Journal of Statistical Planning and Inference 142 (2012) 1980–1990

Contents lists available at SciVerse ScienceDirect

Journal of Statistical Planning and Inference

journal homepage: www.elsevier.com/locate/jspi

Fiducial prediction intervals ☆

C.M. Wang [a,*], Jan Hannig [b], Hari K. Iyer [a,c,1]

[a] Statistical Engineering Division, National Institute of Standards and Technology, Boulder, CO 80305, United States
[b] Department of Statistics and Operations Research, University of North Carolina, Chapel Hill, NC 27599, United States
[c] Department of Statistics, Colorado State University, Fort Collins, CO 80523, United States

ARTICLE INFO

Article history:
Received 3 June 2011
Received in revised form
11 October 2011
Accepted 11 February 2012
Available online 18 February 2012

Keywords:
Fiducial inference
Gamma distribution
Statistical intervals

ABSTRACT

This paper presents an approach for constructing prediction intervals for any given distribution. The approach is based on the principle of fiducial inference. We use several examples, including the normal, binomial, exponential, gamma, and Weibull distributions, to illustrate the proposed procedure.

Published by Elsevier B.V.

1. Introduction

Prediction intervals, used in many practical applications, are statistical intervals that contain, with a specific probability, future realizations of a random variable from a distribution of interest. Hahn and Meeker (1991) provided a summary of methods for constructing statistical intervals, which include prediction intervals, for the normal, binomial, and Poisson distributions. Prediction intervals for the gamma distribution were considered by Hamada et al. (2004), Bhaumik and Gibbons (2006), and Krishnamoorthy et al. (2008). For applications involving the Weibull distribution, prediction intervals based on censored and uncensored data were proposed by Escobar and Meeker (1999), Nelson (2000), Nordman and Meeker (2002), and Krishnamoorthy et al. (2009).

Methods for constructing prediction intervals described in the literature include the pivotal-based method for the normal distribution, the approximate pivotal-based method for the gamma, and the simulation-based method for the Weibull. Lawless and Fredette (2005) described a general method, called pivotal method, that can be used to construct predictions intervals for a wide variety of problems. Let X and Y be random variables, representing the current and future data, from a distribution indexed by a parameter θ. Their approach is based on the random variable

$$U = F_{\hat{\theta}}(Y), \tag{1}$$

where $F_\theta(\cdot)$ is the distribution function of X (and Y) and $\hat{\theta} = \hat{\theta}(X)$ is the maximum likelihood estimator of θ. Prediction intervals for Y are obtained from the distribution function of U, $G_\theta(u)$, which may or may not depend on θ. If $G_\theta(u)$ is free of θ, i.e., U is a pivotal, the resulting prediction intervals are "exact". If $G_\theta(u)$ depends on θ, prediction intervals are obtained

☆ Hannig and Iyer's research was supported in part by the National Science Foundation under Grant No. 0707037 and 1007543.

* Corresponding author.
E-mail address: jwang@boulder.nist.gov (C.M. Wang).
[1] Research Information Fellow at Caterpillar Inc.

0378-3758/$ - see front matter Published by Elsevier B.V.
doi:10.1016/j.jspi.2012.02.021

from an approximate distribution such as $G_{\hat{\theta}}(u)$. In many applications, $G(u)$ is intractable and needs to be approximated by simulation. Lawless and Fredette (2005) also showed that their predictive distributions are optimal in certain settings and illustrated the procedure with several examples.

In this paper we propose an alternative procedure for constructing prediction intervals for a variety of distributions. The procedure is based on the fiducial method described by Hannig et al. (2006) and Hannig (2009a). The fiducial-based methods have been used in other inference problems where classical methods do not lead to satisfactory procedures, for example, see Iyer et al. (2004) and E et al. (2008). Recent research results (Hannig et al., 2006; Hannig, 2009a) and many simulation studies show that fiducial inference is a valid statistical method with good operating characteristics.

The proposed fiducial method is similar to the pivotal method of Lawless and Fredette (2005) in the case of independently and identically distributed (iid) data. In many problems, fiducial prediction intervals can be obtained from the distribution function of

$$U^\star = F_{R_\theta}(Y), \tag{2}$$

where R_θ is a function of observed data and random variables whose distributions are completely known and free of θ. Notice that the similarity between (1) and (2). We use an exponential-distribution example in Section 3 to illustrate this similarity.

The rest of the paper is organized as follows. In the next section we describe the procedure for constructing prediction intervals based on the fiducial inference approach. We show that under regularity conditions, fiducial prediction intervals have asymptotically correct frequentist coverage. In Section 3 we use several examples, including the normal, binomial, exponential, gamma, and Weibull distributions, as well as a simple linear regression, to illustrate the proposed procedure. For the Weibull distribution, we consider only uncensored data in this paper. Summary remarks are provided in the final section.

2. The fiducial approach

In this section we describe a general procedure for constructing prediction intervals for future observations from a population of arbitrary distribution. The procedure we propose is based on the fiducial method described by Hannig et al. (2006) and Hannig (2009a).

We use a simple example throughout this section to illustrate the idea and procedure. Let X_1 and X_2 be a random sample from $\mathrm{N}(\mu,1)$. Suppose, based on these two measurements, we desire a prediction interval on the mean of three future observations from the same population. Let Y_i, $i=1,2,3$, be the random variables denoting the three future observations from $\mathrm{N}(\mu,1)$. We represent the current and future data with the following models:

$$X_1 = \mu + E_1, \tag{3}$$

$$X_2 = \mu + E_2 \tag{4}$$

and

$$Y_i = \mu + E_i^\star, \quad i=1,2,3, \tag{5}$$

where E_1, E_2, and $E_i^\star$ are independent $\mathrm{N}(0,1)$ random errors.

The original fiducial recipe described in Hannig (2009a) is as follows. Let $\boldsymbol{X} \in \mathbb{R}^n$ be a random vector with a distribution indexed by a parameter $\xi \in \Xi$. Assume that the data generating mechanism for $\boldsymbol{X}$ is expressed as

$$\boldsymbol{X} = G(\xi, U), \tag{6}$$

where G is a jointly measurable function and U is a random variable or vector with a completely known distribution independent of any parameters. Eq. (6) is termed the *structural equation*. For the above example, Eqs. (3) and (4) are the structural equations.

We define for any observed value $\boldsymbol{x} \in \mathbb{R}^n$ a set-valued function

$$Q(\boldsymbol{x}, u) = \{\xi : G(\xi, u) = \boldsymbol{x}\}. \tag{7}$$

Assume that $Q(\boldsymbol{x},u)$ is a measurable function of u. The function $Q(\boldsymbol{x},u)$ is an inverse image of the function G.

Let us now assume that a data set was generated using (6) and it has been observed that the sample value is $\boldsymbol{x}$. Clearly, the values of ξ and u used to generate the observed data will satisfy $G(\xi,u) \in \mathbb{R}^n$. This leads to the following definition of a generalized fiducial distribution for ξ:

$$Q(\boldsymbol{x}, U^\star) \,\big|\, \{Q(\boldsymbol{x}, U^\star) \neq \emptyset\}, \tag{8}$$

where $U^\star$ is an independent copy of U.

There are two possible issues with (8). First, Eq. (7) might not be a singleton, in which case we propose to randomly select one of the elements of $Q(\boldsymbol{x},u)$. Second, for many examples $\Pr[Q(\boldsymbol{x}, U^\star) \neq \emptyset] = 0$. In this case, Hannig (2009a) proposes a way of interpreting the conditional distribution (8) that has been theoretically justified by Hannig (2009b).

For the above example, a fiducial distribution for μ is the distribution of $x_1 - E_1$ conditional on $E_1 - E_2 = x_1 - x_2$, which is $\mathrm{N}(\bar{x}, 1/2)$, where x_1 and x_2 are the realized values of X_1 and X_2, and $\bar{x} = (x_1 + x_2)/2$. This is obtained by solving for μ in (3) (inverse image), replacing X_1 in the solution with x_1, and plugging the solution into (4) for conditioning.

The problem of prediction can be easily incorporated into this scenario by modifying the structural equation (6). Let us assume that $\boldsymbol{X}=(\boldsymbol{X}^o, \boldsymbol{X}^p)$, where $\boldsymbol{X}^o$ is observed and $\boldsymbol{X}^p$ is unobserved and needs to be predicted. The structural equation (6) for $\boldsymbol{X}$ is

$$(\boldsymbol{X}^o, \boldsymbol{X}^p)=(G^o(\xi,U), G^p(\xi,U)). \tag{9}$$

Notice that the left-hand side of (9) has both observed values $\boldsymbol{X}^o$ and an unobserved quantity $\boldsymbol{X}^p$ that will be treated together with ξ as the unknown parameters. Again U is a random variable with a fully known distribution independent of any parameters.

Following the generalized fiducial recipe above, the inverse image (7) becomes

$$Q(\boldsymbol{x}^o,u)=\{(\boldsymbol{x}^p,\xi):G^o(\xi,u)=\boldsymbol{x}^o,\ G^p(\xi,u)=\boldsymbol{x}^p\}$$

and the joint fiducial distribution on the parameters and the predicted values is given by

$$Q(\boldsymbol{x}^o, U^\star)\,|\,\{Q(\boldsymbol{x}^o,U^\star)\neq\emptyset\}. \tag{10}$$

Define

$$Q^p(\boldsymbol{x}^o,u)=\{\boldsymbol{x}^p : G^o(\xi,u)=\boldsymbol{x}^o,\ G^p(\xi,u)=\boldsymbol{x}^p \text{ for some } \xi\}$$

and $Q^o(\boldsymbol{x}^o,u)=\{\xi:G^o(\xi,u)=\boldsymbol{x}^o\}$. Notice that $\{u:Q(\boldsymbol{x}^o,u)\neq\emptyset\}=\{u:Q^o(\boldsymbol{x}^o,u)\neq\emptyset\}$ and therefore the condition in (10) is the same as in the definition of the usual generalized fiducial distribution (8) ignoring the prediction part of the problem. The predictive fiducial distribution of $\boldsymbol{X}^p$ is obtained by marginalizing (10) as

$$Q^p(\boldsymbol{x}^o, U^\star)\,|\,\{Q^o(\boldsymbol{x}^o, U^\star)\neq\emptyset\}. \tag{11}$$

If additionally, as in the case of iid data, the structural equation (9) factorizes to

$$\boldsymbol{X}^o=G^o(\xi,U^o) \quad \text{and} \quad \boldsymbol{X}^p=G^p(\xi,U^p), \tag{12}$$

where U^o and U^p are independent, then this recipe is equivalent to first deriving a generalized fiducial distribution from the observed data

$$Q^o(\boldsymbol{x}^o, U^{o,\star})\,|\,\{Q^o(\boldsymbol{x}^o,U^{o,\star})\neq\emptyset\} \tag{13}$$

and then plugging it into the prediction part of the equation. In particular, denote a random variable having the same distribution as the fiducial distribution of ξ by $\tilde{\xi}$. We call $\tilde{\xi}$ a *fiducial quantity* of ξ. Then the predictive fiducial distribution (11) in this situation simplifies to the distribution of

$$\boldsymbol{X}^p=G^p(\tilde{\xi}, U^{p,\star}),$$

where $U^{p,\star}$ is a copy of U^p independent of $\tilde{\xi}$. A simple algebra shows that in this case the density of the predictive distribution becomes

$$\int_{\Xi} f(\boldsymbol{x}^p|\xi)r(\xi|\boldsymbol{x}^o)\,d\xi,$$

where $f(\boldsymbol{x}^p|\xi)$ is the density of the predicted data and $r(\xi|\boldsymbol{x}^o)$ is the fiducial density of (13).

For the above example, let $\tilde{\mu}$ be the fiducial quantity, that is, $\tilde{\mu}\sim \mathrm{N}(\overline{x}, 1/2)$. The predictive distribution (11) is then the distribution of

$$\tilde{Y}_i=\tilde{\mu}+E_i^\star. \tag{14}$$

That is, $\tilde{Y}_i$ is obtained from (5) by replacing μ with its fiducial quantity $\tilde{\mu}$. Prediction limits for Y_i are obtained from the distribution of $\tilde{Y}_i$. Specifically, a fiducial quantity for $\overline{Y}_3$, the mean of Y_1, Y_2, and Y_3, is obtained as

$$\tilde{\overline{Y}}_3=\frac{\tilde{Y}_1+\tilde{Y}_2+\tilde{Y}_3}{3}=\tilde{\mu}+\frac{E_1^\star+E_2^\star+E_3^\star}{3}. \tag{15}$$

Since $\tilde{\mu}\sim \mathrm{N}(\overline{x}, 1/2)$, $\tilde{\overline{Y}}_3$ is normally distributed with mean $\overline{x}$ and variance 5/6. A 95% two-sided prediction interval to contain the mean of the three future observations is then given by

$$(\overline{x}-1.96\sqrt{5/6},\ \overline{x}+1.96\sqrt{5/6}). \tag{16}$$

This is exactly the prediction interval one would get using the frequentist approach.

In this simple example we were able to determine the prediction interval in (16) analytically. In many problems the exact form of (11) cannot be obtained in closed form. In such cases, we will use Monte Carlo techniques to generate a sample having the fiducial predictive distribution. The main idea is to first generate a random sample of $U_1^\star,\ldots,U_m^\star$ from the conditional distribution of

$$U^\star\,|\,\{Q^o(\boldsymbol{x}^o, U^\star)\neq\emptyset\}.$$

Then the random variables

$$X_i^p=G^p(Q^o(\boldsymbol{x}^o,U_i^\star), U_i^\star), \quad i=1,2,\ldots,m$$

are a random sample from the predictive fiducial distribution of (11). Appropriate quantiles from this empirical distribution can be used to compute prediction limits.

2.1. Theoretical result

Theoretical properties of generalized fiducial inference for iid data has been studied in Hannig et al. (2006) and Hannig (2009a,b). They prove that, under regularity conditions, confidence intervals based on the generalized fiducial distribution (13) have asymptotically correct coverage. In particular, they prove a theorem similar to the Bernstein–von Mises theorem for Bayesian posteriors, i.e., if ξ_0 is the parameter used to generate $\boldsymbol{X}^o$ and R_ξ is random with the distribution given by (13), then there is a random variable $T(\boldsymbol{X}^o,\ \xi_0)$ such that $\sqrt{n}(R_\xi - T(\boldsymbol{X}^o,\ \xi_0)) \xrightarrow{\mathcal{D}} \mathrm{N}(0,\Sigma)$ and $\sqrt{n}(T(\boldsymbol{X}^o,\ \xi_0) - \xi_0) \xrightarrow{\mathcal{D}} \mathrm{N}(0,\Sigma)$ as the number of observed values $n \to \infty$. This in particular means that the generalized fiducial distribution is consistent, $R_\xi \xrightarrow{P} \xi_0$.

We now show that if the generalized fiducial distribution is consistent, then the generalized fiducial predictive distribution leads to prediction sets that have asymptotically correct coverage.

Theorem 1. *Let us assume that the structural equation factorizes as in* (12) *with* $G^p(\xi, U^p)$ *continuous in* ξ *at* ξ_0 *for all* U^p *and such that the distribution of* $G^p(\xi_0,\ U^p)$ *is continuous. Further assume that the generalized fiducial distribution* (13) *is consistent. Finally, assume that the prediction sets* $C(\boldsymbol{x}^0)$ *for future data are based on the fiducial predictive distribution* (11) *and have a shape satisfying the Assumption* 3 *of Theorem* 1 *of Hannig (2009a), such as the equal tailed region or one sided interval. Then the prediction set has asymptotically correct frequentist coverage.*

Proof. Let R_ξ be a random variable with the distribution given by (13). By assumption $R_\xi \xrightarrow{P} \xi_0$. Because of the independence of U^p and U^0 in (12) we see that the predictive fiducial distribution (11) is equal to the distribution of $G^p(R_\xi,\ U^{p,*})$, where $U^{p,*}$ is an independent copy of U^p and is independent of R_ξ. By continuity

$$G^p(R_\xi, U^{p,*}) \xrightarrow{P} G^p(\xi_0,\ U^{p,*}).$$

The proof now follows similar arguments as in the proof of Theorem 1 in Hannig (2009a). □

We use several examples to illustrate this procedure.

3. Examples

3.1. One-sample normal distribution

In the first example, we consider a one-sided fiducial prediction interval to contain at least p out of m future observations based on a random sample of size n from a normal distribution $\mathrm{N}(\mu,\sigma^2)$. A fiducial quantity for (μ,σ^2) can be derived from the structural equations based on the minimal sufficient statistics of the model, and is given by

$$R_{(\mu,\sigma^2)}(\overline{x},s^2) = (\overline{x} - sZ^{(1)}/\sqrt{nV/(n-1)},\ (n-1)s^2/V),$$

where $\overline{x}$ and s are the observed sample mean and standard deviation of the random sample; $Z^{(1)} \sim \mathrm{N}(0,1)$; and $V \sim \chi^2_{n-1}$ and is independent of $Z^{(1)}$. Let Y_i be the ith future observation. Then the data-generating mechanism is given by

$$Y_i = \mu + \sigma Z_i^{(2)}, \quad i = 1,\ldots,m, \tag{17}$$

where $Z_i^{(2)}$ are iid $\mathrm{N}(0,1)$ random variable and independent of $Z^{(1)}$. Substituting (μ,σ^2) with $R_{(\mu,\sigma^2)}(\overline{x},s^2)$ into (17), we obtain a fiducial prediction quantity for Y_i as

$$\tilde{Y}_i = \overline{x} - sZ^{(1)}/\sqrt{nV/(n-1)} + sZ_i^{(2)}/\sqrt{V/(n-1)}. \tag{18}$$

Let $Y_{(s)}$ be the sth order statistic of $\{Y_1,\ldots,Y_m\}$. Then the event that at least p out of m future observations will exceed L is equivalent to the event that $Y_{(m-p+1)} > L$. The fiducial prediction quantity for $Y_{(m-p+1)}$ is $\tilde{Y}_{(m-p+1)}$, where $\tilde{Y}_{(s)}$ is the sth order statistic of $\{\tilde{Y}_{(1)},\ldots,\tilde{Y}_{(m)}\}$. Thus, a $1-\alpha$ lower prediction limit L on $Y_{(m-p+1)}$ is the α quantile of the distribution of $\tilde{Y}_{(m-p+1)}$. Since the ordering among

$$(\tilde{Y}_i - \overline{x})/s = -Z^{(1)}/\sqrt{nV/(n-1)} + Z_i^{(2)}/\sqrt{V/(n-1)}, \quad i = 1,\ldots,m$$

is identical to the ordering among $\tilde{Y}_i$, $i = 1,\ldots,m$, the α quantile of the distribution of $\tilde{Y}_{(m-p+1)}$ is given by $\overline{x} - rs$, where r is the α quantile of the distribution of the $(m-p+1)$th order statistic of $Z_i^{(2)}/\sqrt{V/(n-1)} - Z^{(1)}/\sqrt{nV/(n-1)}$, which can be obtained easily by simulation. For example, the following R code will calculate the value of r based on nrun Monte Carlo samples:

```
z1<-rnorm(nrun)/sqrt(n)
V<-sqrt(rchisq(nrun,n-1)/(n-1))
z2i<-matrix(rnorm(m*nrun),byrow=T,ncol=m)
xij<-apply((z2i-z1)/V,1,function(x,q)sort(x)[q],m-p+1)
quantile(xij,alpha)
```

Values of r so obtained are identical to the factors calculated and tabulated by Fertig and Mann (1977), using numerical procedures. Similarly, the two-sided symmetric fiducial prediction intervals can be constructed by use of the $1-\alpha$ quantile of the distribution of the appropriate order statistic of the absolute value of $Z_i^{(2)}/\sqrt{V/(n-1)} - Z^{(1)}/\sqrt{nV/(n-1)}$. These quantiles are identical to the factors calculated and tabulated by Odeh (1990).

3.2. Binomial distribution

Let X be a random variable representing the number of "successes" in n repeated independent Bernoulli trials when p is the probability of success at each trial. If we define the value of the random variable to be 1 if a trial results in success, and 0 otherwise, then we have the following structural equation:

$$X = \sum_{i=1}^{n} \mathrm{I}_{[0,p]}(U_i),$$

where $\mathrm{I}_A(\cdot)$ is the indicator function and U_i, $i=1,\ldots,n$, are independent uniform(0, 1) random variables. Following Hannig (2009a), a fiducial quantity for p is given by

$$\tilde{p} = U_{(x)} + (U_{(x+1)} - U_{(x)})D, \tag{19}$$

where $U_{(s)}$ is the sth order statistic of $\{U_1, \ldots, U_n\}$, and D is a uniform(0, 1) random variable and independent of U_i, $i=1,\ldots,n$. Let Y be the number of successes in a future random sample of size m from the same population. Then we can write

$$Y = \sum_{i=1}^{m} \mathrm{I}_{[0,p]}(U_i^\star), \tag{20}$$

where $U_i^\star$, $i=1,\ldots,m$, are independent uniform(0, 1) random variables. Substituting p in (20) with $\tilde{p}$ of (19), we obtain a fiducial prediction quantity for Y as

$$\tilde{Y} = \sum_{i=1}^{m} \mathrm{I}_{[0,\tilde{p}]}(U_i^\star).$$

Prediction limits on Y can be obtained from the appropriate quantiles of the distribution of $\tilde{Y}$.

We use an example from Hahn and Meeker (1991) to illustrate the above procedure. In a test of 1000 randomly selected integrated circuits,20 nonconforming units were found. An upper 95% prediction bound for the number of nonconforming units in a future lot of 1000 units,which are randomly sampled from the same production process,is desired. This upper bound is found to be 31,calculated by use of the following R code based on 50 000 Monte Carlo samples:

```
n <- 1000 # current sample sizes
x <- 20 # observed nonconforming units
m <- 1000 # future sample sizes
nrun <- 50 000 # Monte Carlo samples
mat1 <- matrix(runif(n*nrun), ncol = n)
mat2 <- apply(mat1, 1, sort)
uvec <- runif(nrun)
fq.p <- mat2[x,]+uvec*(mat2[(x+1),]-mat2[x,]) # FQ for p
Y <- 1 : nrun
for (i in (1 : nrun))
  Y[i] <- sum(runif(m) < fq.p[i]) # realizations from Ỹ
quantile(Y,0.95) # 95% upper bound
```

For comparison, Hahn and Meeker (1991) calculated the upper prediction limit to be 32, based on a hypergeometric-distribution approach.

3.3. Simple linear regression

Suppose

$$Y_i = \beta_0 + \beta_1 x_i + E_i, \quad i = 1, \ldots, n,$$

where β_0 and β_1 are respectively the intercept and the slope of the straight line model relating Y_i to x_i, and E_i are independently normally distributed random errors with mean 0 and unknown variance σ^2. Based on these n pairs of data, a prediction interval to contain the mean of m future values of Y at $x = x_0$ is derived by use of the fiducial approach.

The parameters of the model are β_0, β_1, and σ. Let $\boldsymbol{B} = (B_0\ B_1)^t$ denote the least-squares estimator of $\boldsymbol{\beta} = (\beta_0\ \beta_1)^t$. Then

$$\boldsymbol{B} \sim \mathrm{N}(\boldsymbol{\beta},\ \sigma^2 \boldsymbol{V}), \tag{21}$$

where

$$\boldsymbol{V} = \begin{pmatrix} n & \sum_{i=1}^n x_i \\ \sum_{i=1}^n x_i & \sum_{i=1}^n x_i^2 \end{pmatrix}^{-1} = \frac{1}{D_x} \begin{pmatrix} \sum_{i=1}^n x_i^2 & -\sum_{i=1}^n x_i \\ -\sum_{i=1}^n x_i & n \end{pmatrix}$$

and $D_x = n\sum_{i=1}^n x_i^2 - (\sum_{i=1}^n x_i)^2$. From (21) we can write

$$\boldsymbol{B} = \boldsymbol{\beta} + \sigma \boldsymbol{T\Phi}, \tag{22}$$

where $\boldsymbol{\Phi}$ is a bivariate standard normal random variable, and $\boldsymbol{T}$ is the Cholesky factor of $\boldsymbol{V}$; that is, $\boldsymbol{T}\boldsymbol{T}^t = \boldsymbol{V}$ and is given by

$$\boldsymbol{T} = \frac{1}{\sqrt{\sum_{i=1}^n (x_i - \overline{x})^2}} \begin{pmatrix} \sqrt{\sum_{i=1}^n x_i^2/n} & 0 \\ -\overline{x}/\sqrt{\sum_{i=1}^n x_i^2/n} & \sqrt{\sum_{i=1}^n (x_i-\overline{x})^2/\sum_{i=1}^n x_i^2} \end{pmatrix}.$$

Also, let

$$S^2 = \frac{\sum_{i=1}^n (Y_i - B_0 - B_1 x_i)^2}{n-2}$$

then S^2 is independent of $\boldsymbol{B}$ and

$$V = \frac{(n-2)S^2}{\sigma^2} \sim \chi^2(n-2). \tag{23}$$

Eqs. (22) and (23) are the structural equations that relate the observable random variables $\boldsymbol{B}$ and S^2 to the model parameters $\boldsymbol{\beta}$, σ, and the error processes $\boldsymbol{T\Phi}$ and V, whose distributions are fully known. By solving $\boldsymbol{\beta}$ and σ from these two equations (no conditioning is needed), we obtain the fiducial quantities for σ and $\boldsymbol{\beta}$ as

$$\tilde{\sigma} = \frac{s}{\sqrt{V/(n-2)}},$$

$$\tilde{\boldsymbol{\beta}} = \boldsymbol{b} - \tilde{\sigma}\boldsymbol{T\Phi}, \tag{24}$$

where s and $\boldsymbol{b} = (b_0\ b_1)^t$ are observed values of S and $\boldsymbol{B}$. The mean of m future values of Y at $x = x_0$ is given by

$$\overline{Y}_m^\star = \beta_0 + \beta_1 x_0 + \overline{E}_m^\star, \tag{25}$$

where $\overline{E}_m^\star \sim \mathrm{N}(0,\sigma^2/m)$ and is independent of $\boldsymbol{\Phi}$. Replacing σ and $\boldsymbol{\beta}$ in (25) with their fiducial quantities $\tilde{\sigma}$ and $\tilde{\boldsymbol{\beta}}$, we obtain the fiducial prediction quantity for $\overline{Y}_m^\star$ as

$$\tilde{\overline{Y}}_m^\star = b_0 + b_1 x_0 - \frac{s}{\sqrt{V/(n-2)}}\left[(t_{11}+t_{21}x_0)Z_1 + t_{22}x_0 Z_2 - \frac{1}{\sqrt{m}} Z_3\right],$$

where Z_1, Z_2, and Z_3 are mutually independent standard normal random variables ($\boldsymbol{\Phi} = (Z_1\ Z_2)^t$ and $\overline{E}_m^\star = \sigma Z_3/\sqrt{m}$), and t_{ij} are the (i,j)th elements of $\boldsymbol{T}$. Routine calculation shows that $(t_{11}+t_{21}x_0)Z_1 + t_{22}x_0 Z_2 - Z_3/\sqrt{m}$ is distributed as normal with mean 0 and variance

$$\delta^2 = \frac{1}{m} + \frac{1}{n} + \frac{(x_0-\overline{x})^2}{\sum_{i=1}^n (x_i-\overline{x})^2}.$$

Thus, we can write

$$\tilde{\overline{Y}}_m^\star = b_0 + b_1 x_0 - \delta s T_{n-2},$$

where T_{n-2} is distributed as t distribution with $n-2$ degrees of freedom. A $1-\alpha$ two-sided prediction interval on $\overline{Y}_m^\star$ is obtained from

$$\Pr[L < \tilde{\overline{Y}}_m^\star < U] = 1-\alpha,$$

which produces $(L,U) = b_0 + b_1 x_0 \pm t_{1-\alpha/2,n-2}\delta s$. This prediction interval is identical to the one derived by use of the frequentist approach, for example, see Draper and Smith (1981).

3.4. Exponential distribution

Let X_i, $i = 1,\ldots,n$, be a random sample from the exponential distribution, exponential(λ), with density function

$$f(x) = \lambda e^{-\lambda x} I_{(0,\infty)}(x), \quad \lambda > 0.$$

This example is also used by Lawless and Fredette (2005). Let gamma(α,λ) denote the gamma distribution with density function

$$f(x) = \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x} I_{(0,\infty)}(x), \quad \alpha > 0,\ \lambda > 0.$$

It is known that $\sum_{i=1}^n X_i$ is distributed as gamma(n,λ). Furthermore, $\lambda\sum_{i=1}^n X_i$ is distributed as gamma(n,1). Let $F_n(\cdot)$ be the distribution function of gamma(n,1). Since

$$U = F_n\left(\lambda \sum_{i=1}^n X_i\right) \sim \mathrm{uniform}(0,1)$$

we have the following structural equation:

$$\sum_{i=1}^{n} X_i = \frac{F_n^{-1}(U)}{\lambda}$$

for the model. As a consequence, a fiducial quantity for λ is given by

$$\tilde{\lambda} = \frac{F_n^{-1}(U)}{\sum_{i=1}^{n} x_i}. \quad (26)$$

Let Y_i, $i=1,\ldots,m$, be the future observations from, and $G_\lambda(\cdot)$ be the distribution of, exponential(λ). Since

$$U_i^\star = G_\lambda(Y_i) = 1 - e^{-\lambda Y_i} \sim \text{uniform}(0,1)$$

we can write

$$Y_i = \frac{-\log(1-U_i^\star)}{\lambda}. \quad (27)$$

Substituting λ in (27) with $\tilde{\lambda}$ of (26), we obtain a fiducial prediction quantity for Y_i as

$$\tilde{Y}_i = \frac{-\log(1-U_i^\star)}{F_n^{-1}(U)} \sum_{j=1}^{n} x_j.$$

Prediction limits on functions of Y_i can be obtained from the appropriate quantiles of the corresponding distribution of functions of $\tilde{Y}_i$.

It is easily seen that the distribution of $\tilde{Y}_i$ is identical to the distribution of $G_{\tilde{\lambda}}(Y_i)$ as indicated in (2). Also, for the case of a single future observation, i.e., $m=1$, the distribution function of $\tilde{Y}$ is given by

$$\Pr[\tilde{Y} < y] = \Pr\left[\frac{-\log(1-U^\star)}{F_n^{-1}(U)} < \frac{y}{n\overline{x}}\right].$$

Since

$$\frac{-\log(1-U^\star)}{F_n^{-1}(U)} = \frac{W}{V} = \frac{1}{n} F_{2,2n},$$

where $W \sim$ exponential(1), $V \sim$ gamma(n,1) and are independent, and $F_{2,2n}$ is distributed as an F with degrees of freedom 2 and $2n$ (Lawless and Fredette, 2005), we have

$$\Pr[\tilde{Y} < y] = \int_0^{y/\overline{x}} \left(1+\frac{z}{n}\right)^{-(n+1)} dz = 1 - \left(1+\frac{y}{n\overline{x}}\right)^{-n},$$

which is identical to the distribution derived based on the pivotal method.

Although the closed-form expression of the predictive distribution is available in this example, it is easier, however, to obtain the desired prediction limits by simulation in practice. For example, an upper 95% prediction limit for a single future observation ($m=1$) when $n=30$ and $\overline{x}=1$ can be calculated by use of the following R code based on 500 000 Monte Carlo samples:

```
n <-30 # current sample sizes
xbar <-1 # observed mean
nrun <-500 000 # Monte Carlo samples
U1 <-runif(nrun)
lambda <-qgamma(U1,n)/(n*xbar) # FQ for lambda
U2 <-runif(nrun)
Y <--log(1-U2)/lambda
quantile(Y,0.95) #95% upper bound
```

3.5. Gamma distribution

Let X_i, $i=1,\ldots,n$, be a random sample from gamma(α,λ). Let $F_{\alpha,\lambda}(\cdot)$ be the distribution function of gamma(α,λ). Simple calculation shows that

$$U_i = F_{\alpha,\lambda}(X_i) = F_{\alpha,1}(\lambda X_i) \sim \text{uniform}(0,1).$$

Denote the inverse distribution function $B(U_i,\ \alpha,\lambda) = F_{\alpha,\lambda}^{-1}(U_i)$. Thus, we have the following structural equations:

$$X_i = B(U_i,\ \alpha,\lambda), \quad i=1,\ldots,n \quad (28)$$

for the model. A fiducial distribution of (α,λ) can be obtained based on these equations.

In this model, a closed-form fiducial quantity for (α,λ) is not available. Thus, fiducial prediction quantities for this model must be based on realizations generated from the fiducial distribution of (α,λ). As an example, a fiducial prediction quantity for a future observation Y from gamma(α,λ) can be derived from

$$\tilde{Y} = B(U^\star,\ \tilde{\alpha},\tilde{\lambda}),$$

where $(\tilde{\alpha},\tilde{\lambda})$ is a realization from the fiducial distribution of (α,λ), and $U^\star \sim \text{uniform}(0,1)$. The distribution of $\tilde{Y}$ can be estimated from a large number of realizations by simulation. The key requirement here is the ability to generate realizations from the fiducial distribution of (α,λ).

To this end, Hannig (2009b) derives the density of the fiducial distribution for an iid sample from an absolutely continuous distribution having a distribution function $F_{\boldsymbol{\theta}}(x)$ and density $f_{\boldsymbol{\theta}}(x)$, where $\boldsymbol{\theta} = (\theta_1,\ldots,\theta_p)$ is p-dimensional. If the structural equations are formed using the inverse distribution function, then the fiducial density of $\boldsymbol{\theta}$ is proportional to

$$L(\boldsymbol{\theta}|\boldsymbol{x})J(\boldsymbol{x},\boldsymbol{\theta}),$$

where $L(\boldsymbol{\theta}|\boldsymbol{x}) = \prod_{i=1}^n f_{\boldsymbol{\theta}}(x_i)$ is the likelihood

$$J(\boldsymbol{x},\boldsymbol{\theta}) = \sum_{i_1 < \cdots < i_p} \frac{|\det(\nabla_{\boldsymbol{\theta}} F_{\boldsymbol{\theta}}(x_{i_1}),\ldots,\nabla_{\boldsymbol{\theta}} F_{\boldsymbol{\theta}}(x_{i_p}))|}{f_{\boldsymbol{\theta}}(x_{i_1})\cdots f_{\boldsymbol{\theta}}(x_{i_p})}$$

and $\det(\nabla_{\boldsymbol{\theta}} F_{\boldsymbol{\theta}}(x_{i_1}),\ldots,\nabla_{\boldsymbol{\theta}} F_{\boldsymbol{\theta}}(x_{i_p}))$ is the determinant of the $p \times p$ matrix of gradients of $F_{\boldsymbol{\theta}}(x_{i_j})$, $j=1,\ldots,p$, computed with respect to $\boldsymbol{\theta}$. In the present case of structural equation (28), the likelihood is

$$L(\alpha,\lambda|\boldsymbol{x}) = \frac{\lambda^{n\alpha}\exp(-\lambda\sum_{i=1}^n x_i + (\alpha-1)\sum_{i=1}^n \log x_i)}{\Gamma(\alpha)^n}$$

and the Jacobian term is

$$J(\alpha,\lambda,\boldsymbol{x}) = \sum_{1\le i<j\le n} \frac{\left|\det\begin{pmatrix} \frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_i) & \frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_j) \\ \frac{\partial}{\partial\lambda}F_{\alpha,1}(\lambda x_i) & \frac{\partial}{\partial\lambda}F_{\alpha,1}(\lambda x_j)\end{pmatrix}\right|}{\lambda^{2\alpha}(x_ix_j)^{\alpha-1}e^{-\lambda(x_i+x_j)}/\Gamma(\alpha)^2} = (\lambda\alpha)^{-1}\sum_{1\le i<j\le n} x_ix_j\left|\frac{\Gamma(\alpha+1)\frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_i)}{(\lambda x_i)^\alpha e^{-\lambda x_i}} - \frac{\Gamma(\alpha+1)\frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_j)}{(\lambda x_j)^\alpha e^{-\lambda x_j}}\right|.$$

Thus the fiducial distribution of (α,λ) has a density proportional to

$$\frac{\lambda^{n\alpha-1}e^{-\lambda\sum_{i=1}^n x_i + (\alpha-1)\sum_{i=1}^n \log x_i}}{\alpha\Gamma(\alpha)^n}\sum_{1\le i<j\le n} x_ix_j\left|\frac{\Gamma(\alpha+1)\frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_i)}{(\lambda x_i)^\alpha e^{-\lambda x_i}} - \frac{\Gamma(\alpha+1)\frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_j)}{(\lambda x_j)^\alpha e^{-\lambda x_j}}\right|. \tag{29}$$

We use an example in Hamada et al. (2004) to illustrate the proposed fiducial procedure. In the example, the first breakdown times (in hours) of 20 machines were recorded. Assuming breakdown times are distributed as gamma(α,λ), a 90% one-sided (lower-limit) prediction interval to contain the fifth breakdown time of five similar machines that are to be used simultaneously is desired. Let Y_i, $i=1,2,\ldots,5$, be the breakdown times of five future machines, we are interested in constructing a lower-limit prediction interval for $Y_{(5)}$, where $Y_{(1)} < \cdots < Y_{(5)}$. We use simulation to estimate the distribution of $\tilde{Y}_{(5)}$, the fiducial prediction quantity for $Y_{(5)}$. A single realization of $\tilde{Y}_{(5)}$ may be generated as follows:

1. Obtain a realization $(\tilde{\alpha},\ \tilde{\lambda})$ from the fiducial distribution of (α,λ).
2. Generate five independent uniform $(0,1)$ random deviates $u_1,\ldots,u_5$ and calculate $y_i = B(u_i,\ \tilde{\alpha},\ \tilde{\lambda})$, $i=1,\ldots,5$.
3. Calculate the maximum, $y_{(5)}$, of $y_1,\ldots,y_5$.

To simulate $(\tilde{\alpha},\ \tilde{\lambda})$ from density (29), we implement an importance sampling algorithm. The algorithm produces a weighted sample from (29) and consequently a weighted sample from the predictive distribution of $y_{(5)}$. The details are given in Appendix. Based on a weighted sampling algorithm with effective sample size 9.644×10^6, the 0.1 quantile of these realizations is found to be 74.36 (hours), which is the desired prediction lower limit on the fifth breakdown time. For comparison, Hamada et al. (2004) determined the lower prediction limit to be 71.8 (hours) based on a Bayesian approach. Notice that because we are computing the lower prediction limit, higher values are more desirable.

To evaluate the true frequentist prediction probability of the proposed procedure, we used the MLE estimates of $\alpha = 0.8763, \lambda = 0.0110$ to generate 10 000 data sets. Based on these data sets we have computed the empirical coverage probability associated with the prediction limits using the proposed procedure. The estimated prediction probability of the nominal 10% prediction interval was estimated as 10.11% indicating acceptable coverage. Additionally, we performed a small scale simulation study. In this study we considered Gamma distribution with the scale parameter $\lambda = 1$ and the shape parameter $\alpha = 0.5$, 1, 10, 100 and sample sizes $n=5$, 25, 125. For each of these parameter combinations we simulated 10 000 data sets and found the upper and lower predictive interval for one future observation and the fifth largest future observation. The estimated coverage probabilities of 95% prediction interval were all between 94.71% and 95.69%; the estimated coverage probabilities of 5% prediction interval were all between 4.73% and 5.31% indicating coverage within expected simulation error of the stated value.

3.6. Weibull distribution

Let X_i, $i=1,\ldots,n$, be a random sample from the Weibull distribution, weibull(α,β), with density function

$$f(x)=\alpha\beta(\alpha x)^{\beta-1}e^{-(\alpha x)^\beta}I_{(0,\infty)}(x),\quad \alpha>0,\ \beta>0.$$

Using the distribution function we get

$$U_i=F_{\alpha,\beta}(X_i)=1-e^{-(\alpha X_i)^\beta} \tag{30}$$

is a uniform(0, 1) random variable. From (30) we have the following structural equations:

$$X_i=B(U_i,\alpha,\beta)=\frac{(-\log(1-U_i))^{1/\beta}}{\alpha},\quad i=1,\ldots,n \tag{31}$$

for the model. A fiducial distribution of (α,β) can be obtained based on these equations.

Applying the same method that was used for deriving a fiducial distribution of $(\alpha,\ \lambda)$ in the gamma-distribution case, we obtain the likelihood

$$L(\alpha,\beta|\boldsymbol{x})=\alpha^{n\beta}\beta^n\prod_{i=1}^n(x_i^{\beta-1})e^{-\alpha^\beta\sum_{i=1}^n x_i^\beta}$$

and the Jacobian

$$J(\alpha,\beta,\boldsymbol{x})=\sum_{1\le i<j\le n}\frac{x_ix_j|\log x_i-\log x_j|}{\alpha\beta}.$$

Consequently, the density of the fiducial distribution of (α,β) is proportional to

$$\alpha^{n\beta-1}\beta^{n-1}\prod_{i=1}^n(x_i^{\beta-1})e^{-\alpha^\beta\sum_{i=1}^n x_i^\beta}. \tag{32}$$

We point out that in this case the fiducial distribution is the same as a Bayesian posterior computed using the reference prior $(\alpha\beta)^{-1}$ of Berger et al. (2009). Prediction intervals for the Weibull distribution can be constructed from realizations of the fiducial distribution of (α,β) in the same way as described in the gamma-distribution case.

We use an example in Krishnamoorthy et al. (2009) to illustrate the proposed fiducial procedure. In the example, the vinyl chloride concentration collected from 34 clean upgradient monitoring wells. As argued by Krishnamoorthy et al. (2009), Weibull distribution fits the data well. We computed the 95% one-sided (upper-limit) prediction interval for at least l of m observations at each of r locations; i.e., let Y_{ij}, $i=1,\ldots,r$, $j=1,\ldots,m$, be the future m observations at r locations, we are interested in constructing a upper-limit prediction interval for $Y=\max_{i=1,\ldots,r}Y_{i,(l)}$, where $Y_{i,(1)}<\cdots<Y_{i,(m)}$ for each $i=1,\ldots,r$.

We use simulation to estimate the distribution of $\tilde{Y}$, the fiducial prediction quantity for Y. A single realization of $\tilde{Y}$ may be generated as follows:

1. Obtain a realization $(\tilde{\alpha},\ \tilde{\beta})$ from the fiducial distribution of (α,β).
2. Generate $r\times m$ independent uniform (0, 1) random variables u_{ij} and calculate $y_{ij}=B(u_{ij},\ \tilde{\alpha},\tilde{\beta})$, $i=1,\ldots,r$, $j=1,\ldots,m$.
3. Calculate $y=\max_{i=1,\ldots,r}y_{i,(l)}$.

Using an importance sampling algorithm with effective sample size of 9.774×10^7 we took the 0.95 quantile of the sampled values of y. The computed values for r,m,l combinations taken from Krishnamoorthy et al. (2009) are summarized in Table 1; recall that smaller values are desirable for upper-limit prediction bounds. The computed values are compared to the values obtained by Krishnamoorthy et al. (2009) using a different generalized variable approach. Notice that the prediction bounds computed using the proposed method are all smaller than the prediction bounds of Krishnamoorthy et al. (2009).

To check the true frequentist coverage probability of the proposed procedure, we used the MLE estimates of $\alpha=0.52968,\beta=1.01022$ to generate 10 000 data sets. Based on this we have computed the empirical coverage probability associated with the prediction limits using the proposed procedure for the various r, m, and l values. The estimated

Table 1
95% Upper prediction limits for vinyl chloride data.

r	m	l	Fiducial	Krishnamoorthy et al. (2009)
1	2	1	2.962	2.974
10	2	1	5.480	5.483
10	3	1	3.595	3.618
10	3	2	6.784	6.797

coverage probability of the 95% prediction interval ranged from 0.9502 to 0.9506 indicating excellent coverage. Additionally, we perform a small scale simulation study. In this study we considered Weibull distribution with the scale parameter $\alpha=1$ and the shape parameter $\beta=0.1$, 0.2, 0.5, 1, 2, 5, 10 and sample sizes $n=5$, 25, 125. For each of these parameter combinations we simulated 10 000 data sets and found the predictive interval for the various r,m,l values from Krishnamoorthy et al. (2009). The estimated coverage probabilities of 95% prediction interval were all between 94.72% and 95.17% indicating coverage within expected simulation error of the stated value.

4. Concluding remarks

In this paper we have provided an approach for constructing prediction intervals for any model where a fiducial distribution of the parameters is available. Since a fiducial recipe is available for arbitrary statistical models (Hannig, 2009a), the proposed method provides a general approach for constructing prediction intervals. These fiducial prediction intervals coincide with the exact pivotal-based intervals, when they exist, and possess good statistical properties, otherwise.

This paper did not deal with prediction intervals for censored data. Censored data can be viewed as a special case of discretely observed data. Hannig et al. (2007) presented a fiducial inference on the parameters of discretely observed normal data. Hannig (2009b) studied asymptotic properties of fiducial inference for discretely observed data. Fiducial prediction intervals based on censored failure or life data are currently under investigation and will be reported in a future communication.

Appendix A

A.1. An importance sampling algorithm for fiducial distribution for gamma parameters

Let $f(x)$ be a density. Importance sampling is a standard Monte Carlo computational procedure for approximating quantities of the type $I=\int h(x)f(x)\,dx$ (Robert and Casella, 2004). To approximate the quantity I, we generate a sample $X_1,\dots,X_m$ from an instrumental density $g(x)$. Each of the observation is then associated with an unscaled weight $w_i=f(X_i)/g(X_i)$. This gives us a weighted sample generated from the density $f(x)$ and allows us to estimate

$$I\approx\frac{\sum_{i=1}^m w_i h(X_i)}{\sum_{i=1}^m w_i}.$$

By strong law of large numbers the approximation converges to the quantity of interest, provided that the support of f is included in the support of g.

In this paper we will be interested in computing α quantiles of distributions. This can be done by inverting the approximation

$$\alpha=\int_{-\infty}^{a} f(x)\,dx\approx\frac{\sum_{i=1}^m w_i I_{\{X_i\le a\}}}{\sum_{i=1}^m w_i}.$$

In particular, let $X_{(1)}<\cdots<X_{(m)}$ be the ordered sample with the corresponding weights $w_{(1)},\dots,w_{(m)}$. Here $w_{(i)}$'s are ordered to match the ordering of $X_{(i)}$'s and are not necessarily increasing. Let j be so that $\sum_{i=1}^{j} w_{(i)}\le\alpha\sum_{i=1}^m w_{(i)}<\sum_{i=1}^{j+1} w_{(i)}$ and estimate $a\approx X_{(j)}$.

The quality of the approximations depends on the choice of the instrumental distribution $g(x)$, and consequently the distribution of the weights w. For example, if the sample contains a small number of very large values, importance sampler is effectively using only a small number of observations to compute its estimator leading to large variance. For this reason, it is recommended to select $g(x)$ so that the center of g is roughly in the same place as the center of f but the tails of g are heavier than the tails of f. To measure the effective sample size (ESS) of the weighted sample, Kong et al. (1994) propose using

$$\mathrm{ESS}=\frac{(\sum_{i=1}^m w_i)^2}{\sum_{i=1}^m w_i^2}.$$

To sample from the fiducial distribution (29) we need to find a good instrumental distribution $g(\alpha,\lambda)$. We chose to use the following instrumental distribution. Denote the MLE of α by $\hat{\alpha}$, $c=n\log(n^{-1}\sum_{i=1}^n x_i)-\sum_{i=1}^n\log x_i$, and the density function of gamma(α,λ) by $f_{\alpha,\lambda}(x)$. Let

$$g(\alpha)=0.5f_{n-3,(n-3)/\hat{\alpha}}(\alpha)+0.5f_{\hat{\alpha}c,c}(\alpha).$$

This is a mixture of gamma distributions chosen so that its mean is at the maximum likelihood estimator and the tails are heavier than the tails of the fiducial distribution. Then let

$$g(\lambda|\alpha)=f_{n\alpha,\ \sum x_i}(\lambda).$$

So the instrumental distribution is given by

$$g(\alpha,\lambda) = g(\alpha)\ g(\lambda|\alpha)$$

and the weight is computed as

$$w(\alpha,\lambda) = \frac{\hat{J}(\alpha,\lambda)\Gamma(n\alpha)\left(\prod_{i=1}^{n} x_i\right)^{\alpha-1}}{\alpha\Gamma(\alpha)^n\left(\sum_{i=1}^{n} x_i\right)^{n\alpha} g(\alpha)},$$

where

$$\hat{J}(\alpha,\lambda) = \sum_{1 \le i < j \le n} x_i x_j \left| \frac{\Gamma(\alpha+1)\frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_i)}{(\lambda x_i)^{\alpha} e^{-\lambda x_i}} - \frac{\Gamma(\alpha+1)\frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x_j)}{(\lambda x_j)^{\alpha} e^{-\lambda x_j}} \right|.$$

Finally we remark that the derivative in $\hat{J}(\alpha,\lambda)$ needs to be computed numerically. To increase numerical stability of the numerical derivative in the tails we used the following identity:

$$\begin{aligned}\frac{\Gamma(\alpha+1)\frac{\partial}{\partial\alpha}F_{\alpha,1}(\lambda x)}{(\lambda x)^{\alpha} e^{-\lambda x}} &= \frac{\partial}{\partial\alpha}\frac{\Gamma(\alpha+1)F_{\alpha,1}(\lambda x)}{(\lambda x)^{\alpha} e^{-\lambda x}} - \frac{\Gamma(\alpha+1)F_{\alpha,1}(\lambda x)}{(\lambda x)^{\alpha} e^{-\lambda x}}(\psi(\alpha+1) - \log(\lambda x)) \\ &= -\frac{\partial}{\partial\alpha}\frac{\Gamma(\alpha+1)(1-F_{\alpha,1}(\lambda x))}{(\lambda x)^{\alpha} e^{-\lambda x}} + \frac{\Gamma(\alpha+1)(1-F_{\alpha,1}(\lambda x))}{(\lambda x)^{\alpha} e^{-\lambda x}}(\psi(\alpha+1) - \log(\lambda x))\end{aligned}$$

with $\psi(\alpha)$ is the digamma function. Thus instead of numerically computing $\partial F_{\alpha,1}(\lambda x)/\partial\alpha$, depending on which tail we are in, we numerically compute

$$\frac{\partial}{\partial\alpha}\frac{\Gamma(\alpha+1)F_{\alpha,1}(\lambda x)}{(\lambda x)^{\alpha} e^{-\lambda x}}$$

or

$$\frac{\partial}{\partial\alpha}\frac{\Gamma(\alpha+1)(1-F_{\alpha,1}(\lambda x))}{(\lambda x)^{\alpha} e^{-\lambda x}},$$

which is more stable. The computer code is available from authors open request.

References

Berger, J.O., Bernardo, J.M., Sun, D., 2009. The formal definition of reference priors. Annals of Statistics 37, 905–938.

Bhaumik, D.K., Gibbons, R.D., 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.

Draper, N.R., Smith, H., 1981. Applied Regression Analysis, second ed. Wiley, New York.

E, L., Hannig, J., Iyer, H.K., 2008. Fiducial intervals for variance components in an unbalanced two-component normal mixed linear model. Journal of the American Statistical Association 103, 854–865.

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

Fertig, K.W., Mann, N.R., 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., Meeker, W.Q., 1991. Statistical Intervals: A Guide for Practitioners. Wiley, New York.

Hamada, M., Johnson, V., Moore, L., Wendelberger, J., 2004. Bayesian prediction intervals and their relationship to tolerance intervals. Technometrics 46, 452–459.

Hannig, J., Iyer, H.K., Patterson, P., 2006. Fiducial generalized confidence intervals. Journal of the American Statistical Association 101, 254–269.

Hannig, J., Iyer, H.K., Wang, C.M., 2007. Fiducial approach to uncertainty assessment accounting for error due to instrument resolution. Metrologia 44, 476–483.

Hannig, J., 2009a. On generalized fiducial inference. Statistica Sinica 19, 491–544.

Hannig, J., 2009b. On Asymptotic Properties of Generalized Fiducial Inference for Discretized Data. Technical Report 09/02, Department of Statistics and Operations Research, The University of North Carolina.

Iyer, H.K., Wang, C.M., Matthew, T., 2004. Models and confidence intervals for true values in interlaboratory trials. Journal of the American Statistical Association 99, 1060–1071.

Kong, A., Liu, J., Wong, W., 1994. Sequential imputations and Bayesian missing data problems. Journal of the American Statistical Association 89, 278–288.

Krishnamoorthy, K., Lin, Y., Xia, Y., 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.

Krishnamoorthy, K., Mathew, T., Mukherjee, S., 2008. Normal-based methods for a gamma distribution: prediction and tolerance intervals and stress-strength reliability. Technometrics 50, 69–78.

Lawless, J.F., Fredette, M., 2005. Frequentist prediction intervals and predictive distributions. Biometrika 92, 529–542.

Nelson, W., 2000. Weibull prediction of a future number of failures. Quality and Reliability Engineering International 16, 23–26.

Nordman, D.J., Meeker, W.Q., 2002. Weibull prediction intervals for a future number of failures. Technometrics 44, 15–23.

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.

Robert, C.P., Casella, G., 2004. Monte Carlo Statistical Methods, second ed. Springer, New York.