortunately, $\hat{\beta}(w)$ in (1) is still an unbiased estimator of β for any distribution $\{q_n, n \geq 0\}$ of N . In particular, instead of taking q_n as in (2), we can take $q_n = P_k(1 - P_k)^n$, where P_k is defined below. Under the proposed implementation, when the given budget is sufficiently large, half of the budget is used for estimating P_k and then w is chosen such that the remaining half the budget is used for generating a sample of the unbiased estimator.

To simplify the discussion, assume that there is a known constant $0 < \varepsilon \leq \beta$; for example, if $Z \leq b$ for a constant b , we can take $\varepsilon = 1/b$. Let $\tilde{Z}_1, \tilde{Z}_2, \dots, \tilde{Z}_k$ be a sequence of iid copies of Z , independent of the sequence Z_1, Z_2, \dots , which is used in the construction of the unbiased estimator $\hat{\beta}$ in (1). Define the first two sample moments: $M_1(k) = \frac{1}{k} \sum_{i=1}^k \tilde{Z}_i$ and $M_2(k) = \frac{1}{k} \sum_{i=1}^k \tilde{Z}_i^2$. If $M_1(k) > 0$, define,

$$P_k = 1 - \sqrt{\frac{1}{k} \sum_{i=1}^k (1 - w_k \tilde{Z}_i)^2} \quad \text{with} \quad w_k = \min\left(\frac{1}{k M_1(k)}, \frac{M_1(k)}{M_2(k)}, \varepsilon\right).$$

Otherwise, take $P_k = 1/k$ and $w_k = \varepsilon/k$. The condition $w_k < 2 \frac{M_1(k)}{M_2(k)}$ guarantees that $P_k > 0$. Further, whether $M_1(k) = 0$ or not, $w_k < \beta$ and hence it guarantees that the estimator $\hat{\beta}(w_k)$ (defined by (1)) with $q_n = P_k(1 - P_k)^n$ is an unbiased estimator of β . Note that given $M_1(k)$ and $M_2(k)$, the expected cost to construct $\hat{\beta}(w_k)$ is $k + 1/P_k$ (including the cost to construct P_k), since $\mathbb{E}[N(w_k) | M_1(k), M_2(k)] = 1/P_k - 1$. Theorem 3 states that for large values of k , this total expected cost is approximately $2k$, and conditioned on $M_1(k)$ and $M_2(k)$, the expected time variance product goes to a random variable with mean $4 \frac{\mathbb{V}\text{ar}Z}{(\mathbb{E}Z)^4}$. See Section A.4 for a proof Theorem 3.

Theorem 3 Under the above construction, $\mathbb{E}\hat{\beta}(w_k) = \beta$, and as $k \rightarrow \infty$, $k P_k \rightarrow 1$, *a.s.*, and

$$\left(k + \frac{1}{P_k}\right) \mathbb{V}\text{ar}\left(\hat{\beta}(w_k) | M_1(k), M_2(k)\right) \rightarrow 2 \frac{\mathbb{V}\text{ar}Z}{(\mathbb{E}Z)^4} [1 + \chi_1^2], \text{ a.s.},$$

where χ_1^2 is the square of a standard normal random variable.

To understand Theorem 3, consider the example given in Remark 1. We estimated the expected total cost $k + \mathbb{E}[1/P_k]$ and $\mathbb{V}\text{ar}\hat{\beta}(w_k)$ using 10000 samples of P_k and $\hat{\beta}(w_k)$, respectively, with $k = 10000$. Our simulation results show that the estimated expected time relative variance product is 3969.75, which is approximately equal to $4 \frac{\mathbb{V}\text{ar}Z}{(\mathbb{E}Z)^2} = 4 \times 999 = 3996$, as expected.

4 CONCLUSION

We investigated the theoretical properties of a parametrized family $\{\hat{\beta}(w), w > 0\}$ of unbiased estimators of $1/\mathbb{E}Z$ for a non-negative random variable Z . We studied the variance and the expected time variance product as functions of w and established several asymptotic results. We showed that with an optimal choice of w , the asymptotic performance of the unbiased estimator $\hat{\beta}(w)$ is comparable to the maximum likelihood (biased) ratio estimator. We further proposed an implementable unbiased estimation based on our results. Similar to Theorem 2, our ongoing research establishes a central limit theorem type convergence result for $\hat{\beta}(w_k)$ defined in Section 3, by taking the budget parameter $k \rightarrow \infty$.

ACKNOWLEDGMENTS

The work of the first and the second authors has been supported by the Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS), under grant number CE140100049.

A APPENDIX

To simplify the notation in this section, we use $z_1 := \mathbb{E}Z$ and $z_2 := \mathbb{E}Z^2$. We also write p'_w for the derivative $\frac{dp_w}{dw}$ and p''_w for the second derivative. $\mathcal{E}(\lambda)$ and $\mathcal{G}(p)$ denote respectively the mean $1/\lambda$ exponential distribution and the geometric distribution on non-negative integers and the success probability p .

A.1 Proof of Theorem 1

From the definition, $p'_w = (z_1 - wz_2)/(1 - p_w)$. It follows that $p''_w = -\frac{1}{(1-p_w)^3} [z_2 - z_1^2] < 0$, where the strict inequality holds because $z_2 > z_1^2$, which follows from the assumption that Z is non-degenerate. Therefore, p_w is strictly concave over $(0, 2z_1/z_2)$ and it achieves its maximum value $1 - \sqrt{1 - z_1^2/z_2}$ at $w = z_1/z_2$. From the definition of p_w , it is evident that $\lim_{w \searrow 0} p_w = \lim_{w \nearrow 2z_1/z_2} p_w = 0$.

Recall from (3) that the variance of the estimator is equal to $w^2/p_w^2 - z_1^2$. Its derivative can be written as

$$\frac{d\text{Var}\hat{\beta}(w)}{dw} = 2w \frac{wz_1 - p_w}{p_w^3(1 - p_w)} \quad (6)$$

and the second derivative as

$$\frac{d^2\text{Var}\hat{\beta}(w)}{dw^2} = \frac{2(wz_1 - p_w)}{p_w^3(1 - p_w)^3} (3(wz_1 - p_w) + 2p_w(p_w - w^2z_2)). \quad (7)$$

Using Jensen's inequality, $\mathbb{E}(1 - wZ)^2 > (1 - wz_1)^2$, where the strict inequality holds again because Z is non-degenerate. On the other hand, by Bernoulli's inequality, $\sqrt{\mathbb{E}(1 - wZ)^2} = \sqrt{1 + (-2wz_1 + w^2z_2)}$ is maximized by $1 - wz_1 + w^2z_2/2$. Thus,

$$z_1 - wz_2/2 \leq \frac{p_w}{w} < z_1. \quad (8)$$

Using (8), we have $wz_1 - p_w > 0$ and $p_w - w^2z_2 \geq 2wz_1 - p_w > wz_1$, and hence for all $w \in (0, 2z_1/z_2)$, $\frac{d\text{Var}\hat{\beta}(w)}{dw} > 0$ and $\frac{d^2\text{Var}\hat{\beta}(w)}{dw^2} > 0$, which establishes the convexity of $\text{Var}\hat{\beta}(w)$ over $(0, 2z_1/z_2)$.

We now prove that the expected time variance product is a strictly increasing over $(0, 2z_1/z_2)$. Let $g_1(w) := \mathbb{E}N(w)$, $g_2(w) := \frac{\text{Var}\hat{\beta}(w)}{\beta^2}$ and $g(w) := g_1(w)g_2(w) = \left(\frac{1}{p_w} - 1\right) \left(\frac{w^2z_1^2}{p_w^2} - 1\right)$. Then,

$$\frac{dg_1}{dw}(w) = \frac{wz_2 - z_1}{p_w^2(1 - p_w)}, \quad \text{and} \quad \frac{dg_2}{dw}(w) = 2wz_1^2 \left[\frac{wz_1 - p_w}{p_w^3(1 - p_w)} \right],$$

and hence, we write

$$\begin{aligned} \frac{dg}{dw}(w) &= \frac{2wz_1^2(1 - p_w)(wz_1 - p_w) + (w^2z_1^2 - p_w^2)(wz_2 - z_1)}{p_w^4(1 - p_w)} \\ &= \frac{wz_1 - p_w}{p_w^4(1 - p_w)} [wp_w(z_2 - z_1^2) + z_1(wz_1 - p_w) + wz_1(wz_2 - z_1p_w)] > 0, \end{aligned}$$

where the inequality holds because $z_2 > z_1^2$, $wz_1 > p_w$ and $wz_2 > wz_1^2 > z_1 p_w$. Therefore, $\mathbb{E}N(w) \text{Var} \widehat{\beta}(w)$ is strictly increasing over $(0, 2z_1/z_2)$.

The claims that $\lim_{w \nearrow 2z_1/z_2} \text{Var} \widehat{\beta}(w) = \infty$ and $\lim_{w \nearrow 2z_1/z_2} \mathbb{E}N(w) \text{Var} \widehat{\beta}(w) = \infty$ hold trivially because $\lim_{w \nearrow 2z_1/z_2} p_w = 0$. To complete the proof of the theorem, we can write, by Taylor's theorem, for any $x \in (0, 1)$: $\sqrt{x} = 1 + \frac{(x-1)}{2} - \frac{(x-1)^2}{8} + \frac{(x-1)^3}{16\tilde{x}^{5/2}}$, for some $\tilde{x} \in (x, 1)$. Consequently,

$$\sqrt{\mathbb{E}(1-wZ)^2} = 1 + \frac{\mathbb{E}(1-wZ)^2 - 1}{2} - \frac{(\mathbb{E}(1-wZ)^2 - 1)^2}{8} + R(w),$$

where $R(w) = \frac{(\mathbb{E}(1-wZ)^2 - 1)^3}{16\tilde{x}^{5/2}}$ for some $\tilde{x} \in (\mathbb{E}(1-wZ)^2, 1)$. Since $\tilde{x} \rightarrow 1$ as $w \rightarrow 0$ and $\mathbb{E}(1-wZ)^2 - 1 = w^2 z_2 - 2wz_1$, we have $R(w) = O(w^3)$. Further simplification yields that $p_w = wz_1 - \frac{w^2}{2}(z_2 - z_1^2) + O(w^3)$, and thus, $\frac{p_w^2}{w^2 z_1^2} = 1 - \frac{w}{z_1}(z_2 - z_1^2) + O(w^2)$. Since $1 - p_w = 1 + O(w)$,

$$\frac{1 - p_w}{p_w} = \frac{1}{wz_1} (1 + O(w)), \quad \text{and} \quad \frac{w^2 z_1^2}{p_w^2} - 1 = w \frac{(z_2 - z_1^2)}{z_1} (1 + O(w)). \quad (9)$$

We conclude that both $\text{Var} \widehat{\beta}(w)$ and $\mathbb{E}N(w) \text{Var} \widehat{\beta}(w)$ go to their respective minima as $w \searrow 0$.

A.2 Proof of Theorem 2

Statement (i) follows directly from Theorem 1 and Chebyshev's inequality:

$$\mathbb{P} \left(|\widehat{\beta}(w) - \beta| > \varepsilon \right) \leq \text{Var} \widehat{\beta}(w) / \varepsilon^2 \rightarrow 0, \text{ as } w \searrow 0,$$

for every $\varepsilon > 0$. To prove (ii), consider a decreasing sequence $w_1 > w_2 > \dots$ such that $w_1 \leq z_1/z_2$ and $\lim_{k \rightarrow \infty} w_k = 0$. We construct an almost surely increasing sequence $N_1 \leq N_2 \leq \dots$ such that $N_k \sim \mathcal{G}(p_{w_k})$ and

$$\lim_{k \rightarrow \infty} [w_k N_k] = X_\infty / z_1, \quad a.s., \quad (10)$$

for a random variable $X_\infty \sim \mathcal{E}(1)$. To do so we invoke Theorem 3.1 of (Moka and Juneja 2015). Let $\lambda_k = -\log(1 - p_{w_k})$ and $E_k \sim \mathcal{E}(\lambda_k)$. Then, (Moka and Juneja 2015) says that for each k , there exist a random variable Y_k with cumulative distribution function $G_k(x) = 1 - \left(1 - \frac{\lambda_{k+1}}{\lambda_k}\right) \exp(-\lambda_{k+1}x)$, $x \geq 0$, such that Y_k is independent of E_k , and E_{k+1} has the same distribution as $E_k + Y_k$. Therefore, without loss of generality we assume that there is sequence of independent random variables $Y_k \sim G_k(x)$ such that $E_{k+1} = E_k + Y_k = E_1 + \sum_{i=1}^k Y_i$ for all $k \geq 1$.

Consider the natural filtration $\{\mathcal{F}_k = \sigma(E_1, \dots, E_k), k \geq 0\}$. Since $E_{k+1} = E_k + Y_k$,

$$\begin{aligned} \lambda_{k+1} E_{k+1} - 1 &= \lambda_{k+1} [E_k + Y_k - 1/\lambda_k - \mathbb{E}Y_k] \\ &= \frac{\lambda_{k+1}}{\lambda_k} \lambda_k [E_k - 1/\lambda_k] + \lambda_{k+1} [Y_k - \mathbb{E}Y_k] \\ &\leq \lambda_k E_k - 1 + \lambda_{k+1} [Y_k - \mathbb{E}Y_k], \end{aligned}$$

where the last inequality holds because $\lambda_{k+1} \leq \lambda_k$. We have $\mathbb{E}[\lambda_{k+1} E_{k+1} | \mathcal{F}_k] \leq \lambda_k E_k$ since Y_k is independent of \mathcal{F}_k . Thus, $\{X_k := \lambda_k E_k, k \geq 1\}$ is a supermartingale (with respect to $\{\mathcal{F}_k\}$). In fact the sequence $\{X_k, k \geq 1\}$ is bounded in \mathcal{L}^2 , because $\sup_{k \geq 1} \mathbb{E}X_k^2 = 2$, making it uniformly integrable submartingale.

Thus, $X_\infty = \lim_{k \rightarrow \infty} X_k$ exists *a.s.* (see Theorem 2 in Section 4 of Chapter VII of (Shiryaev 1996)). Since $\mathbb{P}(X_k \leq x) = \mathbb{P}\left(E_k \leq \frac{x}{\lambda_k}\right) = 1 - \exp(-x)$. This implies that $X_\infty \sim \mathcal{E}(1)$.

Let $N_k = \lfloor E_k \rfloor$. Then for all k we have $N_k \sim \mathcal{G}(p_{w_k})$ and $N_k \leq N_{k+1}$. From (8), $\lim_{k \rightarrow \infty} (1 - p_{w_k})^{1/w_k} = \exp(-z_1)$ and hence $\lim_{k \rightarrow \infty} \frac{w_k}{\lambda_k} = 1/z_1$. From the convergence of the sequence X_1, X_2, \dots , we have $\frac{w_k}{\lambda_k} X_k - w_k \leq w_k N_k \leq \frac{w_k}{\lambda_k} X_k$. Therefore, (10) holds.

Now define

$$\widehat{\beta}_k := \frac{w_k}{(1 - p_{w_k})^{N_k} p_{w_k}} \prod_{i=1}^{N_k} (1 - w_k Z_i). \quad (11)$$

From the definitions, $\widehat{\beta}_k$ is identical to $\widehat{\beta}(w_k)$ in distribution. We now conclude the proof Theorem 2 by establishing lower and upper bounds on $\widehat{\beta}_k$ separately. Let b be an upper bound on Z . From the construction of $\widehat{\beta}_k$ given by (11), for all k such that $w_k < 1/b$, we have using (8) that

$$\widehat{\beta}_k \geq \frac{1}{z_1 (1 - p_{w_k})^{N_k}} \exp\left(\sum_{i=1}^{N_k} \log(1 - w_k Z_i)\right) \geq \frac{1}{z_1} \exp\left(N_k w_k (z_1 - w_k z_2 / 2) + \sum_{i=1}^{N_k} \log(1 - w_k Z_i)\right).$$

Using Taylor's theorem, $\log(1 - x) \geq -x - \frac{x^2}{2(1-x)^2}$ for any $x \geq 0$, and thus,

$$\begin{aligned} \widehat{\beta}_k &\geq \frac{1}{z_1} \exp\left(N_k w_k z_1 - N_k \frac{w_k^2 z_2}{2} - w_k \sum_{i=1}^{N_k} Z_i - \sum_{i=1}^{N_k} \frac{w_k^2 Z_i^2}{2(1 - w_k Z_i)^2}\right) \\ &= \frac{1}{z_1} \exp\left(-N_k w_k \frac{1}{N_k} \sum_{i=1}^{N_k} (Z_i - z_1) - \frac{N_k w_k^2}{2} \left(z_2 + \frac{b^2}{(1 - w_k b)^2}\right)\right). \end{aligned} \quad (12)$$

On the other hand, from (11) and (8),

$$\begin{aligned} \widehat{\beta}_k &\leq \frac{w_k}{(1 - p_{w_k})^{N_k} p_{w_k}} \exp\left(-\sum_{i=1}^{N_k} w_k Z_i\right) \leq \frac{w_k}{(1 - w_k z_1)^{N_k} p_{w_k}} \exp\left(-\sum_{i=1}^{N_k} w_k Z_i\right) \\ &= \frac{w_k}{p_{w_k}} \exp\left(-\sum_{i=1}^{N_k} w_k Z_i - N_k \log(1 - w_k z_1)\right) \\ &= \frac{w_k}{p_{w_k}} \exp\left(-w_k N_k \frac{1}{N_k} \sum_{i=1}^{N_k} (Z_i - z_1) + N_k w_k^2 \frac{z_1^2}{2(1 - w_k z_1)^2}\right). \end{aligned} \quad (13)$$

Using the strong law of large numbers and Theorem 1 of (Richter 1965), $\lim_{k \rightarrow \infty} \frac{1}{N_k} \sum_{i=1}^{N_k} (Z_i - z_1) = 0$ almost surely. Further, using (8) and (10), we have that $\lim_{k \rightarrow \infty} \widehat{\beta}_k = \beta$ almost surely. From Taylor's theorem with a Cauchy remainder term, we have almost surely

$$\log(\widehat{\beta}_k) - \log(\beta) = \frac{(\widehat{\beta}_k - \beta)}{\beta} - \frac{(\widehat{\beta}_k - \widehat{X})(\widehat{\beta}_k - \beta)}{\widehat{X}^2} = \frac{(\widehat{\beta}_k - \beta)}{\beta} [1 + o(1)]$$

for a random variable \widehat{X} that takes values between $\widehat{\beta}_k$ and β . Therefore, to complete the proof of (ii), it is enough to show that

$$\frac{1}{\sqrt{w_k}} \log \widehat{\beta}_k \xrightarrow{d} \left(\sqrt{\text{Var} Z X_\infty}\right) \mathcal{N}(0, 1).$$

From (12),

$$\frac{1}{\sqrt{w_k}} \log[z_1 \hat{\beta}_k] \geq -\sqrt{w_k} N_k \frac{1}{\sqrt{N_k}} \sum_{i=1}^{N_k} (Z_i - z_1) - \frac{N_k w_k^{3/2}}{2} \left(z_2 + \frac{b^2}{(1 - w_k b)^2} \right). \quad (14)$$

From (10) and because $w_k \searrow 0$, the second term on the right hand side of the expression goes to zero. Again using (10), we conclude that the right hand side of (14) goes to $\left(\sqrt{\frac{\text{Var} Z}{z_1}} X_\infty\right) \mathcal{N}(0, 1)$ in distribution. From (13),

$$\frac{1}{\sqrt{w_k}} \log[z_1 \hat{\beta}_k] \leq \frac{1}{\sqrt{w_k}} \log \left[\frac{z_1 w_k}{p_{w_k}} \right] - \sqrt{w_k} N_k \frac{1}{\sqrt{N_k}} \sum_{i=1}^{N_k} (Z_i - z_1) + N_k w_k^{3/2} \frac{z_1^2}{2(1 - w_k z_1)^2}.$$

We complete the proof because from (8), as $k \rightarrow \infty$,

$$0 \leq \frac{1}{\sqrt{w_k}} \log \left[\frac{z_1 w_k}{p_{w_k}} \right] \leq \frac{1}{\sqrt{w_k}} \log \left[\frac{1}{\left(1 - \frac{w_k z_2}{2z_1}\right)} \right] = -\sqrt{w_k} \log \left[\left(1 - \frac{w_k z_2}{2z_1}\right)^{1/w_k} \right] \rightarrow 0, \quad a.s.$$

A.3 Proof of Equation (5)

First, observe from the definitions that $\mathbb{E}\tilde{N}(w)\text{Var}\tilde{\beta}(w) \leq \mathbb{E}N(w)\text{Var}\hat{\beta}(w)$. Further using the fact that $\beta = 1/z_1$ and (8),

$$\mathbb{E}\tilde{N}(w) \geq \sum_{n=1}^{\infty} n \frac{p_w(1-p_w)^n}{\sqrt{1+d_w m_1^2 n}} = \mathbb{E}N(w) \sum_{n=1}^{\infty} \frac{1}{\sqrt{1+d_w z_1^2 n}} \frac{np_w(1-p_w)^n}{\mathbb{E}N(w)} \geq \mathbb{E}N(w) \frac{1}{\sqrt{1+d_w z_1^2 \mathbb{E}N(w)^2 / \mathbb{E}N(w)}},$$

where the last inequality holds from Jensen's inequality, because $1/\sqrt{1+ax}$ is a convex function of x for any constant $a > 0$ and $\left(\frac{np_w(1-p_w)^n}{\mathbb{E}N(w)}, n \geq 0\right)$ is a probability distribution. Furthermore, using $\mathbb{E}N(w)^2 = \mathbb{E}N(w) \frac{(2-p_w)}{p_w} \leq 2\mathbb{E}N(w)/p_w$, we write $\mathbb{E}\tilde{N}(w) \geq \mathbb{E}N(w) \frac{1}{\sqrt{1+2d_w z_1^2/p_w}}$. Since the distribution (4) is not the variance minimizing distribution, we have

$$\begin{aligned} \mathbb{E}\tilde{N}(w)\text{Var}\tilde{\beta}(w) &\geq \mathbb{E}\tilde{N}(w)\text{Var}\hat{\beta}(w) \\ &\geq \mathbb{E}N(w)\text{Var}\hat{\beta}(w) \left(\frac{1}{\sqrt{1+2d_w z_1^2/p_w}} \right). \end{aligned}$$

We establish (5) by showing that $\frac{1}{\sqrt{1+2d_w z_1^2/p_w}} = 1 + O(w)$. From the definition of d_w ,

$$\frac{wz_1}{p_w} \sum_{n=0}^{\infty} \frac{p_w(1-p_w)^n}{\sqrt{1+d_w z_1^2 n}} = 1,$$

and thus, using $p_w \leq wz_1$, we write that $1 - wz_1 \leq (1 - p_w) \sum_{n=1}^{\infty} \frac{p_w(1-p_w)^{n-1}}{\sqrt{1+d_w z_1^2 n}} \leq (1 - p_w) \frac{1}{\sqrt{1+d_w z_1^2}}$. Consequently, $1 + d_w z_1^2 \leq \left(\frac{1-p_w}{1-wz_1}\right)^2$. Hence, using (8), $1 + d_w z_1^2 = 1 + O(w^2)$, that is, $d_w = O(w^2)$ and thus, $d_w/p_w = O(w)$ because $p_w = O(w)$. This concludes that $\frac{1}{\sqrt{1+2d_w z_1^2/p_w}} = 1 + O(w)$ and hence establishes (5).

A.4 Proof of Theorem 3

From the assumption, we have $w_k \leq \frac{M_1(k)}{M_2(k)}$. Therefore, similar to (8), we can obtain that

$$0 < w_k M_1(k) - w_k^2 M_2(k)/2 \leq P_k \leq w_k M_1(k). \quad (15)$$

From the definitions of $w_k, M_1(k)$ and $M_2(k)$, it is easy to see that $\lim_{k \rightarrow \infty} k w_k M_1(k) = 1, a.s.$. Using the upper bound in (15), we show that $\delta_k := 1 - \frac{1+w_k^2 z_2 - 2w_k z_1}{1-P_k}$ is lower bounded by $\frac{w_k}{1-P_k} (2z_1 - M_1(k) - w_k z_2)$. Since $\lim_{k \rightarrow \infty} M_1(k) = z_1, a.s.$ and $\lim_{k \rightarrow \infty} k w_k = 1/z_1, a.s.$, for every realization of $\tilde{Z}_1, \tilde{Z}_2, \dots$, there exists a K such that $\delta_k > 0$ for all $k \geq K$, and hence $\text{Var}(\hat{\beta}(w_k) | M_1(k), M_2(k))$ is finite and equal to $\frac{w_k^2}{P_k \delta_k} - \beta^2$. It is now enough to show that

$$k \left(\frac{w_k^2 z_1^2}{P_k \delta_k} - 1 \right) \rightarrow \frac{\text{Var}Z}{z_1^2} [1 + \chi_1^2], a.s., \quad \text{as } k \rightarrow \infty.$$

Write $\left(\frac{w_k^2 z_1^2}{P_k \delta_k} - 1 \right) = \left(\frac{w_k^2 z_1^2}{p_{w_k}^2} - 1 \right) + w_k^2 z_1^2 \left(\frac{1}{P_k \delta_k} - \frac{1}{p_{w_k}^2} \right)$, where $p_{w_k} = 1 - \sqrt{1 + w_k^2 z_2 - 2w_k z_1}$. Using (9) and $\lim_{k \rightarrow \infty} k w_k = 1/z_1$, we have $\lim_{k \rightarrow \infty} k \left(\frac{w_k^2 z_1^2}{p_{w_k}^2} - 1 \right) = \frac{\text{Var}Z}{z_1^2}, a.s.$ By simplifying the terms in $P_k \delta_k - p_{w_k}^2$, we have $\frac{1}{P_k \delta_k} - \frac{1}{p_{w_k}^2} = \frac{1}{P_k \delta_k (1-P_k)} \left(\frac{P_k - p_{w_k}}{p_{w_k}} \right)^2$. Since $P_k - p_{w_k} = w_k (M_1(k) - z_1)(1 + O(w_k))$, we can write $\frac{P_k - p_{w_k}}{p_{w_k}} = \frac{w_k}{p_{w_k}} (M_1(k) - z_1)(1 + O(w))$. Therefore, using $\lim_{k \rightarrow \infty} \frac{w_k^2 z_1^2}{P_k \delta_k (1-P_k)} = 1, a.s.$ and the fact that asymptotically $\sqrt{k}(M_1(k) - z_1)$ has a zero-mean normal distribution with variance $\text{Var}Z$, we complete the proof with the observation that $w_k/p_{w_k} \rightarrow 1/z_1$ as $k \rightarrow \infty$.

REFERENCES

- Andrieu, C., and G. O. Roberts. 2009. "The Pseudo-Marginal Approach for Efficient Monte Carlo Computations". *The Annals of Statistics* 37(2):697–725.
- Asmussen, S., and P. W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*, Volume 57 of *Stochastic Modelling and Applied Probability*. New York: Springer.
- Beaumont, M. A. 2003. "Estimation of Population Growth or Decline in Genetically Monitored Populations". *Genetics* 164(3):1139–1160.
- Blanchet, J. H., N. Chen, and P. W. Glynn. 2015. "Unbiased Monte Carlo Computation of Smooth Functions of Expectations via Taylor Expansions". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 360–367. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Blanchet, J. H., and P. W. Glynn. 2015. "Unbiased Monte Carlo for Optimization and Functions of Expectations via Multi-Level Randomization". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3656–3667. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Carter, L., and E. Cashwell. 1975. *Particle-Transport Simulation with the Monte Carlo Method*. (Technical Report). New Mexico: Los Alamos Scientific Lab.
- Glynn, P. W. 2006. "Chapter 16 Simulation Algorithms for Regenerative Processes". In *Simulation*, edited by S. G. Henderson and B. L. Nelson, Volume 13 of *Handbooks in Operations Research and Management Science*, 477–500. Elsevier.
- Glynn, P. W., and C.-H. Rhee. 2014. "Exact Estimation for Markov Chain Equilibrium Expectations". *Journal of Applied Probability* 51A:377–389.
- Glynn, P. W., and W. Whitt. 1992. "The Asymptotic Efficiency of Simulation Estimators". *Operations Research* 40(3):505–520.
- Jacob, P. E., and A. H. Thiery. 2015. "On Nonnegative Unbiased Estimators". *The Annals of Statistics* 43(2):769–784.
- Lux, I., and L. Koblinger. 1991. *Monte Carlo Particle Transport Methods: Neutron and Photon Calculations*, Volume 102. Boca Raton: CRC Press.
- Lyne, A.-M., M. Girolami, Y. Atchadé, H. Strathmann, and D. Simpson. 2015. "On Russian Roulette Estimates for Bayesian Inference with Doubly-Intractable Likelihoods". *Statistical Science* 30(4):443–467.

- McLeish, D. 2011. “A General Method for Debiasing a Monte Carlo Estimator”. *Monte Carlo Methods and Applications*. 17:301–315.
- Moka, S. B., and S. Juneja. 2015. “Regenerative Simulation for Queuing Networks with Exponential or Heavier Tail Arrival Distributions”. *ACM Transactions on Modeling and Computer Simulation* 25(4):22:1–22:22.
- Møller, J., and R. P. Waagepetersen. 2003. “An Introduction to Simulation-Based Inference for Spatial Point Processes”. In *Spatial Statistics and Computational Methods*, edited by J. Møller, 143–198. New York: Springer.
- Rhee, C.-H., and P. W. Glynn. 2015. “Unbiased Estimation with Square Root Convergence for SDE Models”. *Operations Research* 63(5):1026–1043.
- Richter, W. 1965. “Limit Theorems for Sequences of Random Variables with Sequences of Random Indices”. *Theory of Probability & Its Applications* 10(1):74–84.
- Rubinstein, R. Y., and D. P. Kroese. 2017. *Simulation and the Monte Carlo method*. 3rd ed. Hoboken: John Wiley & Sons.
- Shiryayev, A. N. 1996. *Probability*. 2nd ed. Graduate Texts in Mathematics. New York: Springer-Verlag.
- Walker, S. G. 2011. “Posterior Sampling When the Normalizing Constant is Unknown”. *Communications in Statistics - Simulation and Computation* 40(5):784–792.
- Wei, C., and I. Murray. 2016. “Markov Chain Truncation for Doubly-Intractable Inference”. *arXiv e-prints*:arXiv:1610.05672.

AUTHOR BIOGRAPHIES

SARAT BABU MOKA is a Postdoctoral Research Fellow in The School of Mathematics and Physics at The University of Queensland. He obtained his PhD in System Science from Tata Institute of Fundamental Research in Mumbai. His research interests broadly lie in Applied Probability with a specific focus on Monte Carlo Simulation. In particular, his current research focus is on Spatial Point Processes, Perfect Sampling, Large Deviations Theory, and Variance Reduction techniques in Monte Carlo Simulation. His email address is s.babumoka@uq.edu.au

DIRK P. KROESE is a Professor of Mathematics and Statistics at The University of Queensland. He is the co-author of several influential monographs on simulation and Monte Carlo methods, including *Handbook of Monte Carlo Methods and Simulation and the Monte Carlo Method*, (3rd Edition). Dirk is a pioneer of the well-known Cross-Entropy method—an adaptive Monte Carlo technique, invented by Reuven Rubinstein, which is being used around the world to help solve difficult estimation and optimization problems in science, engineering, and finance. His personal website can be found under <https://people.smp.uq.edu.au/DirkKroese/>. His email address is kroese@maths.uq.edu.au.

SANDEEP JUNEJA is a Professor and Dean at the School of Technology and Computer Science in Tata Institute of Fundamental Research in Mumbai. He is currently on the editorial board of *Stochastic Systems*. Earlier he has been on editorial boards of *Mathematics of Operations Research*, *Management Science* and *ACM TOMACS*. His research interests lie in applied probability including in sequential decision making using bandit methods, financial mathematics, Monte Carlo methods, and game theoretic analysis of queues. His email address is juneja@tifr.res.in.