Proceedings of the 2010 Winter Simulation Conference
B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, eds.

CONFIDENCE INTERVALS FOR QUANTILES AND VALUE-AT-RISK WHEN APPLYING IMPORTANCE SAMPLING

Fang Chu

Department of Information Systems
New Jersey Institute of Technology
Newark, NJ 07102, U.S.A.

Marvin K. Nakayama

Department of Computer Science
New Jersey Institute of Technology
Newark, NJ 07102, U.S.A.

ABSTRACT

We develop methods to construct asymptotically valid confidence intervals for quantiles and value-at-risk when applying importance sampling (IS). We first employ IS to estimate the cumulative distribution function (CDF), which we then invert to obtain a point estimate of the quantile. To construct confidence intervals, we show that the IS quantile estimator satisfies a Bahadur-Ghosh representation, which implies a central limit theorem (CLT) for the quantile estimator and can be used to obtain consistent estimators of the variance constant in the CLT.

1 INTRODUCTION

Consider a random variable X having CDF F. For $0 < p < 1$, the p-quantile of X is defined to be $\xi_p = F^{-1}(p) \equiv \inf\{x : F(x) \geq p\}$. Quantiles are widely used as risk measures in practice. In finance, a quantile is known as a value-at-risk (VaR), and VaRs are often employed to assess the potential loss of a portfolio of risky assets. For example, there is a 1% chance of the loss of the portfolio over a given period of time (e.g., two weeks) exceeding the 0.99-VaR $\xi_{0.99}$. In project planning, one may want to determine a date by which there is a 95% chance that the project completes. This date is then the 0.95-quantile of the project-completion time.

This paper considers estimating ξ_p via simulation. Suppose we generate independent and identically distributed (i.i.d.) samples $X_1, X_2, \ldots, X_n$ from F. The fact that $\xi_p = F^{-1}(p)$ suggests estimating the p-quantile by $\hat{\xi}_{p,n} = F_n^{-1}(p)$, where F_n is the empirical distribution function, which assigns mass $1/n$ to each sample X_i. When the simulation process is based on generating i.i.d. samples from F, we call the method *crude Monte Carlo*.

We can indicate the error in the point estimate of a quantile by constructing a confidence interval. This is typically accomplished by first showing that the quantile estimator satisfies a central limit theorem (CLT) and then replacing the variance constant in the CLT with a consistent estimator to obtain a confidence interval. Section 2.3.3 of Serfling (1980) establishes a CLT for the quantile estimator when applying crude Monte Carlo.

However, applying crude Monte Carlo to estimate a quantile may result in a large confidence interval, especially when the quantile is extreme (i.e., p is close to 0 or 1). Different variance-reduction techniques may be used to address this issue. In particular, importance sampling (IS) (Glynn and Iglehart 1989), when properly applied, can reduce variance by orders of magnitude in rare-event simulations; see Heidelberger (1995). For rare-event problems, effective application of IS often requires changing the probabilistic dynamics of the system to increase the occurrence of the rare events of interest, such as extreme portfolio losses in a VaR example, and then recovering an unbiased estimator by multiplying by a correction factor known as the likelihood ratio.

Previous work on applying IS to estimate a quantile does not provide a direct method to construct confidence intervals. Glynn (1996) establishes CLTs for quantile estimators obtained by inverting various CDF estimators when applying IS, and Glasserman, Heidelberger, and Shahabuddin (2000) prove a CLT for a quantile estimator from a combination of IS and stratification. However, neither of these papers provides a consistent estimator of the variance constant κ_p^2 in the CLT for the IS quantile estimator. It turns out that $\kappa_p = \psi_p / f(\xi_p)$, where ψ_p^2 is the variance constant in the CLT satisfied by the IS estimator of the CDF evaluated at ξ_p and $f(\xi_p)$ is the density function (when it exists) of the original CDF F evaluated at the (unknown) quantile. In the case of crude Monte Carlo, Bloch and Gastwirth (1968), Bofinger (1975), and Babu (1986) provide consistent estimators of $f(\xi_p)$, but their proofs of consistency do not generalize when applying IS.

In this article, we provide consistent estimators of ψ_p and ϕ_p when applying IS, allowing the construction of asymptotically valid confidence intervals for ξ_p. We accomplish this by first showing that the IS quantile estimator satisfies Ghosh's (1971) weaker form of a Bahadur (1966) representation, which we call a Bahadur-Ghosh representation. The Bahadur-Ghosh representation not only leads to a consistent estimator of $f(\xi_p)$ but also implies a CLT for the IS quantile estimator under weaker conditions than those used in Glynn (1996) and Glasserman, Heidelberger, and Shahabuddin (2000). (We recently discovered that, independent of our work, the work of Sun and Hong 2010 establishes, under a stronger set of conditions, a stronger form of a Bahadur representation for the IS quantile estimator, which they use to prove a CLT. However, the particular representation they derive does not permit estimating $f(\xi_p)$.)

The rest of the article has the following organization. Section 2 reviews methods for estimating quantiles and the Bahadur-Ghosh representation for crude Monte Carlo. Section 3 establishes the Bahadur-Ghosh representation for IS quantile estimators and presents our methods to construct confidence intervals. Section 4 contains an empirical study of the finite-sample behavior of the confidence intervals for two stochastic models. We provide some concluding remarks in Section 5. All the proofs of our results can be found in Chu and Nakayama (2010).

2 QUANTILE ESTIMATION AND BAHADUR-GHOSH REPRESENTATION FOR CRUDE MONTE CARLO

For crude Monte Carlo, estimating quantiles typically entails first generating i.i.d. samples $X_1, X_2, \ldots, X_n$ from distribution F. Then, the empirical CDF F_n is defined as

$$F_n(x) = \frac{1}{n}\sum_{i=1}^{n} I(X_i \leq x), \tag{1}$$

where $I(A)$ is the indicator function of a set A, which assumes value 1 on A and 0 on its complement. Inverting F_n results in the p-quantile estimator $\hat{\xi}_{p,n} = F_n^{-1}(p)$. An equivalent way of computing $\hat{\xi}_{p,n}$ is to first sort the n samples in ascending order as $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$, and then set $\hat{\xi}_{p,n} = X_{(\lceil np \rceil)}$, where $\lceil \cdot \rceil$ is the round-up function.

Consider the following heuristic argument. When the sample size n is large, we have $F_n \approx F$, so $\hat{\xi}_{p,n} \approx \xi_p$. Thus, since $p = F(\xi_p)$, a Taylor approximation shows

$$p \approx F(\hat{\xi}_{p,n}) \approx F(\xi_p) + f(\xi_p)(\hat{\xi}_{p,n} - \xi_p) \approx F_n(\xi_p) + f(\xi_p)(\hat{\xi}_{p,n} - \xi_p)$$

because $F_n(\xi_p) \approx F(\xi_p)$. Hence, $\hat{\xi}_{p,n} \approx \xi_p - (F_n(\xi_p) - p)/f(\xi_p)$.

Making the above heuristic argument rigorous, Bahadur (1966) proves the following holds under the assumption that $f(\xi_p) > 0$ and the second derivative of F is bounded in a neighborhood of ξ_p:

$$\hat{\xi}_{p,n} = \xi_p - \frac{F_n(\xi_p) - p}{f(\xi_p)} + R_n, \tag{2}$$

where almost surely (a.s.),

$$R_n = O(n^{-3/4}(\log n)^{1/2}(\log\log n)^{1/4}) \text{ as } n \to \infty. \tag{3}$$

By "$Y_n = O(g(n))$ a.s." we mean that there exists a set Ω_0 such that $P(\Omega_0) = 1$ and for each $\omega \in \Omega_0$, there exists a constant $B(\omega)$ such that $|Y_n(\omega)| \leq B(\omega)g(n)$, for n sufficiently large. Equations (2) and (3) are known as a *Bahadur representation*. To understand the implications of this result, let $N(a, b^2)$ denote a normal random variable with mean a and variance b^2, and let $\Rightarrow$ denote convergence in distribution (Billingsley 1995, Section 25). It is well known (e.g., Serfling 1980, Section 2.3.3) that $\sqrt{n}(\hat{\xi}_{p,n} - \xi_p) \Rightarrow N(0, p(1-p)/f^2(\xi_p))$ as $n \to \infty$, and $\sqrt{n}(p - F_n(\xi_p))/f(\xi_p)$ has the same weak limit since $F_n(\xi_p)$ is the sample average of i.i.d. indicator functions $I(X_i \leq \xi_p)$, each with mean p. But the Bahadur representation goes further, showing the difference of the two quantities vanishes a.s., and it provides the rate at which this occurs. Thus, the Bahadur representation also sheds light onto why a quantile estimator, which is *not* a sample average, satisfies a CLT.

Ghosh (1971) establishes a weaker form of the Bahadur representation in (2) and (3). Requiring only that $f(\xi_p) > 0$, he shows that $\hat{\xi}_{p_n,n} = F_n^{-1}(p_n)$ with $p_n = p + O(n^{-1/2})$ satisfies

$$\hat{\xi}_{p_n,n} = \xi_p - \frac{F_n(\xi_p) - p_n}{f(\xi_p)} + R'_n \tag{4}$$

with

$$\sqrt{n}R'_n \Rightarrow 0 \text{ as } n \to \infty. \tag{5}$$

We call (4) and (5) a *Bahadur-Ghosh representation*, which suffices for most applications, including ours. It is easy to see that the Bahadur-Ghosh representation also implies that $\hat{\xi}_{p,n}$ satisfies a CLT (see Theorem 10.3 of David and Nagaraja 2003).

3 QUANTILE ESTIMATION USING IMPORTANCE SAMPLING

Because crude Monte Carlo is sometimes inefficient for estimating quantiles, particularly extreme quantiles, we now consider using IS to estimate ξ_p, as in Glynn (1996). Let F_* be another CDF such that F is absolutely continuous with respect to F_* (p. 422 of Billingsley 1995). Define E_* to be expectation under CDF F_*. Also, let $L(t) = F(dt)/F_*(dt)$ be the likelihood ratio at t. Then we can write

$$F(x) = \int I(t \le x)\, F(dt) = \int I(t \le x)\, L(t)\, F_*(dt) = E_*\left[I(X \le x)\, L(X)\right].$$

This then motivates estimating the CDF F via IS as follows. Generate i.i.d. samples $X_1, \ldots, X_n$ from CDF F_* and the IS estimator of F is then

$$\tilde{F}_n(x) = \frac{1}{n}\sum_{i=1}^{n} I(X_i \le x) L(X_i). \tag{6}$$

(We will sometimes denote $L(X)$ and $L(X_i)$ by L and L_i, respectively, in the following.) Inverting $\tilde{F}_n$ results in the IS quantile estimator. The following theorem shows that the IS quantile estimator satisfies a Bahadur-Ghosh representation.

Theorem 1. *Suppose $f(\xi_p) > 0$, and suppose there exists $\varepsilon > 0$ and $\delta > 0$ such that $E_*[I(X < \xi_p + \delta)L^{2+\varepsilon}] < \infty$. Then $\tilde{\xi}_{p_n,n} = \tilde{F}_n^{-1}(p_n)$ with $p_n = p + O(n^{-1/2})$ satisfies*

$$\tilde{\xi}_{p_n,n} = \xi_p - \frac{\tilde{F}_n(\xi_p) - p_n}{f(\xi_p)} + \tilde{R}_n \tag{7}$$

with

$$\sqrt{n}\tilde{R}_n \Rightarrow 0 \text{ as } n \to \infty. \tag{8}$$

As we noted in Section 1, Sun and Hong (2010) establish that the IS quantile estimator satisfies an a.s. Bahadur representation as in (2) and (3) under stronger assumptions. They further assume that the density f is positive and continuously differentiable in a neighborhood of ξ_p and that the likelihood ratio $L(x)$ is bounded in a neighborhood of ξ_p. Moreover, they do not examine the case of perturbed p_n, which is essential in our approach for estimating $1/f(\xi_p)$, an important component in constructing a confidence interval for ξ_p.

It is straightforward to show that Theorem 1 implies $\tilde{\xi}_{p,n} = \tilde{F}_n^{-1}(p)$ is a consistent estimator of ξ_p. Moreover, $\tilde{\xi}_{p,n}$ satisfies the following CLT.

Theorem 2. *If the conditions in Theorem 1 hold, then*

$$\frac{\sqrt{n}}{\kappa_p}(\tilde{\xi}_{p,n} - \xi_p) \Rightarrow N(0,1) \tag{9}$$

as $n \to \infty$, where $\kappa_p = \psi_p \phi_p$ with $\phi_p = 1/f(\xi_p)$ and

$$\psi_p^2 = E_*\left[I(X \le \xi_p)L^2\right] - p^2. \tag{10}$$

We allow for $F_* \equiv F$, in which case IS reduces to crude Monte Carlo, so Theorems 1 and 2 generalize previous results for crude Monte Carlo. Thus, the variance constant κ_p^2 in (9) has the same basic form of $\kappa_p = \psi_p \phi_p$ for both IS and crude Monte Carlo. The value of ψ_p differs for the two methods (and equals $\sqrt{p(1-p)}$ for crude Monte Carlo), but ϕ_p does not change. As a consequence, when estimating a quantile, choosing a change of measure for IS focuses on trying to reduce ψ_p.

If we have consistent estimates of ψ_p and ϕ_p, then the CLT in Theorem 2 provides a way to construct a confidence interval for ξ_p when applying IS. It can be shown that ψ_p^2 in (10) is the variance constant in the CLT for $\tilde{F}_n(\xi_p)$, and a natural estimator for ψ_p^2 is

$$\tilde{\psi}_{p,n}^2 = \left(\frac{1}{n}\sum_{i=1}^{n} I(X_i \le \tilde{\xi}_{p,n})L_i^2\right) - p^2. \tag{11}$$

To derive an estimator for ϕ_p, we first note that $\frac{d}{dp}F^{-1}(p) = 1/f(\xi_p) = \phi_p$ by the chain rule of differentiation, and we will estimate ϕ_p using finite-difference estimators (e.g., Section 7.1 of Glasserman 2004). Let $c \neq 0$ be any constant, and define

$$\tilde{\phi}_{p,n,1}(c) = \frac{\tilde{F}_n^{-1}(p+cn^{-1/2}) - \tilde{F}_n^{-1}(p)}{cn^{-1/2}}, \tag{12}$$

$$\tilde{\phi}_{p,n,2}(c) = \frac{\tilde{F}_n^{-1}(p+cn^{-1/2}) - \tilde{F}_n^{-1}(p-cn^{-1/2})}{2cn^{-1/2}}. \tag{13}$$

Thus, for $c > 0$ (resp., $c < 0$), $\tilde{\phi}_{p,n,1}(c)$ is a forward (resp., backward) finite-difference estimator, and $\tilde{\phi}_{p,n,2}(c)$ is a central finite-difference estimator. In addition, we can define other estimators of ϕ_p through weighted combinations of the previous finite-difference estimators. Let $c_1,\ldots,c_r$ and $w_1,\ldots,w_r$ be any nonzero constants (some possibly negative) with $\sum_{j=1}^r w_j = 1$. Then we define combined estimators of ϕ_p as

$$\bar{\phi}_{p,n,i}(c_1,\ldots,c_r) = \sum_{j=1}^{r} w_j\, \tilde{\phi}_{p,n,i}(c_j), \text{ for } i = 1,2. \tag{14}$$

The following theorem shows that $\tilde{\psi}_{p,n}$ and all our estimators of ϕ_p are consistent. We can thus consistently estimate $\kappa_p = \psi_p\phi_p$ in (9) by taking the product of the consistent estimators of ψ_p and ϕ_p. In addition, the CLT in (9) still holds when κ_p is replaced by its consistent estimator.

Theorem 3. *Assume the conditions of Theorem 1 hold. Then for any nonzero constants c and $c_1,\ldots,c_r$,*

$$\tilde{\phi}_{p,n,i}(c) \quad \Rightarrow \quad \phi_p, \tag{15}$$

$$\bar{\phi}_{p,n,i}(c_1,\ldots,c_r) \quad \Rightarrow \quad \phi_p, \tag{16}$$

as $n \to \infty$, for $i = 1,2$. Moreover, $\tilde{\psi}_{p,n} \Rightarrow \psi_p$ as $n \to \infty$, and

$$\frac{\sqrt{n}}{\tilde{\kappa}_{p,n}}(\tilde{\xi}_{p,n} - \xi_p) \Rightarrow N(0,1) \tag{17}$$

as $n \to \infty$, with $\tilde{\kappa}_{p,n} = \tilde{\psi}_{p,n}\tilde{\phi}_{p,n,i}(c)$ or $\tilde{\kappa}_{p,n} = \tilde{\psi}_{p,n}\bar{\phi}_{p,n,i}(c_1,\ldots,c_r)$ for $i = 1,2$.

Hong (2009), Liu and Hong (2009) and Fu, Hong, and Hu (2009) develop consistent estimators for derivatives of quantiles with respect to certain model parameters, but their methods do not apply for estimating ϕ_p (nor when using IS). Bloch and Gastwirth (1968) and Bofinger (1975) provide estimators of ϕ_p analogous to $\tilde{\phi}_{p,n,i}(c)$, $i = 1,2$, in (12) and (13) for crude Monte Carlo. Babu (1986) considers combining estimators of ϕ_p as in (14) for crude Monte Carlo. However, the consistency proofs of the last three papers mentioned do not generalize to IS as they rely on expressing each sample X_i as $X_i = F^{-1}(U_i)$ with $U_i \sim \text{unif}[0,1]$. Thus, in Chu and Nakayama (2010) we establish (15) and (16) via a different approach based on the Bahadur-Ghosh representation in (7) and (8).

The right tail of $\tilde{F}_n$ in (6) may not behave as a proper CDF since it is possible that $\lim_{x\to\infty}\tilde{F}_n(x) = a$ with $a < 1$ or $a > 1$. To address this issue, Glynn (1996) also proposes another IS estimator of the CDF:

$$\tilde{F}_n'(x) = 1 - \frac{1}{n}\sum_{i=1}^{n} I(X_i > x)L(X_i), \tag{18}$$

which can be more effective when estimating the p-quantile when $p \approx 1$. The following two theorems, in which primed variables replace non-primed ones from before, show that quantile estimators based on inverting $\tilde{F}_n'$ satisfy a Bahadur-Ghosh representation and a CLT.

Theorem 4. *Suppose $f(\xi_p) > 0$, and suppose there exists $\varepsilon > 0$ and $\delta > 0$ such that $E_*[I(X > \xi_p - \delta)L^{2+\varepsilon}] < \infty$. Then $\tilde{\xi}_{p_n,n}' = \tilde{F}_n'^{-1}(p_n)$ with $p_n - p = O(n^{-1/2})$ satisfies the Bahadur-Ghosh representation in (7) and (8).*

Theorem 5. *Under the conditions of Theorem 4, $\tilde{\xi}_{p,n}' = \tilde{F}_n'^{-1}(p)$ satisfies the CLTs in (9) and (17), where*

$$\psi_p'^2 = E_*\left[I(X > \xi_p)L^2\right] - (1-p)^2, \tag{19}$$

$$\tilde{\psi}_{p,n}'^2 = \left(\frac{1}{n}\sum_{i=1}^{n} I(X_i > \tilde{\xi}_{p,n}')L_i^2\right) - (1-p)^2. \tag{20}$$

Glynn (1996) and Glasserman, Heidelberger, and Shahabuddin (2000) establish CLTs analogous to (9), but they do not consider the CLT in (17) with estimated variance. Their proofs are based on the Berry-Esséen theorem (p. 33 of Serfling 1980), thus requiring the likelihood ratio to have a finite third moment, which is stronger than our assumptions in Theorems 2 and 5. Sun and Hong (2010) also prove the CLT in (9) (but not (17)) under the stronger conditions mentioned earlier after Theorem 1.

We now explain how to construct a $100(1-\alpha)\%$ confidence interval for ξ_p. First generate i.i.d. pairs (X_i, L_i), $i = 1, \ldots, n$, using F_*. Then sort $X_1, X_2, \ldots, X_n$ in ascending order as $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$, where $X_{(i)}$ denotes the ith smallest value, and let $L^{(i)} = L(X_{(i)})$. The algorithm now depends on whether we use the IS CDF estimator $\tilde{F}_n$ in (6) or $\tilde{F}'_n$ in (18). Glynn (1996) provides the following algorithms to invert these IS CDF estimators. If we work with $\tilde{F}_n$, then $\tilde{F}_n^{-1}(q) = X_{(i_q)}$ for $0 < q < 1$, where i_q is the smallest integer for which $\sum_{i=1}^{i_q} L^{(i)} \geq qn$. The p-quantile estimator is $\tilde{\xi}_{p,n} = \tilde{F}_n^{-1}(p)$. We then compute $\tilde{\kappa}_{p,n}$ in (17) in this case by taking the product of $\bar{\phi}_{p,n,i}(c_1, \ldots, c_r)$ in (14) (for $i = 1$ or 2) and $\tilde{\psi}_{p,n}$ from (11). (Note that (12) and (13) are special cases of (14) with $r = 1$, so we only consider (14).) Alternatively, if we instead use $\tilde{F}'_n$ as the IS CDF estimator, then $\tilde{F}_n'^{-1}(q) = X_{(i'_q)}$, where i'_q is the greatest integer such that $\sum_{i=i'_q}^{n} L^{(i)} \geq n(1-q)$. In this case the p-quantile estimator is $\tilde{\xi}_{p,n} = \tilde{F}_n'^{-1}(p)$, and we compute $\tilde{\kappa}_{p,n}$ by taking the product of $\bar{\phi}_{p,n,i}(c_1, \ldots, c_r)$ in (14) and $\tilde{\psi}'_{p,n}$ from (20), where (14) (and (12) and (13)) now uses $\tilde{F}_n'^{-1}$. Finally, in either case an asymptotically valid $100(1-\alpha)\%$ confidence interval for ξ_p is $(\tilde{\xi}_{p,n} \pm z_{1-\alpha/2}\tilde{\kappa}_{p,n}/\sqrt{n})$, where $z_\beta = \Phi^{-1}(\beta)$ and Φ is the CDF of a $N(0,1)$ random variable.

4 EMPIRICAL STUDY

In the previous sections, we established the asymptotic validity of confidence intervals for ξ_p as the sample size $n \to \infty$. However, in practice, the parameter n has to be finite, so we ran experiments to study how the sample size n and the "smoothing parameter" c used in the finite-difference estimators in (12) and (13) affect the coverage of the resulting intervals.

We ran our experiments on two stochastic models: a normal distribution and a stochastic activity network, as specified in the subsequent sections. For both models we consider estimating the p-quantile for $p \approx 1$, so we work with the CDF estimator $\tilde{F}'_n$ in (18) to obtain a quantile estimator and confidence intervals.

4.1 Normal Distribution

The first set of experiments involves estimating the p-quantile ξ_p of a standard normal random variable X, so F is the standard normal CDF Φ. We might think of X as representing the loss in value of a portfolio over the next two weeks, so ξ_p is the $100p\%$ VaR. We obtain the IS distribution F_* by exponentially tilting F with tilting parameter θ, defined by $F_*(dx) = e^{\theta x - \zeta(\theta)}F(dx)$, where $\zeta(\theta) = \ln(E[e^{\theta X}]) = \theta^2/2$ is the cumulant generating function of Φ. It is straightforward to show that F_* is also normal with unit variance and mean $\zeta'(\theta) = \theta$, the derivative of $\zeta(\theta)$. The likelihood ratio is $L(x) = F(dx)/F_*(dx) = \exp(-\theta x + \frac{\theta^2}{2})$. To choose a value for θ, consider the following approximation applied by Glynn (1996): $P(X > x) \approx \exp(-x\theta_x + \zeta(\theta_x))$ for $x \gg 0$, where θ_x is the root of the equation $\zeta'(\theta_x) = x$, so $\theta_x = x$. Since we are interested in x satisfying $P(X > x) = 1 - p$ (i.e., the p-quantile), we arrive at the equation $-\theta^2 + \theta^2/2 = \ln(1-p)$. Solving for θ gives $\theta = \sqrt{-2\ln(1-p)} = \zeta'(\theta)$ as the mean of F_*.

4.2 Stochastic Activity Network

The second model we consider is a stochastic activity network (SAN); SANs are often employed to model the project completion time in project planning. In this experiment, we use a simple SAN previously studied by Hsu and Nelson (1990). The SAN consists of 5 activities, whose durations $A_1, \ldots, A_5$ are i.i.d. exponential random variables with mean 1. There are 3 paths in the SAN, and let $D_1 = \{1,2\}$, $D_2 = \{1,3,5\}$, and $D_3 = \{4,5\}$, where D_j is the set of activities on path j. The length of the jth path is denoted by $T_j = \sum_{i \in D_j} A_i$. Let m_j be the number of activities on the jth path, which is also the mean of T_j. Set $X = \max\{T_1, T_2, T_3\}$ as the length of the longest path, and we want to estimate the p-quantile of X. As noted by Hsu and Nelson (1990), the CDF of X is given by, for $x \geq 0$,

$$F(x) = 1 + (3 - 3x - x^2/2)e^{-x} + (-3 - 3x + x^2/2)e^{-2x} - e^{-3x},$$

whose density $f(x)$ is positive for all $x \geq 0$.

We now describe how we use IS to estimate ξ_p. Our change of measure for IS is a modification of an approach developed in Juneja, Karandikar, and Shahabuddin (2007). The basic idea is to use a mixture of three distributions, each defined by exponentially tilting one path length T_j and not changing the distributions of the durations of activities not on that path. Specifically, define f_i to be the density function of A_i, so $f_i(t) = e^{-t}$ for $t \geq 0$ for each $i = 1, \ldots, 5$. Define f_i^θ to be the exponentially tilted version of f_i under tilting parameter θ, so $f_i^\theta(t) = e^{\theta t - \chi_i(\theta)} f_i(t)$, where

$\chi_i(\theta) = \ln E[e^{\theta A_i}] = -\ln(1-\theta)$ is the cumulant generating function of A_i, which exists for $\theta < 1$. It is simple to show that f_i^{θ} is the density of an exponential with rate $1-\theta$. For each $j = 1,2,3$, define a probability measure P_j such that each A_i has density $f_i^{\theta_j}$ when $i \in D_j$ and density f_i when $i \notin D_j$, where θ_j denotes P_j's tilting parameter, which we specify later. The A_i, $i = 1,\ldots,5$, are mutually independent under each measure P_j. Now define the IS measure P_* to be the mixture of the P_j using positive weights α_j (specified later) satisfying $\sum_{j=1}^{3} \alpha_j = 1$; i.e., $P_*(B) = \sum_{j=1}^{3} \alpha_j P_j(B)$ for any event B. The likelihood ratio is then

$$L = \left[\sum_{j=1}^{3} \alpha_j \exp\left(\theta_j T_j - \zeta_j(\theta_j)\right)\right]^{-1}, \tag{21}$$

where $\zeta_j(\theta) = \sum_{i \in D_j} \chi_i(\theta) = -m_j \ln(1-\theta)$ is the cumulant generating function of T_j.

To compute the tilting parameter θ_j used with measure P_j, we apply an idea outlined by Glynn (1996). The approach is based on large-deviations theory, which suggests that under certain conditions,

$$P(T_j > x) \approx \exp(-x\theta_x + \zeta_j(\theta_x)) \tag{22}$$

for $x \gg E[T_j] = m_j$, where θ_x is the root of the equation $\zeta_j'(\theta_x) = x$ and prime denotes derivative, so $\zeta_j'(\theta) = m_j/(1-\theta)$. Since we are interested in the p-quantile, we equate the right side of (22) to $1-p$. This yields $-\zeta_j'(\theta)\theta + \zeta_j(\theta) = \ln(1-p)$, and we take θ_j to be its root. Also, we get $\zeta'(\theta_j) = m_j/(1-\theta_j)$ as a (crude) approximation for the p-quantile of T_j when p is close to 1.

We now modify a heuristic in Juneja, Karandikar, and Shahabuddin (2007) to obtain the mixture weights α_j used to define P_*. We want the variance in (19) to be small, and the idea is to select the α_j to minimize an upper bound for an approximation to the second moment $b \equiv E_*[I(X > \xi_p)L^2]$ from (19). Since $\zeta_j'(\theta_j)$ is roughly equal to the p-quantile of T_j and since $X = \max_j T_j$, we first approximate ξ_p via $\bar{\xi}_p \equiv \max_j \zeta_j'(\theta_j)$. This then leads to approximating b by $E_*[I(X > \bar{\xi}_p)L^2]$, which we now want to bound from above. Let $K_j = \exp(-\theta_j \bar{\xi}_p + \zeta_j(\theta_j))$, so (21) implies $L \le K_j/\alpha_j$ for $\theta_j > 0$ when $T_j > \bar{\xi}_p$. Hence, since $\{X > \bar{\xi}_p\} = \cup_{j=1}^{3}\{T_j > \bar{\xi}_p\}$, we get

$$E_*[I(X > \bar{\xi}_p)L^2] \le \left(\max_{j=1,2,3} \frac{K_j}{\alpha_j}\right)^2. \tag{23}$$

We then choose α_j to minimize our upper bound in (23), subject to $\sum_{j=1}^{3} \alpha_j = 1$, which results in $\alpha_j = K_j / \sum_{s=1}^{3} K_s$.

4.3 Choosing the Smoothing Parameter c

Recall the estimators of ϕ_p in (12) and (13), where specifying different values of c results in different estimators of ϕ_p and hence different confidence intervals. We now discuss some recommendations to select c for the central finite-difference estimator in (13). We consider two criteria for choosing c. One is to minimize the coverage error of the resulting confidence interval for ξ_p; the other minimizes the mean-square error (MSE) of the estimator of ϕ_p. In the case of crude Monte Carlo, Hall and Sheather (1988) and Bofinger (1975) have carried out asymptotic analyses regarding these two issues, which, when applied in our context, suggest selecting c as large as possible when n is large; see Chu and Nakayama (2010) for further details.

When combining different values of c as in (14), we can use a recommendation on p. 384 of Glasserman (2004) developed for reducing the bias of finite-difference estimators of a derivative of a mean. The suggestion is to combine $r = 2$ values c_1 and c_2 in (14) with $c_2 = 2c_1$, $w_1 = 4/3$ and $w_2 = -1/3$.

4.4 Discussion of Empirical Results

In all our experiments we let the sample size $n = 100 \times 4^j$ for $0 \le j \le 4$. Also, we varied the smoothing parameter between $c = 0.025$ and $c = 1$, using different values in different experiments. Two boundary conditions arise from (13) and (12) that govern the allowable values of (c,n) pairs. First we need $0 \le p + cn^{-1/2} \le 1$ to ensure $\tilde{F}_n'^{-1}$ is evaluated at a value corresponding to a probability. Also, we require $|n \times c/\sqrt{n}| \ge 1$ so that $\tilde{F}_n'^{-1}(p + cn^{-1/2})$ and $\tilde{F}_n'^{-1}(p)$ in (12) return different order statistics. In all cases we constructed confidence intervals having nominal level $1-\alpha = 0.9$, and we estimated coverages and average half-widths using $m = 10^4$ independent replications.

Figures 1–3 plot coverage as a function of n for different values of c when using the central finite-difference estimator of ϕ_p from (13). Figure 1 presents the results when applying crude Monte Carlo to the normal distribution (left) and the SAN (right), both for $p = 0.95$. Figure 2 shows the results when applying IS to the normal distribution for $p = 0.95$ (left) and $p = 0.99$ (right). Figure 3 is for applying IS to the SAN model for $p = 0.95$ (left) and $p = 0.99$ (right). The figures show the coverage levels are converging to the nominal level as n grows for each fixed c, demonstrating the asymptotic validity of our confidence intervals. When n is large, larger values of c seem to result in better coverage. This agrees with the recommendations in Section 4.3 to choose large c when applying crude Monte Carlo, so the same suggestion also may be appropriate when applying IS. When n is small, larger c also leads to larger coverage but not necessarily closer to the nominal level of 0.9. Thus, the recommendations for choosing large c require large sample sizes n.

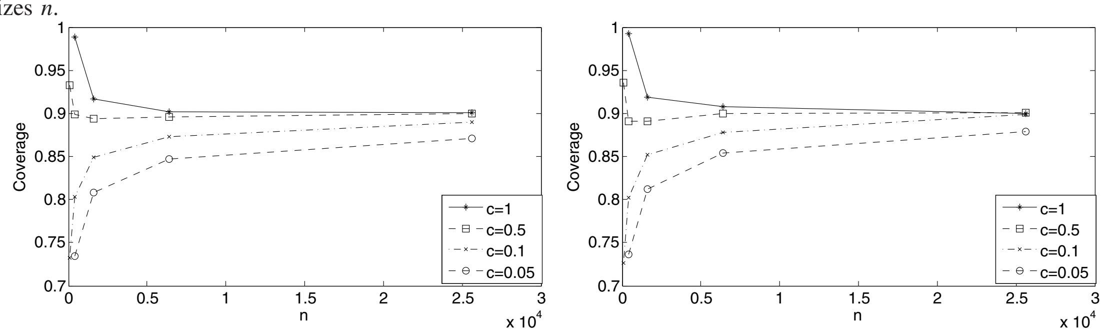

Figure 1: Coverage for CMC for the normal (left) and for the SAN (right) for $p = 0.95$.

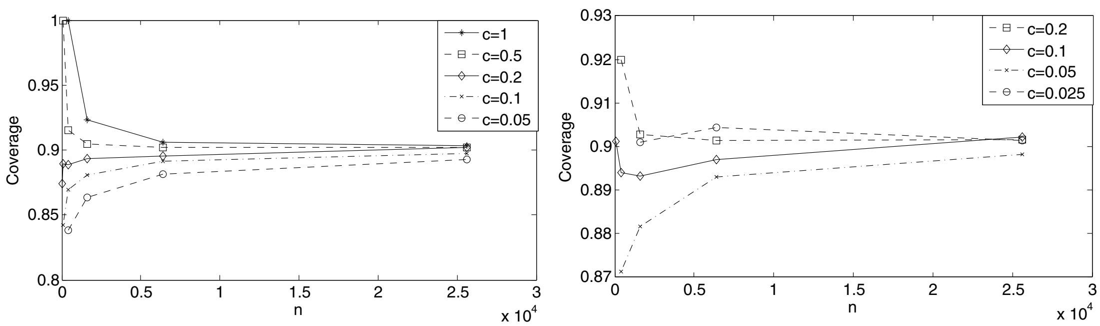

Figure 2: Coverage for IS for the normal when $p = 0.95$ (left) and $p = 0.99$ (right).

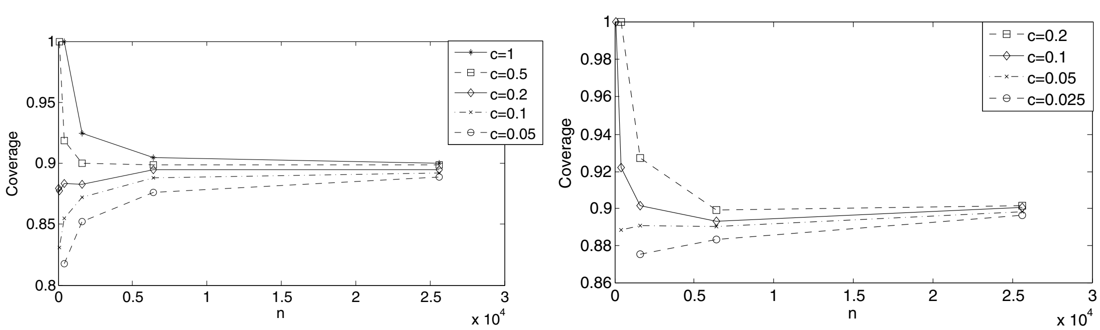

Figure 3: Coverage for IS for the SAN when $p = 0.95$ (left) and $p = 0.99$ (right).

Tables 1–6 provide additional results from our experiments for $c = 0.1$ and $c = 0.2$. Tables 1 and 4 are for the normal model and the SAN, respectively, when applying crude Monte Carlo. Tables 2 and 3 are for the normal model with IS for $p = 0.95$ and 0.99, respectively, and Tables 5 and 6 contain results for the SAN when applying IS for $p = 0.95$ and 0.99, respectively. In each table the first column gives the sample size n. The next three columns give the results for the central, forward and backward finite-difference estimators of ϕ_p from (13) and (12), abbreviated as CFD, FFD and BFD, respectively, for $c = 0.1$. Columns 5–7 present the same for $c = 0.2$. Each sample size n corresponds to two rows: the first gives the coverage and the second provides the average half-width of the confidence intervals. Overall, BFD seems to produce the worst coverage levels of the three estimators. Also, CFD slightly outperforms FFD for most cases, which complements analysis in Section 7.1 of (Glasserman 2004) showing that central finite-difference estimators of the derivative of a mean have asymptotically smaller MSE than forward estimators. In addition, the confidence intervals for IS have smaller half-widths than those for crude Monte Carlo, indicating IS reduces variance.

Table 1: Coverages (and average half-width) for CMC with $c = 0.1$ and $c = 0.2$ for the Normal when $p = 0.95$

	$c = 0.1$			$c = 0.2$					
n	CFD	FFD	BFD	CFD	FFD	BFD	ϕ_p	Batching	Mean
100	0.733	0.659	0.623	0.801	0.778	0.690	0.903	0.629	0.894
	(0.324)	(0.345)	(0.303)	(0.331)	(0.378)	(0.285)	(0.348)	(0.331)	(0.164)
400	0.804	0.749	0.725	0.848	0.828	0.786	0.900	0.642	0.898
	(0.169)	(0.174)	(0.164)	(0.171)	(0.184)	(0.159)	(0.174)	(0.171)	(0.082)
1600	0.849	0.822	0.807	0.894	0.865	0.843	0.899	0.833	0.899
	(0.087)	(0.088)	(0.086)	(0.087)	(0.091)	(0.084)	(0.087)	(0.092)	(0.041)
6400	0.873	0.854	0.847	0.895	0.878	0.866	0.900	0.882	0.903
	(0.043)	(0.044)	(0.043)	(0.043)	(0.044)	(0.043)	(0.044)	(0.047)	(0.021)

Table 2: Coverages (and average half-width) for IS with $c = 0.1$ and $c = 0.2$ for the Normal when $p = 0.95$

	$c = 0.1$			$c = 0.2$					
n	CFD	FFD	BFD	CFD	FFD	BFD	ϕ_p	Batching	Mean
100	0.842	0.820	0.742	0.890	0.911	0.774	0.881	0.887	0.253
	(0.119)	(0.120)	(0.129)	(0.125)	(0.145)	(0.104)	(0.119)	(0.185)	(0.661)
400	0.870	0.855	0.819	0.889	0.899	0.839	0.894	0.889	0.393
	(0.061)	(0.062)	(0.064)	(0.061)	(0.067)	(0.057)	(0.061)	(0.074)	(0.636)
1600	0.881	0.874	0.854	0.894	0.898	0.866	0.901	0.896	0.525
	(0.031)	(0.031)	(0.032)	(0.031)	(0.031)	(0.030)	(0.031)	(0.034)	(0.477)
6400	0.892	0.887	0.876	0.895	0.899	0.883	0.902	0.902	0.614
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.015)	(0.016)	(0.017)	(0.337)

Table 3: Coverages (and average half-width) for IS with $c = 0.1$ and $c = 0.2$ for the Normal when $p = 0.99$

	$c = 0.1$			$c = 0.2$					
n	CFD	FFD	BFD	CFD	FFD	BFD	ϕ_p	Batching	Mean
100	1.000	1.000	0.727	0.999	1.000	0.644	0.874	0.867	0.085
	(0.487)	(0.898)	(0.076)	(0.256)	(0.449)	(0.062)	(0.104)	(0.174)	(0.833)
400	0.920	0.961	0.813	1.000	1.000	0.763	0.895	0.880	0.184
	(0.059)	(0.072)	(0.045)	(0.285)	(0.530)	(0.039)	(0.054)	(0.065)	(0.762)
1600	0.903	0.930	0.857	0.925	0.969	0.828	0.901	0.898	0.302
	(0.028)	(0.031)	(0.025)	(0.030)	(0.036)	(0.023)	(0.027)	(0.030)	(0.669)
6400	0.901	0.917	0.881	0.909	0.938	0.865	0.904	0.907	0.415
	(0.014)	(0.014)	(0.013)	(0.014)	(0.015)	(0.012)	(0.014)	(0.015)	(0.526)

We also constructed confidence intervals using the exact value of ϕ_p rather than estimating it to evaluate the effects of having to estimate ϕ_p on coverage levels. The columns in the tables labeled "ϕ_p" contain these results. For smaller sample sizes, using the exact ϕ_p generally seems to produce better coverages than when estimating it, demonstrating the degradation in the quality of the intervals from estimating ϕ_p.

We also applied batching as an alternative approach to construct confidence intervals. This method divides all the data into $b \geq 2$ batches of size n/b and computes a quantile estimate from each batch. It then produces an overall point estimate and confidence interval from the sample average and sample variance of the b i.i.d. quantile estimates from the batches using a critical point from a t-distribution with $b-1$ degrees of freedom; e.g., see p. 491 of Glasserman (2004). The columns labeled "Batching" show coverage levels for $b = 10$ batches. For large n, batching produces slightly wider intervals on average. In terms of coverage, it is not clear which of batching and finite-difference estimation of ϕ_p is better when applying IS. However, for crude Monte Carlo, coverage with batching is clearly inferior when n is small. This is because accurate quantile estimation with crude Monte Carlo typically requires large sample sizes (Avramidis and Wilson 1998), and batching effectively reduces the sample size by a factor of b. Thus, the quantile estimate from each batch is usually inaccurate with crude Monte Carlo, leading to poor coverage.

Table 4: Coverages (and average half-width) for CMC with $c = 0.1$ and $c = 0.2$ for the SAN when $p = 0.95$

	$c = 0.1$			$c = 0.2$					
n	CFD	FFD	BFD	CFD	FFD	BFD	ϕ_p	Batching	Mean
100	0.726	0.662	0.608	0.802	0.791	0.681	0.910	0.561	0.894
	(0.881)	(0.966)	(0.796)	(0.905)	(1.069)	(0.742)	(0.952)	(0.912)	(0.278)
400	0.802	0.751	0.722	0.844	0.831	0.722	0.898	0.660	0.896
	(0.464)	(0.482)	(0.447)	(0.468)	(0.510)	(0.427)	(0.476)	(0.455)	(0.140)
1600	0.852	0.815	0.806	0.874	0.863	0.806	0.897	0.837	0.900
	(0.236)	(0.241)	(0.231)	(0.237)	(0.248)	(0.227)	(0.238)	(0.249)	(0.070)
6400	0.878	0.857	0.854	0.891	0.886	0.854	0.906	0.888	0.901
	(0.118)	(0.120)	(0.117)	(0.119)	(0.122)	(0.116)	(0.119)	(0.128)	(0.035)

Table 5: Coverages (and average half-width) for IS with $c = 0.1$ and $c = 0.2$ for the SAN when $p = 0.95$

	$c = 0.1$			$c = 0.2$					
n	CFD	FFD	BFD	CFD	FFD	BFD	ϕ_p	Batching	Mean
100	1.000	1.000	0.726	0.890	0.911	0.774	0.871	0.846	0.798
	(0.403)	(0.440)	(0.366)	(0.125)	(0.145)	(0.104)	(0.400)	(0.532)	(0.710)
400	0.922	0.964	0.794	0.889	0.899	0.839	0.891	0.879	0.835
	(0.207)	(0.217)	(0.196)	(0.061)	(0.067)	(0.057)	(0.207)	(0.234)	(0.356)
1600	0.900	0.931	0.845	0.894	0.898	0.866	0.896	0.894	0.864
	(0.104)	(0.106)	(0.102)	(0.031)	(0.031)	(0.030)	(0.104)	(0.114)	(0.178)
6400	0.899	0.815	0.871	0.895	0.899	0.883	0.898	0.902	0.883
	(0.052)	(0.053)	(0.051)	(0.016)	(0.016)	(0.015)	(0.052)	(0.057)	(0.089)

We also constructed confidence intervals for the mean $\mu = E[X] = E_*[XL]$ as another benchmark for comparison. When applying IS, we used the same change of measure as when estimating the quantiles. For IS on the SAN, coverage levels for the mean are close to the nominal level. However, coverages are poor for the normal with IS, and we now explore this issue. Estimating μ with IS entails averaging i.i.d. copies of XL, where $X \sim F_*$ and $L = L(X)$ is the likelihood ratio. The mean of F_* is $\theta = \sqrt{-2\ln(1-p)}$, so $\theta = 2.4477$ when $p = 0.95$ and $\theta = 3.0349$ when $p = 0.99$, both of which are quite far from the original mean $\mu = 0$. (The variance of F_* remains at 1.) For large sample sizes n, Edgeworth expansions show that the coverage level when estimating a mean is largely affected by skewness for one-sided confidence intervals and by skewness and kurtosis for two-sided intervals; see pp. 50 and 72–73 of Hall (1992). We calculate the skewness of XL under F_* to be $E_*[(XL)^3]/(E_*[(XL)^2])^{3/2} = \exp(3\theta^2/2)(-8\theta^3 - 6\theta)/(1+\theta^2)^{3/2}$, which works out to -5.7×10^4 when $p = 0.95$ and -7.4×10^6 when $p = 0.99$, indicating the distributions are quite asymmetric with much heavier left tails. The kurtosis of XL is $E_*[(XL)^4]/(E_*[(XL)^2])^2 = \exp(4\theta^2)(3 + 54\theta^2 + 81\theta^4)/(1+\theta^2)^2$, which is 1.7×10^{12} when $p = 0.95$ and 7.1×10^{17} when $p = 0.99$. The huge values for skewness and kurtosis seem to explain the poor coverage when estimating the mean of the normal using IS. Further experiments with even larger sample sizes (not shown) indicate that coverages are converging to the nominal level at a very slow rate.

Table 6: Coverages (and average half-width) for IS with $c = 0.1$ and $c = 0.2$ for the SAN when $p = 0.99$

	$c = 0.1$			$c = 0.2$					
n	CFD	FFD	BFD	CFD	FFD	BFD	ϕ_p	Batching	Mean
100	1.000	1.000	0.704	1.000	1.000	0.614	0.873	0.811	0.889
	(5.005)	(9.703)	(0.307)	(2.549)	(4.851)	(0.246)	(0.445)	(0.659)	(0.965)
400	0.922	0.962	0.800	1.000	1.000	0.741	0.897	0.876	0.896
	(0.253)	(0.317)	(0.189)	(3.246)	(6.330)	(0.162)	(0.232)	(0.270)	(0.486)
1600	0.902	0.933	0.847	0.927	0.974	0.815	0.899	0.892	0.900
	(0.119)	(0.134)	(0.105)	(0.128)	(0.161)	(0.096)	(0.117)	(0.129)	(0.243)
6400	0.893	0.910	0.899	0.933	0.850	0.866	0.893	0.894	0.899
	(0.059)	(0.063)	(0.060)	(0.067)	(0.053)	(0.055)	(0.059)	(0.064)	(0.122)

Finally, Table 7 presents results from applying the combined estimator of ϕ_p from (14). We combined $r = 2$ values of c, using the strategy described at the end of Section 4.3. However, in terms of coverage, combining demonstrates no clear improvement, and perhaps even a slight degradation for small n.

Table 7: Coverages (and average half-width) with CFD ($c = 0.1$) and combined ($r = 2$, $c_1 = 0.1$, $c_2 = 0.2$, $w_1 = 4/3$ and $w_2 = -1/3$) for the SAN when $p = 0.95$

n	CFD	Combined
100	0.831	0.798
	(0.403)	(0.397)
400	0.855	0.835
	(0.207)	(0.206)
1600	0.872	0.864
	(0.104)	(0.104)
6400	0.888	0.883
	(0.052)	(0.052)

5 CONCLUDING REMARKS

In this paper we developed asymptotically valid confidence intervals for quantiles when applying IS. To do this we provided a consistent estimator of the asymptotic variance κ_p^2 appearing in the CLT that the quantile estimator satisfies. It turns out that $\kappa_p = \psi_p \phi_p$, and we gave consistent estimators of ψ_p and $\phi_p = 1/f(\xi_p)$. Our estimators of ϕ_p are finite-difference estimators, and their consistency can be shown by exploiting a Bahadur-Ghosh representation for the IS quantile estimator, which we also establish.

In Chu and Nakayama (2010) we extend these results to a general framework for variance-reduction techniques (VRTs), allowing the construction of asymptotically valid confidence intervals when applying any VRT within our framework. Our framework specifies conditions on the VRT estimator of the CDF, and we show the framework encompasses antithetic variates, control variates, and a combination of IS and stratification.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation under grant number CMMI-0926949. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

Avramidis, A. N., and J. R. Wilson. 1998. Correlation-induction techniques for estimating quantiles in simulation. *Operations Research* 46:574–591.

Babu, G. J. 1986. Efficient estimation of the reciprocal of the density quantile function at a point. *Statistics and Probability Letters* 4:133–139.

Bahadur, R. R. 1966. A note on quantiles in large samples. *Annals of Mathematical Statistics* 37:577–580.

Billingsley, P. 1995. *Probability and measure*. Third ed. New York: John Wiley & Sons.

Bloch, D. A., and J. L. Gastwirth. 1968. On a simple estimate of the reciprocal of the density function. *Annals of Mathematical Statistics* 39:1083–1085.

Bofinger, E. 1975. Estimation of a density function using order statistics. *Australian Journal of Statistics* 17:1–7.

Chu, F., and M. K. Nakayama. 2010. Confidence intervals for quantiles when applying variance-reduction techniques. *submitted*.

David, H. A., and H. Nagaraja. 2003. *Order statistics*. Third ed. Hoboken, NJ: Wiley.

Fu, M. C., L. J. Hong, and J.-Q. Hu. 2009. Conditional Monte Carlo estimation of quantile sensitivities. *Management Science* 55:2019–2027.

Ghosh, J. K. 1971. A new proof of the Bahadur representation of quantiles and an application. *Annals of Mathematical Statistics* 42:1957–1961.

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

Glasserman, P., P. Heidelberger, and P. Shahabuddin. 2000. Variance reduction techniques for estimating value-at-risk. *Management Science* 46:1349–1364.

Glynn, P. W. 1996. Importance sampling for Monte Carlo estimation of quantiles. In *Mathematical Methods in Stochastic Simulation and Experimental Design: Proceedings of the 2nd St. Petersburg Workshop on Simulation*, 180–185: Publishing House of St. Petersburg University, St. Petersburg, Russia.

Glynn, P. W., and D. L. Iglehart. 1989. Importance sampling for stochastic systems. *Management Science* 35:1367–1393.

Hall, P. 1992. *The bootstrap and Edgeworth expansions*. New York: Springer.

Hall, P., and S. J. Sheather. 1988. On the distribution of a Studentized quantile. *Journal of the Royal Statistical Society B* 50:381–391.

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

Hong, L. J. 2009. Estimating quantile sensitivities. *Operations Research* 57:118–130.

Hsu, J. C., and B. L. Nelson. 1990. Control variates for quantile estimation. *Management Science* 36:835–851.

Juneja, S., R. Karandikar, and P. Shahabuddin. 2007. Asymptotics and fast simulation for tail probabilities of maximum of sums of few random variables. *ACM Transactions on Modeling and Computer Simulation* 17:article 2, 35 pages.

Liu, G., and L. J. Hong. 2009. Kernel estimation of quantile sensitivities. *Naval Research Logistics* 56:511–525.

Serfling, R. J. 1980. *Approximation theorems of mathematical statistics*. New York: John Wiley & Sons.

Sun, L., and L. J. Hong. 2010. Asymptotic representations for importance-sampling estimators of value-at-risk and conditional value-at-risk. *Operations Research Letters*:to appear.

AUTHOR BIOGRAPHIES

FANG CHU received a Ph.D. in Information Systems from New Jersey Institute of Technology. His research interests are in the area of simulation and modeling, particularly in financial applications.

MARVIN K. NAKAYAMA is a professor in the Department of Computer Science at the New Jersey Institute of Technology. He received his Ph.D. in operations research from Stanford University. He won second prize in the 1992 George E. Nicholson Student Paper Competition sponsored by INFORMS and is a recipient of a CAREER Award from the National Science Foundation. He is the Simulation Analysis and Stochastic Modeling Area Editor for *ACM Transactions on Modeling and Computer Simulation* and the Simulation Area Editor for *INFORMS Journal on Computing*.