Submitted to: Proceedings of the 2007 INFORMS Simulation Society Research Workshop

ASYMPTOTIC ROBUSTNESS OF ESTIMATORS IN RARE-EVENT SIMULATION

Jose H. Blanchet

Department of Statistics
Harvard University, USA

Peter W. Glynn

Department of Management Science and Engineering
Standford University, USA

Pierre L'Ecuyer

IRISA, Rennes, and DIRO, Université de Montreal
C.P. 6128, Succ. Centre-Ville
Montréal (Québec), Canada, H3C 3J7

Werner Sandmann

Otto-Friedrich-Universität Bamberg
Feldkirchenstrasse 21, D-96045 Bamberg, Germany

Bruno Tuffin

IRISA/INRIA, Campus de Beaulieu
35042 Rennes Cedex, France

ABSTRACT

The asymptotic robustness of estimators in the context of rare-event simulation is often qualified by properties such as bounded relative error (BRE) and logarithmic efficiency (LE), also called asymptotic optimality. These properties guarantee that the estimator's relative error remains bounded or does not increase too fast, respectively, when the rare events becomes rarer. Their work-normalized versions take the computational work into account. Other recently introduced characterizations of estimators include bounded normal approximation, which implies a uniform Berry-Esseen bound as a function of the rarity parameter. However, these properties do not suffice to ensure that moments of order higher than 1 are well estimated. For example, they do not guarantee that the variance of the empirical variance remains under control as a function of the rarity parameter.

We introduce generalizations of the BRE and LE properties that take care of this limitation. They are named bounded relative moment of order k (BRM-k) and logarithmic efficiency of order k (LE-k). These properties are of interest for various estimators, including the empirical mean and the empirical variance. Work-normalized versions are also defined. As an illustration, we study the hierarchy between these properties (and a few others) for a model of highly-reliable Markovian system (HRMS) where the goal is to estimate the failure probability of the system.

1 INTRODUCTION

Rare-event simulation is a key tool in several areas such as reliability, telecommunications, finance, insurance, and computational physics, among others (Bucklew 2004, Heidelberger 1995, Juneja and Shahabuddin 2006). In typical rare-event settings, the Monte Carlo method is not viable unless special "acceleration" techniques are used to make the important rare events occur frequently enough. The two main families of techniques for doing that are splitting (Ermakov and Melas 1995, L'Ecuyer, Demers, and Tuffin 2007) and importance sampling (IS) (Glynn and Iglehart 1989, Juneja and Shahabuddin 2006).

Suppose we want to estimate a positive quantity $\gamma = \gamma(\varepsilon)$ that depends on a *rarity* parameter $\varepsilon > 0$. We assume that γ is a monotone (strictly) increasing function of ε and that $\lim_{\varepsilon\to 0^+} \gamma(\varepsilon) = 0$. We have a family of estimators $Y = Y(\varepsilon)$ taking their values in $[0,\infty)$, such that $\mathbb{E}[Y(\varepsilon)] = \gamma(\varepsilon)$ for each $\varepsilon > 0$. In applications, $\gamma(\varepsilon)$ can be a performance measure defined as a mathematical expectation, and some model parameters are defined as functions of ε in a convenient way. For example, in queuing systems, the service time and inter-arrival time distributions and the buffer sizes might depend on ε, while in Markovian reliability models, the failure rates and repair rates might be functions of ε. The convergence speed of $\gamma(\varepsilon)$ toward 0 may depend on how the model is parameterized, but the robustness properties introduced in this paper do not depend on this speed; they depend only on the convergence speed of some moments of $Y(\varepsilon)$ relative to that of $\gamma(\varepsilon)$.

We may want to compute a confidence interval on $\gamma(\varepsilon)$ based on i.i.d. replicates of $Y(\varepsilon)$. To do this via the classical central limit theorem (CLT), we need reliable estimators for both the mean $\gamma(\varepsilon)$ and the variance $\sigma^2(\varepsilon) = \mathbb{E}[(Y(\varepsilon) - \gamma(\varepsilon))^2]$. We want these estimators to remain robust in the sense that their relative error remains bounded (or grows only very slowly) when $\varepsilon \to 0$. Under the (unrealistic) assumption that the width of the confidence interval can be computed by using the exact variance, the relative width remains bounded if the relative error $\sigma(\varepsilon)/\gamma(\varepsilon)$ is bounded. An estimator with the latter property is said to have *bounded relative error* (BRE) (e.g., Heidelberger 1995). To estimate this relative width properly, we need a robust estimator of $\sigma^2(\varepsilon)$ relative to $\gamma^2(\varepsilon)$; this involves the fourth moment of $Y(\varepsilon)$. In rare-event settings, reliable (relative) mean and variance estimators are typically difficult to obtain. The variance is also often more difficult to estimate than the mean (relative to the mean).

More generally, a CLT based on an Edgeworth-type expansion would require reliable estimates of higher relative moments. As another example, if we want to compare the efficiencies of alternative mean estimators, we want the error on the variance to be significantly smaller than the variance itself; that is, a small relative error of the empirical variance. Robustness of relative moment estimators can be important for other applications as well.

In the remainder of this extended abstract, we define such robustness characterizations, briefly examine some of their properties, and give examples. The details will be provided in the (forthcoming) full paper.

We use the following notation. For a function $f : (0,\infty) \to \mathbb{R}$, we say that $f(\varepsilon) = o(\varepsilon^d)$ if $f(\varepsilon)/\varepsilon^d \to 0$ as $\varepsilon \to 0$; $f(\varepsilon) = O(\varepsilon^d)$ if $|f(\varepsilon)| \leq c_1\varepsilon^d$ for some constant $c_1 > 0$ for all ε sufficiently small; $f(\varepsilon) = \underline{O}(\varepsilon^d)$ if $|f(\varepsilon)| \geq c_2\varepsilon^d$ for some constant $c_2 > 0$ for all ε sufficiently small; and $f(\varepsilon) = \Theta(\varepsilon^d)$ if $f(\varepsilon) = \underline{O}(\varepsilon^d)$ and $f(\varepsilon) = O(\varepsilon^d)$. We use the shorthand notation $Y(\varepsilon)$ to refer to the family of estimators $\{Y(\varepsilon), \varepsilon > 0\}$. We also write "$\to 0$" to mean "$\to 0^+$."

2 ROBUSTNESS PROPERTIES OF THE RELATIVE MOMENTS

Bounded relative moments. The *relative moment of order k* of the estimator $Y(\varepsilon)$ is defined as

$$m_k(\varepsilon) = \mathbb{E}[Y^k(\varepsilon)]/\gamma^k(\varepsilon). \tag{1}$$

We say that $Y(\varepsilon)$ has *bounded relative moment of order k* (BRM-*k*) if

$$\limsup_{\varepsilon \to 0} m_k(\varepsilon) < \infty. \tag{2}$$

Since $m_2(\varepsilon) = \mathbb{E}[Y^2(\varepsilon)]/\gamma^2(\varepsilon) = 1 + \sigma^2(\varepsilon)/\gamma^2(\varepsilon)$, it is easily seen that BRE is equivalent to BRM-2. The BRE property means that the expected width of a confidence interval on $\gamma(\varepsilon)$ based on i.i.d. replicates of $Y(\varepsilon)$ and the classical central-limit theorem (CLT) decreases at least as fast as $\gamma(\varepsilon)$ when $\varepsilon \to 0$.

Using Jensen's inequality, we can show that $m_k(\varepsilon)$ is nondecreasing in k, so that BRM-k implies BRM-k' for $1 \leq k' < k$. We also have that for any positive integers k, ℓ, m, and any non-negative random variable $X(\varepsilon)$, if $Y(\varepsilon) = X^\ell(\varepsilon)$ is BRM-mk, then $Y'(\varepsilon) = X^{m\ell}(\varepsilon)$ is BRM-k.

Work-normalized measures. The BRM-*k* property does not take the computational work into account. For example, if an estimator $\tilde{Y}(\varepsilon)$ has $m_2(\varepsilon) = 1/\varepsilon$, we can turn it into an estimator $Y(\varepsilon)$ with the BRM-*k* property simply by defining $Y(\varepsilon)$ as the average of $\lceil 1/\varepsilon \rceil$ i.i.d. copies of $\tilde{Y}(\varepsilon)$. However, the computing cost of $Y(\varepsilon)$ increases without bound as $\varepsilon \to 0$. A less trivial example of an estimator that is BWRM-2 but not BRM-2 is exhibited in Cancela, Rubino, and Tuffin (2005). In that example, $t(\varepsilon) = O(\varepsilon)$ but $m_2(\varepsilon) = \underline{O}(\varepsilon^{-1})$.

To take the work into account, we introduce the *work-normalized relative moment of order k*, defined as $t(\varepsilon)m_k(\varepsilon)$, where $t(\varepsilon)$ is the expected computational time to generate $Y(\varepsilon)$. We say that $Y(\varepsilon)$ has *bounded work-normalized relative moment of order k* (BWRM-*k*) if

$$\limsup_{\varepsilon \to 0} t(\varepsilon)m_k(\varepsilon) < \infty. \tag{3}$$

Logarithmic efficiency. The estimator $Y(\varepsilon)$ has *logarithmic efficiency of order k* (LE-*k*) if

$$\lim_{\varepsilon \to 0} \frac{\ln \mathbb{E}[Y^k(\varepsilon)]}{k \ln \gamma(\varepsilon)} = 1. \tag{4}$$

One intuitive interpretation of LE-*k* could be that when $\gamma^k(\varepsilon)$ converges to zero exponentially fast, $\mathbb{E}[Y^k(\varepsilon)]$ also converges exponentially fast and at the same exponential rate. This is the best possible rate; it cannot converge at a faster rate because Jensen's inequality ensures that $\mathbb{E}[Y^k(\varepsilon)] - \gamma^k(\varepsilon) \geq 0$. LE-2 is the usual definition of logarithmic efficiency, also called asymptotic optimality. In general, LE-*k* is weaker than BRM-*k*, but there are situations where the two are equivalent. The following examples illustrate the two possibilities. There are several rare-event applications where $\gamma(\varepsilon)$ decreases exponentially fast with ε and where practical BRM-2 estimators are not readily available (e.g., in queueing and finance), but where estimators with the (weaker) LE-2 property, and often with the LE-*k* property for all *k*, have been constructed by exploiting the theory of large deviations (Asmussen 2002, Glasserman 2004, Heidelberger 1995, Juneja and Shahabuddin 2006, Siegmund 1976).

Example 1 Suppose that $\gamma(\varepsilon) = \exp[-\eta/\varepsilon]$ for some constant $\eta > 0$ and that our estimator has $\sigma^2(\varepsilon) = q(1/\varepsilon)\exp[-2\eta/\varepsilon]$ for some polynomial function q. Then, the LE property is easily verified, whereas BRE does not hold because $m_2(\varepsilon) = q(1/\varepsilon) \to \infty$ when $\varepsilon \to 0$.

Example 2 Suppose that $\gamma^k(\varepsilon) = q_1(\varepsilon) = \varepsilon^{t_1} + o(\varepsilon^{t_1})$ and $\mathbb{E}[Y^k(\varepsilon)] = q_2(\varepsilon) = \varepsilon^{t_2} + o(\varepsilon^{t_2})$. That is, both converge to 0 as a polynomial in ε. Clearly, $t_2 \leq t_1$, because $\mathbb{E}[Y^k(\varepsilon)] - \gamma^k(\varepsilon) \geq 0$. We have BRM-$k$ if and only if (iff) $q_2(\varepsilon)/q_1(\varepsilon)$ remains bounded when $\varepsilon \to 0$, iff $t_2 = t_1$. On the other hand, $-\ln q_1(\varepsilon) = -\ln(\varepsilon^{t_1}(1+o(1))) = -t_1\ln(\varepsilon) - \ln(1+o(1))$ and similarly for $q_2(\varepsilon)$ and t_2. Then,

$$\lim_{\varepsilon\to 0}\frac{\ln\mathbb{E}[Y^k(\varepsilon)]}{k\ln\gamma(\varepsilon)} = \lim_{\varepsilon\to 0}\frac{t_2\ln\varepsilon}{t_1\ln\varepsilon} = \frac{t_2}{t_1}.$$

Thus, LE-k holds iff $t_2 = t_1$, which means that BRM-k and LE-k are equivalent in this case.

Other robustness properties We quickly mention other related properties. We say that $Y(\varepsilon)$ has the *bounded normal approximation* (BNA) property if

$$\limsup_{\varepsilon\to 0}\mathbb{E}\left[|Y(\varepsilon)-\gamma(\varepsilon)|^3\right]/\sigma^3(\varepsilon) < \infty. \quad (5)$$

This property implies that $\sqrt{n}|F_n(x) - \Phi(x)|$ remains bounded as a function of ε, i.e., that the Berry-Esseen bound on the approximation error of F_n by the normal distribution remains in $O(n^{-1/2})$ uniformly in ε. BNA is equivalent to BRM-3 if $\sigma(\varepsilon) = \Theta(\gamma(\varepsilon))$, but not otherwise.

Two other robustness properties are introduced in Tuffin (2004), in the context of the application of IS to an HRMS model, under the name of "well estimated mean and variance." They are equivalent to none of the other properties defined here.

Robustness of the empirical variance An important special case is the stability of the empirical variance as an estimator of the true variance $\sigma^2(\varepsilon)$. Let $X_1(\varepsilon),\ldots,X_n(\varepsilon)$ be an i.i.d. sample of $X(\varepsilon)$, where $n \geq 2$. The empirical mean and empirical variance are $\bar{X}_n(\varepsilon) = (X_1(\varepsilon)+\cdots+X_n(\varepsilon))/n$ and

$$S_n^2(\varepsilon) = \frac{1}{n-1}\sum_{i=1}^{n}(X_i(\varepsilon)-\bar{X}_n(\varepsilon))^2.$$

If we take $Y(\varepsilon) = S_n^2(\varepsilon)$ in our framework of the previous subsections, we obtain definitions of the robustness properties for $S_n^2(\varepsilon)$ as an estimator of $\sigma^2(\varepsilon)$.

It can be proved that if $\sigma^2(\varepsilon) = \Theta(\gamma^2(\varepsilon))$ (where $\gamma(\varepsilon) = \mathbb{E}[X(\varepsilon)]$), then BRE-$2k$ for $X(\varepsilon)$ implies BRM-k for $S_n^2(\varepsilon)$, for any $k \geq 1$. These two properties are not equivalent in general, however. In particular, BRE for $S_n^2(\varepsilon)$ is not equivalent to BRM-4 for $X(\varepsilon)$. Indeed, $\mathrm{Var}[S_n^2(\varepsilon)]/\sigma^4(\varepsilon) = \Theta\left(\mathbb{E}[(X(\varepsilon)-\gamma(\varepsilon))^4]/\sigma^4(\varepsilon)\right)$, which differs in general from $\Theta\left(\mathbb{E}[X^4(\varepsilon)]/\gamma^4(\varepsilon)\right)$.

3 EXAMPLE: A RELIABILITY MODEL

We consider an HRMS with c types of components and n_i components of type i, for $i = 1,\ldots,c$, all those components being subject to failures and repairs (Shahabuddin 1994). The system is modeled by a continuous time Markov chain $(Y(t) = (Y_1(t),\ldots,Y_c(t)), t \geq 0\}$, where $Y_i(t)$ is the number of failed type-i components at time t. We suppose that the state space is partitioned in two subsets $\mathscr{U}$ and $\mathscr{F}$ of up and failed states, where $\mathscr{U}$ is a decreasing set (i.e., if $x \in \mathscr{U}$ and $x \geq y$, then $y \in \mathscr{U}$) that contains the state $\mathbf{0} = (0,\ldots,0)$ in which all the components are operational. Failure and repair rates of individual components are independent and given by $\lambda_i(x) = a_i(x)\varepsilon^{b_i(x)}$ and $\mu_i(x) = \Theta(1)$ for type-i components when the current state is x, where $a_i(x) > 0$ is a real number and $b_i(x) \geq 1$ an integer. The parameter $\varepsilon \ll 1$ represents the rarity of failures; failure propagation is allowed (which may depend on ε) as well as grouped repairs (independent of ε).

Our goal is to estimate $\gamma(\varepsilon) = \mathbb{P}[\tau_{\mathscr{F}} < \tau_{\mathbf{0}}]$, where $\tau_{\mathscr{F}} = \inf\{j > 0 : X_j \in \mathscr{F}\}$ and $\tau_{\mathbf{0}} = \inf\{j > 0 : X_j = \mathbf{0}\}$. We further assume that from $\mathbf{0}$, the failures having probability $\Theta(1)$ do not directly drive to $\mathscr{F}$, since otherwise $\gamma = \Theta(1)$ is not a rare event probability. Shahabuddin (1994) shows that for this model, there is a number $r > 0$ such that $\gamma(\varepsilon) = \Theta(\varepsilon^r)$. To study this probability, we can limit ourselves to the canonically embedded discrete-time Markov chain (DTMC) $\{X_j,\ j \geq 0\}$, with transition matrix $\mathbf{P}$, is defined by $X_j = Y(\xi_j)$ for $j = 0,1,2,\ldots$, where $\xi_0 = 0$ and $0 < \xi_1 < \xi_2 < \cdots$ are the jump times of the CTMC. We use $\mathbb{P}$ to denote the corresponding measure on the sample paths of the DTMC. Our final assumptions are that the DTMC is irreducible and that at least one repairman is active whenever a component is failed.

Naive Monte Carlo simulation is inefficient in this case (the relative error of this estimator increases toward infinity when $\varepsilon \to 0$) and something else must be done to obtain a viable estimator. Several IS schemes have been proposed in the literature for this HRMS model; see, e.g., Cancela, Rubino, and Tuffin (2002), Nakayama (1996), Shahabuddin (1994). They all pertain to a class $\mathscr{I}$ of measures $\mathbb{P}^*$ defined by a transition probability matrix $\mathbf{P}^*$ with the following property: whenever $\mathbf{P}(x,y) = \Theta(\varepsilon^d)$, then $\mathbf{P}^*(x,y) = \Theta(\varepsilon^\ell)$ for $\ell \leq d$. This means that the probability of a transition under the new probability transition matrix is never significantly smaller than under the original one. From now on, we assume that $\mathbf{P}^*$ satisfies this property. For example, the so-called *simple failure biasing* (SFB), is such that for states $x \in \mathscr{F} \cup \{\mathbf{0}\}$, we have $\mathbf{P}^*(x,y) = \mathbf{P}(x,y)$ for all states y, i.e., the transition probabilities are unchanged. For any other state x, a fixed probability ρ is assigned to the set of all failure transitions, and a probability $1-\rho$ is assigned to the set of all repair transitions. In each of these two subsets, the individual probabilities are taken proportionally to the original ones.

For a given sample path ending at step $\tau = \min(\tau_{\mathscr{F}},\ \tau_{\mathbf{0}})$, the likelihood ratio is

$$L = L(X_0,\ldots,X_\tau) = \frac{\mathbb{P}[(X_0,\ldots,X_\tau)]}{\mathbb{P}^*[(X_0,\ldots,X_\tau)]} = \prod_{j=1}^{\tau}\frac{\mathbf{P}(X_{j-1},X_j)}{\mathbf{P}^*(X_{j-1},X_j)}$$

and the corresponding (unbiased) IS estimator of $\gamma(\varepsilon)$ is given by $Y(\varepsilon) = 1_{\{\tau_{\mathscr{F}} < \tau_{\mathbf{0}}\}} L(X_0,\ldots,X_\tau)$.

A characterization of the IS schemes for the HRMS model that satisfy the BRE property was obtained by Nakayama (1996) and the equivalence between BRE and LE for this model was mentioned without proof in Heidelberger (1995). Our first result generalizes this. Indeed, in our HRMS framework, with a measure in $\mathscr{I}$, we can prove that the three properties BRM-*k*, BWRM-*k*, and LE-*k* are equivalent. These three properties are also equivalent if we replace $Y(\varepsilon)$ by its *g*th empirical moment or by its empirical variance.

Another interesting remark is that for any IS measure in $\mathscr{I}$, $\sigma^2(\varepsilon) = \underline{O}(\gamma^2(\varepsilon))$, and even more, $\mathbb{E}[(Y(\varepsilon)-\gamma(\varepsilon))^k] = \underline{O}(\gamma^k(\varepsilon))$ in general.

We have also been able to characterize BRM-*k* for the *g*th empirical moment in the HRMS framework. In particular, we have gives characterizations for BRM-*k* (of the mean), as well as BRE and LE for the empirical variance. Define Δ_m as the set of paths $(x_0,\cdots,x_n)$ such that $\tau_{\mathscr{F}} < \tau_{\mathbf{0}}$ and $\mathbb{P}\{(X_0,\cdots,X_\tau) = (x_0,\cdots,x_n)\} = \Theta(\varepsilon^m)\}$ and let s_m be the integer such that $\mathbb{E}[Y^m(\varepsilon)] = \Theta(\varepsilon^{s_m})$ with $s_m \le mr$. For an IS measure $\mathbb{P}^* \in \mathscr{I}$, we have BRM-*k* for the *g*th empirical moment if and only if for all integers *m* such that $r \le m < ks_g$ and all $(x_0,\cdots,x_n) \in \Delta_m$,

$$\mathbb{P}^*\{(X_0,\cdots,X_\tau) = (x_0,\cdots,x_n)\} = \Theta(\varepsilon^\ell)$$

for some $\ell \le k(mg - s_g)/(kg-1)$. This result means that a path cannot be too rare under the IS measure $\mathbb{P}^*$ to verify BRM-*k* of the *g*th moment. Special cases of this result were obtained under the same conditions by Nakayama (1996) for BRE and by Tuffin (1999), Tuffin (2004) for BNA.

In the specific case of empirical mean and variance, we have BRM-*k* for $Y(\varepsilon)$ if and only if for all integers *m* such that $r \le m < kr$ and all $(x_0,\cdots,x_n) \in \Delta_m$,

$$\mathbb{P}^*\{(X_0,\cdots,X_\tau) = (x_0,\cdots,x_n)\} = \Theta(\varepsilon^\ell)$$

for $\ell \le k(m-r)/(k-1)$. We also have BRM-*k* for $Y^2(\varepsilon)$ if and only if the same condition holds with $\ell \le k(2m - s_2)/(2k-1)$. We have BRM-*k* for the empirical variance if and only if $\ell \le k(2m-s)/(2k-1)$ with *s* the integer such that $\sigma^2(\varepsilon) = \Theta(\varepsilon^s)$ (with $s_2 = s$ iff $\sigma^2(\varepsilon) = \Theta(\mathbb{E}[Y^2(\varepsilon)])$).

Some relationships between measures of robustness were proved by Tuffin (2004) for a more restricted class of IS measures such that for each failure transition not starting from $\mathbf{0}$, whenever $\mathbf{P}(x,y) = \Theta(\varepsilon^d)$, then $\mathbf{P}^*(x,y) = \Theta(\varepsilon^\ell)$ with $\ell < d$. It was shown that BNA implies BRE, but that the converse is not true. Now, for the same class of measures, we are also able to prove that BRE for the empirical variance implies BNA, but that the converse is not true.

ACKNOWLEDGMENTS

This research has been supported by NSERC-Canada grant No. ODGP0110050 and a Canada Research Chair to the third author, as well as EuroFGI Network of Excellence and INRIA's cooperative research initiative RARE to the fifth author.

REFERENCES

Asmussen, S. 2002. Large deviations in rare events simulation: Examples, counterexamples, and alternatives. In *Monte Carlo and Quasi-Monte Carlo Methods 2000*, ed. K.-T. Fang, F. J. Hickernell, and H. Niederreiter, 1–9. Berlin: Springer-Verlag.

Bucklew, J. A. 2004. *Introduction to rare event simulation*. New York: Springer-Verlag.

Cancela, H., G. Rubino, and B. Tuffin. 2002. MTTF estimation by Monte Carlo methods using Markov models. *Monte Carlo Methods and Applications* 8 (4): 312–341.

Cancela, H., G. Rubino, and B. Tuffin. 2005. New measures of robustness in rare event simulation. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 519–527.

Ermakov, S. M., and V. B. Melas. 1995. *Design and analysis of simulation experiments*. Dordrecht, The Netherlands: Kluwer Academic.

Glasserman, P. 2004. *Monte Carlo methods in financial engineering*. New York: Springer-Verlag.

Glynn, P. W., and D. L. Iglehart. 1989. Importance sampling for stochastic simulations. *Management Science* 35:1367–1392.

Heidelberger, P. 1995. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation* 5 (1): 43–85.

Juneja, S., and P. Shahabuddin. 2006. Rare event simulation techniques: An introduction and recent advances. In *Simulation*, ed. S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science, 291–350. Amsterdam, The Netherlands: Elsevier. Chapter 11.

L'Ecuyer, P., V. Demers, and B. Tuffin. 2007. Rare-events, splitting, and quasi-Monte Carlo. *ACM Transactions on Modeling and Computer Simulation* 17 (2). To appear.

Nakayama, M. K. 1996. General conditions for bounded relative error in simulations of highly reliable Markovian systems. *Advances in Applied Probability* 28:687–727.

Shahabuddin, P. 1994. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science* 40 (3): 333–352.

Siegmund, D. 1976. Importance sampling in the Monte Carlo study of sequential tests. *The Annals of Statistics* 4:673–684.

Tuffin, B. 1999. Bounded normal approximation in simulations of highly reliable Markovian systems. *Journal of Applied Probability* 36 (4): 974–986.

Tuffin, B. 2004, September. On numerical problems in simulations of highly reliable Markovian systems. In *Proceedings of the 1st International Conference on Quantitative Evaluation of SysTems (QEST)*, 156–164. University of Twente, Enschede, The Netherlands: IEEE CS Press.

AUTHOR BIOGRAPHIES

JOSE H. BLANCHET is Assistant Professor of Statistics at Harvard University. Jose holds a M.Sc. in Engineering-Economic Systems and Operations Research and a Ph.D. in Management Science and Engineering, both from Stanford University. He also holds two B.Sc. degrees: one in Actuarial Science and another one in Applied Mathematics from ITAM (Mexico). Jose worked for two years as an analyst in Protego Financial Advisors, a leading investment bank in Mexico. He has research interests in applied probability, computational finance, performance engineering, queueing theory, risk management, rare-event analysis, statistical inference, stochastic modelling, and simulation.

PETER W. GLYNN received his Ph.D. in Operations Research from Stanford University in 1982. He then joined the faculty of the University of Wisconsin at Madison, where he held a joint appointment between the Industrial Engineering Department and Mathematics Research Center, and courtesy appointments in Computer Science and Mathematics. In 1987, he returned to Stanford, where he is now the Thomas Ford Professor of Engineering in the Department of Management Science and Engineering. He also has a courtesy appointment in the Department of Electrical Engineering. He is a member of INFORMS and a fellow of the Institute of Mathematical Statistics. His e-mail address is `glynn@stanford.edu`, and his web page is `www.stanford.edu/dept/MSandE/people/faculty/glynn`.

PIERRE L'ECUYER is Professor in the Département d'Informatique et de Recherche Opérationnelle, at the Université de Montréal, Canada. He holds the Canada Research Chair in Stochastic Simulation and Optimization. His main research interests are random number generation, quasi-Monte Carlo methods, efficiency improvement via variance reduction, sensitivity analysis and optimization of discrete-event stochastic systems, and discrete-event simulation in general. He is currently Associate/Area Editor for *ACM TOMACS*, *ACM TOMS*, and *Statistical Computing*. He obtained the *E. W. R. Steacie* fellowship in 1995-97, a *Killam* fellowship in 2001-03, and became an INFORMS Fellow in 2006. His recent research articles are available on-line from his web page: `http://www.iro.umontreal.ca/~lecuyer`.

WERNER SANDMANN received his diploma degree and his PhD both in Computer Science from the University of Bonn (Germany) in 1998 and 2004, respectively. From 1998 to 2003 he was a research and teaching assistant in the Computer Science Department of the University of Bonn. Since 2004 he is assistant professor of Computer Science in the Department of Information Systems and Applied Computer Science of the University of Bamberg (Germany). His research interests are in stochastic modeling and performance evaluation of computer and communication systems. His e-mail address is `werner.sandmann@wiai.uni-bamberg.de`, and his web page is `http://www.uni-bamberg.de/index.php?id=sandmann&L=1`.

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. He spent 8 months as a postdoc at Duke University in 1999. His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems, and developing new Internet-pricing schemes. His web page is `http://www.irisa.fr/armor/lesmembres/Tuffin/Tuffin_en.htm`.