Monte Carlo Estimation for Guaranteed-Coverage Non-normal Tolerance Intervals

Huifen Chen

Bruce W. Schmeiser

School of Industrial Engineering, Purdue University

West Lafayette, Indiana 47907-1287, U.S.A.

We propose a Monte Carlo sampling algorithm for estimating guaranteed-coverage tolerance factors for non-normal continuous distributions with known shape but unknown location and scale. The algorithm is based on reformulating this root-finding problem as a quantile-estimation problem.

The reformulation leads to a geometrical interpretation of the tolerance-interval factor. For arbitrary distribution shapes, we analytically and empirically investigate various relationships among tolerance-interval coverage, confidence, and sample size.

KEY WORDS: Quantile, Reliability, Root Finding, Stochastic Approximation.

1   INTRODUCTION

Let $\{X_1, X_2, ..., X_n\}$ be a random sample from the distribution F_X of a continuous random variable X with known shape but unknown location and unknown scale. That is, F_X has all known standardized moments but the mean μ and variance σ^2 are unknown. Define $\bar{X} = n^{-1} \sum_{i=1}^{n} X_i$ and $S^2 = (n-1)^{-1} \sum_{i=1}^{n} (X_i - \bar{X})^2$, the sample mean and sample variance of $\{X_1, X_2, ..., X_n\}$. We consider the guaranteed-coverage tolerance interval $I(k^*)$ for X where $I(k^*)$ is $[\bar{X} - k^* S, \infty)$ for lower one-sided, $(-\infty, \bar{X} + k^* S]$ for upper one-sided, and $[\bar{X} - k^* S, \bar{X} + k^* S]$ for two-sided intervals. The constant tolerance factor k^* is defined so that with $100\gamma\%$ confidence the random tolerance interval covers the proportion α of the distribution, i.e.,

$$\Pr_{\bar{X},S}\{ \Pr_X\{ X \in I(k^*) \} \geq \alpha \} = \gamma. \tag{1}$$

Here the future X is assumed to be independent of the sample statistics $\bar{X}$ and S. The value of k^* depends on sample size $n \in \{2, 3, ...\}$, coverage $\alpha \in (0, 1)$, confidence $\gamma \in (0, 1)$, and the distribution shape of X. For a single application, a single interval is computed from observed values of $\bar{X}$ and S^2; the probability that a future observation X lies in the interval is random but should be at least α in $100\gamma\%$ of many applications.

If μ and σ^2 were known, one could solve an equation for the unknown k^* to obtain a tolerance interval having 100% confidence. For example, the 100%-confidence α-coverage lower one-sided tolerance interval $[F_X^{-1}(1 - \alpha), \infty)$ is obtained by solving "$F_X(\mu - k^*\sigma) = 1 - \alpha$" for $k^* = [\mu - F_X^{-1}(1 - \alpha)]/\sigma$ and substituting into $[\mu - k^*\sigma, \infty)$. Our interest, however, is when μ and σ^2 are unknown but are estimated by $\bar{X}$ and S^2.

Based on $\bar{X}$ and S^2, the tolerance interval $I(k^*)$ is used to predict future behavior. In computer simulation of a manufacturing system, X might be the throughput of a single future shift and the observed data $X_1, X_2, ..., X_n$ might be the n simulated shift throughputs. An α proportion of future throughput is predicted to be in the interval with confidence γ. In reliability, based on product test results $X_1, X_2, ..., X_n$, a system is designed at the tolerance bounds to ensure, with confidence γ, that system reliability is at least α; i.e., at least $100\alpha\%$ of the systems built will not fail. In quality control, a contract might specify constants n, k^* and c so that a lot is accepted if a subset $X_1, X_2, ..., X_n$ yields a value of $\bar{X} - k^*S$ at least c units (lower specification limit); these constants can be chosen using tolerance-interval logic to guarantee that a particular lot containing $100(1 - \alpha)\%$ defective items (defined as $X < c$) is accepted with probability $1 - \gamma$.

Despite the broad range of applications, most guaranteed-coverage tolerance-interval literature assumes normally distributed X's, e.g., Wald and Wolfowitz (1946), Guttman (1970), Aitchison and Dunsmore (1975), and Eberhardt et al. (1989). The one-sided tolerance factor for the normal distribution is

$$k^* = t_{n-1,\gamma}\left(\sqrt{n}z_\alpha\right) / \sqrt{n}, \tag{2}$$

where $t_{\nu,\gamma}(\lambda)$ is the γ^{th} quantile of the noncentral t distribution with ν degrees of freedom

and noncentrality parameter λ, and z_α is the α^{th} quantile of the standard normal. The normal two-sided tolerance factor k^* can be computed by solving the equation

$$\int_{-\infty}^{\infty} \Pr\{\ \chi_{n-1}^2 \geq \frac{(n-1)v^2}{k^{*2}}\ \}\ \sqrt{\frac{n}{2\pi}}\ \exp\{\frac{-nu^2}{2}\}\ du\ =\ \gamma,$$

where χ_ν^2 is chi-square distributed with ν degrees of freedom, and v satisfies $\Phi(u+v) - \Phi(u-v) = \alpha$, where Φ is the standard normal distribution function. Odeh and Owen (1980) provide tables for one-sided and two-sided tolerance factors for normal distributions.

Some non-normal literature exists. Aitchison and Dunsmore (1975) and Patel (1986) also discuss different forms of tolerance intervals for binomial, Poisson, exponential, gamma, and some other standard populations. Guenther (1985) provides an extensive discussion of distribution-free tolerance intervals. Wald (1942) develops maximum-likelihood tolerance limits through asymptotic theory. Rao, Subrahmaniam, and Owen (1972) study the effect of non-normality on two-sided tolerance limits with control of both tails using the distribution shape dependent on third and fourth moments (α_3, α_4); they consider, however, only a small range of (α_3, α_4) around the normal shape. Jilek (1981) and Jilek and Ackermann (1989) provide a bibliography of the existing tolerance-region literature.

We focus on lower one-sided guaranteed-coverage tolerance intervals. Calculating the factor k^* in the upper one-sided guaranteed-coverage tolerance-interval $(-\infty, \bar{X} + k^* S]$ is a variation of the lower one-sided problem; in Appendix A we show that the upper one-sided factor with coverage α and confidence γ is the negative of the lower one-sided factor with coverage $1 - \alpha$ and confidence $1 - \gamma$. The factor k^* in the two-sided interval $[\bar{X} - k^* S,\ \bar{X} + k^* S]$ can be found by a modified algorithm, as discussed in Section 6.

In Section 2, we define the research problem. In Section 3, we propose a Monte Carlo Quantile Estimation (QE) algorithm for the lower one-sided guaranteed-coverage tolerance factor k^* based on reformulating k^* as a distribution quantile. Algorithmic convergence speed is also discussed and compared to that of stochastic approximation. In Section 4, we study the behavior of the lower one-sided factor k^* as a function of n, α, γ, and distribution shape. In Section 5, we show two examples of using tolerance intervals in acceptance sampling and

reliability. In Section 6, we show that the two-sided tolerance factor is also a distribution quantile and hence can be solved by modifying the QE algorithm.

2   PROBLEM DEFINITION

Let $F_X(\cdot)$ denote the distribution function from which the future observation X and the independent sample $\{X_1, X_2, ..., X_n\}$ are drawn. We assume that X is continuous (i.e., there is no point with positive probability mass), the shape of F_X is known, and the mean μ and standard deviation $\sigma > 0$ are unknown. We want to find the lower one-sided guaranteed-coverage tolerance factor k^*, given sample size n, α, γ and distribution shape, such that Equation 1, i.e.,

$$\Pr_{\bar{X},S}\{ \Pr_X\{ X \geq \bar{X} - k^* S \} \geq \alpha \} = \gamma, \tag{3}$$

is satisfied. For convenience, we define the function

$$g(k) = \Pr_{\bar{X},S}\{ \Pr_X\{ X \geq \bar{X} - kS \} \geq \alpha \}, \tag{4}$$

the confidence that the interval $[\bar{X} - kS, \infty)$ contains at least the proportion α of the measurements. Then this problem is to solve the equation $g(k^*) = \gamma$ for k^*.

We diverge from past practice. Despite some distribution-free approaches that are appropriate for large sample sizes, past practice has focused on building tables of tolerance-interval constants as a function of n, α, and γ (see, e.g., the comprehensive survey by Patel (1986)). The form of the tolerance constant differs depending on distributional assumptions, with our k^* corresponding to the common normality assumption. The analytical and computational effort to build a new table is often substantial. Although computed numerically, the resulting tables are often accurate to several digits. A practitioner uses such tables by assuming a distribution family and reading the appropriate constant from the appropriate table. The limitations to this approach are that (1) tables might not exist for the assumed distribution, that (2) existing tables might not be readily available, and that (3) all tables are incomplete, forcing the user to approximate the desired constant using interpolation or extrapolation.

Our divergence from past practice is motivated by our work with Thiokol Corporation, which uses tolerance intervals in its reliability program. Sample sizes are typically small, coverage probabilities are usually high, and many distribution shapes are possible. For example, typically for Thiokol applications $n < 10$, coverage $\alpha > .99999$, and distribution shape is measured using the third and fourth standardized moments. Making available tables to practicing engineers is not feasible, because the necessary tables would be huge.

Our approach is to avoid tables by computing the tolerance-interval constant k^* in real time as it is needed. Such an approach avoids the limitations of tables by assuming that engineers have interactive computing readily available. Thus, rather than provide tables, we provide a black-box algorithm that returns to the user the desired constant as a function of n, α, γ, and distribution shape.

We describe here a Monte Carlo sampling algorithm that estimates k^* to about two digits within a few seconds. (The specific tradeoff between computing time and accuracy depends upon n, α, γ, and distribution shape.) Precision beyond two digits is available by devoting more computing time. Providing an estimate of the standard error of the the Monte Carlo estimate of k^* is important, and straightforward.

Precision beyond two digits is seldom needed. The argument lies at three levels. First, the four- or five-digit accuracy shown in tables is oftenn far beyond the modeling error caused by not knowing the distribution shape precisely. Second, two-digit accuracy is sufficient even if the distribution shape is known. The first derivative of confidence with respect to the tolerance-interval factor, $g'(k^*)$, is small. An analysis that provides a confidence of $\gamma = 0.901$ when the nominal confidence is $\gamma = 0.900$ is not a bad analysis. Third, a Monte Carlo estimate of k^* is sometimes too low and sometimes too high, resulting in an average confidence that is close to the nominal confidence.

3   METHOD

Here we discuss solving the equation $g(k^*) = \gamma$ for the tolerance factor k^*, where $g(k) = \Pr_{\bar{X},S}\{ \Pr_X\{ X \geq \bar{X} - kS \} \geq \alpha \}$ as defined in Equation 4. We show three properties of

the root-finding function g, propose an efficient quantile estimation (QE) algorithm using these properties, and compare Algorithm QE to stochastic approximation algorithms, which are designed for general root-finding problems.

In our root-finding problem, three properties of g are useful:

1. $g : \Re \rightarrow [0,1]$ is a continuous nondecreasing function and strictly increasing in the set $\{k : g(k) \in (0,1)\}$,

2. for $0 < \gamma < 1$, equation $g(k) = \gamma$ has a unique solution, k^*, and

3. g does not depend on the population mean μ or standard deviation σ.

These three properties are straightforward. If the value of k increases, the value of $\bar{X} - kS$ decreases. Therefore, $g(k)$, the probability of having coverage at least α, increases. In the limits, $g(-\infty) = 0$ and $g(\infty) = 1$. Continuity follows since F_X has no mass points, so the first property holds. For the second property, existence of the root follows from continuity and the intermediate value theorem; uniqueness follows from g being increasing. To show the third property, let $Y = (X - \mu)/\sigma$ and $Y_i = (X_i - \mu)/\sigma$ for $i = 1, 2, ..., n$. Then

$$\Pr\nolimits_X\{\, X \geq \bar{X} - kS \,\} \;=\; \Pr\nolimits_Y\Big\{\, Y \geq \frac{\sum_{i=1}^{n} Y_i}{n} - k \sqrt{\frac{\sum_{i=1}^{n}[Y_i - \frac{\sum_{j=1}^{n} Y_j}{n}]^2}{n-1}} \,\Big\}$$
$$=\; \Pr\nolimits_Y\{\, Y \geq \bar{Y} - k\, S_Y \,\},$$

where $\bar{Y}$ and $S_Y{}^2$ are the sample mean and sample variance of $Y_1, ..., Y_n$, respectively. Hence,

$$g(k) \;=\; \Pr\nolimits_{\bar{Y}, S_Y}\{\, \Pr\nolimits_Y\{\, Y \geq \bar{Y} - k\, S_Y \,\} \geq \alpha \,\}$$

and the third property holds.

Despite not depending on μ or σ, g can be easily computed only for special cases, such as the normal distribution, since g depends upon the joint distribution function of $\bar{X}$ and S. However, $g(k)$ can be estimated by Monte Carlo simulation experiments using any arbitrary values of μ and σ^2.

We propose interpreting k^* as a quantile, which allows application of our Quantile Es-

timation (QE) algorithm. Then we show that QE is asymptotically more efficient than general-purpose stochastic-approximation algorithms.

3.1 Quantile Estimation Algorithm (QE)

A natural approach to solving for k^* in Equation 3 would be to invert g, defined in Equation 4. By Property 1, the inverse function g^{-1} always exists for the domain $(0, 1)$. However, it is easy to compute g^{-1} only for special cases, such as when X is normally distributed.

Nevertheless, we always can simplify g to reformulate k^* as a distribution quantile. We show this in Result 1.

Result 1 *Define the random variable $K = [\ \bar{X} - F_X^{-1}(1 - \alpha)\ ]\ /\ S$. Then $k^* = F_K^{-1}(\gamma)$, where $F_K(\cdot)$ is the distribution function of K.*

Proof:
$$
\begin{aligned}
g(k) &= \Pr_{\bar{X}, S}\{\ \Pr_X\{X \geq \bar{X} - kS\} \geq \alpha\ \} \\
&= \Pr_{\bar{X}, S}\{\ \Pr_X\{X < \bar{X} - kS\} \leq 1 - \alpha\ \} \\
&= \Pr_{\bar{X}, S}\{\ \bar{X} - kS\ \leq\ F_X^{-1}(1 - \alpha)\ \} \\
&= \Pr_{\bar{X}, S}\{\ \frac{\bar{X} - F_X^{-1}(1 - \alpha)}{S}\ \leq\ k\ \} \\
&= F_K(k).
\end{aligned}
$$

Hence $g^{-1}(\cdot) = F_K^{-1}(\cdot)$. Therefore, $k^* = g^{-1}(\gamma) = F_K^{-1}(\gamma)$, the γ^{th} quantile of F_K.

Algorithm QE estimates k^* for lower one-sided tolerance intervals $[\bar{X} - k^*\, S,\ \infty)$ by generating m Monte Carlo independent realizations of K. The estimate is

$$
\hat{k}^* = \omega\ k_{(\ \lfloor (m+1)\gamma \rfloor\ )}\ +\ (1 - \omega)\ k_{(\ \lceil (m+1)\gamma \rceil\ )}, \tag{5}
$$

the convex combination of the $\lfloor (m + 1)\gamma \rfloor^{\text{th}}$ and $(\lceil (m + 1)\gamma \rceil)^{\text{th}}$ order statistics, with the weight $\omega = \lceil (m + 1)\gamma \rceil - (m + 1)\gamma$ chosen to reduce the first-order bias of the quantile estimate. (The $\lfloor a \rfloor$ is the biggest integer that is no larger than a and $\lceil a \rceil$ is the smallest integer that is no less than a.) See Avramidis (1993) for other possibilities. Neither $\hat{k}^*$ nor k^* depends on μ or σ.

Algorithm QE(m): Given n, α, γ, and distribution shape, estimate k^*.

    Step 0. Initialize $i = 1$.

    Step 1. Independently generate a random sample $\{x_1, x_2, ..., x_n\}$ from the population X with any arbitrary values of μ and σ.

    Step 2. Compute the sample mean $\bar{x}$ and standard deviation s from the sample.

    Step 3. Compute $k_i = [\bar{x} - F_X^{-1}(1 - \alpha)]/s$.

    Step 4. If $i < m$, set $i \leftarrow i + 1$ and go to Step 1.

    Step 5. Compute $\hat{k}^*$ from $k_1, k_2, ..., k_m$ using Equation 5.

Given $\hat{k}^*$ from QE, a practitioner can form a lower one-sided tolerance interval $[\bar{X} - \hat{k}^* S, \infty)$ using observed values of $\bar{X}$ and S from real-world data.

3.2   Efficiency of Algorithms QE and Stochastic Approximation

We show here that the QE algorithm always converges at rate $O(m^{-1/2})$, the best that stochastic approximation algorithms can achieve. Furthermore, QE has no algorithmic parameter. Hence our QE is easier to apply and asymptotically more efficient.

The asymptotic distribution of the QE estimate $\hat{k}^*$, a combination of order statistics based on m independent realizations of random variable K, is (Lehmann (1983), p. 394)

$$\sqrt{m}(\hat{k}^* - k^*) \xrightarrow{D} N(0, \frac{\gamma(1 - \gamma)}{f_K^2(k^*)}), \tag{6}$$

where $f_K(\cdot)$ is the density function of K. Hence QE always converges at rate $O(m^{-1/2})$.

Stochastic approximation is a classical Monte Carlo approach first proposed by Robbins and Monro (1951) for root-finding problems when the function value $g(\cdot)$ is difficult to compute. There are several variations (e.g., Kesten (1958), Venter (1967), Andradóttir (1992), Polyak and Juditsky (1992)). All are iterative methods requiring only an ability to estimate $g(k)$. Each has several algorithmic parameters (initial point, step size, etc.), which strongly affect the speed of convergence.

Stochastic approximation achieves its best asymptotic distribution, that of Equation 6, when the optimal step size, which depends on $g'(k^*)$ (Sacks (1958), Venter (1967)), is chosen at each iteration. However, $g'(k^*)$ is unknown since k^* is unknown. Hence QE is asymptotically more efficient than stochastic approximation.

Variations of QE can improve performance. Because the normal distribution yields a fast solution via the noncentral t distribution, the normal-distribution estimator can be used as a control variate. A second variation is to sample K dependently using, for example, Latin hypercube sampling. We do not pursue these variations here because simple sampling has been adequate for our needs.

4 ANALYSIS

The lower one-sided tolerance factor k^* is a function of the parameter values n, α, γ, and distribution shape. In Subsection 4.1 we discuss properties of k^* for symmetric distribution shape and for infinite sample size. In Subsection 4.2 we discuss the sensitivity of k^* to the parameter values n, α, γ, and distribution shape.

4.1 Symmetric Distribution Shape and Infinite Sample Size

Here we show that for symmetric distributions the value of k^* with coverage α and confidence γ is the negative value of k^* with coverage $1 - \alpha$ and confidence $1 - \gamma$. We also show that, as the sample size n goes to infinity, k^* goes to $[\mu - F_X^{-1}(1 - \alpha)]/\sigma$.

Result 2 *Let $k_{n,\alpha,\gamma}^*$ denote the tolerance factor for the random variable X such that the coverage is α, confidence level is γ, and the sample size is n. If the distribution of X is symmetric, then for $0 < \alpha, \gamma < 1$ and $n \in \{2, 3, ...\}$*

$$k_{n,\alpha,\gamma}^* = -k_{n,1-\alpha,1-\gamma}^*. \tag{7}$$

The proof is in Appendix B.

9

The limiting value of k^* as the sample size n goes to infinity, for given values of α and γ, is sometimes useful as a bound, as an initial guess, or as an approximation when n is large. Odeh and Owen (1980) propose the normal-case limiting value, which is generalized to every continuous distribution in our Result 3.

Result 3 *For every* $0 < \alpha, \gamma < 1$ *and distribution* F_X, $\lim_{n \to \infty} k^*_{n,\alpha,\gamma} = [\mu - F_X^{-1}(1-\alpha)] / \sigma$.

When $n \to \infty$, $\bar{X}$ converges in probability to μ and S converges in probability to σ. Therefore, Slutsky's theorem implies from Result 1 that the random variable K converges in distribution to the constant $[\mu - F_X^{-1}(1-\alpha)]/\sigma$. Therefore, all quantiles $g^{-1}(q), 0 < q < 1$, converge to this same constant, yielding Result 3.

This limiting value is a function of only α and distribution shape. As always, it is not a function of μ or σ. In addition, the limiting value is not a function of the confidence γ, since the limiting joint distribution of $(S, \bar{X})$ is degenerate at (σ, μ).

4.2 Sensitivity Analysis

We show here that the lower one-sided factor k^* is an increasing function of α and of γ, but that k^* is not necessarily a monotonic function of n. The distribution shape can affect the values of k^* substantially.

To measure distribution shape, we use the skewness α_3 and kurtosis α_4, the third and fourth standardized moments. For any specified point (α_3, α_4), we choose the unique corresponding Johnson distribution. The Johnson family, proposed by Johnson (1949), includes three transformations of the standard normal distribution. Let X and Z denote the Johnson and standard normal random variables, respectively. The three transformations are:

$$
\begin{aligned}
S_L : \quad & Z = \eta + \delta \ln(\frac{X - \xi}{\lambda}), && \lambda(\, X - \xi) \geq 0, \\
S_B : \quad & Z = \eta + \delta \ln(\frac{X - \xi}{\xi + \lambda - X}), && 0 \leq X - \xi \leq \lambda, \\
S_U : \quad & Z = \eta + \delta \sinh^{-1}(\frac{X - \xi}{\lambda}), && -\infty < X < \infty.
\end{aligned}
$$

The constants ξ and λ are location and scale parameters, respectively; η and δ are the shape parameters. The second transformation, S_B, provides a bounded random variable X; the third transformation, S_U, results in an unbounded X. For lognormal distributions, S_L, the range is bounded below if $\lambda > 0$ and bounded above if $\lambda < 0$. DeBrota et al. (1989) develop two public-domain software packages; VISIFIT and FITTR1. VISIFIT allows visual fitting to a desired density shape. FITTR1 fits Johnson distributions to data using any of several criteria. We use the numerical routines of Hill, Hill, and Holder (1976) to find the Johnson distribution having desired moments μ, σ, α_3, and α_4. We arbitrarily use $\mu = 0$ and $\sigma = 1$, since k^* is not a function of μ or σ.

Tables 1 and 2 show values of k^* for thirty-six design points: $n \in \{2, 10, 30, \infty\}$, $\alpha \in \{.001, .5, .99\}$, and $\gamma \in \{.001, .5, .99\}$. The normal-distribution results in Table 1 are computed numerically. The $(\alpha_3, \alpha_4) = (4, 30)$ Johnson-distribution results in Table 2 are estimates using the QE algorithm based on 500,000 independent Monte Carlo samples of size n; only significant digits are shown, based on standard errors estimated using Schmeiser et al. (1990), and Hashem and Schmeiser (1994).

Table 1: Tolerance Factors for the Normal Distribution

		n			
α	γ	2	10	30	∞
.001	.001	-2365	-8.93	-5.15	-3.09
.001	.5	-4.53	-3.21	-3.13	-3.09
.001	.99	-0.97	-1.85	-2.28	-3.09
.5	.001	-225	-1.36	-0.62	0.00
.5	.5	0.00	0.00	0.00	0.00
.5	.99	22.49	0.89	0.45	0.00
.99	.001	0.15	1.08	1.49	2.33
.99	.5	3.38	2.41	2.35	2.33
.99	.99	186	5.07	3.45	2.33

These two tables illustrate four points that are true in general: (1) The tolerance factor k^* increases as the coverage α increases. (2) The tolerance factor k^* increases as the confidence γ increases. (3) The sensitivity to n is least when $\alpha \approx 0.5$ and $\gamma \approx 0.5$, with $k^* = 0$ in

Table 2: Tolerance Factors for the Johnson S_B with Skewness 4 and Kurtosis 30

		n			
α	γ	2	10	30	∞
.001	.001	-1.9E+4	-74.9	-32	-8.61
.001	.5	-27.4	-12.6	-10.3	-8.61
.001	.99	-1.5	-2.7	-3.7	-8.61
.5	.001	-400	-1.7	-0.24	0.33
.5	.5	0.30	0.36	0.35	0.33
.5	.99	15	0.89	0.62	0.33
.99	.001	0.43	0.41	0.38	0.74
.99	.5	1.35	0.98	0.86	0.74
.99	.99	67	1.94	1.34	0.74

symmetric cases such as the normal. (4) The value of k^* is sensitive to the non-normality.

The behavior of k^* is not always monotonically decreasing as n increases. For the normal distribution in Table 1, it is true that k^* moves monotonically to the limiting value; for the usual case of large confidence values γ, k^* decreases. For the Johnson distribution in Table 2, however, k^* does not move monotonically to its limiting value in the fifth and seventh rows. If we expanded Table 2, the seventh-row values of k^* would be 0.40, 0.44 and 0.49 when n is 60, 100 and 200, respectively. Hence, the value of k^* decreases with n until some value of n between 30 and 40, where it starts increasing toward the limiting value 0.74. The reason for this non-monotonic behavior is that the shape of the $(S, \bar{X})$ joint distribution changes as n increases. We later study this effect using a geometric interpretation (see Figure 4).

To interpret the behavior of the tolerance factor k^*, we now view the problem geometrically. Consider the straight line $L : \bar{x} = k^* s + F_X^{-1}(1 - \alpha)$ in the sample plane of $(S, \bar{X})$ with given values of n, α, γ and (α_3, α_4), where $\bar{x}$ and s are the realizations of $\bar{X}$ and S, respectively. The $\bar{x}$-axis intercept $F_X^{-1}(1 - \alpha)$ is determined only by distribution shape and α. The slope is k^*, determined so that γ is the probability of the random point $(S, \bar{X})$ lying on or below L. The value of k^* is not necessarily positive; negative values of k^* occur when α or γ, or both, are small.

Figure 1 is a scatter plot of one hundred independent observations $(s, \bar{x})$ from standard

Figure 1: Plot of Lines L in the Sample Plane of $(S, \bar{X})$ for $\alpha = 0.01, 0.5$ and 0.99 with Standard Normal Data, $n = 2$, and $\gamma = 0.5$

normal samples of size $n = 2$. For $\gamma = 0.5$, the three lines corresponding to $\alpha = 0.01, 0.5, 0.99$ illustrate the change of k^* with change of α. (The slopes are computed numerically.) As α increases, the intercept on the $\bar{x}$-axis moves down so the slope of line L goes up in order to keep half of the observations $(s, \bar{x})$ below line L.

Figure 1 also illustrates that k^* increases with γ, although the change of γ with α fixed is not plotted. As γ increases, the line L pivots counterclockwise at $(0, F_X^{-1}(1 - \alpha))$ to increase the proportion of the observations below L; hence, the slope of L increases.

Figures 2 and 3 illustrate k^* increasing and decreasing, respectively, with n. One hundred independent observations $(s, \bar{x})$ are plotted for both $n = 2$ and $n = 30$ for the Johnson S_B population with skewness 4 and kurtosis 30. Lines are shown for $n = 2, 30$ and ∞, and for $\alpha = 0.5$. (The slopes are estimated using QE with $m = 500,000$ samples of size n.) The only difference between Figure 2 and Figure 3 is that the value of γ changes from 0.001 to 0.99. As n increases, the slope of line L passing through point $(0, F_X^{-1}(1 - \alpha))$ goes closer to the limiting value of k^* as the joint distribution shrinks toward the point (σ, μ). Since a proportion γ of the points lies below the line, the larger value of γ has the larger slope k^*.

Figure 4 shows that k^* does not necessarily change monotonically with n. Three sets of

Figure 2: Plot of Line L in Sample Plane of $(S, \bar{X})$ from Johnson S_B Distribution for $n = 2$, 30, and ∞, where $\alpha = 0.5$, and $\gamma = 0.001$ (When $n = 2$, the slope is -400 and the Line L is hidden in the $\bar{x}$-axis.)

Figure 3: Plot of Line L in Sample Plane of $(S, \bar{X})$ from Johnson S_B Distribution for $n = 2$, 30, and ∞, where $\alpha = 0.5$, and $\gamma = 0.99$

thirty points $(S, \bar{X})$ from the same Johnson S_B distribution are shown for each of $n = 2$, $n = 30$, and $n = 200$. The four lines (for $n = 2, 30, 200, \infty$) correspond to $\alpha = 0.99$ and $\gamma = 0.1$. As n increases, the slope first decreases and then increases to the limiting k^*. The graph shows that nonmonotonic behavior occurs because the joint distribution of $(S, \bar{X})$ changes shape as it shrinks to (σ, μ). In this case, the changing shape effect dominates for small sample sizes and the shrinking effect dominates for large sample sizes.

Figure 4: Plot of Line L in Sample Plane of $(S, \bar{X})$ from Johnson S_B Distribution for $n = 2$, 30, 200, and ∞, where $\alpha = 0.99$, and $\gamma = 0.1$

5   EXAMPLES

We discuss two examples of tolerance intervals used in acceptance sampling and reliability, respectively. Example 1 assumes the normal population and Example 2 assumes the Johnson population. Comparing these two examples, we see that the value of k^* is sensitive to non-normality.

Example 1:   The acceptance sampling problem is to decide whether a lot should be

15

accepted or rejected, based on an inspected sample of size n from the lot. Given $n = 10$ and the lower specification limit 0, suppose the sampling plan is:

Accept the lot if $\bar{x} - k^* s \geq 0$ and   reject otherwise,

based on an inspected sample $\{x_1, ..., x_{10}\}$. The constant k^* controls the proportion of defective items (i.e., below 0) in accepted lots no more than 1%, with confidence 99%. Suppose the population is normally distributed. Then the acceptance sampling procedure is as follows.

1. Compute k^* with $n = 10$, $\alpha = 1 - 0.01 = 0.99$, $\gamma = 0.99$ and the normal population. Table 1 gives $k^* = 5.07$.

2. Obtain a sample $\{x_1, ..., x_{10}\}$ from the lot. Compute $\bar{x} = \sum_{i=1}^{10} x_i / 10 = 388$, and $s = [\sum_{i=1}^{10} (x_i - \bar{x})^2 / 9]^{1/2} = 200$. Then $\bar{x} - k^* s = 388 - 5.07 \cdot 200 = -626$.

3. Reject the lot since $\bar{x} - k^* s = -626 < 0$.

Example 2: Consider the reliability of solid-fuel rocket engines. Characteristic X is the difference between the strength of the rocket case and the pressure when this rocket is fired. The lower specification limit is therefore zero and the unknown reliability is $r = \Pr_X\{X \geq 0\}$. Suppose $n = 10$, $\gamma = 99\%$, and X has a Johnson S_B distribution with skewness 4 and kurtosis 30. Based on a sample $\{x_1, ..., x_{10}\}$, let $k^* = (\bar{x} - 0)/s$. Given this observed k^*, n, and γ, find the coverage α satisfying Equation 3. Then with 99% confidence, the reliability r is said to be at least α (Owen and Hua (1977)). Notice that α and k^* here are random, but the tolerance limit is fixed at 0. An engineer may perform the reliability program as follows.

1. Observe $\{x_1, ..., x_{10}\}$ from 10 fired test rockets. Compute $\bar{x} = \sum_{i=1}^{10} x_i / 10 = 388$, and $s = [\sum_{i=1}^{10} (x_i - \bar{x})^2 / 9]^{1/2} = 200$. Let $k^* = (\bar{x} - 0)/s = 388/200 = 1.94$.

2. Compute the value of α satisfying Equation 3 with $k^* = 1.94$, $n = 10$, $\gamma = 0.99$, and Johnson S_B shape. Table 2 gives $\alpha = 0.99$.

3. Conclude, with 99% confidence, that the system reliability is at least $\alpha = 99\%$.

6 TWO-SIDED TOLERANCE FACTORS

The factor k^* in the two-sided guaranteed-coverage tolerance interval $[\bar{X} - k^* S, \ \bar{X} + k^* S]$ is also a quantile of an observable random variable. Hence Algorithm QE can be modified to solve for k^*.

The two-sided factor k^* satisfies Equation 1, i.e.,

$$\Pr_{\bar{X},S}\{\Pr_X\{\bar{X} - k^* S \leq X \leq \bar{X} + k^* S\} \geq \alpha\} = \gamma. \tag{8}$$

Let $v_\alpha(\bar{X}) = v$ be the random variable satisfying $F_X(\bar{X} + v) - F_X(\bar{X} - v) = \alpha$. Then the event "$\Pr_X\{\bar{X} - k^* S \leq X \leq \bar{X} + k^* S\} \geq \alpha$" in Equation 8 is equivalent to the event "$k^* S \geq v_\alpha(\bar{X})$". Hence, Equation 8 can be rewritten as

$$\Pr_{\bar{X},S}\{ v_\alpha(\bar{X})/S \ \leq \ k^* \} = \gamma.$$

Define the random variable $K = v_\alpha(\bar{X})/S$, which again does not depend on the population mean μ or standard deviation σ. Then k^* is the γ^{th} quantile of the distribution of K, which can be observed via realizations of $\bar{X}$ and S.

The modification of Algorithm QE for two-sided tolerance intervals estimates k^* from order statistics in Equation 5 based on m independent realizations of the observable K. Analogous to the QE Algorithm for lower one-sided tolerance factors in Section 3, Step 3 is changed to

Step 3: Compute $k_i = v_\alpha(\bar{x})/s$, where $v_\alpha(\bar{x}) = v$ satisfies $F_X(\bar{x} + v) - F_X(\bar{x} - v) = \alpha$.

The new Step 3 requires numerical root finding for v.

ACKNOWLEDGMENTS

17

APPENDIX A: UPPER ONE-SIDED TOLERANCE FACTORS

Upper one-sided tolerance intervals are closely related to lower one-sided intervals. Let $k_{n,\alpha,\gamma}^U$ and $k_{n,\alpha,\gamma}^L$ denote the factors for upper and lower one-sided intervals, respectively, such that the coverage is α, confidence is γ and the sample size is n. (Notice that the previous notation k^* or $k_{n,\alpha,\gamma}^*$ is now $k_{n,\alpha,\gamma}^L$.) Then

$$\Pr_{\bar{X},S}\{\Pr_X\{X \leq \bar{X} + k_{n,\alpha,\gamma}^U\ S\} \geq \alpha\} = \gamma$$

implies that

$$\Pr_{\bar{X},S}\{\Pr_X\{X \geq \bar{X} + k_{n,\alpha,\gamma}^U\ S\} \leq 1 - \alpha\} = \gamma,$$

and therefore

$$\Pr_{\bar{X},S}\{\Pr_X\{X \geq \bar{X} + k_{n,\alpha,\gamma}^U\ S\} \geq 1 - \alpha\} = 1 - \gamma.$$

Hence,

$$k_{n,\alpha,\gamma}^U = -k_{n,1-\alpha,1-\gamma}^L\ . \tag{9}$$

To estimate the upper one-sided tolerance factor k^* with coverage α and confidence γ, we can estimate the lower one-sided tolerance factor with coverage $1 - \alpha$ and confidence $1 - \gamma$ and then change the sign. The limiting value of the upper one-sided guaranteed-coverage tolerance factor is then, from Equation 9 and Result 3,

$$\lim_{n\to\infty} k_{n,\alpha,\gamma}^U = [\ F_X^{-1}(\alpha) - \mu\ ]\ /\ \sigma,$$

for all positive values of γ. If the distribution of X is symmetric, then from Equation 9 and Result 2,

$$k_{n,\alpha,\gamma}^U = k_{n,\alpha,\gamma}^L\ ,$$

as noticed in Patel (1986) for the normal case.

APPENDIX B: PROOF FOR RESULT 2

Here we prove that if X is symmetrically distributed, then the lower one-sided tolerance factor satisfies $k^*_{n,\alpha,\gamma} = -k^*_{n,1-\alpha,1-\gamma}$.

Because changing μ does not affect the tolerance factor, without loss of generality set $\mu = 0$, so that the distribution of X is symmetric at zero. Then, the samples $(x_1, ..., x_n)$ and $(-x_1, ..., -x_n)$ are equally likely. Hence their sample statistics $(\bar{x}, s^2)$ and $(-\bar{x}, s^2)$, respectively, are equally likely. Given $n \in \{2, 3, ...\}$ and $0 < \alpha, \gamma < 1$, then $k^*_{n,\alpha,\gamma}$ satisfies

$$\Pr_{\bar{X},S}\{\Pr_X\{X \geq \bar{X} - k^*_{n,\alpha,\gamma}\, S\} \geq \alpha\} = \gamma.$$

Also, from Equation 9,

$$\Pr_{\bar{X},S}\{\, \Pr_X\{X \leq \bar{X} - k^*_{n,\alpha,\gamma}\, S\} \geq 1 - \alpha\,\} = 1 - \gamma.$$

Because the distribution of X is symmetric at zero,

$$\Pr_{\bar{X},S}\{\, \Pr_X\{X \geq -\bar{X} + k^*_{n,\alpha,\gamma}\, S\} \geq 1 - \alpha\,\} = 1 - \gamma.$$

Because $(\bar{X}, S)$ and $(-\bar{X}, S)$ are equally likely,

$$\Pr_{\bar{X},S}\{\, \Pr_X\{X \geq \bar{X} + k^*_{n,\alpha,\gamma}\, S\} \geq 1 - \alpha\,\} = 1 - \gamma.$$

Hence,

$$k^*_{n,\alpha,\gamma} = -k^*_{n,1-\alpha,1-\gamma}.$$

REFERENCES

Aitchison, J., and Dunsmore, I. R. (1975). *Statistical Prediction Analysis*, Cambridge: Cambridge University Press.

Andradóttir, S. (1992). An Empirical Comparison of Stochastic Approximation Methods

for Simulation Optimization. *Proceedings of the First Industrial Engineering Research Conference*, ed. Klutke, G., Mitta, D. A., Nnaji, B. O., and Seiford, L. M., 471–475. Institute of Industrial Engineers, Norcross, Georgia.

Avramidis, A. N. (1993). *Variance Reduction Techniques for Simulation with Applications to Stochastic Networks*, Ph.D. Thesis, School of Industrial Engineering, Purdue University, West Lafayette, Indiana.

DeBrota, D. J., Dittus, R. S., Swain, J. J., Roberts, S. D., Wilson, J. R., and Venkatraman, S. (1989). Modeling Input Processes with Johnson Distributions. *Proceedings of the 1989 Winter Simulation Conference*, ed. MacNair, E. A., Musselman, K. J., Heidelberger, P., 308–318. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Eberhardt, K. R., Mee, R. W., and Reeve, C. P. (1989). Computing Factors for Exact Two-Sided Tolerance Limits for a Normal Distribution. *Communications in Statistics— Simulation and Computation* **18**, 397–413.

Guenther, W. C. (1985). Two-Sided Distribution-Free Tolerance Intervals and Accompanying Sample Size Problems. *Journal of Quality Technology* **17**, 40–43.

Guttman, I. (1970). *Statistical Tolerance Regions: Classical and Bayesian*, London: Charles Griffin and Co.

Hashem, S., and Schmeiser, B. (1994). Algorithm 727: Quantile Estimation Using Overlapping Batch Statistics. *ACM Transactions on Mathematical Software* **20**, 100-102.

Hill, I. D., Hill, R., and Holder, R. L. (1976). Algorithm AS 99. Fitting Johnson Curves by Moments. *Applied Statistics* **25**, 180–189.

Jilek, M. (1981). A Bibliography of Statistical Tolerance Regions. *Statistics* **12**, 441–456.

Jilek, M. and Ackermann, H. (1989). A Bibliography of Statistical Tolerance Regions, II. *Statistics* **20**, 165–172.

Johnson, N. L. (1949). Systems of Frequency Curves Generated by Methods of Translation. *Biometrika* **36**, 149–176.

Kesten, H. (1958). Accelerated Stochastic Approximation. *Annals of Mathematical Statistics* **29**, 41–59.

Lehmann, E. L. (1983). *Theory of Point Estimation*, New York: John Wiley.

Odeh, R. E., and Owen, D. B. (1980). *Tables for Normal Tolerance Limits, Sampling Plans,*

and Screening, New York: Marcel Dekker Inc.

Owen, D. B., and Hua, Tsunshung A. (1977). Tables of Confidence Limits on the Tail Area of the Normal Distribution. *Communications in Statistics—Simulation and Computation* **6**, 285–311.

Patel, J. K. (1986). Tolerance Intervals—a Review. *Communications in Statistics—Theory and Methods* **15**, 2719–2762.

Polyak, B. T., and Juditsky, A. B. (1992). Acceleration of Stochastic Approximation by Averaging. *SIAM Journal on Control and Optimization* **30**, 838–855.

Rao, J. N. K., Subrahmaniam, K., Owen, D. B. (1972). Effect of Non-Normality on Tolerance Limits Which Control Percentages in Both Tails of Normal Distribution. *Technometrics* **14**, 571–575.

Robbins, H., and Monro, S. (1951). A Stochastic Approximation Method. *Annals of Mathematical Statistics* **22**, 400–407.

Sacks, J. (1958). Asymptotic Distributions of Stochastic Approximation Procedures. *Annals of Mathematical Statistics* **29**, 373–405.

Schmeiser, B. W., Avramidis, T. N., and Hashem, S. (1990). Overlapping Batch Statistics. *Proceedings of the 1990 Winter Simulation Conference*, ed. Balci, O., Sadowski, R., and Nance, R., 395–398. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Venter, J. H. (1967). An Extension of the Robbins-Monro Procedure. *Annals of Mathematical Statistics* **38**, 181–190.

Wald, A. (1942). Setting of Tolerance Limits When the Sample Is Large. *Annals of Mathematical Statistics* **13**, 389–399.

Wald, A., and Wolfowitz, J. (1946). Tolerance Limits for a Normal Distribution. *Annals of Mathematical Statistics* **17**, 208–215.