2017

A general algorithm for computing simultaneous prediction intervals for the (log)-location-scale family of distributions

Yimeng Xie
Virginia Tech

Yili Hong
Virginia Tech

Luis A. Escobar
Louisiana State University

William Q. Meeker
Iowa State University, wqmeeker@iastate.edu

A general algorithm for computing simultaneous prediction intervals for the (log)-location-scale family of distributions

Abstract

Making predictions of future realized values of random variables based on currently available data is a frequent task in statistical applications. In some applications, the interest is to obtain a two-sided simultaneous prediction interval (SPI) to contain at least *k* out of *m* future observations with a certain confidence level based on *n* previous observations from the same distribution. A closely related problem is to obtain a one-sided upper (or lower) simultaneous prediction bound (SPB) to exceed (or be exceeded) by at least *k* out of *m* future observations. In this paper, we provide a general approach for computing SPIs and SPBs based on data from a particular member of the (log)-location-scale family of distributions with complete or right censored data. The proposed simulation-based procedure can provide exact coverage probability for complete and Type II censored data. For Type I censored data, our simulation results show that our procedure provides satisfactory results in small samples. We use three applications to illustrate the proposed simultaneous prediction intervals and bounds.

Keywords

Censored data, coverage probability, k out of m, lognormal, simulation, Weibull

Disciplines

Statistics and Probability

Comments

This is an Accepted Manuscript of an article published by Taylor & Francis as Xie, Yimeng, Yili Hong, Luis A. Escobar, and William Q. Meeker. "A general algorithm for computing simultaneous prediction intervals for the (log)-location-scale family of distributions." *Journal of Statistical Computation and Simulation* 87, no. 8 (2017): 1559-1576. DOI: 10.1080/00949655.2016.1277426. Posted with permission.

A General Algorithm for Computing Simultaneous Prediction Intervals for the (Log)-Location-Scale Family of Distributions

Yimeng Xie and Yili Hong
Department of Statistics
Virginia Tech
Blacksburg, VA 24061

Luis A. Escobar
Department of Experimental Statistics
Louisiana State University
Baton Rouge, LA 70803

William Q. Meeker
Department of Statistics
Center for Nondestructive Evaluation
Iowa State University
Ames, IA 50011

Abstract

Making predictions of future realized values of random variables based on currently available data is a frequent task in statistical applications. In some applications, the interest is to obtain a two-sided simultaneous prediction interval (SPI) to contain at least k out of m future observations with a certain confidence level based on n previous observations from the same distribution. A closely related problem is to obtain a one-sided upper (or lower) simultaneous prediction bound (SPB) to exceed (or be exceeded) by at least k out of m future observations. In this paper, we provide a general approach for computing SPIs and SPBs based on data from a particular member of the (log)-location-scale family of distributions with complete or right censored data. The proposed simulation-based procedure can provide exact coverage probability for complete and Type II censored data. For Type I censored data, our simulation results show that our procedure provides satisfactory results in small samples. We use three applications to illustrate the proposed simultaneous prediction intervals and bounds.

Key Words: Censored Data; Coverage Probability; k out of m; Lognormal; Simulation; Weibull.

1 Introduction

1.1 Motivation

Prediction intervals are used to quantify the uncertainty associated with future realized values of random variables. In predicting future outcomes, one might be interested in point predictions. Often, however, the focus is on whether the future observations will fall within a prediction interval (PI) or conforming to a one-sided prediction bound (PB) obtained from the available data and a pre-specified confidence level.

In some applications, it is desirable to obtain a two-sided simultaneous prediction interval (SPI) or a one-sided simultaneous prediction bound (SPB) for at least k out of m future observations, where $1 \leqslant k \leqslant m$. For example, Fertig and Mann (1977) consider time to failure of turbine nozzles subject to a certain load. The company had manufactured 50 nozzles. Based on the failure times in a life test of 10 of those nozzles, they obtained a 95% lower prediction bound to be exceeded by at least 90% of the remaining 40 nozzles (i.e., 36 out of 40). In another study, Mann and Fertig (1973) use failure times (in hours) based on a life test of aircraft components to obtain an SPI to contain the failure times of all 10 future components.

SPIs and SPBs are related to tolerance intervals (TIs) and tolerance bounds (TBs), respectively. A tolerance interval is an interval that contains at least a specified proportion (β) of the distribution with a specified degree of confidence of $100(1-\alpha)\%$. SPIs/SPBs are used when the interest is prediction of a small number of future observations. TIs/TBs are used when the interest is to make a statement about a distribution.

Much research has been done for statistical prediction of a single future observation. Details and additional references can be found in Mee and Kushary (1994) and Escobar and Meeker (1999). There has been some work on SPIs/SPBs to contain at least k out of m future observations. Those procedures, however, have been developed only for specific distributions (e.g., normal and Weibull distributions). Hence, it is desirable to have a general approach to compute SPIs/SPBs for a general class of distributions. For highly reliable products, it may be hard to obtain failures for all tested units in a limited time period even under adverse working conditions. Therefore, censoring is common in reliability applications. In this paper, we develop a general algorithm to compute SPIs and SPBs for the location-scale family and the log-location-scale family of distributions. The proposed procedures can be used with complete or censored data and can be extended, in an approximate manner, to other distributions.

1.2 Literature Review and Contributions of This Work

There is some previous work on the computing of SPIs/SPBs to contain/bound at least k out of m future observations. Danziger and Davis (1964) described and provided tables of coverage probabilities for non-parametric SPIs to contain k out of m future observations and corresponding one-sided SPBs. Hahn (1969) considered the special case of $k = m$ for a normal distribution. Hahn (1969) gave the factors to calculate two-sided SPIs. One-sided SPBs were considered in Hahn (1970). Fertig and Mann (1977) presented factors for constructing one-sided SPBs to contain at least k out of m future observations for a normal distribution. Odeh (1990) provided a method for computing k out of m two-sided SPIs for a normal distribution. Due to computational limitations, these papers provided tables of factors for a limited number of combinations of n, k, m and for some specified confidence levels. In the area of environmental monitoring, some articles considered the use of SPIs/SPBs for at least k out of m future observations at p locations. Davis and McNichols (1987) studied this type of problem for one-sided prediction bounds for the normal distribution. This method was implemented in the STATCALC 3.0 software (Krishnamoorthy 2016). Beran (1990) gave theoretical results on the coverage properties of the prediction regions based on simulation. Kushary (1996) constructed prediction intervals for location-scale families based on Type II censored data. Lawless (2003) provided a general approach for finding confidence limits and prediction limits for distributions in the location-scale family. Bhaumik and Gibbons (2006) developed an approximate upper SPB for the gamma distribution. Wu and Lu (2007) developed prediction intervals for ordered observations from the logistic distribution based on censored data. Bhaumik et al. (2008) constructed a one-sided SPB for left-censored data from the normal distribution. Krishnamoorthy, Lin, and Xia (2009) constructed one-sided upper prediction bounds for the Weibull distribution based on generalized pivotal quantities. Wang, Hannig, and Iyer (2012) presented an approach for constructing prediction intervals for any given distribution based on the principle of fiducial inference.

In this paper, we provide a unified approach for SPIs/SPBs for the location-scale (e.g., the smallest extreme value, the normal, and the largest extreme value distributions) and the related log-location-scale family of distributions (e.g., the Weibull, the lognormal, and the Fréchet distributions). The proposed methods are exact (except for Monte Carlo error) for complete and Type II censored data. We use simulation to study the coverage properties for the proposed interval/bound procedures with Type I censored data, which is also common in life tests. The proposed methods will be useful for practitioners who may need a generic algorithm for computing SPIs/SPBs for different distributions in the (log)-location-scale family of distributions.

1.3 Overview

The rest of this paper is organized as follows. Section 2 introduces the data and model setting for the problem. Section 3 gives the formal definition of the proposed SPI/SPB procedure and shows the connection between SPI/SPBs and TI/TBs. Section 4 proposes a general algorithm to compute an SPI or an SPB, followed by illustrative examples. Section 5 describes simulation studies that evaluate the performance of the proposed procedure for Type I censored data. Section 6 illustrates the use of the proposed method with applications. Section 7 contains concluding remarks and some discussion about related extensions and applications of the methods.

2 Data, Model, and Maximum Likelihood Estimation

2.1 Data

We consider situations in which there are n independent observations. The data consist of: (a) r exact observations and (b) a set of $(n-r)$ right-censored observations at x_c, where x_c is larger or equal to the maximum of the exact observations. Three important special cases of these data structure are: (a) complete data, when $r = n$; (b) Type II censored data, when r $(2 \leqslant r \leqslant n)$ is pre-specified and x_c is equal to the maximum of the exact observations. Note that in this case x_c is random; (c) Type I censored data, when x_c is pre-specified and x_c is greater than or equal to the largest observation. Note that in the case of Type I censoring, r $(1 \leqslant r \leqslant n)$ is random. When $r = 0$, the maximum likelihood (ML) estimate does not exist. Note that for Type I censoring, the ML estimate exists when $r = 1$ because the proportion failing at the fixed censoring time gives extra information.

To be precise, let $\boldsymbol{X} = (X_1, \ldots, X_n)$ denote the random variables for the observations from the n units, where $-\infty < X_i < \infty, i = 1, \ldots, n$. Define

$$\delta_i = \begin{cases} 1, & \text{if } X_i \text{ is an exact observation} \\ 0, & \text{if } X_i \text{ is a right-censored observation,} \end{cases} \tag{1}$$

to indicate the censoring status. For Type I and Type II censoring, the data are $x_i = \min(X_i, x_c)$ and δ_i, $i = 1, \ldots, n$. The observed values are denoted by $\boldsymbol{x} = (x_1, \ldots, x_n)$. This data structure is general and includes data from reliability and lifetime studies with right-censored data from a positive response. For log-location-scale distributions, all the components of $\boldsymbol{X}$ take positive values.

Table 1: The pdfs and cdfs of different commonly-used members of the standard location-scale and log-location-scale distributions.

Location-Scale	Log-Location-Scale	pdf $\phi(x)$	cdf $\Phi(x)$
Normal	Lognormal	$\dfrac{\exp(-x^2/2)}{\sqrt{2\pi}}$	$\int_{-\infty}^{x} \phi(w)\, dw$
Logistic	Loglogistic	$\dfrac{\exp(x)}{[1+\exp(x)]^2}$	$\dfrac{\exp(x)}{1+\exp(x)}$
Largest extreme value	Fréchet	$\exp[-x-\exp(-x)]$	$\exp[-\exp(-x)]$
Smallest extreme value	Weibull	$\exp[x-\exp(x)]$	$1-\exp[-\exp(x)]$

2.2 Model

To construct an SPI or an SPB for a set of future observations, we use a statistical model to describe the population of interest. In this paper, we assume the observations have a distribution in the location-scale or log-location-scale family of distributions. In a location-scale distribution, the location parameter μ and the scale parameter σ are typically unknown and need to be estimated. The probability density function (pdf) and the cumulative distribution function (cdf) of a location-scale distribution are

$$f(x) = \frac{1}{\sigma}\phi\left(\frac{x-\mu}{\sigma}\right) \quad \text{and} \quad F(x) = \Phi\left(\frac{x-\mu}{\sigma}\right),$$

respectively. The definitions of the standard pdf $\phi(\cdot)$ and cdf $\Phi(\cdot)$ functions for the different members of this family are given in Table 1.

The pdf and cdf of the log-location-scale family are

$$f(t) = \frac{1}{\sigma t}\phi\left[\frac{\log(t)-\mu}{\sigma}\right] \quad \text{and} \quad F(t) = \Phi\left[\frac{\log(t)-\mu}{\sigma}\right],$$

respectively. The Weibull, lognormal, Fréchet, and log-logistic distributions are members of the log-location-scale family. For these distributions, σ is a shape parameter and $\exp(\mu)$ is a scale parameter. In the reminder of this paper, however, we will refer to μ and σ as location and scale parameters, respectively.

This paper focuses on the construction of SPIs and SPBs containing at least k of m future observations $\boldsymbol{Y} = (Y_1, \ldots, Y_m)$ from a previously sampled population. The sample data are denoted by $\boldsymbol{X}$. The assumptions are that $\boldsymbol{Y}$ and $\boldsymbol{X}$ are independent random samples from the same distribution.

2.3 Maximum Likelihood Estimation

We use maximum likelihood (ML) to estimate the unknown parameters (μ, σ). Under the independent and identically distributed (i.i.d.) assumptions in Section 2.2, the likelihood of the right censored data has the form

$$L(\mu, \sigma) = \mathcal{C} \prod_{i=1}^{n} [f(x_i; \mu, \sigma)]^{\delta_i} [1 - F(x_i; \mu, \sigma)]^{1-\delta_i},$$

where $\mathcal{C}$ is a constant that does not depend on μ or σ, $f(x_i; \mu, \sigma)$ is the assumed pdf, and $F(x_i; \mu, \sigma)$ is the corresponding cdf. Here, δ_i is the censoring indicator as defined in (1). The ML estimates are the values of μ and σ that maximize the likelihood function. In general, there is no closed-form expression for the ML estimates, which are denoted by $(\widehat{\mu}, \widehat{\sigma})$. Consequently, numerical methods are used to find the ML estimates.

2.4 Distribution Specification

When using the (log)-location-scale family of distributions to describe data, a critical step is to choose the distribution that best fits the data. In some applications, knowledge of the data-generating mechanism (e.g., physics of failure) will suggest the appropriate distribution. For example if failure is due to fracture from fatigue in ductile materials (e.g., most metal alloys), the lognormal distribution is known to provide an appropriate model because of the cumulative damage type of mechanism (e.g., Chapter 11 of Meeker and Escobar 1998). On the other hand, the Weibull distribution is a good approximation to the distribution of minima based on extreme value theory (e.g., Chapter 5 of Castillo 1988). For example, if failure is due to fracture from fatigue in a brittle material (such as ceramics), failure is determined by the weakest part of the ceramic structure and the Weibull distribution provides an appropriate model for time to fracture. Beyond such physical motivation, for a new dataset, probability plots with nonparametric simultaneous confidence band (e.g., Nair 1984) are useful to help select an appropriate distribution. Alternatively, one can do a goodness of fit test such as the Kolmogorov-Smirnov test to assess the statistical significance (e.g., Barr and Davidson 1973 and Dufour and Maag 1978).

3 Simultaneous Prediction Intervals and Bounds

3.1 Two-sided Simultaneous Prediction Intervals

This section shows how to construct an SPI $[L(\boldsymbol{x}, 1-\alpha),\ U(\boldsymbol{x}, 1-\alpha)]$ that will contain at least k out of m independent future observations from the sampled distribution, with a specified confidence level $1-\alpha$. Conditioning on the observed data $\boldsymbol{X}=\boldsymbol{x}$, the conditional coverage probability (CP) of the interval $[L(\boldsymbol{x},1-\alpha),\ U(\boldsymbol{x},1-\alpha)]$ with nominal confidence level $1-\alpha$ is

$$\begin{aligned}\mathrm{CP}(\boldsymbol{\theta}|\boldsymbol{X}=\boldsymbol{x}) &= \Pr\{\text{at least } k \text{ of } m \text{ values lie in } [L(\boldsymbol{x},1-\alpha),\ U(\boldsymbol{x},1-\alpha)]\,|\boldsymbol{X}=\boldsymbol{x}\} \\ &= \sum_{j=k}^{m}\binom{m}{j}p^{j}(1-p)^{m-j}, \qquad (2)\end{aligned}$$

where $\boldsymbol{\theta}=(\mu,\sigma)$ is the vector of unknown parameters and

$$p=\Pr\{\text{a future observation is in } [L(\boldsymbol{x},1-\alpha),\ U(\boldsymbol{x},1-\alpha)]\,|\boldsymbol{X}=\boldsymbol{x}\}.$$

The conditional CP is unobservable because it depends on the unknown parameters and varies from sample to sample because it depends on the data. To evaluate the prediction interval procedure, we use the unconditional CP

$$\mathrm{CP}(\boldsymbol{\theta})=\mathrm{E}_{\boldsymbol{X}}\left[\sum_{j=k}^{m}\binom{m}{j}p^{j}(1-p)^{m-j}\right],$$

where the expectation is taken with respect to the joint distribution of the data $\boldsymbol{X}$.

Because the $(Y_i-\widehat{\mu})/\widehat{\sigma}, i=1,\dots,m$ are pivotal quantities, one can construct a two-sided $100(1-\alpha)\%$ SPI to contain at least k out of m future observations with the following form

$$[\,\widehat{\mu}+u_L(k,m;\alpha)\widehat{\sigma},\ \widehat{\mu}+u_U(k,m;\alpha)\widehat{\sigma}\,],$$

for the location-scale family of distributions. Here $u_L(k,m;\alpha)$ and $u_U(k,m;\alpha)$ are factors to be chosen so that the SPI will have CP equal to $1-\alpha$. For notational simplicity, we let $u_L=u_L(k,m;\alpha)$ and $u_U=u_U(k,m;\alpha)$. The following gives the result to determine the factors (u_L,u_U).

Result 1 *For two-sided* $100(1-\alpha)\%$ *SPI, the factors* (u_L,u_U) *satisfy the equation*

$$1-\alpha=\int_{0}^{\infty}\int_{-\infty}^{\infty}\sum_{j=k}^{m}\binom{m}{j}\left[\Phi(a)-\Phi(b)\right]^{j}\left[1-\Phi(a)+\Phi(b)\right]^{m-j}f_{\boldsymbol{Z}}(z_1,z_2)\,dz_1dz_2, \qquad (3)$$

where $a=z_1+u_Uz_2$, $b=z_1+u_Lz_2$, $\boldsymbol{Z}=(Z_1,Z_2)$, $f_{\boldsymbol{Z}}(z_1,z_2)$ *is the joint pdf of* $Z_1=(\widehat{\mu}-\mu)/\sigma$ *and* $Z_2=\widehat{\sigma}/\sigma$, *and* $\Phi(\cdot)$ *is the standard cdf of* X. *Note that (3) can be written as*

$$1-\alpha=\mathbf{E}_{\boldsymbol{Z}}\left[\sum_{j=k}^{m}\binom{m}{j}\left[\Phi(A)-\Phi(B)\right]^{j}\left[1-\Phi(A)+\Phi(B)\right]^{m-j}\right], \qquad (4)$$

where $A=Z_1+u_UZ_2$, $B=Z_1+u_LZ_2$, *and* $\mathbf{E}_{\boldsymbol{Z}}(\cdot)$ *is the expectation with respect to the joint distribution of* $\boldsymbol{Z}$.

The proof of **Result 1** is given in Appendix A. For distributions in the log-location-scale family, the corresponding two-sided $100(1-\alpha)\%$ SPI to contain at least k out of m future observations has the form $[\exp(\widehat{\mu}+u_L\widehat{\sigma}),\ \exp(\widehat{\mu}+u_U\widehat{\sigma})]$. Thus, (3) can still be used to obtain a prediction interval for distributions in the log-location-scale family.

The details of the computations are given in Section 4. For complete or Type II censored data from the location-scale/log-location-scale family of distributions, the quantities $\boldsymbol{Z}$ are pivotal quantities (e.g., Lawless 2003, pages 217 and 262). That is, the distribution of $\boldsymbol{Z}$ does not depend on unknown parameters, which implies that the procedure based on (3) will provide an exact CP for the SPIs/SPBs.

For Type I censoring, the pivotal property of $\boldsymbol{Z}$ no longer holds but we have the following result.

Result 2 *Under Type I censoring, the joint distribution of* $Z_1 = (\widehat{\mu}-\mu)/\sigma$ *and* $Z_2 = \widehat{\sigma}/\sigma$ *depends only on the expected fraction failing* $p_f = \Phi[(x_c-\mu)/\sigma]$.

The proof of **Result 2** is given in Appendix A.2. In practice, p_f is usually unknown but one can use $\widehat{p}_f$ to approximate p_f. Thus, the quantities $\boldsymbol{Z}$ can be treated as approximate pivotal quantities. Thus we can still use (3) to obtain the asymptotically correct SPIs under Type I censoring, and other types of non-informative censoring.

3.2 One-sided Simultaneous Prediction Bounds

There are similar CP statements for one-sided simultaneous prediction bounds. In particular, for a one-sided lower simultaneous prediction bound, the conditional CP is

$$\begin{aligned} \text{CP}_L(\boldsymbol{\theta}|\boldsymbol{X}=\boldsymbol{x}) &= \Pr\left[\text{at least } k \text{ of } m \text{ values are larger than } L(\boldsymbol{x}, 1-\alpha)|\boldsymbol{X}=\boldsymbol{x}\right] \\ &= \sum_{j=k}^{m}\binom{m}{j}p_L^j(1-p_L)^{m-j}, \end{aligned} \tag{5}$$

where $p_L = \Pr\left[\text{a single future observation is larger than } L(\boldsymbol{x}, 1-\alpha)|\boldsymbol{X}=\boldsymbol{x}\right]$.

The unconditional CP is

$$\text{CP}_L(\boldsymbol{\theta}) = \mathbf{E}_{\boldsymbol{X}}\left[\sum_{j=k}^{m}\binom{m}{j}p_L^j(1-p_L)^{m-j}\right].$$

For the location-scale family of distributions, a one-sided lower simultaneous prediction bound to be exceeded by at least k out of m future observations can be expressed as $L(\boldsymbol{x}, 1-\alpha) = \widehat{\mu} + u'_L(k,m;\alpha)\widehat{\sigma}$, where $u'_L(k,m;\alpha)$ is a factor to be chosen so that the interval will give a CP of $1-\alpha$. Let $u'_L = u'_L(k, m; \alpha)$ and note that u'_L satisfies the equation

$$1-\alpha = \int_0^\infty \int_{-\infty}^\infty \sum_{j=k}^{m} \binom{m}{j} [1-\Phi(b)]^j [\Phi(b)]^{m-j} \; f_{\boldsymbol{Z}}(z_1, z_2)\, dz_1 dz_2$$
$$= \mathbf{E}_{\boldsymbol{Z}} \left[\sum_{j=k}^{m} \binom{m}{j} [1-\Phi(B)]^j [\Phi(B)]^{m-j} \right], \qquad (6)$$

where $b = z_1 + u'_L z_2$ and $B = Z_1 + u'_L Z_2$. When $k = m$, one obtains the lower prediction bound to contain all m new additional observations.

Similarly, for a one-sided upper simultaneous prediction bound, the conditional CP is

$$\mathrm{CP}_U(\boldsymbol{\theta}|\boldsymbol{X}=\boldsymbol{x}) = \Pr\left[\text{at least } k \text{ of } m \text{ values are less than } U(\boldsymbol{x}, 1-\alpha) | \boldsymbol{X}=\boldsymbol{x}\right]$$
$$= \sum_{j=k}^{m} \binom{m}{j} p_U^j (1-p_U)^{m-j}, \qquad (7)$$

where

$$p_U = \Pr\left[\text{a single future observation is less than } U(\boldsymbol{x}, 1-\alpha) | \boldsymbol{X}=\boldsymbol{x}\right].$$

The unconditional CP is

$$\mathrm{CP}_U(\boldsymbol{\theta}) = \mathbf{E}_{\boldsymbol{X}} \left[\sum_{j=k}^{m} \binom{m}{j} p_U^j (1-p_U)^{m-j} \right].$$

A one-sided upper simultaneous prediction bound to exceed at least k out of m future observations for the location-scale family of distributions is $U(\boldsymbol{x}, 1-\alpha) = \widehat{\mu} + u'_U(k, m; \alpha)\widehat{\sigma}$, where $u'_U(k, m; \alpha)$ is a factor to be chosen so that the interval will give a CP equal to $1-\alpha$. Let $u'_U = u'_U(k, m; \alpha)$ and note that u'_U satisfies the equation

$$1-\alpha = \int_0^\infty \int_{-\infty}^\infty \sum_{j=k}^{m} \binom{m}{j} [\Phi(a)]^j [1-\Phi(a)]^{m-j} \; f_{\boldsymbol{Z}}(z_1, z_2)\, dz_1 dz_2$$
$$= \mathbf{E}_{\boldsymbol{Z}} \left[\sum_{j=k}^{m} \binom{m}{j} [\Phi(A)]^j [1-\Phi(A)]^{m-j} \right], \qquad (8)$$

where $a = z_1 + u'_U z_2$ and $A = Z_1 + u'_U Z_2$.

For the log-location-scale family of distributions, the lower and upper SPBs have the form $L(\boldsymbol{x}, 1-\alpha) = \exp(\widehat{\mu} + u'_L\widehat{\sigma})$ and $U(\boldsymbol{x}, 1-\alpha) = \exp(\widehat{\mu} + u'_U\widehat{\sigma})$, respectively. The factors u'_L and u'_U are obtained as solutions of (6) and (8), respectively.

3.3 Relationship with Tolerance Interval/Bound

SPIs/SPBs quantify the uncertainty in future samples from a previously-sampled distribution. TIs/TBs cover a specific probability content of a previously-sampled distribution. There is an interesting relationship between these two kinds of intervals/bounds. The following result states the asymptotic equivalence of SPIs/SPBs to contain at least k out of m future observations and TIs/TBs to contain at least a fraction $\beta = k/m$ of the distribution.

Result 3 *When $k \to \infty$, $m \to \infty$, and $0 < (k/m) = \beta < 1$, the two-sided SPI in (3) converges to a TI that, with confidence $100(1-\alpha)\%$, will contain at least a proportion β of the distribution. This property also holds for the one-sided SPBs given in (6) and (8).*

The proof is given in Appendix A.3.

4 Computations of SPIs/SPBs

In this section, we introduce a general procedure for finding the factors so that the two-sided SPIs and one-sided SPBs will have the correct CP. The computing procedure requires solving equations (4), (6), and (8). In general, there is no closed-form expression for the solution of these equations. The exact distribution of $\boldsymbol{Z}$ can be complicated, especially with censored data. Therefore, we use Monte Carlo simulation to obtain the distribution of $\boldsymbol{Z}$ and evaluate the expectation based on the simulated samples.

4.1 Complete and Type II Censored Data

The two-sided SPI for complete or Type II censored data can be obtained from the following algorithm.

Algorithm 1:

1. Draw a complete or Type II censored sample of size n from a (log)-location-scale family of distributions with $(\mu, \sigma) = (0, 1)$. Detailed discussion on efficient simulation of censored samples can be found in Meeker and Escobar (1998, Section 4.13).

2. Repeat step 1 B_1 times and compute ML estimates $(\widehat{\mu}_l^*, \widehat{\sigma}_l^*)$ for each simulated sample, $l = 1, \ldots, B_1$.

   Note that for data with Type II censoring, the proportion of censoring is known to be r/n. The distribution of pivotal quantities depends only on sample size n and number of failures r. To save computing time, these $(\widehat{\mu}_l^*, \widehat{\sigma}_l^*)$ values are stored and used to compute all the SPIs and SPBs for the particular censoring specification (n, r) as shown below.

3. For every (u_L, u_U), in a collection of chosen values, compute

$$\mathrm{CP}^*(u_L, u_U) = \frac{1}{B_1}\sum_{l=1}^{B_1}\left\{\sum_{j=k}^{m}\binom{m}{j}p_l(u_L, u_U)^j[1 - p_l(u_L, u_U)]^{m-j}\right\}, \tag{9}$$

   where $p_l(u_L, u_U) = \Phi(\widehat{\mu}_l^* + u_U\widehat{\sigma}_l^*) - \Phi(\widehat{\mu}_l^* + u_L\widehat{\sigma}_l^*)$ and $u_L < u_U$.

4. Find (u_L, u_U) such that $\mathrm{CP}^*(u_L, u_U) = 1 - \alpha$.

Note that the choice of $(\mu, \sigma) = (0, 1)$ in Step 1 above is justified because for the Type II censored and complete data case, the **Algorithm 1** procedure does not depend on unknown parameters due to the pivotal property of $\boldsymbol{Z}$.

Finding (u_L, u_U) such that $\mathrm{CP}^*(u_L, u_U) = 1-\alpha$ is a two-dimensional root-finding problem and there are multiple solutions. An additional constraint on u_L and u_U is needed for a unique solution. For symmetric distributions, $u_L = -u_U$ is an appropriate constraint and leads to two-sided SPIs with equal error probabilities in both tails. For non-symmetric distributions, the two-sided SPI with equal error probabilities in both tails is appealing from a practical point of view. The computation, however, is more complicated. Detailed discussion of the computation is given in Section 4.2.

For one-sided SPBs, modifications to the algorithm are needed. Specifically, for the lower SPB, replace (9) with

$$\mathrm{CP}_L^*(u_L') = \frac{1}{B_1}\sum_{l=1}^{B_1}\left\{\sum_{j=k}^{m}\binom{m}{j}p_l(u_L')^j[1 - p_l(u_L')]^{m-j}\right\},$$

where $p_l(u_L') = 1-\Phi(\widehat{\mu}_l^*+u_L'\widehat{\sigma}_l^*)$. Then find the unique value of u_L' such that $\mathrm{CP}_L^*(u_L') = 1-\alpha$. For the upper SPB, replace (9) with

$$\mathrm{CP}_U^*(u_U') = \frac{1}{B_1}\sum_{l=1}^{B_1}\left\{\sum_{j=k}^{m}\binom{m}{j}p_l(u_U')^j[1 - p_l(u_U')]^{m-j}\right\},$$

where $p_l(u_U') = \Phi(\widehat{\mu}_l^* + u_U'\widehat{\sigma}_l^*)$. Then find the unique value of u_U' such that $\mathrm{CP}_U^*(u_U') = 1-\alpha$. For one-sided prediction bounds, numerical root finding methods such as the bi-section method can be used to find the factors for SPBs based on the CP curve $[(1-\alpha)$ versus u_L' or u_U', respectively$]$ for desired confidence levels.

4.2 Two-sided SPI with Equal Error Probabilities in Both Tails

In applications, even involving a non-symmetric distribution, it is preferable to have a two-sided prediction interval with equal error probabilities in both tails. For this purpose, we define the error probability for each tail as the tail error probability of the corresponding one-sided bound. Therefore, the equal error probabilities in both tails implies that $\mathrm{CP}_L(u_L) = \mathrm{CP}_U(u_U)$. Except for the special case of $k = 1$ (i.e., a prediction interval for exactly one new observation), combining a one-sided lower $100(1-\alpha_1)\%$ prediction bound and a one-sided upper $100(1-\alpha_2)\%$ prediction bound will not provide a two-sided $100(1-\alpha_1-\alpha_2)\%$ SPI. Thus, a special procedure for a two-sided SPI with equal error probabilities in both tails is needed. For a given confidence level $1-\alpha$, we can obtain u_L and u_U by solving numerically the equations

$$\mathrm{CP}(u_L, u_U) = 1 - \alpha \quad \text{and} \quad \mathrm{CP}_L(u_L) - \mathrm{CP}_U(u_U) = 0. \tag{10}$$

To find the solutions to (10), based on the values of $\mathrm{CP}(u_L, u_U)$ and $\mathrm{CP}_L(u_L) - \mathrm{CP}_U(u_U)$ on a grid values of u_L and u_U, one can find numerically the contour lines of the functions $\mathrm{CP}(u_L, u_U)$ and $\mathrm{CP}_L(u_L) - \mathrm{CP}_U(u_U)$. Then the solution to (10) is the intersecting point of the contour line $\mathrm{CP}(u_L, u_U) = 1-\alpha$ and $\mathrm{CP}_L(u_L) - \mathrm{CP}_U(u_U) = 0$. Alternatively, one can re-express the two-sided CP as a function of the one-sided error probability to reduce the dimension of root-finding, and then find the common error probability that gives the desired two-sided CP. Illustration of this method is given in Section 4.5.

4.3 Type I Censored Data

For Type I censored data, the statistics Z_1 and Z_2 are no longer pivotal and their distributions depend on the fraction failing p_f. One can use $\widehat{p}_f$ to approximate p_f. Therefore, for Type I censoring, we use the following algorithm.

Algorithm 2:

1. For the observed Type I data, calculate the ML estimates $(\widehat{\mu}, \widehat{\sigma})$ and then compute estimated expected fraction failing $\widehat{p}_f = \Phi[(x_c - \widehat{\mu})/\widehat{\sigma}]$.

2. Draw a censored sample of size n from the (log)-location-scale family of distributions with $(\mu, \sigma) = (0, 1)$ and censoring time is $x_c = \Phi^{-1}(\widehat{p}_f)$.

3. Follow steps 2 to 4 in **Algorithm 1**.

If we fix the censoring time x_c and let sample size n increase to infinity, $\widehat{p}_f$ will converge to p_f. Thus the CP of SPIs/SPBs computed by **Algorithm 2** will approach the nominal confidence level as n increases. Also, **Algorithm 2** can be viewed as a "parametric bootstrap" approach. In Section 5, we study finite sample CPs for SPIs and SPBs obtained using **Algorithm 2**.

4.4 Controlling Monte Carlo Errors

For the number of Monte Carlo samples needed to give the desired precision (i.e., how to choose B_1), we quantify the Monte Carlo error in estimating the CP. For the two-sided SPI procedure, we have the following result.

Result 4 *Let* $g(u_L, u_U) = \sum_{j=k}^{m} \binom{m}{j} p_l(u_L, u_U)^j [1 - p_l(u_L, u_U)]^{m-j}$ *and let* I_o *be the Fisher information matrix evaluated at* $\boldsymbol{\theta} = (0, 1)$. *By the delta method, we have*

$$\mathrm{Var}[\mathrm{CP}^*(u_L, u_U)] = \frac{1}{B_1} \left(\frac{\partial g}{\partial \widehat{\mu}_l^*}, \frac{\partial g}{\partial \widehat{\sigma}_l^*} \right) I_o^{-1} \left(\frac{\partial g}{\partial \widehat{\mu}_l^*}, \frac{\partial g}{\partial \widehat{\sigma}_l^*} \right)' \Bigg|_{(\widehat{\mu}_l^*, \widehat{\sigma}_l^*) = (0,1)} .$$

The proof of **Result 4** is in Appendix A.4. The result also shows that the variance of estimated CP $\mathrm{CP}^*(u_L, u_U)$ is a function of $u_L, u_U, k, m, n, x_c, \Phi(\cdot)$ and B_1 for Type I censored data, and is a function of u_L, u_U, $k, m, n, r, \Phi(\cdot)$ and B_1 for Type II censored data. We have calculated the variance of $\mathrm{CP}^*(u_L, u_U)$ for different combinations of (k, m, x_c, r) and different distributions over a large grid of factors (u_L, u_U). If the expected number of failures is not too small (> 10), $B_1 = 10{,}000$ is roughly enough for $\mathrm{CP}^*(u_L, u_U)$ to be accurate at the first two decimals. We suggest more samples (100,000 or more) if the expected number of failures is small (< 10). Similar results hold for one-sided SPBs. In the literature, variance reduction techniques, such as the control variates variance method, are available. **Result 4**, however, shows that we can control the Monte Carlo error well by using large B_1. With today's computing power, the computation time is usually not a burden.

4.5 Illustrations of SPI/SPB Computing

• **Illustration A**: Upper SPB for Type II Censoring and Complete Data. For purpose of illustration, we generate the CP curve for a one-sided upper SPB for at least 4 out of 5 future observations from a previously sampled Weibull distribution. The sample size is $n = 20$ and we consider the Type II censored configurations corresponding to $r = 5, 10, 15$, and 20 (complete data case). The number of simulations B_1 is set to be 100,000 so that the results are stable (i.e., negligible Monte Carlo error). Figure 1 shows the CP as a function of

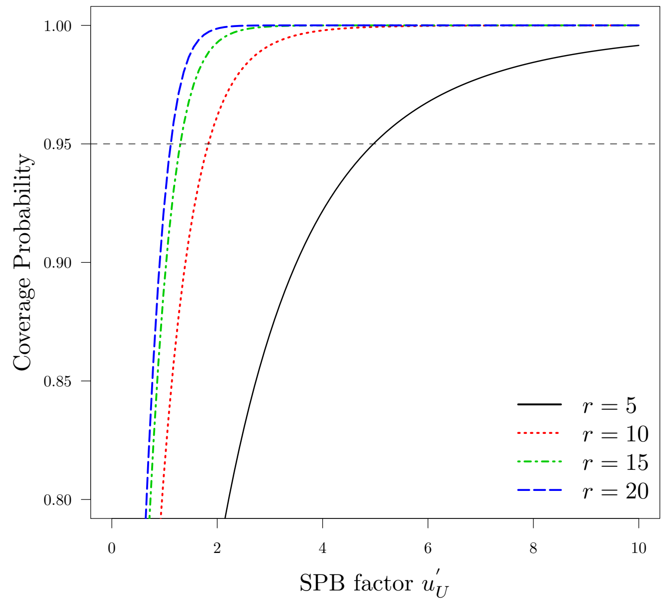

Figure 1: CP curves of one-sided upper SPBs for $n = 20, r = 5, 10, 15$ and $20, k = 4$, and $m = 5$ based on **Algorithm 1**.

u'_U and r. For a desired coverage level, say $1 - \alpha = 0.95$ and a specific value of r, the value of u'_U is determined from the CP curve corresponding to the specified r value.

• **Illustration B**: Two-sided SPI with Equal Error Probabilities in Both Tails. Here the objective is to construct a two-sided SPI from a previously sampled Weibull distribution. Again, B_1 is chosen to be 100,000. The contour plot of the CP as a function of u_L and u_U is shown in Figure 2(a). To obtain the prediction interval with equal error probability in each tail, we solve the equations in (10). Figure 2(b) shows the contour lines of the two equations. The upper and lower limits u_U and u_L of the 95% SPI with equal error probabilities in both tails are the coordinates of the intersection point of the two non-linear curves in Figure 2(b). In Figure 2(b), the coordinates $(u_L, u_U) = (-6.46, 3.82)$ produce both the 0.95 overall CP and the equal error probabilities in both tails.

5 Simulation Study for Type I censoring

This section studies the CP properties of the simulation-based procedure proposed in Section 4.3 in small samples, with Type I censored samples from a Weibull with $(\mu, \sigma) = (0, 1)$.

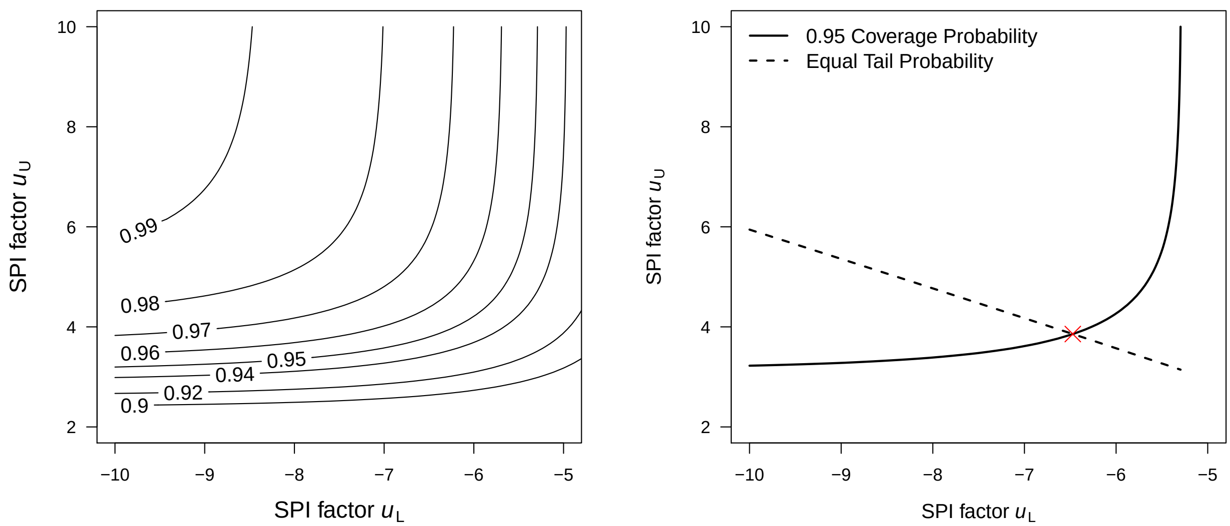

(a) Coverage Probability Plot (b) 0.95 Coverage and Equal Error Probability Plot

Figure 2: (a) Contour plot of the CP as a function of u_L and u_U for $n = 20, r = 8, k = m = 3$, based on 100,000 simulations. (b) The contour lines of equation (10).

5.1 Simulation Setting

For the Type I censored data case, the procedure properties will depend on unknown parameters through the expected fraction failing. The CP of the SPIs/SPBs, however, will converge to the nominal confidence level as the sample size increases to infinity. Here we study the effect of the expected number of failures $n_f = n \times p_f$ on the CP of the SPIs/SPBs in small samples. Similar simulation designs can be found in Vander Weil and Meeker (1990) and Jeng and Meeker (2000).

5.2 Simulation Algorithm

In **Algorithm 2**, we calculated the SPIs/SPBs based on the ML estimates $(\widehat{\mu}, \widehat{\sigma})$, which are determined from the observed data. To evaluate the performance of **Algorithm 2**, we simulate the data many times and average over the results. The detailed simulation plan is as follows.

1. For $l = 1$ to B_2 do steps 2-4:

2. Simulate $\boldsymbol{X} = (X_1, X_2, \ldots, X_n)$ with the pre-determined censoring time from the

Weibull distribution with parameters $(\mu, \sigma) = (0, 1)$. Then, calculate the ML estimates of (μ, σ) and the corresponding estimated expected fraction failing.

3. Use **Algorithm 2** to obtain the SPIs/SPBs. For example, we obtain the one-sided upper SPB by computing u'_U.

4. Evaluate the conditional CP of the interval in step 3. The conditional CP is defined in (2), (5), and (7) for SPI, lower SPB, and Upper SPB, respectively. Note that this evaluation is conditional on the particular dataset obtained in Step 2. It is computationally much more efficient to evaluate the unconditional CP of a prediction interval by averaging these conditional CP, as opposed to averaging the binary outcomes of whether particular intervals cover or not.

5. Average the B_2 conditional CP to estimate the CP of the procedure.

Because the focus is on the CP for small sample sizes, we simulate datasets with the expected number of failures equal to $n_f = 5, 7, 10, 25$ and the expected fraction failing equal to $p_f = 0.25$. Here we chose $B_2 = 500$.

5.3 Simulation Results

Figure 3 displays the estimated actual CP versus the nominal confidence level for the one-sided lower and upper SPBs, and the two-sided SPI. Figure 3 shows that there are some deviations from the nominal CP when the expected number of failures n_f is small (around 10). The estimated actual CP is close to the nominal confidence level when n_f is large enough (e.g., around 25). In the case of $n_f = 25$, the corresponding line is nearly the same as the identity line. When n_f is large, the observed data tends to have more failures, thus the estimates are more accurate and the SPIs/SPBs have better CP. We also note that the two-sided SPI tends to perform better than one-sided SPBs when n_f is small. As indicated earlier, we used $(\mu, \sigma) = (0, 1)$ in the simulation. For other values of (μ, σ), the simulation results are similar because they depend on the expected number of failures. Overall, **Algorithm 2** provides satisfactory results for Type I censoring in finite samples when the expected number of failures is at least 5.

From **Result 2**, we know that the accuracy of SPI/SPB obtained from **Algorithm 2** depends on how close the estimated expected fraction failing is to the expected fraction failing. To avoid the interaction and confounding of sample size and expected fraction failing, we fix the expected fraction failing ($p_f = 0.25$) in the simulation study and use different expected number of failures ($n_f = 5, 7, 15, 25$).

Note that the CP of the SPI tends to be closer to the nominal confidence levels when compared to the CP of SPBs. It is possible that the asymmetry of the underlying distribution is related to this fact. In general, the performance of the SPI under Type I censoring depends on the performance of the "approximate pivotal quantities." When the expected number of failures is relatively large, the CP values are very close to the nominal confidence levels.

We also performed additional simulations for $p_f = 0.10$ and 0.50, using different values of k and m, and distributions other than the Weibull distribution. The results are presented in the online supplementary materials. Overall, the results are similar and consistent with results presented here.

6 Applications

In this section, we use three examples to illustrate the applicability of the proposed procedure.

6.1 Nozzle Failure Time Data

This example is adapted from the application described in Fertig and Mann (1977). They wanted to compute a 95% lower prediction bound (they called it a "warranty period") of the failure times of at least 36 or 40 out of 40 nozzles. They provided the sample mean and sample standard derivation of the logarithm of failure times (which they assumed to have normal distribution) of 10 nozzles, which are $\widehat{\mu} = 3.850$ and $\widehat{\sigma} = 0.034$, respectively. Applying **Algorithm 1**, we found that the lower SPBs to be exceeded by at least $k = 36$ and $k = 40$ out of $m = 40$ nozzles are, respectively, 43.35 and 40.96 hours (based on 100,000 Monte Carlo trials).

6.2 Aircraft Component Failure Time Data

Mann and Fertig (1973) describes a study yielding ten failure times out of 13 aircraft components that were tested. The failure times were 0.22, 0.50, 0.88, 1.00, 1.32, 1.33, 1.54, 1.76, 2.50, and 3.00 hours. The three right censored observations occurred at 3.00 hours. Both Mann and Fertig (1973) and Hsieh (1996) state that it is reasonable to assume a Weibull model for the data. The Weibull probability plot with 95% simultaneous confidence band (Nair 1984) in Figure 4(a) corroborates the adequacy of the Weibull model. Based on Figure 4(b), the lognormal distribution, however, is also suitable to describe the failure-time distribution of the aircraft component. In addition, the p-value (based on $10,000$ Monte Carlo samples) of Kolmogorov-Smirnov test for censored samples is 0.9998 for a Weibull fit;

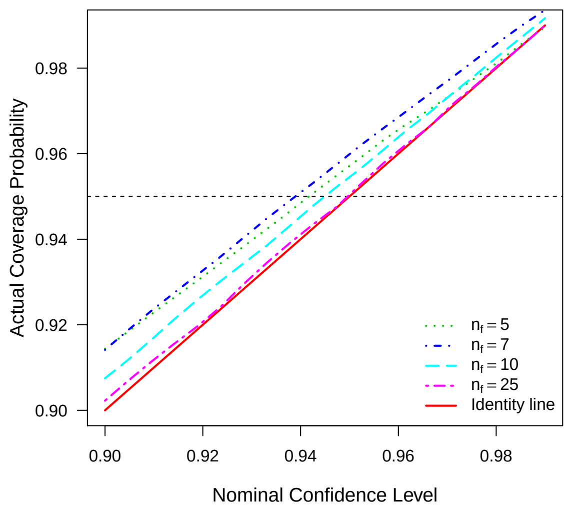

(a) Lower SPB

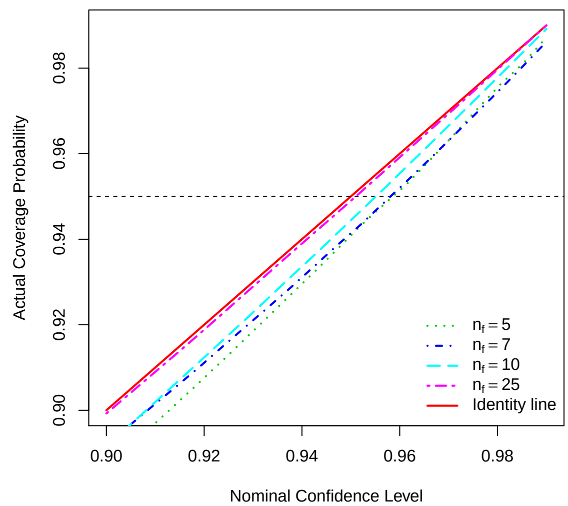

(b) Upper SPB

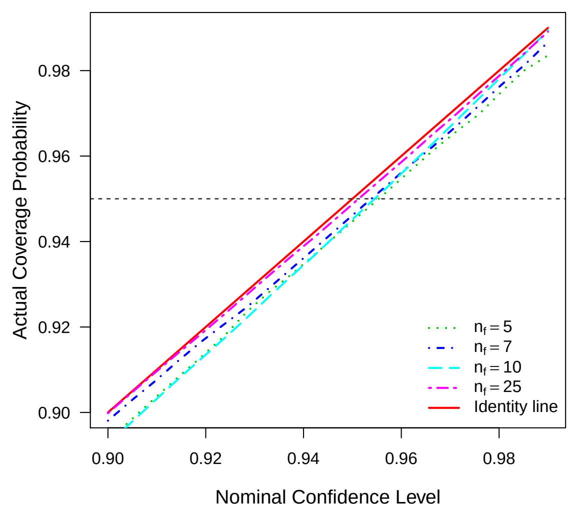

(c) Two-sided SPI

Figure 3: Estimated actual CP versus nominal confidence level for a Weibull model with fixed $p_f = 0.25$, when $k = 4$ and $m = 5$. (a) Lower SPB. (b) Upper SPB. (c) Two-sided SPI.

0.9834 for a lognormal fit. Using **Algorithm 1** one obtains 95% lower SPBs of the failure times of all 10 future aircraft components, which are 0.003 hours and 0.04 hours for the Weibull and lognormal distributions, respectively. Also we found that the 95% upper SPBs are 39.789 hours and 107.465 hours for the Weibull and lognormal distributions, respectively. The large difference is due to the implied extrapolation, especially into the upper tail of the failure-time distribution.

6.3 Vinyl Chloride Data

This application uses data consisting of 34 vinyl chloride concentrations (in $\mu g/L$) from clean upgradient ground-water monitoring wells. The data were given in Bhaumik and Gibbons (2006). The probability plot in Bhaumik and Gibbons (2006) indicates that the gamma distribution fits the data well. Figures 5(a) and 5(b) indicate that the Weibull and lognormal distributions also provide good fit to the vinyl chloride data. Furthermore, p-value of Kolmogorov-Smirnov test (based on 10,000 Monte Carlo samples) is 0.3867 for the Weibull distribution and 0.8417 for the lognormal distribution. Bhaumik and Gibbons (2006) wanted to obtain a 95% upper SPB to exceed at least $k = 1$ out of $m = 2$ future observations. For the gamma distribution, the 95% upper SPB is 2.931 $\mu g/L$. Using **Algorithm 1**, for the Weibull distribution, the 95% upper SPB is $\exp(0.635 + 0.464 \times 0.99) = 2.989$ $\mu g/L$; for the lognormal distribution, the 95% upper SPB is $\exp(0.092 + 0.829 \times 1.120) = 2.773$ $\mu g/L$. For this application, the 95% upper SPBs for the gamma, Weibull, and lognormal distributions are close to each other. This is because extrapolation is not required to construct this interval.

7 Conclusion and Areas for Future Research

In this paper, we propose a general method for computing simultaneous two-sided prediction intervals for at least k out of m future observations as well as the corresponding one-sided bounds for the (log)-location-scale family of distributions. For the Type II censored or complete data cases, the method provides a procedure with CP equal to the nominal confidence level (ignoring Monte Carlo error that can be made arbitrarily small). For Type I censored data, the approximate procedure provides coverage probabilities that are close to the nominal confidence level if the expected number of failures is not too small.

The procedures in this paper can also be extended to data involving multiple censoring or random censoring. With complete data, the extension of the proposed methods to the regression case is straightforward because the pivotal properties still hold (Lawless 2003,

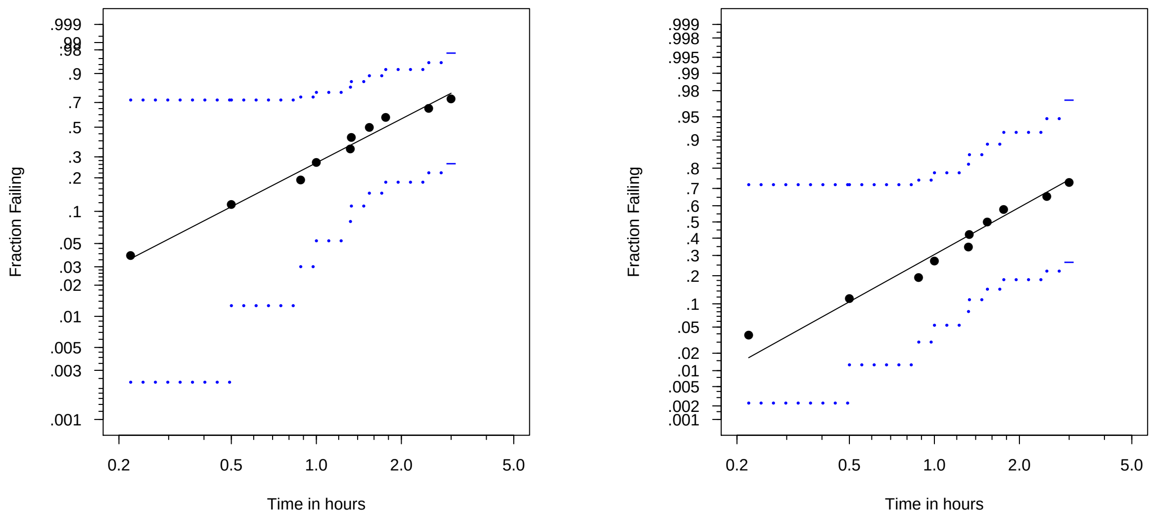

(a) Weibull (b) lognormal

Figure 4: Probability plots for the aircraft component data. The solid line shows the ML estimate of the cdf and the dotted lines show the 95% simultaneous confidence band. (a) Weibull fit. (b) Lognormal fit.

Appendix E4). As long as the pivotal property holds, the proposed procedure can be easily extended to give exact prediction intervals. When the pivotal property no longer holds (e.g., with regression and censoring, even if the censoring is Type II), the approximate pivotal approach can be applied. Our simulation results show that the CP for Type I censored data is satisfactory even when the sample size is small. Bias correction methods in estimating the CP, however, can be considered in future research.

Acknowledgments

The authors thank the editor, an associate editor, and two referees, for their valuable comments and suggestions that helped us to improve this paper. The authors acknowledge Advanced Research Computing at Virginia Tech for providing computational resources. The work by Hong was partially supported by the National Science Foundation under Grant CNS-1565314 to Virginia Tech.

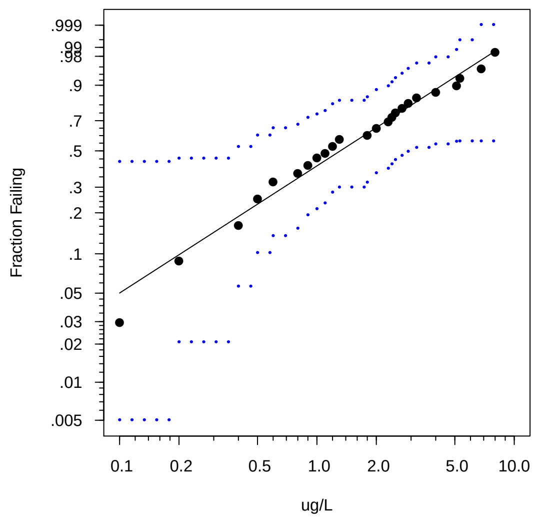

(a) Weibull

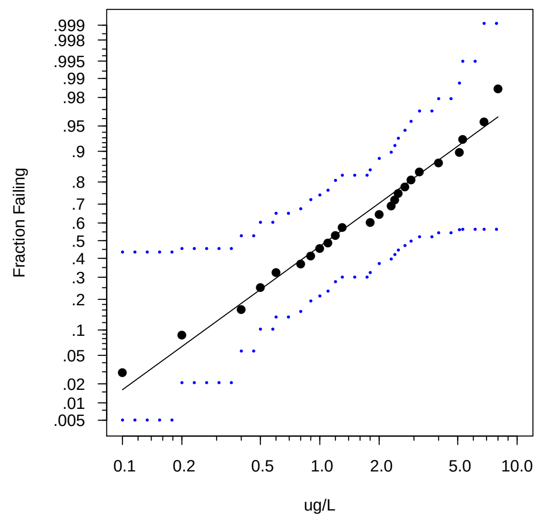

(b) lognormal

Figure 5: Probability plots for the vinyl chloride data. The solid line shows the ML estimate of the cdf and the dotted lines show the 95% simultaneous confidence band. (a) Weibull fit. (b) Lognormal fit.

A Technical Details

A.1 Proof of Result 1

Let D_j be the event that exactly j of the future observations $\boldsymbol{Y}$ lie in the prediction interval $[\widehat{\mu} + u_L\widehat{\sigma},\ \widehat{\mu} + u_U\widehat{\sigma}]$. To compute $\Pr(D_j)$, first we compute the conditional probability $\Pr(D_j|\widehat{\mu}, \widehat{\sigma})$ and then average this conditional probability over the sampling distribution of the ML estimates $(\widehat{\mu}, \widehat{\sigma})$.

First we proceed to compute $\Pr(D_j|\widehat{\mu}, \widehat{\sigma})$. Define the indicator variables

$$I_j = \begin{cases} 1 & \text{if } Y_j \in [\widehat{\mu} + u_L\widehat{\sigma},\ \widehat{\mu} + u_U\widehat{\sigma}] \\ 0 & \text{otherwise,} \end{cases}$$

where $j = 1, \ldots, m$.

The I_j variables are independent and identically distributed (iid) because the Y_j are iid. The I_j are Bernoulli(p) distributed where the p parameter is given in (11). Consequently, the number of future observations, say $S = \sum_{j=1}^{m} I_j$, contained by the conditional prediction interval $[\widehat{\mu} + u_L\widehat{\sigma},\ \widehat{\mu} + u_U\widehat{\sigma}]$ is Binomial(m, p) distributed.

The parameter p is

$$\begin{aligned}
p &= P(I_j = 1|\widehat{\mu}, \widehat{\sigma}) = \Pr\left(\widehat{\mu} + u_L\widehat{\sigma} \le Y_j \le \widehat{\mu} + u_U\widehat{\sigma}\right) \\
&= \Pr\left(Y_j \le \widehat{\mu} + u_U\widehat{\sigma}\right) - \Pr\left(Y_j \le \widehat{\mu} + u_L\widehat{\sigma}\right) \\
&= \Phi\left(\frac{\widehat{\mu} - \mu + u_U\widehat{\sigma}}{\sigma}\right) - \Phi\left(\frac{\widehat{\mu} - \mu + u_L\widehat{\sigma}}{\sigma}\right) \\
&= \Phi\left(\frac{\widehat{\mu} - \mu}{\sigma} + u_U\frac{\widehat{\sigma}}{\sigma}\right) - \Phi\left(\frac{\widehat{\mu} - \mu}{\sigma} + u_L\frac{\widehat{\sigma}}{\sigma}\right) \\
&= \Phi(a) - \Phi(b) \qquad (11)
\end{aligned}$$

where $a = z_1 + u_U z_2$, $b = z_1 + u_L z_2$, with z_1 and z_2 being realizations of $Z_1 = (\widehat{\mu} - \mu)/\sigma$ and $Z_2 = \widehat{\sigma}/\sigma$, respectively. The value of p is the same for all the variables $I_j,\ j = 1, \ldots, m$, because its value does not depend on the variable Y_j chosen to do the probability computation in (11). Note that $\Phi(\cdot)$ denotes the standard cdf of a location-scale distribution. Thus

$$\Pr(D_j \mid \widehat{\mu}, \widehat{\sigma}) = \Pr(S = j) = \binom{m}{j} p^j (1 - p)^{m-j}$$

and the unconditional probability for D_j is

$$\Pr(D_j) = \int_0^\infty \int_{-\infty}^\infty \binom{m}{j} p^j (1 - p)^{m-j} f_{(L,S)}(\widehat{\mu}, \widehat{\sigma})\, d\widehat{\mu} d\widehat{\sigma}$$

where $f_{(L,S)}(\widehat{\mu}, \widehat{\sigma})$ is the sampling distribution of $(\widehat{\mu}, \widehat{\sigma})$.

Define M to be the number of future observations contained by the prediction interval $[\widehat{\mu} + u_L\widehat{\sigma},\ \widehat{\mu} + u_U\widehat{\sigma}]$. Then the probability that the prediction interval contains at least k out of m future observations is

$$\begin{aligned}
\Pr(M \ge k) &= \sum_{j=k}^{m} \Pr(D_j) \\
&= \int_0^\infty \int_{-\infty}^\infty \sum_{j=k}^{m} \binom{m}{j} [\Phi(a) - \Phi(b)]^j \left[1 - \Phi(a) + \Phi(b)\right]^{m-j} f_{(L,S)}(\widehat{\mu}, \widehat{\sigma})\, d\widehat{\mu} d\widehat{\sigma} \\
&= \mathbf{E}\left[\sum_{j=k}^{m} \binom{m}{j} [\Phi(A) - \Phi(B)]^j \left[1 - \Phi(A) + \Phi(B)\right]^{m-j}\right].
\end{aligned}$$

Using (3), (u_L, u_U) can be chosen (selected/computed) to ensure that CP is equal to $(1 - \alpha)$.

A.2 Proof of Result 2

The loglikelihood of the right censored data has the following form,

$$l(\mu, \sigma) = \sum_{i=1}^{n} \left\{\delta_i \log f(x_i; \mu, \sigma) + (1 - \delta_i) \log[1 - F(x_i; \mu, \sigma)]\right\},$$

up to a constant. Based on the loglikelihood function, the ML estimators $(\widehat{\mu}, \widehat{\sigma})$ satisfies the following score equations:

$$\sum_{i=1}^{n}\left\{-\delta_i \frac{\phi'\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)}{\phi\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)} + (1-\delta_i)\frac{\phi\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)}{\left[1-\Phi\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)\right]}\right\} = 0, \tag{12}$$

$$\sum_{i=1}^{n}\left\{-\delta_i \left[1 + \frac{\phi'\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)}{\phi\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)}\right] + (1-\delta_i)\frac{\phi\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)}{1-\Phi\left(\frac{x_i-\widehat{\mu}}{\widehat{\sigma}}\right)}\right\} = 0. \tag{13}$$

For Type I censored data, x_c is fixed. Hence, $\delta_i = I(x_i \le x_c) = I[(x_i-\mu)/\sigma \le (x_c-\mu)/\sigma]$, where $I(\cdot)$ is an indicator function.

Notice that

$$\frac{x_i-\widehat{\mu}}{\widehat{\sigma}} = \frac{x_i-\mu}{\sigma}\left(\frac{\widehat{\sigma}}{\sigma}\right)^{-1} - \frac{\widehat{\mu}-\mu}{\sigma}\left(\frac{\widehat{\sigma}}{\sigma}\right)^{-1},$$
$$\frac{x_c-\mu}{\sigma} = \Phi^{-1}(p_f).$$

Because the distribution of $(x_i-\mu)/\sigma$ does not depend on (μ,σ), the left side of score equations are the function of $Z_1 = (\widehat{\mu}-\mu)/\sigma$, $Z_2 = \widehat{\sigma}/\sigma$ and p_f. Therefore the distribution of Z_1, Z_2 depends only on p_f.

In **Algorithm 2**, we use $\widehat{p}_f$ to approximate p_f. If we fix the censoring time x_c and let sample size n increases to infinity, $\widehat{p}_f$ will converge to p_f, Z_1 and Z_2 are asymptotic pivotal statistics, thus the CP of SPIs/SPBs computed by **Algorithm 2** will approach the nominal confidence level.

A.3 Proof of Result 3

Let $Y_1, \ldots, Y_{m_i}$ be a random sample of size m_i from a continuous, monotone increasing, distribution function $F(y)$. Denote the order statistics in the sample by $Y_{(1)} < \cdots < Y_{(k_i)} < \cdots < Y_{(m_i)}$. Let y_β be the β quantile of the distribution $F(y)$. Under some mild assumptions, $Y_{(k_i)}$ converges in probability and distribution to y_β [see Arnold, Balakrishnan, and Nagaraja (2008, page 3) for a proof]. Denoting by $F_{(k_i)}(x)$ the distribution of $Y_{(k_i)}$, we have $F_{(k_i)}(y) = \sum_{j=k_i}^{m_i}\binom{m_i}{j}[F(y)]^j[1-F(y)]^{m_i-j}$, and

$$\lim_{m_i\to\infty} F_{(k_i)}(y) = I_{\{y\ge y_\beta\}} = \begin{cases} 1 & \text{if } y > y_\beta \\ 0 & \text{otherwise,} \end{cases}$$

where $I_{\{y\ge y_\beta\}}$ is the indicator function of the set $\{y \ge y_\beta\}$.

Consider a two-sided prediction interval $[B,\ A]$, like the one in (4), to contain, with confidence $100(1-\alpha)\%$, at least k_i out of m_i future observations from the same population. Then

$$\begin{aligned}1-\alpha &= \mathbf{E}_{\boldsymbol{Z}}\left[\sum_{j=k_i}^{m_i}\binom{m_i}{j}\left[\Phi(A)-\Phi(B)\right]^j\left[1-\Phi(A)+\Phi(B)\right]^{m_i-j}\right] \\ &= \mathbf{E}_{\boldsymbol{Z}}\left(F_{(k_i)}\left\{F^{-1}\left[\Phi(A)-\Phi(B)\right]\right\}\right). \end{aligned} \tag{14}$$

Taking limits on both sides of (14), we get

$$\begin{aligned}1-\alpha &= \mathbf{E}_{\boldsymbol{Z}}\left(\lim_{m_i\to\infty}F_{(k_i)}\left\{F^{-1}\left[\Phi(A)-\Phi(B)\right]\right\}\right) \\ &= \mathbf{E}_{\boldsymbol{Z}}\left(I_{\left\{F^{-1}[\Phi(A)-\Phi(B)]>y_\beta\right\}}\right) = \Pr\left\{\left[\Phi(A)-\Phi(B)\right]>\beta\right\}. \end{aligned} \tag{15}$$

The expression in (15) shows that in the limit, as m_i increases, the prediction interval $[B,\ A]$ is equivalent to a two-sided tolerance interval to contain, with confidence $100(1-\alpha)\%$, at least a proportion β of the population.

For a one-sided upper or low prediction bound, a similar result can be obtained. The same type of argument can be used to show that a two-sided SPI $[B,\ A]$ with equal error probability in both tails, to contain, with confidence $100(1-\alpha)\%$, at least k_i out of m_i future observation, is (in the limit) equivalent to the two-sided TI below.

$$1-\alpha = \Pr\left\{\left[\Phi(A)-\Phi(B)\right]>\beta\right\},\ \text{and}\ \Pr\left\{\left[1-\Phi(B)\right]>\beta\right\} = \Pr\left[\Phi(A)>\beta\right].$$

Note, that the constraint $\Pr\left\{\left[1-\Phi(B)\right]>\beta\right\} = \Pr\left[\Phi(A)>\beta\right]$ has the effect of equating the lower and upper error rates of the two-sided $[B,\ A]$ TI.

A.4 Proof of Result 4

In this section, we obtain the variance of the estimated CP $\mathrm{CP}^*(u_L, u_U)$ given in (9). For each simulated sample $l=1,\dots,B_1$, ML estimates $\widehat{\boldsymbol{\theta}}^* = (\widehat{\mu}_l^*, \widehat{\sigma}_l^*)$ are independent. Moreover,

$$\widehat{\boldsymbol{\theta}}^* \dot{\sim} \mathrm{N}\left[\boldsymbol{\theta}, I^{-1}(\boldsymbol{\theta})\right],$$

where $I(\boldsymbol{\theta})$ is the Fisher information matrix. In the simulation, we take $\boldsymbol{\theta}=(0,1)'$. Denote the corresponding Fisher information matrix by I_o. A discussion on the calculation of Fisher information matrix with censored data can be found in Appendix B.3 of Escobar and Meeker (1998). They showed that I_o is only the function of n, x_c, and $\Phi(\cdot)$ for Type I censored data; I_o is the function of n, r, and $\Phi(\cdot)$ for Type II censored data.

Denote $g(u_L, u_U) = \sum_{j=k}^{m} \binom{m}{j} p_l(u_L, u_U)^j [1 - p_l(u_L, u_U)]^{m-j}$, therefore we have

$$\text{Var}[\text{CP}^*(u_L, u_U)] = \frac{1}{B_1} \text{Var}[g(u_L, u_U)].$$

The first derivatives of $g(u_L, u_U)$ with respect to $(\widehat{\mu}_l^*, \widehat{\sigma}_l^*)$ are

$$\frac{\partial g}{\partial \widehat{\mu}_l^*} = m \binom{m-1}{k-1} [p_l(u_L, u_U)]^{k-1} [1 - p_l(u_L, u_U)]^{m-k} \frac{\partial p_l(u_L, u_U)}{\partial \widehat{\mu}_l^*},$$
$$\frac{\partial g}{\partial \widehat{\sigma}_l^*} = m \binom{m-1}{k-1} [p_l(u_L, u_U)]^{k-1} [1 - p_l(u_L, u_U)]^{m-k} \frac{\partial p_l(u_L, u_U)}{\partial \widehat{\sigma}_l^*},$$

where

$$\frac{\partial p_l(u_L, u_U)}{\partial \widehat{\mu}_l^*} = \phi(\widehat{\mu}_l^* + u_U \widehat{\sigma}_l^*) - \phi(\widehat{\mu}_l^* + u_L \widehat{\sigma}_l^*),$$
$$\frac{\partial p_l(u_L, u_U)}{\partial \widehat{\sigma}_l^*} = \phi(\widehat{\mu}_l^* + u_U \widehat{\sigma}_l^*) u_U - \phi(\widehat{\mu}_l^* + u_L \widehat{\sigma}_l^*) u_L.$$

According to Delta method, we have

$$\text{Var}[g(u_L, u_U)] = \left(\frac{\partial g}{\partial \widehat{\mu}_l^*}, \frac{\partial g}{\partial \widehat{\sigma}_l^*} \right) I_o^{-1} \left(\frac{\partial g}{\partial \widehat{\mu}_l^*}, \frac{\partial g}{\partial \widehat{\sigma}_l^*} \right)' \Bigg|_{(\widehat{\mu}_l^*, \widehat{\sigma}_l^*) = (0,1)}.$$

Therefore the variance of estimated CP $\text{CP}^*(u_L, u_U)$ is a function of $(u_L, u_U, k, m, n, x_c, \Phi(\cdot), B_1)$ for Type I censored data, and a function of $(u_L, u_U, k, m, n, r, \Phi(\cdot), B_1)$ for Type II censored data. Similar results can be found for one-sided SPBs.

References

Arnold, B. C., N. Balakrishnan, and H. Nagaraja (2008). *A First Course in Order Statistics (Classics edition)*. New York: Wiley.

Barr, D. R. and T. Davidson (1973). A Kolmogorov-Smirnov test for censored samples. *Technometrics 15*, 739–757.

Beran, R. (1990). Calibrating prediction regions. *Journal of the American Statistical Association 85*, 715–723.

Bhaumik, D. K. and R. D. Gibbons (2006). One-sided approximate prediction intervals for at least p of m observations from a Gamma population at each of r locations. *Technometrics 48*, 112–119.

Bhaumik, D. K., S. Santra, S. Aryal, and R. Gibbons (2008). One-sided simultaneous prediction limits for left-censored normal random variables. *Sankhyā: The Indian Journal of Statistics 70*, 248–266.

Castillo, E. (1988). *Extreme Value Theory in Engineering.* New York: Academic Press.

Danziger, L. and S. A. Davis (1964). Tables of distribution-free tolerance limits. *Annals of Mathematical Statistics 35*, 1361–1365.

Davis, C. B. and R. J. McNichols (1987). One-sided intervals for at least p of m observations from a normal population on each of r future occasions. *Technometrics 29*, 359–370.

Dufour, R. and U. Maag (1978). Distribution results for modified Kolmogorov-Smirnov statistics for truncated or censored. *Technometrics 20*, 29–32.

Escobar, L. A. and W. Q. Meeker (1998). Fisher information matrices with censoring, truncation, and explanatory variables. *Statistica Sinica 8*, 221–237.

Escobar, L. A. and W. Q. Meeker (1999). Statistical prediction based on censored life data. *Technometrics 41*, 113–124.

Fertig, K. W. and N. R. Mann (1977). One-sided prediction intervals for at least p out of m future observations from a normal population. *Technometrics 19*, 167–177.

Hahn, G. J. (1969). Factors for calculating two-sided prediction intervals for samples from a normal distribution. *Journal of the American Statistical Association 65*, 878–888.

Hahn, G. J. (1970). Additional factors for calculating prediction intervals for samples from a normal distribution. *Journal of the American Statistical Association 65*, 1668–1676.

Hsieh, H. K. (1996). Prediction intervals for Weibull observations, based on early-failure data. *IEEE Transactions on Reliability 45*, 666–670.

Jeng, S.-L. and W. Q. Meeker (2000). Comparisons of Weibull distribution approximate confidence intervals procedures for type I censored data. *Technometrics 42*, 135–148.

Krishnamoorthy, K. (2016). STATCALC 3.0, avaliable at http://www.ucs.louisiana.edu/∼kxk4695/StatCalc.htm. Webpage.

Krishnamoorthy, K., Y. Lin, and Y. Xia (2009). Confidence limits and prediction limits for a Weibull distribution based on the generalized variable approach. *Journal of Statistical Planning and Inference 139*, 2675–2684.

Kushary, D. (1996). Construction of prediction intervals in location-scale families. *Biometrical Journal 38*, 865–878.

Lawless, J. F. (2003). *Statistical Models and Methods for Lifetime Data, 2nd edition.* New York: John Wiley and Sons.

Mann, N. R. and K. W. Fertig (1973). Tables for obtaining Weibull confidence bounds and tolerance bounds based on best linear invariant estimates of parameters of the extreme-value distribution. *Technometrics 15*, 87–101.

Mee, R. W. and D. Kushary (1994). Prediction limits for the Weibull distribution utilizing simulation. *Computational Statistics & Data Analysis 17*, 327–336.

Meeker, W. Q. and L. A. Escobar (1998). *Statistical Methods for Reliability Data*. John Wiley and Sons.

Nair, V. N. (1984). Confidence bands for survival functions with censored data: a comparative study. *Technometrics 26*, 265–275.

Odeh, R. E. (1990). Two-sided prediction intervals to contain at least k out of m future observations from a normal distribution. *Technometrics 32*, 203–216.

Vander Weil, S. A. and W. Q. Meeker (1990). Accuracy of approx confidence bounds using censored Weibull regression data from accelerated life tests. *IEEE Transactions on Reliability 39*, 346–351.

Wang, C. M., J. Hannig, and H. K. Iyer (2012). Fiducial prediction intervals. *Journal of Statistical Planning and Inference 142*, 1980–1990.

Wu, T.-H. and H.-L. Lu (2007). Prediction intervals for an ordered observation from the logistic distribution based on censored samples. *Journal of Statistical Computation and Simulation 77*, 389–405.