Proceedings of the 1992 Winter Simulation Conference
ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson

EFFICIENT METHODS FOR GENERATING SOME EXPONENTIALLY TILTED RANDOM VARIATES

Marvin K. Nakayama

IBM Research Division
T. J. Watson Research Center
Yorktown Heights, NY 10598, U.S.A.

ABSTRACT

Exponentially tilted distributions often arise as importance sampling distributions which are derived using large deviations theory. In this paper we present simple and efficient methods for generating some exponentially tilted random variates when the input distribution is either a Weibull or a positive normal. In particular, our methods are acceptance-rejection algorithms, and we prove that the expected number of iterations tends to 1 as the tilting parameter increases to infinity. We also provide empirical results from using our proposed techniques.

1 INTRODUCTION

Exponential tilting (also known as exponential twisting or shifting) transforms a given distribution into a new one. Tilted distributions are often used in importance sampling schemes derived with the aid of large deviations theory; e.g., see Siegmund (1976), Cottrell, Fort, and Malgouvres (1983), Dupuis and Kushner (1987), Asmussen (1985), Parekh and Walrand (1989), and Sadowski and Bucklew (1990). This previous research showed that when using importance sampling, an exponentially tilted distribution minimizes the variance of the resulting estimator over some class of possible changes of measure.

While exponential changes of measure have been studied from a theoretical standpoint, there has not been substantial research that investigates how to apply them in practice. In particular, techniques for generating exponentially tilted random variates need to be developed.

For certain distributions, exponential tilting only alters the parameter values of the input distribution. This holds for the exponential family, which includes the exponential, gamma, and normal distributions. However, when exponential tilting is applied to other distributions, the resulting distribution is not the same as the initial one (and often not one of the "standard" distributions). For example, this situation occurs with the Weibull and positive (or truncated) normal distributions. In this paper we describe acceptance-rejection algorithms for generating exponentially tilted random variates based on these two input distributions.

The proposed acceptance-rejection schemes have several desirable properties. First, the only inputs to our two algorithms are the parameters of the original distributions and the tilting parameter. This is advantageous as the exponentially tilted distributions themselves depend on the input distributions' moment generating functions and there are no closed form expressions for these quantities. Also, the majorizing densities in both cases are from "standard" distributions for which many fast and simple variate generation techniques are available. Furthermore, our algorithms become more efficient as we increase the tilting. In particular, we prove that the expected number of iterations in our procedures tends to 1 as the amount of tilting increases to infinity.

The rest of the paper is organized as follows. In Section 2, we give a brief review of exponentially tilted distributions. We discuss our acceptance-rejection algorithm for generating exponentially tilted Weibull random variates in Section 3, and Section 4 contains the same for the exponentially tilted positive normal distribution. We present some empirical results in Section 5, and conclude with Section 6.

2 EXPONENTIALLY TILTED DISTRIBUTIONS

Let random variable $X \geq 0$ have distribution function F and mean μ. Define

$$\phi(\theta) = E\left[e^{-\theta X}\right] = \int_0^\infty e^{-\theta x} F(dx),$$

which is the moment generating function of X. Let $\Lambda = \{\theta : \phi(\theta) < \infty\}$. For $\theta \in \Lambda$, the exponentially tilted distribution corresponding to F is given by

$$T_\theta(dx) = \frac{e^{-\theta x}F(dx)}{\phi(\theta)}$$

for $x > 0$, and 0 otherwise. Thus, exponential tilting shifts increasingly more mass of the distribution towards 0 as θ gets larger.

The mean of the distribution T_θ is

$$\nu(\theta) = \frac{\int_0^\infty xe^{-\theta x}F(dx)}{\phi(\theta)}.$$

Using the notation $\phi'(\theta) = \frac{d}{d\theta}\phi(\theta)$, if $\theta \in \Lambda$, then $\nu(\theta) = -\phi'(\theta)/\phi(\theta)$, assuming that we can interchange the order of the derivative and integral operators.

3 WEIBULL DISTRIBUTION

The Weibull distribution is often used for modeling the time to complete a task or the time to failure of a piece of equipment; see Law and Kelton (1991), pp. 333–335. Depending on the choice of the shape parameter, the distribution has either an increasing or decreasing failure rate.

Suppose that the random variable X has a Weibull distribution with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$. Its distribution function is given by

$$F(x) = 1 - \exp\{-(\beta x)^\alpha\}$$

for $x > 0$ and 0 otherwise, and its density is

$$f(x) = \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha\}$$

for $x > 0$ and 0 otherwise. The mean of a Weibull is $\Gamma(1/\alpha)/(\alpha\beta)$, where $\Gamma(\cdot)$ is the gamma function, and

$$\phi(\theta) = \int_0^\infty \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\}dx$$

is its moment generating function. The density of the exponentially tilted Weibull distribution is

$$t_\theta(x) = \frac{\alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\}}{\phi(\theta)}$$

for $x > 0$ and 0 otherwise.

Assume $\theta > 0$. Since $\exp\{-(\beta x)^\alpha\} \leq 1$ for all $x \geq 0$,

$$t_\theta(x) \;\leq\; g_\theta(x) \;\equiv\; \frac{\alpha\beta^\alpha x^{\alpha-1}e^{-\theta x}}{\phi(\theta)}.$$

Hence, we will use g_θ as a majorizing function in an acceptance-rejection algorithm. To this end, define

$$c_\theta \;\equiv\; \int_0^\infty g_\theta(x)dx \;=\; \frac{\alpha\beta^\alpha\Gamma(\alpha)}{\theta^\alpha\phi(\theta)}.$$

Note that $c_\theta \geq 1$ for all $\theta > 0$. Thus, the majorizing density is given by

$$h_\theta(x) = \frac{g_\theta(x)}{c_\theta} = \frac{\theta^\alpha x^{\alpha-1}e^{-\theta x}}{\Gamma(\alpha)}$$

for $x > 0$ and 0 otherwise, which is a gamma density with shape parameter α and scale parameter θ. Our acceptance-rejection algorithm to generate exponentially tilted Weibull random variates is as below.

Algorithm for Generating Exponentially Tilted Weibull Variates

1. Generate $Y \sim \text{gamma}(\alpha, \theta)$.
2. Generate $U \sim \text{uniform}(0, 1)$, independent of Y.
3. If $U \leq \exp\{-(\beta Y)^\alpha\}$, then return $X = Y$. Otherwise, reject (Y, U) and return to step 1.

There are many fast and simple methods for generating the gamma random variate needed in step 1 of our algorithm; e.g., see Devroye (1986). In addition, note that in the algorithm, we do not need to know the value of $\phi(\theta)$. This is desirable since there is no closed form expression for $\phi(\theta)$, and we would otherwise have to evaluate it numerically.

The following result shows that the expected number of iterations in our algorithm converges to 1 as the tilting parameter increases to infinity.

Theorem 1 *For an exponentially tilted Weibull distribution, $c_\theta \to 1$ as $\theta \to \infty$.*

Proof. We need to show that $\theta^\alpha\phi(\theta) \to \alpha\beta^\alpha\Gamma(\alpha)$ as $\theta \to \infty$. To this end, note that

$$\theta^\alpha\phi(\theta) = A(\theta) + B(\theta),$$

where

$$A(\theta) = \theta^\alpha \int_0^{(2\alpha\log\theta)/\theta} \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\}dx$$

and

$$B(\theta) = \theta^\alpha \int_{(2\alpha\log\theta)/\theta}^\infty \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\}dx.$$

First observe that

$$\begin{aligned} B(\theta) &\leq \theta^\alpha \exp\left\{-\theta\frac{2\alpha\log\theta}{\theta}\right\} \\ &\quad \cdot \int_{(2\alpha\log\theta)/\theta}^\infty \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha\}dx \\ &= \theta^{-\alpha}\exp\left\{-\left(\frac{2\alpha\beta\log\theta}{\theta}\right)^\alpha\right\} \\ &\to 0 \end{aligned}$$

as $\theta \to \infty$.

Thus, we must show that $A(\theta) \to \alpha\beta^\alpha\Gamma(\alpha)$ as $\theta \to \infty$. Note that $0 \le e^{-\theta x} \le 1$ for all $x \ge 0$ implies $\alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\} \le f(x)$ for all $x \ge 0$, where f is the density function of the Weibull. Since $\int_0^\infty f(x)dx = 1$ and $\alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\} \to 0$ as $\theta \to \infty$ for all $x \ge 0$, we have $\phi(\theta) \to 0$ as $\theta \to \infty$ by the dominated convergence theorem. Hence,

$$\int_0^{(2\alpha\log\theta)/\theta} \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\}dx \to 0$$

as $\theta \to \infty$, so

$$\lim_{\theta\to\infty} A(\theta) = \lim_{\theta\to\infty} \frac{G(\theta)}{-\alpha\theta^{-\alpha-1}}$$

by L'Hopital's rule, where

$$G(\theta) = \frac{d}{d\theta}\int_0^{(2\alpha\log\theta)/\theta} \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\}dx.$$

We now want to evaluate the limit of $G(\theta)$. To do this, suppose we have a function $h(x_1, x_2)$ that has anti-derivative $H(x_1, x_2)$ with respect to x_2, and $H(x_1, x_2)$ is differentiable in x_1. Then, letting $k(\theta)$ be some differentiable function of θ, we obtain

$$\begin{aligned}
&\frac{d}{d\theta}\int_0^{k(\theta)} h(\theta, x)dx \\
&= \frac{d}{d\theta}[H(\theta, k(\theta)) - H(\theta, 0)] \\
&= H_1(\theta, k(\theta)) + h(\theta, k(\theta))k'(\theta) - H_1(\theta, 0) \\
&= \int_0^{k(\theta)} \frac{d}{d\theta}h(\theta, x)dx + h(\theta, k(\theta))k'(\theta),
\end{aligned}$$

where we use the notation $H_1(x_1, x_2) = \frac{\partial}{\partial x_1}H(x_1, x_2)$ and $k'(\theta) = \frac{d}{d\theta}k(\theta)$. Thus, in our setting, we have $h(\theta, x) = \alpha\beta^\alpha x^{\alpha-1}\exp\{-(\beta x)^\alpha - \theta x\}$ and $k(\theta) = (2\alpha\log\theta)/\theta$, so

$$\lim_{\theta\to\infty} A(\theta) \equiv \lim_{\theta\to\infty}[C(\theta) + D(\theta)],$$

where

$$C(\theta) = \frac{\theta^{\alpha+1}}{\alpha}\int_0^{(2\alpha\log\theta)/\theta} \alpha\beta^\alpha x^\alpha \exp\{-(\beta x)^\alpha - \theta x\}dx$$

and

$$\begin{aligned}
D(\theta) &= -\frac{\theta^{\alpha+1}}{\alpha}\alpha\beta^\alpha\left(\frac{2\alpha\log\theta}{\theta}\right)^{\alpha-1}\frac{2\alpha(1-\log\theta)}{\theta^2} \\
&\quad \cdot \exp\left\{-\left(\frac{2\alpha\beta\log\theta}{\theta}\right)^\alpha - \theta\frac{2\alpha\log\theta}{\theta}\right\}.
\end{aligned}$$

Note that $D(\theta) \to 0$ as $\theta \to \infty$, and

$$\begin{aligned}
C(\theta) &\ge \beta^\alpha\theta^{\alpha+1}\exp\left\{-\left(\frac{2\alpha\beta\log\theta}{\theta}\right)^\alpha\right\} \\
&\quad \cdot \int_0^{(2\alpha\log\theta)/\theta} x^\alpha e^{-\theta x}dx \\
&= \beta^\alpha\exp\left\{-\left(\frac{2\alpha\beta\log\theta}{\theta}\right)^\alpha\right\} \\
&\quad \cdot \gamma(\alpha+1, 2\alpha\log\theta),
\end{aligned}$$

where $\gamma(\cdot,\cdot)$ is the incomplete gamma function (see Gradshteyn and Ryzhik 1981, p. 940) and is given by $\gamma(\lambda, z) = \int_0^z e^{-t}t^{\lambda-1}$. For large values of $|z|$,

$$\begin{aligned}
\gamma(\lambda, z) &= \Gamma(\lambda) - z^{\lambda-1}e^{-z} \\
&\quad \cdot \left[\sum_{m=0}^{M-1}\frac{(-1)^m\Gamma(1-\lambda+m)}{z^m\Gamma(1-\lambda)} + O\left(|z|^{-M}\right)\right]
\end{aligned}$$

for $M = 1, 2, \ldots$ (see Gradshteyn and Ryzhik 1981, p. 942). Hence, for large θ,

$$\begin{aligned}
&C(\theta) \\
&\ge \beta^\alpha\exp\left\{-\left(\frac{2\alpha\beta\log\theta}{\theta}\right)^\alpha\right\} \\
&\quad \cdot \left[\Gamma(\alpha+1) - \frac{(2\alpha\log\theta)^\alpha}{\theta^{2\alpha}}\left(1 + O\left(\frac{1}{2\alpha\log\theta}\right)\right)\right] \\
&\to \beta^\alpha\Gamma(\alpha+1)
\end{aligned}$$

as $\theta \to \infty$. Thus, $\lim_{\theta\to\infty}\theta^\alpha\phi(\theta) \ge \beta^\alpha\Gamma(\alpha+1)$. Finally, $c_\theta = \alpha\beta^\alpha\Gamma(\alpha)/(\theta^\alpha\phi(\theta)) \ge 1$ and $\alpha\Gamma(\alpha) = \Gamma(\alpha+1)$ imply that $\theta^\alpha\phi(\theta) \to \alpha\beta^\alpha\Gamma(\alpha)$ as $\theta \to \infty$, proving our result. ■

4 POSITIVE NORMAL DISTRIBUTION

Suppose that Z is a normally distributed random variable with mean 0 and variance σ^2, and define $X = |Z|$. Then X has a positive (or truncated) normal distribution with scale parameter σ. The density of X is given by

$$f(x) = \frac{2}{\sigma\sqrt{2\pi}}\exp\left\{-\frac{x^2}{2\sigma^2}\right\}$$

for $x \ge 0$ and 0 otherwise. Its moment generating function is

$$\begin{aligned}
\phi(\theta) &= \int_0^\infty \frac{2}{\sigma\sqrt{2\pi}}\exp\left\{-\frac{x^2}{2\sigma^2} - \theta x\right\}dx \\
&= \exp\left\{\frac{\sigma^2\theta^2}{2}\right\}\int_{\sigma^2\theta}^\infty \frac{2}{\sigma\sqrt{2\pi}}\exp\left\{-\frac{y^2}{2\sigma^2}\right\}dy. \quad (1)
\end{aligned}$$

Differentiating $\phi(\theta)$ and evaluating at $\theta = 0$ gives us that the mean of a positive normal is $\mu = 2\sigma/\sqrt{2\pi}$.

The density of the exponentially tilted positive normal distribution is

$$t_\theta(x) = \frac{2}{\phi(\theta)\sigma\sqrt{2\pi}} \exp\left\{-\frac{x_2}{2\sigma^2} - \theta x\right\}$$

for $x \geq 0$ and 0 otherwise.

Assume $\theta \geq 0$. Since $\exp\left\{-\frac{x^2}{2\sigma^2}\right\} \leq 1$ for all $x \geq 0$,

$$t_\theta(x) \leq g_\theta(x) \equiv \frac{2}{\phi(\theta)\sigma\sqrt{2\pi}} e^{-\theta x}.$$

Hence, we use g_θ as a majorizing function in an acceptance-rejection algorithm. Define

$$c_\theta \equiv \int_0^\infty g_\theta(x)dx = \frac{2}{\phi(\theta)\sigma\sqrt{2\pi}}.$$

Note that $c_\theta \geq 1$ for all $\theta \geq 0$. Our majorizing density is given by $h_\theta(x) = \theta e^{-\theta x}$, which is an exponential density with parameter θ. Thus, our acceptance-rejection algorithm to generate exponentially tilted positive normal random variates is as follows:

Algorithm for Generating Exponentially Tilted Positive Normal Variates

1. Generate $Y \sim$ exponential(θ).
2. Generate $U \sim$ uniform(0, 1), independent of Y.
3. If $U \leq \exp\{-Y^2/(2\sigma^2)\}$, then return $X = Y$. Otherwise, reject (Y, U) and return to step 1.

We can easily generate the exponential random variate in step 1 by using inversion. In the above algorithm, we do not need to know the value of $\phi(\theta)$. This is desirable since there is no closed form expression for $\phi(\theta)$, and we would otherwise have to evaluate it numerically.

The next theorem shows that the expected number of iterations in our algorithm converges to 1 as the tilting parameter increases to infinity.

Theorem 2 *For an exponentially tilted positive normal distribution, $c_\theta \to 1$ as $\theta \to \infty$.*

Proof. We need to show that $\theta\phi(\theta) \to 2/(\sigma\sqrt{2\pi})$ as $\theta \to \infty$. Recall our expression for $\phi(\theta)$ given in (1), and note that

$$\int_{\sigma^2\theta}^\infty \frac{2}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{y^2}{2\sigma^2}\right\} dy = P\{|N(0,\sigma^2)| > \sigma^2\theta\} \to 0$$

as $\theta \to \infty$, where $N(0, \sigma^2)$ denotes a normal random variable with mean 0 and variance σ^2. Also, $\theta^{-1}\exp\{-(\sigma^2\theta^2)/2\} \to 0$ as $\theta \to \infty$. Thus, using L'Hopital's rule, we obtain

$$\lim_{\theta\to\infty} \theta\phi(\theta) = \lim_{\theta\to\infty} \frac{\int_{\sigma^2\theta}^\infty \frac{2}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{y^2}{2\sigma^2}\right\} dy}{\theta^{-1}\exp\left\{-\frac{\sigma^2\theta^2}{2}\right\}} = \lim_{\theta\to\infty} \frac{2\sigma}{(\theta^{-2}+\sigma^2)\sqrt{2\pi}} = \frac{2}{\sigma\sqrt{2\pi}},$$

proving the result.

5 EMPIRICAL RESULTS

We now present some empirical results generated using the algorithms presented in the previous sections. The goal of the experiments was to determine how various choices of the tilting parameter θ affect the mean of the tilted distribution and the efficiency of our acceptance-rejection procedures (as measured by the expected number of iterations). Since there are no closed form expressions for these quantities, we estimated them by replicating our algorithms 10,000 times, and we provide 95% confidence intervals for the estimates. We let $\hat{\nu}(\theta)$ denote the estimate for the mean of the tilted distribution with tilting parameter θ, and let $\hat{c}_\theta$ be the estimate of the expected number of iterations in the acceptance-rejection scheme.

First we consider the Weibull distribution with parameters α and β. We varied the values of α and β so as to keep the mean of the original distribution fixed at μ. We selected $\mu = 1$ in Table 1 and $\mu = 10^3$ in Table 2, and varied α between 0.5 and 2.0. To keep the mean fixed, β must decrease as α increases.

Table 1: Results for Weibull Distribution with Original Mean 1

α	β	θ	$\hat{\nu}(\theta)$	$\hat{c}_\theta$
0.5	2.0	10^{-1}	$0.716 \pm 4\%$	$4.3 \pm 2\%$
0.5	2.0	1	$0.237 \pm 3\%$	$1.9 \pm 1\%$
0.5	2.0	10	$0.039 \pm 3\%$	$1.3 \pm 1\%$
0.5	2.0	10^2	$0.005 \pm 3\%$	$1.1 \pm 1\%$
0.5	2.0	10^3	$0.001 \pm 3\%$	$1.0 \pm 0\%$
1.0	1.0	10^{-1}	$0.908 \pm 2\%$	$10.9 \pm 2\%$
1.0	1.0	1	$0.500 \pm 2\%$	$2.0 \pm 1\%$
1.0	1.0	10	$0.092 \pm 2\%$	$1.1 \pm 1\%$
1.0	1.0	10^2	$0.010 \pm 2\%$	$1.0 \pm 0\%$
1.0	1.0	10^3	$0.001 \pm 2\%$	$1.0 \pm 0\%$
2.0	0.89	10^{-1}	$0.965 \pm 1\%$	$170.1 \pm 2\%$
2.0	0.89	1	$0.766 \pm 1\%$	$3.8 \pm 2\%$
2.0	0.89	10	$0.192 \pm 1\%$	$1.1 \pm 0\%$
2.0	0.89	10^2	$0.020 \pm 1\%$	$1.0 \pm 0\%$
2.0	0.89	10^3	$0.002 \pm 1\%$	$1.0 \pm 0\%$

We now examine the results in Table 1. (All of the following observations also apply to Table 2, showing their robustness.) For each fixed value of α and β, the average number of iterations in our acceptance-rejection scheme diminishes as θ grows, which agrees with Theorem 1. In addition, the estimated mean of the exponentially tilted distribution decreases as θ increases. As α gets larger, we need to tilt the distribution more in order to make the new mean smaller. This arises from the fact that as $\alpha \to \infty$, the Weibull becomes degenerate at $1/\beta$. Thus, more of the mass of the (original) distribution is tending toward larger values as α increases (and β decreases). Finally, for the larger values of θ, $\hat{\nu}(\theta)$ is approximately α/θ, which can be explained as follows. Note that $\hat{c}_\theta$ is almost 1 for large θ, implying that the tilted distribution is close to the majorizing distribution. (This is essentially what Theorem 1 states.) The majorizing distribution is a gamma(α, θ), which has mean α/θ, thus showing the desired property for large θ.

Table 2: Results for Weibull Distribution with Original Mean 10^3

α	β	θ	$\hat{\nu}(\theta)$	$\hat{c}_\theta$
0.5	2.0×10^{-3}	10^{-4}	$721.3 \pm 4\%$	$4.3 \pm 2\%$
0.5	2.0×10^{-3}	10^{-3}	$236.7 \pm 3\%$	$1.9 \pm 1\%$
0.5	2.0×10^{-3}	10^{-2}	$39.3 \pm 3\%$	$1.3 \pm 1\%$
0.5	2.0×10^{-3}	10^{-1}	$4.6 \pm 3\%$	$1.1 \pm 1\%$
0.5	2.0×10^{-3}	1	$0.5 \pm 3\%$	$1.0 \pm 0\%$
1.0	1.0×10^{-3}	10^{-4}	$915.3 \pm 2\%$	$10.9 \pm 2\%$
1.0	1.0×10^{-3}	10^{-3}	$506.0 \pm 2\%$	$2.0 \pm 1\%$
1.0	1.0×10^{-3}	10^{-2}	$90.9 \pm 2\%$	$1.1 \pm 1\%$
1.0	1.0×10^{-3}	10^{-1}	$10.0 \pm 2\%$	$1.0 \pm 0\%$
1.0	1.0×10^{-3}	1	$1.0 \pm 2\%$	$1.0 \pm 0\%$
2.0	8.9×10^{-4}	10^{-4}	$966.0 \pm 1\%$	$169.8 \pm 2\%$
2.0	8.9×10^{-4}	10^{-3}	$769.0 \pm 1\%$	$3.8 \pm 2\%$
2.0	8.9×10^{-4}	10^{-2}	$190.8 \pm 1\%$	$1.1 \pm 0\%$
2.0	8.9×10^{-4}	10^{-1}	$20.1 \pm 1\%$	$1.0 \pm 0\%$
2.0	8.9×10^{-4}	1	$2.0 \pm 1\%$	$1.0 \pm 0\%$

Table 3 contains the results from generating exponentially tilted random variates when the input distribution is a positive normal with parameter σ. We let σ take on the values 1, 10, and 100. The same types of observations which we made before also apply in this setting. In particular, as θ increases, $\hat{c}_\theta$ approaches 1 (which agrees with Theorem 2) and $\hat{\nu}(\theta)$ converges to 0. Also, $\hat{\nu}(\theta)$ is approximately $1/\theta$ for large θ since our majorizing distribution is an exponential(θ), which has mean $1/\theta$.

Table 3: Results for Positive Normal Distribution

σ	μ	θ	$\hat{\nu}(\theta)$	$\hat{c}_\theta$
1.0	0.80	10^{-2}	$0.80 \pm 2\%$	$78.8 \pm 2\%$
1.0	0.80	10^{-1}	$0.76 \pm 2\%$	$8.6 \pm 2\%$
1.0	0.80	1	$0.52 \pm 2\%$	$1.5 \pm 1\%$
1.0	0.80	10	$0.10 \pm 2\%$	$1.0 \pm 0\%$
1.0	0.80	100	$0.01 \pm 2\%$	$1.0 \pm 0\%$
10	8.0	10^{-2}	$7.60 \pm 1.5\%$	$8.6 \pm 1.9\%$
10	8.0	10^{-1}	$5.19 \pm 1.7\%$	$1.5 \pm 1.6\%$
10	8.0	1	$0.98 \pm 1.9\%$	$1.0 \pm 0.2\%$
10	8.0	10	$0.10 \pm 2.0\%$	$1.0 \pm 0.0\%$
10	8.0	100	$0.01 \pm 2.0\%$	$1.0 \pm 0.0\%$
100	80	10^{-2}	$53.22 \pm 1.7\%$	$1.5 \pm 1.2\%$
100	80	10^{-1}	$9.87 \pm 1.9\%$	$1.0 \pm 0.2\%$
100	80	1	$1.01 \pm 2.0\%$	$1.0 \pm 0.0\%$
100	80	10	$0.10 \pm 2.0\%$	$1.0 \pm 0.0\%$
100	80	100	$0.01 \pm 2.0\%$	$1.0 \pm 0.0\%$

6 CONCLUSIONS

We have proposed some simple and efficient algorithms for generating exponentially tilted random variates for two different input distributions. An area for future research is to develop variate generation schemes for other tilted distributions.

ACKNOWLEDGMENTS

Part of this work is based on material from the author's Ph.D. dissertation which was written while he was a student in the Operations Research Department at Stanford University. The author would like to thank his dissertation advisor Peter W. Glynn for his guidance. The research at Stanford was supported by IBM under SUR-SST contract 12480042. Also, the author would like to express gratitude to Randy Nelson and Perwez Shahabuddin for their helpful comments.

REFERENCES

Asmussen, S. 1985. Conjugate processes and the simulation of ruin problems. *Stochastic Processes and Their Applications* 20:213–230.

Cottrell, M., J. C. Fort, and G. Malgourves. 1983. Large deviations and rare events in the study of stochastic algorithms. *IEEE Transactions on Automatic Control* 28:907–920.

Devroye, L. 1986. *Non-uniform random variate generation.* New York: Springer-Verlag.

Dupuis, P. and J. Kushner. 1987. Stochastic systems with small noise, analysis and simulation; a phase locked loop example. *SIAM Journal of Applied Mathematics* 47:643–661.

Gradshteyn, I. S. and I. M. Ryzhik. 1981. *Tables of integrals, series, and products*. New York: Academic Press.

Law, A. M. and W. D. Kelton. 1991. *Simulation modelling and analysis*. 2d ed. New York: Academic Press.

Parekh, S. and J. Walrand. 1989. A quick simulation method for excessive backlog in networks of queues. *IEEE Transactions on Automatic Control* 34:54–66.

Sadowski, J. S. and J. A. Bucklew. 1990. On large deviations theory and asymptotically efficient Monte Carlo estimation. *IEEE Transactions on Data Theory* 36:579–588.

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

AUTHOR'S BIOGRAPHY

MARVIN K. NAKAYAMA is currently a postdoctoral fellow at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York. He received a B.A. in Mathematics/Computer Science from University of California, San Diego and a M.S. and Ph.D. in Operations Research from Stanford University. His research interests include simulation output analysis, gradient estimation, and rare event simulation.