

This is the peer reviewed version of the following article: Zachary Zimmer, DoHwan Park and Thomas Mathew, Tolerance limits under zero-inflated lognormal and gamma distributions, Comp and Math Methods. 2020;e1113, doi: <https://doi.org/10.1002/cmm4.1113>, which has been published in final form at <https://doi.org/10.1002/cmm4.1113>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Tolerance limits under zero-inflated lognormal and gamma distributions

Zachary Zimmer | DoHwan Park | Thomas Mathew 

Department of Mathematics and Statistics,
University of Maryland, Baltimore County

Correspondence

Thomas Mathew, Department of
Mathematics and Statistics, University of
Maryland, Baltimore County.
Email: mathew@umbc.edu

The computation of upper tolerance limits is investigated for the zero-inflated lognormal distribution and the zero-inflated gamma distribution, with or without covariates. The methodologies investigated consist of a fiducial approach and bootstrap approaches, including the bias corrected and accelerated bootstrap and a bootstrap-calibrated delta method. Based on estimated coverage probabilities, it is concluded that overall, the bootstrap-calibrated delta method is to be preferred for computing the upper tolerance limit. Two applications are also discussed; the first application is on the analysis of data on health care expenditures, and the second application is on testing the safety of body armor.

KEYWORDS

BCa bootstrap, bootstrap calibration, bootstrap-calibrated delta method, delta method, fiducial quantity

1 | INTRODUCTION

The lognormal distribution and the gamma distribution are widely used for the modeling and analysis of positively skewed data that are also positive. In some practical applications, the sample will include a certain proportion of zeros, in addition to positive values arising from skewed distributions such as the lognormal and the gamma. A common scenario where zeros are expected is in the context of workplace exposure assessment studies where the zeros correspond to nonexposed workers, and the positive values represent exposure measurements that are typically skewed. The proportion of zeros is an unknown parameter since the population proportion of nonexposed workers is usually unknown. A distribution that includes zeros in addition to positive values is referred to as a zero-inflated distribution. Such a distribution is also commonly encountered in the context of health care expenditures since a certain proportion of patients may not incur any cost during a given period. A problem that has been rather well addressed in the literature for the zero-inflated lognormal distribution is inference on the mean, and various solutions have been proposed and numerically compared, including solutions based on a *fiducial approach*; see Hasan and Krishnamoorthy¹ for further background information and for the fiducial approach for inference on the mean and percentiles of the distribution.

This article is on the zero-inflated lognormal and gamma distributions and addresses the problem of obtaining upper tolerance limits, that is, upper confidence limits for specific percentiles; see Krishnamoorthy and Mathew.² Our setup is as follows. Let Y be a nonnegative random variable corresponding to a zero-inflated population, that is, Y assumes positive continuous values along with a positive mass at zero. If π is the proportion of zeros in the population, then the distribution function of Y , say $G(y)$, is given by

$$G(y) = \begin{cases} \pi & \text{for } y = 0 \\ \pi + (1 - \pi)F(y) & \text{for } y > 0, \end{cases} \quad (1)$$

where $F(y)$ is the distribution function of Y when it is positive. It is possible that the distribution function $F(y)$, as well as the proportion π will depend on covariates. When the log-normal distribution of Y depends on covariates, we shall assume the log-normal regression model, that is, the mean of $\ln(Y)$ has a linear regression structure. When the gamma distribution for Y depends on covariates, we model using a log-link function for the mean. That is, the logarithm of the mean has a linear regression structure; see Lawless.³

We shall investigate the construction of upper tolerance limits in the set up (1) when the distribution function $F(y)$ corresponds to either the lognormal or the gamma distribution. In the absence of covariates, the tolerance limit problem is addressed in Hasan and Krishnamoorthy¹ for the zero-inflated lognormal distribution. The authors conclude that a solution based on the fiducial approach, and an approximate solution that is also easy to compute, are both satisfactory. Our goal is to compute upper tolerance limits when the zero-inflated distributions also depend on covariates, as mentioned earlier. For the zero-inflated lognormal distribution, in addition to the fiducial approach, we shall also pursue a bias corrected and accelerated bootstrap (BCa), and also a bootstrap-calibrated delta method, and compare them based on estimated coverage probabilities. Here the bootstrapping is done parametrically. The delta method, being based on asymptotic normality, is not expected to maintain the coverage probability unless the sample size is large. The bootstrap calibration is meant to provide a correction so that accurate coverage probabilities can be expected when the sample size is not large. In the absence of covariates, the fiducial approach turned out to be quite satisfactory for the zero-inflated lognormal distribution, agreeing with the conclusion in Hasan and Krishnamoorthy.¹ In the same scenario, and also for the set up with covariates, the fiducial approach turned out to be quite satisfactory. In addition, the bootstrap-calibrated delta method also provided satisfactory results; however, the BCa bootstrap may fall short with respect to coverage probabilities.

For the zero-inflated gamma distribution, we explored two bootstrap approaches and the bootstrap-calibrated delta method in order to compute an upper tolerance limit. If ξ_p denotes the p th percentile of the zero-inflated gamma distribution, so that the upper tolerance limit is an upper confidence limit for ξ_p , we investigated the BCa bootstrap for the distribution of $\hat{\xi}_p$, and for the distribution of $\hat{\xi}_p/\xi_p$, where $\hat{\xi}_p$ is the maximum likelihood estimator of ξ_p . Bootstrapping the latter quantity is considered in Zhang and An.⁴ Numerical results showed that the latter provides better coverage probabilities, but the coverages can still be somewhat unsatisfactory in some cases. On the other hand, the bootstrap-calibrated delta method continues to provide satisfactory coverages whether or not covariates are present. For the zero-inflated gamma distribution it maybe possible to explore the fiducial approach after applying a cube-root normal approximation, as done in Krishnamoorthy and Wang⁵ and Krishnamoorthy et al;⁶ however, we have not pursued this approach here.

After discussing the methodology and presenting the numerical results on the coverage probabilities, we have applied the proposed methods to two examples. The first example is taken from Callahan et al⁷ and is on the diagnostic test charges for a group of 40 patients. Our second example deals with the analysis of a set of body armor data where the zero-inflated gamma distribution with covariates is appropriate. Data analysis based on the two examples has also brought out the differences among the competing methods for computing upper tolerance limits. Overall, it appears that the fiducial approach and the bootstrap-calibrated delta method provide satisfactory methodology for computing upper tolerance limits for the zero-inflated distributions considered, with or without covariates.

2 | UPPER TOLERANCE LIMITS FOR THE ZERO-INFLATED LOGNORMAL DISTRIBUTION

Let Y be a random variable following the zero-inflated lognormal distribution, and suppose we have a random sample of size n from the distribution. Among the n observations in the sample, let n_0 denote the number of zero observations. If π denotes the unknown proportion of zero observations in the population, we then have $n_0 \sim \text{Binomial}(n, \pi)$. For $Y > 0$, Y follows a log-normal distribution. Here we assume that π does not depend on covariates.

2.1 | Zero-inflated log-normal distribution without covariates

We note that when $Y > 0$, $W = \ln(Y) \sim N(\mu, \sigma^2)$. Let ξ_p denote the p th percentile of the distribution. Here it is assumed that $\pi < p$. Then we have $\xi_p = \mu + z_{(p-\pi)/(1-\pi)} \sigma$, where z_γ denotes the 100γ th percentile of the standard normal distribution. We want to compute a $100(1 - \alpha)\%$ upper confidence limit for ξ_p . We shall first explain the fiducial approach very briefly, omitting the details which are available elsewhere; see, for example, Hasan and Krishnamoorthy.¹

Let Y_1, Y_2, \dots, Y_n denote n independent observations from the zero-inflated lognormal distribution, and let n_0 denote the number of zero observations. Thus when $Y_i > 0$, $W_i = \ln(Y_i) \sim N(\mu, \sigma^2)$. Let \bar{W} and S^2 denote the mean and variance among the W_i s, and let \bar{w} and s^2 denote the corresponding observed values. Then a set of approximate fiducial quantities for π , μ , and σ , say $\tilde{\pi}$, $\tilde{\mu}$, and $\tilde{\sigma}$, respectively, are given by

$$\begin{aligned}\tilde{\pi} &= \text{Beta}\left(n_0 + \frac{1}{2}, n - n_0 + \frac{1}{2}\right) \\ \tilde{\mu} &= \bar{w} + \frac{Z_1}{U/\sqrt{n - n_0 - 1}} \frac{s}{\sqrt{n - n_0}} \\ \tilde{\sigma}^2 &= \frac{(n - n_0 - 1)s^2}{U^2},\end{aligned}\quad (2)$$

where $Z_1 \sim N(0, 1)$ is independent of $U^2 \sim \chi_{n-n_0-1}^2$, and χ_r^2 denotes the chi-square distribution with $\text{df} = r$. An approximate fiducial quantity for ξ_p , say $\tilde{\xi}_p$ is then given by

$$\tilde{\xi}_p = \tilde{\mu} + z_{[(p-\tilde{\pi})/(1-\tilde{\pi})]} \tilde{\sigma}. \quad (3)$$

The $100(1 - \alpha)$ th percentile of $\tilde{\xi}_p$, computed after keeping the observed data fixed, gives the required upper tolerance limit for the zero-inflated lognormal distribution.

We shall now discuss the implementation of a bootstrap-calibrated delta method for computing the upper tolerance limit, that is, an upper confidence limit for ξ_p . We note that the implementation of the delta method requires the computation of the derivatives of $\xi_p = \mu + z_{[(p-\pi)/(1-\pi)]} \sigma$ with respect to π , μ , and σ . However, there is no analytic form for the derivative with respect to π . We computed all the required derivatives numerically using the numderiv package in R (2015). It turned out that the delta method based upper tolerance limit had coverages below the nominal level, implying that in order to get a coverage probability close to $1 - \alpha$, the nominal level to be used has to be more than $1 - \alpha$. We estimated the required nominal level by the bootstrap, resulting in a bootstrap-calibrated solution. The following algorithm summarizes the steps required to obtain the bootstrap-calibrated upper tolerance limit:

Algorithm 1

1. Based on the given sample from the zero-inflated log-normal distribution, estimate the MLEs of π , μ and σ , denoted as $\hat{\pi}$, $\hat{\mu}$ and $\hat{\sigma}$, respectively. Furthermore, let $\hat{\xi}_p = \hat{\mu} + z_{[(p-\hat{\pi})/(1-\hat{\pi})]} \hat{\sigma}$ denote the MLE of $\xi_p = \mu + z_{[(p-\pi)/(1-\pi)]} \sigma$, and let $SE(\hat{\xi}_p)$ denote its asymptotic standard error, computed by the delta method.
 2. Generate B parametric bootstrap samples $(Y_{1j}^*, Y_{2j}^*, \dots, Y_{nj}^*)$, $j = 1, 2, \dots, B$, from the zero inflated log-normal distribution with the estimated parameters $\hat{\pi}$, $\hat{\mu}$, and $\hat{\sigma}$. Compute the MLEs of π , μ , and σ from the j th bootstrap sample, and denote them as $\hat{\pi}_j^*$, $\hat{\mu}_j^*$, and $\hat{\sigma}_j^*$, respectively.
 3. Let $\hat{\xi}_{pj}^*$ denote the MLE of ξ_p computed using the j th bootstrap sample, $j = 1, 2, \dots, B$. Also let $SE(\hat{\xi}_{pj}^*)$ denote its asymptotic standard error.
 4. For a grid of values of $1 - \alpha^*$, and for $j = 1, 2, \dots, B$, let $\hat{\xi}_{pj}^* + z_{1-\alpha^*} SE(\hat{\xi}_{pj}^*)$ be an upper confidence limit for ξ_p based on a nominal confidence level of $1 - \alpha^*$.
 5. For each $1 - \alpha^*$, compute the proportion of the B upper confidence limits obtained in step 4, which exceeds the MLE $\hat{\xi}_p$ of ξ_p .
 6. Choose a value of $1 - \alpha^*$, say $1 - \hat{\alpha}$, so that the above proportion is equal to $1 - \alpha$.
 7. Now go back to the original sample, and compute the upper confidence limit for ξ_p based on a nominal confidence level of $1 - \hat{\alpha}$; the upper confidence limit is given by $\hat{\xi}_p + z_{1-\hat{\alpha}} SE(\hat{\xi}_p)$.
-

In addition to the fiducial solution and the solution based on the bootstrap-calibrated delta method, we also considered the BCa bootstrap for the interval estimation of ξ_p . Later we shall assess the performance of the resulting upper tolerance limits based on their estimated coverage probabilities.

2.2 | Zero-inflated log-normal distribution with covariates

Next we consider the model with covariates. However, we proceed under the assumption that the probability π for an observation Y to be zero is free of covariates. Furthermore, in the presence of covariates, we assume that the positive

values of Y follow a log-regression model. In other words, for $Y > 0$, we assume

$$\ln(Y) \sim N(\mathbf{x}'\boldsymbol{\beta}, \sigma^2), \quad (4)$$

where \mathbf{x} is a $q \times 1$ covariate vector, and $\boldsymbol{\beta}$ is an unknown parameter vector. We shall once again consider a sample of size n that includes n_0 observations equal to zero. Let \mathbf{x}_i ($i = 1, 2, \dots, n - n_0$) denote the values of the covariate vector corresponding to the $n - n_0$ non-zero values in the sample, and let X be an $(n - n_0) \times q$ matrix whose i th row is \mathbf{x}_i' . Under the model (4) for the non-zero observations, let $\hat{\boldsymbol{\beta}}$ denote the least squares estimator of $\boldsymbol{\beta}$ and S^2 denote the residual mean square, where we assume that the matrix X has rank q . Furthermore, let $\hat{\boldsymbol{\beta}}_{\text{obs}}$ and s^2 denote the observed values of $\hat{\boldsymbol{\beta}}$ and S^2 , respectively. Approximate fiducial quantities for $\boldsymbol{\beta}$ and σ^2 , say $\tilde{\boldsymbol{\beta}}$ and $\tilde{\sigma}^2$, respectively, can be obtained similar to the fiducial quantities given in the previous subsection. These are given by

$$\begin{aligned} \tilde{\boldsymbol{\beta}} &= \hat{\boldsymbol{\beta}}_{\text{obs}} - \sqrt{n - n_0 - q} \frac{S}{U} (X'X)^{-1/2} \mathbf{z} \text{ and } \tilde{\sigma}^2 = (n - n_0 - q) s^2 / U^2, \\ \text{where } \mathbf{z} &= \frac{1}{\sigma} (X'X)^{1/2} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \sim N(\mathbf{0}, I_q) \text{ and } U^2 = (n - n_0 - q) \frac{S^2}{\sigma^2} \sim \chi_{n-n_0-q}^2. \end{aligned} \quad (5)$$

Suppose an upper tolerance limit is required for the zero-inflated log-regression model at a specified covariate vector \mathbf{x}_0 . The p th percentile of this distribution is given by $\xi_p = \mathbf{x}_0' \boldsymbol{\beta} + z_{[(p-\pi)/(1-\pi)]} \sigma$. A fiducial quantity for ξ_p is given by

$$\tilde{\xi}_p = \mathbf{x}_0' \tilde{\boldsymbol{\beta}} + z_{[(p-\tilde{\pi})/(1-\tilde{\pi})]} \tilde{\sigma}, \quad (6)$$

where $\tilde{\boldsymbol{\beta}}$ and $\tilde{\sigma}$ are given in Equation (5) and $\tilde{\pi}$ is given in Equation (2). The $100(1 - \alpha)$ th percentile of $\tilde{\xi}_p$ gives the required upper tolerance limit.

3 | ZERO-INFLATED GAMMA DISTRIBUTION

We now consider the case of the zero-inflated gamma distribution, that is, the non-zero observations come from a gamma distribution. Furthermore, when the gamma distribution depends on covariates, we assume a log-link function for the gamma mean, that is, the natural logarithm of the mean has a linear regression structure; see Farewell and Prentice,⁸ Manning et al.,⁹ and Tu.¹⁰ We note that there is no analytic expression for the percentiles of the distribution, so the computations have to be carried out numerically. We have derived upper confidence limits for the p th percentile ξ_p using the BCa bootstrap for the distribution of $\hat{\xi}_p$, and also bootstrapping the distribution of $\hat{\xi}_p / \xi_p$, where $\hat{\xi}_p$ is the MLE of ξ_p ; the latter bootstrap approach is used in Zhang and An.⁴

4 | NUMERICAL RESULTS

In this section, we report numerical results on the estimated coverage probabilities of the upper confidence limits for ξ_p in order to assess the accuracy of the proposed solutions. All coverage probabilities are reported for $p = 0.95$, and computed using 5000 simulated samples from the respective distributions. Furthermore, wherever bootstrapping is involved in our simulations, we have used 300 bootstrap samples. We recall that the bootstrapping is done parametrically. Additionally, for implementing the fiducial approach, we generated 50 000 values of the relevant fiducial quantity.

We start with the log-normal distribution. Table 1 gives the coverage probabilities under the model without covariates. The coverage probabilities are reported for the upper confidence limits based on the BCa bootstrap applied to the MLE $\hat{\xi}_p$, the fiducial approach, and the bootstrap-calibrated delta method, for two choices of the lognormal parameters: $\mu = 2.1$, $\sigma = 1$ and $\mu = 3.5$, $\sigma = 3$.

When $\mu = 2.1$ and $\sigma = 1$, all the three approaches appear to be satisfactory in terms of maintaining the coverage probability. However, for the parameter choice $\mu = 3.5$, $\sigma = 3$, the performance of the BCa bootstrap approach is not satisfactory; the coverage probabilities are somewhat less than the nominal level. In the same set-up, the bootstrap-calibrated delta method also results in poor coverage probability when the sample size is small ($n = 20$); the coverage probabilities are more satisfactory in the other cases. The fiducial approach exhibits satisfactory performance in all scenarios considered for simulation. The satisfactory performance of the fiducial approach is already noted in Hasan and Krishnamoorthy.¹

TABLE 1 Estimated coverage probabilities of the 95% upper confidence limits for the 95th percentile of the a zero-inflated log-normal distribution; (i) BCa bootstrap, (ii) Fiducial approach, and (iii) Bootstrap-calibrated delta method

<i>n</i>	π	$\mu = 2.1, \sigma = 1$			$\mu = 3.5, \sigma = 3$		
		(i)	(ii)	(iii)	(i)	(ii)	(iii)
20	0.5	0.9450	0.9484	0.9474	0.8944	0.9456	0.9062
50	0.3	0.9422	0.9500	0.9448	0.9352	0.9488	0.9374
50	0.5	0.9496	0.9476	0.9452	0.9298	0.9474	0.9404
80	0.2	0.9510	0.9520	0.9474	0.9310	0.9520	0.9440
80	0.3	0.9426	0.9468	0.9444	0.9418	0.9408	0.9424
80	0.4	0.9480	0.9480	0.9456	0.9430	0.9464	0.9456
80	0.5	0.9514	0.9482	0.9470	0.9310	0.9454	0.9470
80	0.6	0.9472	0.9502	0.9472	0.9334	0.9524	0.9390
80	0.7	0.9468	0.9454	0.9504	0.9392	0.9484	0.9452
150	0.5	0.9504	0.9480	0.9468	0.9412	0.9440	0.9460

TABLE 2 Estimated coverage probabilities of the 95% upper confidence limits for the 95th percentile of a zero-inflated log-normal regression corresponding to $x = 0.25$ and $z = 1$ in (7); (i) BCa bootstrap, (ii) Fiducial approach, and (iii) Bootstrap-calibrated delta method

<i>n</i>	π	$\sigma = 1$			$\sigma = 3$		
		(i)	(ii)	(iii)	(i)	(ii)	(iii)
100	0.4	0.9276	0.9336	0.9454	0.9312	0.9278	0.9406
100	0.5	0.9280	0.9296	0.9450	0.9368	0.9222	0.9394
200	0.4	0.9356	0.9408	0.9492	0.9414	0.9360	0.9438
200	0.5	0.9402	0.9340	0.9502	0.9426	0.9352	0.9504

We now consider a log-regression model for Y , given by

$$\ln(Y) \sim N(2.1 + 1.4x - z, \sigma^2), \quad (7)$$

where x takes the values 0, 0.25, 0.5, and 0.75 with equal frequency. Thus we consider sample sizes that are multiples of four. Furthermore, when $x = 0$ or 0.5, we assume that z takes the value 0, and when x assumes the values 0.5 and 0.75, we assume that z takes the value 1. Coverage probabilities are listed in Table 2 for the upper tolerance limits corresponding to $x = 0.25$ and $z = 1$, and for the choices $\sigma = 1$ and $\sigma = 3$.

We note that in spite of the somewhat large sample sizes in Table 2, only the bootstrap-calibrated delta method provides satisfactory coverages; both the BCa bootstrap and the fiducial approach seem to fall short. More extensive simulations appear to be necessary before firm conclusions can be drawn.

Tables 3 and 4, respectively, provide estimated coverage probabilities under the gamma distribution without and with covariates. In Table 3, the scale parameter is chosen as 1.2 and 0.1. The shape parameter was chosen as $\exp(-0.2) = 0.67032$ and $\exp(0.2) = 1.491825$. For the scenario with covariates considered in Table 4, we assume that the mean μ and the probability π both depend on covariates according to the following functions:

$$\ln(\mu) = 1.2 + 1.8x - z, \quad \text{logit}(\pi) = -1.09861 - 2.19722z, \quad (8)$$

where the choices of x and z are as specified for the model (7). Furthermore, for Table 4, we have chosen the shape parameter to be $\exp(0.2) = 1.491825$. Three methodologies are considered in Tables 3 and 4 for computing an upper confidence limit for ξ_p : the BCa bootstrap, bootstrapping the distribution of $\hat{\xi}_p/\xi_p$, and the bootstrap-calibrated delta method. Here $\hat{\xi}_p$ is once again the MLE of ξ_p .

The results indicate that overall, the methodology based on bootstrapping $\hat{\xi}_p/\xi_p$ provides more accurate performance compared to the BCa percentile bootstrap. However, the bootstrap-calibrated delta method appears to be the most accurate methodology.

TABLE 3 Estimated coverage probabilities of the 95% upper confidence limits for the 95th percentile of a zero-inflated gamma distribution; (i) BCa bootstrap, (ii) Bootstrapping $\hat{\xi}_p/\xi_p$, and (iii) Bootstrap-calibrated delta method

<i>n</i>	π	Shape parameter	Scale parameter = 1.2			Scale parameter = 0.1		
			(i)	(ii)	(iii)	(i)	(ii)	(iii)
40	0.2	1.491825	0.9324	0.9394	0.9460	0.9214	0.9360	0.9452
40	0.3	1.491825	0.9236	0.9510	0.9464	0.9316	0.9362	0.9408
40	0.4	0.670320	0.9358	0.9302	0.9486	0.9232	0.9314	0.9448
40	0.4	1.491825	0.9206	0.9344	0.9462	0.9242	0.9298	0.9468
40	0.5	0.670320	0.9236	0.9304	0.9514	0.9222	0.9318	0.9434
40	0.5	1.491825	0.9106	0.9356	0.9478	0.9292	0.9398	0.9450
80	0.4	1.491825	0.9248	0.9400	0.9444	0.9388	0.9452	0.9472
80	0.5	1.491825	0.9316	0.9510	0.9476	0.9284	0.9430	0.9442
80	0.5	0.670320	0.9316	0.9396	0.9454	0.9400	0.9420	0.9418
80	0.6	1.491825	0.9286	0.9504	0.9458	0.9286	0.9432	0.9458
150	0.4	1.491825	0.9384	0.9410	0.9454	0.9374	0.9414	0.9424
240	0.4	1.491825	0.9410	0.9494	0.9506	0.9432	0.9376	0.9506

<i>n</i>	<i>z</i>	<i>x</i>	(i)	(ii)	(iii)
100	1	0.25	0.9354	0.9412	0.9494
100	0	0.25	0.9310	0.9398	0.9452
200	1	0.25	0.9400	0.9474	0.9470
200	0	0.25	0.9398	0.9456	0.9486

TABLE 4 Estimated coverage probabilities of the 95% upper confidence limits for the 95th percentile of a zero-inflated gamma distribution with shape parameter $\exp(0.2)$ and the covariates specified in Equation (8); (i) BCa bootstrap, (ii) Bootstrapping $\hat{\xi}_p/\xi_p$, and (iii) Bootstrap-calibrated delta method

5 | EXAMPLES

We shall present two examples where we apply our tolerance limit computation. The first example is on a health expenditure study reported in Callahan et al.⁷ The same dataset is also analyzed in Hasan and Krishnamoorthy,¹ Tian,¹¹ and Zhou and Tu.¹² A zero-inflated lognormal distribution is appropriate for these data. Our second example deals with the analysis of a set of body armor data where the zero-inflated gamma distribution with covariates is appropriate.

For both the examples, we implemented the fiducial approach by generating 50 000 values of the fiducial quantity, and bootstrap approaches (including the calibrated delta method) were implemented using 1000 parametric bootstrap samples.

5.1 | Health expenditure data

The data are from a study reported in Callahan et al.⁷ and deal with diagnostic test charges for a group of 40 patients. Among the 40 patients, 10 patients had no diagnostic tests performed at all, and incurred no costs; thus we have zero-inflated data. The same example is also considered in Hasan and Krishnamoorthy,¹ and earlier by Tian¹¹ and Zhou and Tu¹² where a zero-inflated lognormal distribution is used to analyze the data. For this example, we have computed an upper tolerance limit corresponding to $p = 0.95$ and 95% confidence level. Thus 95% or more of the patients in the corresponding population will have their diagnostic test charges below such a tolerance limit, with 95% confidence.

Based on the data, the MLEs are $\hat{\mu} = 6.8535$, $\hat{\sigma}^2 = 1.8696$, and $\hat{\pi} = 0.25$. The upper tolerance limits corresponding to $p = 0.95$ and 95% confidence level are 15202.53 by the fiducial approach and 14217.27 by the bootstrap-calibrated delta

TABLE 5 Upper tolerance limits for the body armor data set for the threat T1 when $p = 0.95$ and a 95% confidence level

Environmental condition	Drop number	Bootstrapping $\hat{\xi}_p/\xi$	Bootstrap calibrated delta method
C1	1	14.222	14.512
C1	2	24.035	23.829
C1	3	8.475	8.543
C2	1	12.173	10.526
C2	2	20.991	21.052
C2	3	7.278	6.597
C3	1	21.674	21.656
C3	2	35.589	35.507
C3	3	12.425	12.976

method. Even though both approaches had satisfactory coverage probabilities, it turns out that the latter provides a smaller upper tolerance limit.

5.2 | A body armor example

In this example, the zero-inflated gamma distribution is applicable and covariates are present. The application deals with the testing of body armor for stab resistance. For an overview of testing practices and background information on this problem, we refer to the technical report by Cavallaro¹³ and to the documents from the National Institute of Justice,^{14,15} National Research Council,^{16,17} and the National Science Academy.¹⁸ The main purpose of the collection and analysis of the relevant data is to assess the performance of new materials, or to evaluate the performance of an existing material line, in order to ensure that there are no changes in the performance. The data consist of stab depths collected based on established testing protocols. The response variable Y is a penetration depth associated with the different types of threat, say X , when the threat penetrated the armor. If the threat did not penetrate the armor, then Y has a value of 0. The material was tested against several different threats and environmental conditions to ensure that the armor was safe. We fit a zero-inflated gamma regression model to analyze and compute the tolerance limits. The smaller the penetration depth is, the safer the armor. Consequently smaller values of the upper tolerance limit are indicative of the armor being safe. Here X represents covariates, and some details are provided later in the article. Armor testing based on upper tolerance limits for binary response data is discussed in Zimmer et al.¹⁹ Due to the confidential nature of the application, the actual data cannot be made available, and only brief details are given here.

To test the stab resistance, the penetration depth resulting from a simulated threat is recorded after it is dropped onto the body armor. If the threat does not pierce the armor, then a value of zero is recorded. The data consisted of 124 data points collected based on three covariates: threat type, sample condition, and drop number. There were two types of threats, labeled here as T1 and T2. The different threats used in this test are surrogates for an actual threat that would be used. The samples were treated under three environmental conditions, labeled as C1, C2, and C3. Each of these conditions are based upon governmental regulations to ensure that the armor maintains its effectiveness in all environments. Lastly on each sample the threat was dropped three times, labeled as S1, S2, and S3; their placement was done at different locations to ensure that the measurements are independent. These drops are also governed by regulations to ensure that consistent testing is performed. The drop heights are regulated to ensure that the armor's stab resistance would provide protection against the force from an actual attacker. We use a zero-inflated gamma regression model to analyze the data; we assume that the proportion π (for an observation to be zero) also depends on the covariates, and this is modeled using logistic regression. If μ denotes the gamma mean, then the fitted regression models for μ and π are given by

$$\begin{aligned}\ln(\mu) &= 1.7808 + 0.2606T2 + 0.4941S2 - 0.5861S3 - 0.1252C2 + 0.4289C3 \\ \text{logit}(\pi) &= -0.6233 + 1.693T2 - 3.3291S2 - 0.3326S3 + 1.448C2 - 2.318C3.\end{aligned}\quad (9)$$

The adequacy of the gamma regression model is assessed using the methodology given in Lin et al,²⁰ and the goodness of fit P -value from the Kolmogorov Smirnov test was 0.989 for the gamma regression. The goodness of fit was assessed using a suitably modified version of the `cumres` function in the R `gof` package. Clearly, the model fit is excellent.

Table 5 gives the upper tolerance limits based on bootstrapping the pivot statistic $\hat{\xi}_p/\xi_p$, and based on the bootstrap calibration of the delta method for the threat type T1, under different environmental conditions and drop numbers. Both methods give comparable results. We also note that there are significant differences among the environmental conditions, and also among the drop numbers. We see that the upper tolerance limit is the largest corresponding to the second drop number, and smallest for the third drop number. Due to the confidential nature of the data, the reasons for the differences noted cannot be explained here.

6 | DISCUSSION

In this article, we have investigated various methodologies for computing upper tolerance limits for zero-inflated log-normal and gamma distributions. The models considered include models without covariates and regression models with covariates. For the zero-inflated lognormal distribution, the methodologies investigated consist of bootstrap approaches, a fiducial approach, and a bootstrap-calibrated delta method. For the zero-inflated gamma distribution, the methodologies investigated consist of bootstrap approaches and a bootstrap-calibrated delta method. The fiducial approach has found numerous applications in the recent literature. We refer to Hannig²¹ for a detailed discussion of the approach. The concept of a fiducial interval was introduced earlier in Weerahandi²² under the name *generalized confidence interval*. However, for the present problems, based on estimated coverage probabilities, our recommendation is the fiducial approach and the bootstrap-calibrated delta method; our overall recommendation is in favor of the latter.

When performance in terms coverage probabilities is satisfactory, the upper tolerance limits can be compared based on their expected values; it is certainly desirable to have smaller expected values. Based on the examples considered, there is some indication that the bootstrap-calibrated delta approach may provide upper tolerance limits having a smaller expected value, at least for the zero-inflated lognormal distribution. A detailed numerical study is necessary in order to draw confirm conclusions. Here we have not undertaken such a numerical study. It should also be noted that the bootstrap-calibrated delta method is applicable to other distributions, for example, a three-parameter gamma distribution. The method was indeed investigated in Zimmer et al²³ for computing upper tolerance limits for a normal mixture distribution, and the resulting solution exhibited satisfactory performance in terms of coverage probabilities.

R codes are available from the authors for computing the tolerance limits proposed in this article.

ACKNOWLEDGEMENTS

The authors are grateful to two reviewers for several helpful suggestions.

ORCID

Thomas Mathew  <https://orcid.org/0000-0003-4152-3628>

REFERENCES


1. Hasan MS, Krishnamoorthy K. Confidence intervals for the mean and a percentile based on zero-inflated lognormal data. *J Stat Simulat Comput*. 2018;88:1499-1514.
2. Krishnamoorthy K, Mathew T. *Statistical Tolerance Regions: Theory, Applications, and Computation*. New York, NY: Wiley; 2009.
3. Lawless JF. *Statistical Model and Methods for Lifetime Data*. 2nd ed. New York, NY: Wiley; 2003.
4. Zhang J, An W. Assessing circular error probable when the errors are elliptical normal. *J Stat Comput Simulat*. 2012;82:565-586.
5. Krishnamoorthy K, Wang X. Fiducial confidence limits and prediction limits for a gamma distribution: censored and uncensored cases. *Environmetrics*. 2016;27:479-493.
6. Krishnamoorthy K, Mathew T, Mukherjee S. Normal-based methods for a gamma distribution: prediction and tolerance intervals and stress-strength reliability. *Technometrics*. 2008;50:69-78.
7. Callahan CM, Kesterson JG, Tierney WM. Association of symptoms of depression with diagnostic test charges. *Ann Internal Med*. 1997;126:426-432.
8. Farewell VT, Prentice RL. A study of distributional shape in life testing. *Technometrics*. 1977;19:69-75.
9. Manning WG, Basu A, Mullahy J. *Modeling costs with generalized gamma regression*. 2002. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.127.2182>.
10. Tu W. Zero-inflated data. In: El-Shaarawi AH, Piegorisch WW, eds. *Encyclopedia of Environmetrics*. Vol 4. Chichester, England: Wiley; 2006:2387-2391.
11. Tian L. Inferences on the mean of zero-inflated lognormal data: the generalized variable approach. *Stat Med*. 2005;24:3223-3232.
12. Zhou XH, Tu W. Confidence Intervals for the mean of diagnostic test charge data containing zeros. *Biometrics*. 2000;56:1118-1125.

13. Cavallaro PV. Soft body armor: an overview of materials, manufacturing, testing, and ballistic impact dynamics. NUWC-NPT Technical Report 12057, 2011.
14. National Institute of Justice. *Stab Resistance of Personal Body Armor. NIJ Standard 0115.00*. Washington, DC: National Institute of Justice; 2000.
15. National Institute of Justice. *Ballistic Resistance of Body Armor. NIJ Standard 0101.06*. Washington, DC: National Institute of Justice; 2008.
16. National Research Council. *Phase I Report on Review of the Testing of Body Armor Materials for use by the U.S. Army*. Washington, DC: The National Academies Press; 2009.
17. National Research Council. *Testing of Body Armor Materials: Phase III*. Washington, DC: National Academies Press; 2012.
18. National Science Academy. *Testing of Body Armor Materials for use by the U.S. Army—Phase II: Letter Report*. Washington, DC: The National Academies Press; 2010.
19. Zimmer Z, Park D, Mathew T. Pointwise and simultaneous tolerance limits under logistic regression. *Technometrics*. 2014;56:282-290.
20. Lin DY, Wei LJ, Ying Z. Model checking techniques based on cumulative residuals. *Biometrics*. 2002;58:1-12.
21. Hannig J. On generalized fiducial inference. *Stat Sinica*. 2009;19:491-544.
22. Weerahandi S. Generalized confidence intervals. *J Am Stat Assoc*. 1993;88:899-905.
23. Zimmer Z, Park D, Mathew T. Tolerance limits under normal mixtures: application to the evaluation of nuclear power plant safety and to the assessment of circular error probable. *Comput Stat Data Anal*. 2016;103:304-315.

AUTHOR BIOGRAPHIES



Zachary Zimmer received a Bachelors in Civil Engineering and Statistics at the University of Illinois at Urbana-Champaign in 2002, a Masters in Statistics from Virginia Commonwealth University in 2006, and a PhD in Statistics from the University of Maryland Baltimore County (UMBC) in 2014. His research interests have focused on methodologies to compute tolerance intervals.



DoHwan Park's research has been concentrated on both theory and applications, comprising primarily methodological developments motivated by biological and medical applications. She is interested in topics including survival and longitudinal data analysis and have continued to work on related problems to date. She has also worked on high-dimensional hypothesis testing problems which have played a crucial role in statistics and related areas such as bioinformatics and genomics.



Thomas Mathew received his PhD in Statistics from the Indian Statistical Institute in 1983. He is currently a professor of statistics at the University of Maryland Baltimore County (UMBC). His research interests have focused on both theoretical and applied statistics. He is the co-author of two books: one of hypothesis testing in mixed effects models, and a second book of statistical tolerance intervals. He is a Fellow of the American Statistical Association and the Institute of Mathematical Statistics.

How to cite this article: Zimmer Z, Park DH, Mathew T. Tolerance limits under zero-inflated lognormal and gamma distributions. *Comp and Math Methods*. 2020;e1113. <https://doi.org/10.1002/cmm4.1113>