Applied Stochastic Models in Business and Industry

WILEY

RESEARCH ARTICLE OPEN ACCESS

Simulated Exact Confidence Intervals: With Applications to Censored Exponential Reliability Data

Bo Henry Lindqvist | Gunnar Taraldsen

Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway

Correspondence: Bo Henry Lindqvist (bo.lindqvist@ntnu.no)

Received: 20 May 2024 | **Revised:** 7 October 2024 | **Accepted:** 9 October 2024

Keywords: confidence distribution | exact confidence bound | exponential distribution | fiducial inference | hybrid type II censoring | life testing | type I censoring

ABSTRACT

A method for constructing exact simulated confidence intervals is presented, valid for situations with both discrete and continuous observations. The idea of the method is to invert a data generating function, which needs not be monotone, and where special attention is taken when the data generating function contains jumps. The method is applied to obtain exact confidence intervals for certain types of censored data from exponential distributions. The censoring schemes under study are earlier treated in the literature, and a comparison to these approaches is considered. The connection to fiducial inference is discussed, and a difference in the paradigm of obtaining intervals for the parameter is studied.

1 | Introduction

Confidence intervals serve as indispensable tools in statistical inference, providing insights into the precision of parameter estimates. The motivation for the present paper comes from reliability engineering, where assessments of reliability are often based on limited samples. Thus, traditional methods for constructing confidence intervals, based on asymptotic theory, may be inappropriate in many applications. There has, on the other hand, been a certain interest in constructing exact confidence intervals, mainly by means of computer simulations. This is what the present paper is about.

Before turning to simulations, it should be noted 'en passant' that in reliability engineering it is well known how to derive an exact confidence interval for the parameter of an exponential distribution in the case of type II censoring [1]. Recall that type II censoring of a sample of lifetimes means censoring at the event of the rth failure for a given $r < n$. This is a simple example of the more general concept of life testing, for which there is a rich literature, and from which the main examples of the present paper are taken. Life testing is indeed in the core of reliability engineering, and has motivated much of the truly original and profound work by Nozer Singpurwalla throughout several decades, nicely documented in his monograph [2]. While Nozer Singpurwalla was indeed known as a dedicated and convinced Bayesian, it will be clear that our approach below is mostly within the frequentistic paradigm. We will, however, consider some connections to fiducial inference [3], which may be viewed as a bridge to Bayesian statistics [4].

Now suppose that T is a statistic with distribution, which depends on a single (real) unknown parameter θ. If T is stochastically increasing in θ, then it is in principle straightforward to construct an exact or conservative confidence interval for θ. For example, an upper α-confidence point when t is the observed value of T is obtained by solving the equation $F_\theta(t) = \alpha$ where F_θ is the cumulative distribution function (cdf) with respect to θ. The method is often known as "pivoting the cdf," see Casella and Berger [5].

This method has been used in the reliability literature for constructing exact confidence intervals for various types of censored exponential data [6–9]. The mentioned papers derive confidence intervals based on first calculating the distribution of the maximum likelihood estimator of the mean value of the underlying exponential distribution. While the three first of the cited papers *assumed* the stochastic monotonicity of the estimator, this property was later proven for fairly general settings of this kind by Balakrishnan and Iliopoulos [9].

In the present paper, we discuss how appropriate generalizations and modifications of the idea of "pivoting the cdf" can be used to derive exact confidence intervals more generally. The point of departure is the papers by Bølviken and Skovlund [10] and Lillegård and Engen [11], which also were the basis for the conference paper [12].

The basic assumption in these papers is that there is a random vector $\mathbf{U}$ with a known distribution, and a function $\tau(\mathbf{u},\theta)$ such that T under parameter θ has the same distribution as $\tau(\mathbf{U},\theta)$, for each possible θ. The function τ is then what in modern terminology is called a *data generating function* [13]. Lillegård and Engen [11] treated the case when for each possible observation t of T and each $\mathbf{u}$ there is a unique solution $\hat{\theta}(\mathbf{u},t)$ for θ of the equation $\tau(\mathbf{u},\theta)=t$. This excludes, however, the possibility of having a $\tau(\mathbf{u},\theta)$ with jumps or a $\tau(\mathbf{u},\theta)$ which is not everywhere increasing. The method in Bølviken and Skovlund [10] basically includes the above possibilities, but is often more difficult to use than the method of Lillegård and Engen [11]. We shall in Section 2 show how one may in a certain sense unify and generalize the ideas of [10, 11].

As indicated above, the approach of the paper may also be connected to fiducial inference. In a series of papers in the 1920s and 1930s, for example, [3], Sir Ronald Fisher introduced the concept of a fiducial distribution for a parameter, proposing to use this instead of the Bayesian posterior distribution for interval estimation of a parameter. Fisher's proposal was rather extensively discussed in the statistical literature of the 1950s and 1960s. Lindley [14] gave the necessary and sufficient conditions for reconciliation of the fiducial with the Bayesian inference in the one-parameter case. Savage [15] viewed the fiducial probability as "a bold attempt to make the Bayesian omelet without breaking the Bayesian egg." The topic became seemingly of less interest in the last part of the 20th century. In the last 10–15 years there has been, however, a renewed interest in the fiducial idea, particularly promoted by papers by Hannig [13, 16] and collaborators, but also by a wider group of researchers, see for example, [17] for further information. We will see in Sections 4 and 6 how the theory of fiducial intervals fits into and compares to the approach of the present paper.

The structure of the following sections is as follows. In Section 2, we present the basic theoretical results, including an example showing the well-known Clopper-Pearson intervals [18] to be a special case. In Section 3, we show how the main results are used in practice to derive confidence bounds of guaranteed coverage, including conditions for these bounds being exact. The connection to fiducial inference is considered in Section 4, to be exemplified later in Section 6. Section 5 shows in some detail how the general results from Section 3 are applied to type I censoring of exponential lifetimes. The corresponding use of fiducial inference is then considered in Section 6. As a further example, Section 7 considers certain censoring schemes known from the literature, named hybrid censoring. The relevant data generating function is presented. Section 8 contains some concluding remarks, while an Appendix A ends the paper.

2 | The Main Results

As discussed in Section 1, consider a situation where there is an unknown real parameter θ for which an observable random variable T is at hand for making inference. Let $\tau(\mathbf{U},\theta)$ be a data generating function corresponding to T. Recall from Section 1 that this means that T has the same distribution as $\tau(\mathbf{U},\theta)$ for each θ, where $\mathbf{U}$ has a known distribution. As a simple example, let $T=\sum_{i=1}^{n}X_i$ where the X_i are iid and $N(\theta,1)$ where θ is unknown. Then if $\mathbf{U}=(U_1,\dots,U_n)$ is a vector of independent standard normal variables, we can define $\tau(\mathbf{u},\theta)=n\theta+\sum_{i=1}^{n}u_i$.

Even if we do not essentially require $\tau(\mathbf{u},\cdot)$ to be non-decreasing in θ, we shall at least for intuition assume that large values of T are associated with large values of θ. For example this is natural if T is an estimator for θ.

Theorem 1. *With definitions as above, define* $\overline{\theta}(\mathbf{u},t)=\sup\{\theta:\tau(\mathbf{u},\theta)\le t\}$ *and let* $a(t)$ *be a left continuous function such that* $P_{\mathbf{U}}(\overline{\theta}(\mathbf{U},t)\ge a(t))\le\alpha$ *for all* t*, where* $0<\alpha<1$*. Then* $a(T)$ *is an upper* α *confidence bound (possibly conservative) for* θ*, that is,* $P_\theta(\theta\ge a(T))\le\alpha$ *for all* α.

Proof. Define $a^{-1}(\theta)=\sup\{t:a(t)\le\theta\}$. Then

$$\begin{aligned}P_\theta(\theta\ge a(T)) &\le_{(i)} P_\theta\left(T\le a^{-1}(\theta)\right)\\ &=_{(ii)} P_{\mathbf{U}}\left(\tau(\mathbf{U},\theta)\le a^{-1}(\theta)\right)\\ &\le_{(iii)} P_{\mathbf{U}}\left(\overline{\theta}\left(\mathbf{U},a^{-1}(\theta)\right)\ge\theta\right)\\ &\le_{(iv)} P_{\mathbf{U}}\left(\overline{\theta}\left(\mathbf{U},a^{-1}(\theta)\right)\ge a\left(a^{-1}(\theta)\right)\right)\\ &\le_{(v)} \alpha.\end{aligned}$$

Here (i) follows by the definition of $a^{-1}(\theta)$. The equality (ii) and the inequality (iii) follow by the respective definitions of $\tau(\mathbf{U},\theta)$ and $\overline{\theta}(\mathbf{u},t)$. Finally, (iv) follows since left continuity of $a(t)$ implies $a\left(a^{-1}(\theta)\right)\le\theta$, while (v) follows from the assumption made for $a(t)$. □

A lower α confidence bound for θ now can be obtained similarly, here formulated as a corollary to the above theorem:

Corollary 1. *Define* $\underline{\theta}(\mathbf{u},t)=\inf\{\theta:\tau(\mathbf{u},\theta)\ge t\}$ *and let* $b(t)$ *be a right continuous function such that* $P_{\mathbf{U}}(\underline{\theta}(\mathbf{U},t)\le b(t))\le\alpha$ *for all* t*, where* $0<\alpha<1$*. Then* $b(T)$ *is a lower* α *confidence bound (possibly conservative) for* θ*, that is,* $P_\theta(\theta\le b(T))\le\alpha$ *for all* θ.

The distributions of $\overline{\theta}(\mathbf{U},t)$ and $\underline{\theta}(\mathbf{U},t)$ are so called confidence distributions [19, 20]. The following result is of particular interest in applications, and is an easy consequence of the above results.

Corollary 2. *Suppose* $\tau(\mathbf{u},\theta)$ *for any fixed* $\mathbf{u}$ *is strictly increasing in* θ *(but may make jumps as a function of* θ*).*

Then $\underline{\theta}(\mathbf{u}, t) = \overline{\theta}(\mathbf{u}, t)$ *for all* $\mathbf{u}, t$. *If in this case,* $\overline{\theta}(\mathbf{U}, t)$ *has a continuous distribution, and* $a(t)$ *and* $b(t)$ *are chosen with equality in their respective definitions in Theorem 1 and Corollary 1, then the conclusions of Theorem 1 and Corollary 1 hold with equality, that is,* $P_\theta(\theta \geq a(T)) = \alpha$, $P_\theta(\theta \leq b(T)) = \alpha$.

The main examples of the paper, for example, Section 5, satisfy the conditions of Corollary 2. The following is an example where these conditions do not hold. This will be typical for situations with discrete observations.

2.1 | Example: Clopper-Pearson Interval

Let T be binomially distributed with success probability θ, where confidence limits for θ are sought based on an observation $T = t$ with $t \in \{0, 1, \dots, n\}$. The purpose of this example is to show that Theorem 1 and Corollary 1 lead to the well-known Clopper-Pearson interval [18]. To see this, let

$$\tau(\mathbf{U}, \theta) = \sum_{i=1}^{n} I(U_i < \theta)$$

be the data generating function for T, where $\mathbf{U} = (U_1, \dots, U_n)$ is a vector of independent uniform variables on $[0, 1]$. Then for $t \in \{0, 1, \dots, n\}$,

$$\underline{\theta}(\mathbf{U}, t) = U_{(t)}, \quad \overline{\theta}(\mathbf{U}, t) = U_{(t+1)},$$

where $U_{(1)}, \dots, U_{(n)}$ are the order statistics of $\mathbf{U}$, where we put $U_{(0)} = 0, U_{(n+1)} = 1$ by convention. Thus using Theorem 1 and Corollary 1 we may put $a(n) = 1$, $b(0) = 0$, and let $a(t)$ for $t < n$ and $b(t)$ for $t > 0$ be the solutions to, respectively, $P\left(U_{(t+1)} \geq a(t)\right) = \alpha$ and $P\left(U_{(t)} \leq b(t)\right) = \alpha$. Now it is known [21, p. 63]. that the order statistics of the standard uniform distribution are beta-distributed, namely

$$U_{(t)} \sim Beta(t, n+1-t); \quad t = 0, 1, \dots, n$$

Thus, letting $B(x; r, s)$ be the cumulative distribution function of $\text{Beta}(r, s)$, we have

$$a(t) = B(1-\alpha; t+1, n-t), \qquad b(t) = B(\alpha; t, n+1-t)$$

It is known (see e.g., Thulin [22]) that these correspond precisely to the confidence limits derived by Clopper and Pearson [18].

3 | Simulated Confidence Limits

In the same way as in the cited papers [10, 11], the results of the previous section are particularly designed for Monte Carlo computations of upper and lower confidence bounds. For this, suppose t is the observed value of T. Then we simulate a large number of independent realizations of $\mathbf{U}$ and calculate the corresponding $\overline{\theta}(\mathbf{U}, t)$ and $\underline{\theta}(\mathbf{U}, t)$. Now for a given α, the upper αth quantile of the empirical distribution of $\overline{\theta}(U, t)$ is an approximate upper αth quantile for θ, while similarly the lower αth quantile of the distribution of $\underline{\theta}(U, t)$ is an approximate lower αth quantile for θ.

More precisely, by careful modification of the ideas of Bølviken and Skovlund [10] and Lillegård and Engen [11], we have the following result:

Theorem 2. *Let the situation be as in Section 2. Let* $\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_m$ *be independent realizations of* $\mathbf{U}$ *and let* $\overline{\theta}_{(1)} \leq \overline{\theta}_{(2)} \leq \dots \leq \overline{\theta}_{(m)}$ *be the ordering of the values* $\overline{\theta}(\mathbf{U}_j, t)$ $(j = 1, \dots, m)$ *where* t *is the observed value of* T. *Then* $\overline{\theta}_{(m-k+1)}$ *is a (possibly conservative) exact upper* $(k/(m+1))$*-confidence bound for* θ.

Proof. Following [11], the clue is here that the $\overline{\theta}(\mathbf{U}_j, t)$ for $j = 1, \dots, m$, together with $\overline{\theta}(\mathbf{U}_0, t)$ with $\mathbf{U}_0$ being the underlying (latent) vector $\mathbf{U}$ which led to the observed value of T, form a sample of $m + 1$ independent and identically distributed variables. □

A similar result holds for lower bounds:

Corollary 3. *Let the situation be as in Section 2. Let* $\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_m$ *be independent realizations of* $\mathbf{U}$ *and let* $\underline{\theta}_{(1)} \leq \underline{\theta}_{(2)} \leq \dots \leq \underline{\theta}_{(m)}$ *be the ordering of the values* $\underline{\theta}(\mathbf{U}_j, t)$ $(j = 1, \dots, m)$ *where* t *is the observed value of* T. *Then* $\underline{\theta}_{(k)}$ *is a (possibly conservative) exact lower* $(k/(m+1))$*-confidence bound for* θ.

The following is a consequence of the two results above in view of Corollary 2:

Corollary 4. *Suppose* $\tau(\mathbf{u}, \theta)$ *for any fixed* $\mathbf{u}$ *is strictly increasing in* θ *(but may make jumps as a function of* θ*). Assume further that* $\overline{\theta}(\mathbf{U}, t)$ *has a continuous distribution. Let* $\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_m$ *be independent realizations of* $\mathbf{U}$ *and let* $\overline{\theta}_{(1)} \leq \overline{\theta}_{(2)} \leq \dots \leq \overline{\theta}_{(m)}$ *be the ordering of the values* $\overline{\theta}(\mathbf{U}_j, t)$ $(j = 1, \dots, m)$ *where* t *is the observed value of* T. *Then* $\overline{\theta}_{(k)}$ *and* $\overline{\theta}_{(m-k+1)}$ *are, respectively, exact lower and upper* $(k/(m+1))$ *bounds for* θ. *Thus* $\left(\overline{\theta}_{(k)}, \overline{\theta}_{(m-k+1)}\right)$ *is an exact upper* $(1 - 2k/(m+1))$*-confidence interval for* θ.

4 | Connection to Fiducial Inference

The fiducial approach as formalized by Hannig [13] takes a data generating function of the form $\tau(\mathbf{U}, \theta)$ (e.g., as in Section 2 above), as the basis for definition of a *fiducial distribution* for a parameter θ. In the case where the equation

$$t = \tau(\mathbf{u}, \theta) \tag{1}$$

can be inverted to solve uniquely for θ, with say,

$$\theta = \hat{\theta}(\mathbf{u}, t),$$

then the distribution of $\hat{\theta}(\mathbf{U}, t)$ is the simplest form of a fiducial distribution for θ.

The inverse to Equation (1) does, however, not necessarily exist. There are mainly two possible reasons for this: Either, there are more than one θ satisfying (1) for some values of t and $\mathbf{u}$, or there is no θ satisfying $t = \tau(\mathbf{u}, \theta)$.

The former situation is typically seen for discrete observations, for example as in the Clopper-Pearson example, and when $T = 0$ is observed in the application of Section 5. There are several possible solutions to defining a fiducial distribution in such cases. In the two cases mentioned above, our approach considers "upper" and "lower" solutions to (1) given by, respectively, $\overline{\theta}(\mathbf{u}, t)$ and $\underline{\theta}(\mathbf{u}, t)$.

For the latter situation, Hannig [16] suggested ignoring the values of $\mathbf{u}$ for which there is no solution for θ for an observed t. Thus, more precisely, the suggestion is to use the distribution of $\mathbf{U}$ conditional on the event $\{t = \tau(\mathbf{U}, \theta)$ for some $\theta\}$. The rationale for this choice is that we know (assuming our model to be true) that the observed data t have been generated from some fixed (but unknown) θ_0 and $\mathbf{u}_0$ with $t = \tau(\mathbf{u}_0, \theta_0)$. Values of $\mathbf{u}$ for which (1) does not have a solution for θ could not be the true $\mathbf{u}_0$. Hence only the values of $\mathbf{u}$ for which there is a solution should be considered in the definition of the generalized fiducial distribution.

We will see in the example of the next section how this conditioning may influence the coverage probability of the fiducial intervals.

5 | Application to Type I Censored Exponentially Distributed Lifetimes

Let $X_1, \ldots, X_n$ be the potential lifetimes of n items put on test at time $t = 0$, assumed to be iid exponentially distributed with unknown hazard rate $\theta > 0$. Suppose the lifetimes are subject to a type I censoring at a given time $c > 0$, that is, failure times exceeding c are not observed and hence our observations are $Y_i = \min(X_i, c)$ for $i = 1, \ldots, n$ (see e.g., [1]).

Let

$$R = \sum_{i=1}^{n} I(Y_i = X_i),$$

that is, the number of observed failures among the n items. Then it is straightforward to show (e.g., Meeker and Escobar [1]) that the maximum likelihood estimator of θ based on the Y_i is given as

$$T = \frac{R}{S}, \tag{2}$$

where $S = \sum_{i=1}^{n} Y_i$. Statistical inference about θ based on the asymptotic distribution of T is well established [1]. But as concluded, for example, by Bartholomew [23], the asymptotic sampling theory is inadequate in this case unless the sample size is very large. Thus, we shall consider the approach of Section 2 in order to do exact inference for θ.

Let $\mathbf{U} = (U_1, \ldots, U_n)$ be iid exponentially distributed with expected value 1. In this case, for $\theta > 0$, U_i/θ is exponentially distributed with hazard rate θ and is hence distributed as X_i, $i = 1, \ldots, n$. Now T, when the true parameter is θ, is distributed as $\tau(\mathbf{U}, \theta)$ where

$$\tau(\mathbf{u}, \theta) = \frac{\sum_{i=1}^{n} I(u_i/\theta < c)}{\sum_{i=1}^{n} \min(u_i/\theta, c)}, \tag{3}$$

defined for $\mathbf{u} = (u_1, \ldots, u_n)$ where $u_i > 0$ for $i = 1, \ldots, n$.

Let $u_{(1)} < u_{(2)} < \ldots, u_{(n)}$ be the ordered values of the components of $\mathbf{u}$. Then it is clear that, for fixed $\mathbf{u}$, the denominator of $\tau(\mathbf{u}, \theta)$ is continuous and strictly decreasing in θ for $\theta \geq u_{(1)}/c$, and is constant equal to nc for $0 < \theta < u_{(1)}/c$. Moreover, the numerator is non-increasing in θ, or more precisely, piecewise constant on intervals of the form $[u_{(i)}/c, u_{(i+1)}/c)$, making upward jumps at each $\theta = u_{(i)}$. It follows that, for fixed $\mathbf{u}$, $\tau(\mathbf{u}, \theta)$ is strictly increasing in θ for $\theta \geq u_{(1)}/c$, and is constant for $0 < \theta < u_{(1)}/c$.

Thus, if the observed value of T is $t > 0$, then by Corollary 4, the simulation recipe of Section 3 leads to exact upper and lower confidence bounds for θ. Figure 1 shows the function $\tau(\mathbf{u}, \theta)$ for an example where $n = 2$.

We now turn to the derivation of $\overline{\theta}(\mathbf{u}, t)$. Suppose we have observed n units until time c, giving the value t for T in (2). Then we draw m independent samples (where m is a large number) $\mathbf{u} = (u_1, \ldots, u_n)$ from the unit exponential distribution. For each $\mathbf{u}$, we solve the equation $\tau(\mathbf{u}, \theta) = t$ for θ, to obtain a value for $\overline{\theta}(\mathbf{u}, t)$ *if* there is a solution; *otherwise* we let $\overline{\theta}(\mathbf{u}, t)$ be the value of θ such that $\tau(\mathbf{u}, \theta-) < t < \tau(\mathbf{u}, \theta+)$, in case there is a jump of $\tau(\mathbf{u}, \theta)$ at θ. We refer again to Figure 1 for an illustration with $n = 2$. The figure shows that for a given $\mathbf{u}$, and observation $t > 0$, there is either a crossing of the solid curve, or there is a crossing of a dotted vertical line corresponding to a jump of $\tau(\mathbf{u}, \theta)$. In any case, there corresponds a unique value for $\overline{\theta}(\mathbf{u}, t)$.

For $t = 0$, on the other hand, there is no such unique crossing, which leads to different values for $\overline{\theta}(\mathbf{u}, t)$ and $\underline{\theta}(\mathbf{u}, t)$, namely $\overline{\theta}(\mathbf{u}, 0) = u_{(1)}/c$ and $\underline{\theta}(\mathbf{u}, 0) = 0$. Here, and for the case of general n, only $b(0) < 0$ would satisfy the requirement for a lower confidence bound. We, therefore, concentrate on the upper bound which should be based on $\overline{\theta}(\mathbf{u}, 0) = u_{(1)}/c$. Thus, the upper bound $a(0)$ is found by solving $P_{\mathbf{U}}(U_{(1)} \geq c \;\; a(0)) = \alpha$. This is readily seen to give the following upper bound for the hazard rate,

$$a(0) = \frac{-\log \alpha}{nc},$$

which in fact is a standard result (e.g., Meeker and Escobar [1, p. 168]).

An explicit general expression for $\overline{\theta}(\mathbf{u}, t)$ based on $\tau(\mathbf{u}, \theta)$ in (3) is given by equation (A1) in the Appendix A.

Note that if c tends to infinity, then the numerator of (3) tends to the constant n, which suggests to use the (maximum likelihood) estimator n/S. This statistic may of course be used as the basis for a confidence interval also for finite c, and will then have the data generating function

$$\tau_S(\mathbf{u}, \theta) = n\left(\sum_{i=1}^{n} \min(u_i/\theta, c)\right)^{-1}$$

It follows from the above that τ_S is constant equal to $1/c$ for θ between 0 and $u_{(1)}/c$ and is continuous and strictly increasing for $\theta > u_{(1)}/c$. Exact confidence bounds can then be found from this.

We may also use the numerator R of T as the basis for a confidence interval. The inference for this case is clear from the fact that R

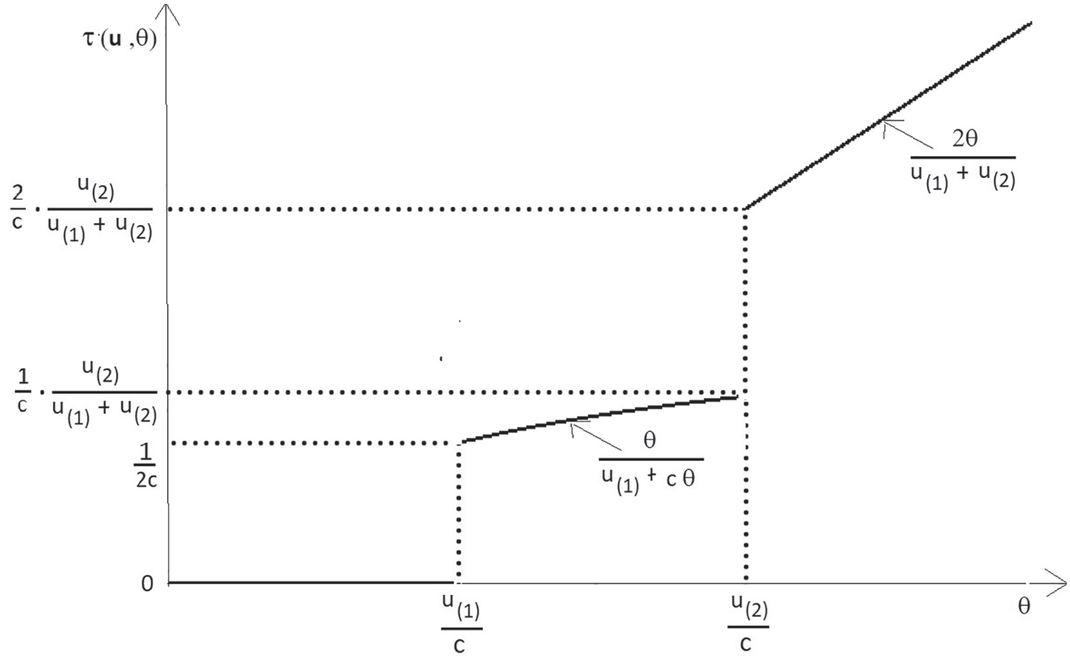

FIGURE 1 | The function $\tau(\mathbf{u}, \theta)$ (solid line) for Type I censored exponential variables with $n = 2$.

is binomially distributed with parameters n and $p = 1 - e^{-c/\theta}$. Thus, as already seen, the use of the machinery in Section 2 leads to the Clopper-Pearson interval for p which then in the end needs to be transformed to an interval for θ. The data generating function is in this case

$$\tau_R(\mathbf{u}, \theta) = \sum_{i=1}^{n} I(u_i/\theta < c)$$

6 | The Fiducial Distribution for Exponential Type I Censoring

Let the setting be as in the previous section. Recall the ideas of fiducial inference as described in Section 4. In the present section, we consider the construction of fiducial intervals for the parameter θ.

Consider first the case $n = 2$ illustrated in Figure 1. It is seen from the figure that for a given t and simulated $\mathbf{u} = (u_1, u_2)$ there are solutions for θ in (1) only if

$$\begin{array}{lll} \text{for } t = 0: & \text{when any } \theta \leq u_{(1)}/c \text{ is a solution, or} & \\ \text{for } \frac{1}{2c} < t < \frac{1}{c} \cdot \frac{u_{(2)}}{u_{(1)}+u_{(2)}}: & \text{when } \hat{\theta} = \frac{t u_{(1)}}{1-ct} \text{ is the unique solution, or} & (4) \\ \text{for } t > \frac{2}{c} \cdot \frac{u_{(2)}}{u_{(1)}+u_{(2)}}: & \text{when } \hat{\theta} = \frac{t(u_{(1)}+u_{(2)})}{2} \text{ is the unique solution.} & \end{array}$$

For other t, there is no solution.

As an example, suppose $c = 1$, and that we have observed one failure at time 0.5 and one censoring at time 1. Thus, we have $R = 1, Y_1 = 0.5, Y_2 = 1$, so that $S = 3/2$ and $T = 2/3$ in the notation of Section 5. With $t = 2/3$ we thus conclude from Figure 1 that

$$\bar{\theta}(\mathbf{u}, t) = 2u_{(1)} \qquad \text{if } \; u_{(2)} \geq 2u_{(1)}$$
$$\bar{\theta}(\mathbf{u}, t) = u_{(2)} \qquad \text{if } \; u_{(2)} < 2u_{(1)}$$

On the other hand, in view of (4), we have the fiducial distribution given by

$$\hat{\theta}(\mathbf{u}, t) = 2u_{(1)} \qquad \text{conditional on } \; u_{(2)} \geq 2u_{(1)} \tag{5}$$

We now simulated 100.000 pairs (U_1, U_2) of independent standard exponentially distributed variables, and plotted in Figure 2 the histograms of respectively $\bar{\theta}(\mathbf{u}, 2/3)$ and the corresponding fiducial version using (5). It is seen that the fiducial distribution in this case is more concentrated on low values. For example, a confidence interval obtained by deleting 2.5% in each end of the histogram gives for the left histogram (0.025, 2.84), while the right one gives the more "optimistic" one (0.017, 2.44).

In order to obtain a simulation recipe for simulation from the fiducial distribution for general n, we generalized the list in (4) by using (A1) in the Appendix A. As we saw in Figure 1, for a given $t > 0$ we will generally have no solution of $\tau(\mathbf{u}, \theta) = t$ if $\mathbf{u}$ satisfies

$$\frac{(i-1)\ u_{(i)}}{\left(\sum_{k=1}^{i} u_{(k)} + (n-i)u_{(i)}\right)c} < t \leq \frac{i\ u_{(i)}}{\left(\sum_{k=1}^{i} u_{(k)} + (n-i)u_{(i)}\right)c};\ i = 2, \ldots, n \tag{6}$$

In order to obtain the fiducial distribution, for a given t one hence needs to condition on $\mathbf{u}$ not giving t in such intervals.

In order to check the coverage probability of a fiducial interval as compared to the exact intervals as treated in this paper, we did a small simulation study for one given setting, letting $n = 10$, $c = 1$ and $\theta = 1$. We simulated 10.000 data sets, giving 10.000 values

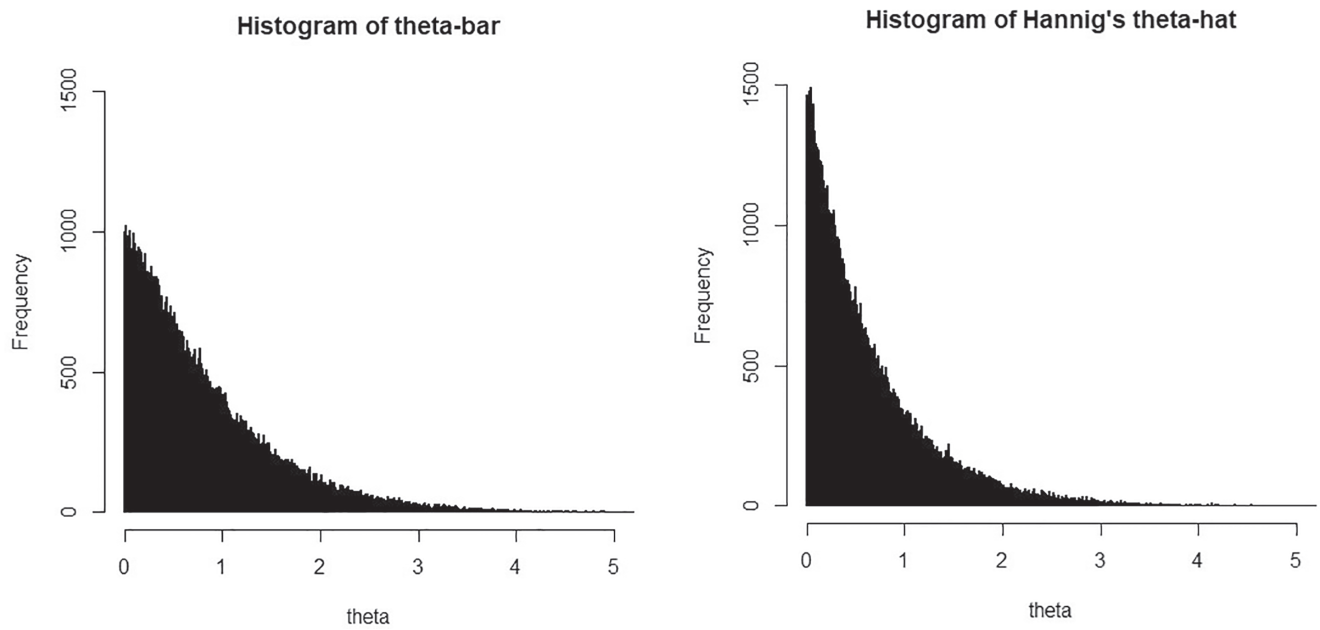

FIGURE 2 | Histograms of simulations of $\overline{\theta}(\mathbf{u}, t)$ (left) and the fiducial counterpart $\widehat{\theta}(\mathbf{u}, t)$ (right).

for the estimate t for θ. For each of them we simulated $m = 1000$ vectors $\mathbf{u}$ of 10 independent standard exponential variables. This gave us 1000 values for $\overline{\theta}(\mathbf{u}, t)$, from which we found lower and upper confidence bounds by excluding 2.5% of the observations in each end of the ordered set of values (see Section 3). For the resulting interval, we found a lower bound less than 1 in 97.54% of the cases, while the upper bound was above 1 in 97.48% of the cases. This of course corresponds very well to the claimed exactness of these bounds.

For the fiducial case, we again simulated 10.000 data sets, and hence 10.000 values of t, and then for each, we simulated 1000 *non-excluded* $\mathbf{u}$, using (6), and made the interval by taking away 2.5% values in each end of the ordered values. In the same manner as above, we now obtained a lower bound less than 1 in 97.58% of the cases, and an upper bound larger than 1 in 95.87% of the cases. The conclusion of fiducial intervals being more "optimistic," is thus seemingly still the case, as it was in the simple example above with $n = 2$. This example indicates that generalized fiducial distributions are not necessarily confidence distributions. It should be noted, however, that generalized fiducial intervals under standard conditions has the correct asymptotic coverage [16].

7 | Hybrid Censoring

While we in Section 5 considered estimation of the exponential *hazard* θ, it seems that much of the corresponding literature considers estimation of the exponential *expectation*, which in our notation is $1/\theta$. In this case the maximum likelihood estimator is undefined if there are no uncensored observations, that is, if $R = 0$ in our notation. It has, therefore, been common to condition on the event $R \geq 1$ [7, 23]. This may of course also be done in our estimation of exponential hazard since an estimate $\theta = 0$ seems uninteresting.

In our approach, it is seen from (3) that $R \geq 1$ is equivalent to $\{U_{(1)} < c\theta\}$. Conditioning on the event $R \geq 1$ would hence imply a modification of the data generating function $\tau(\mathbf{u}, \theta)$ in a way where the distribution of $\mathbf{U}$ depends on θ. This is however not compatible with $\tau(\mathbf{U}, \theta)$ being a data-generating function, where the requirement is that $\mathbf{U}$ has a distribution not depending on θ.

Childs et al. [7] studied various aspects of so-called Type-I and Type-II hybrid sampling for the exponential distribution. With the notation in Section 5, *Type-I hybrid sampling* means terminating the life testing experiment at the time $\min\left(X_{(r)}, c\right)$ for a given integer $r \geq 1$, thus avoiding to run the experiment until time c if an a priori given number r of failures have already happened. On the other hand, there is still a problem when no failures are observed until time c. Thus, for this case, [7] considered conditional distributions given at least one failure.

In order to avoid the problem with no failures, Childs et al. [7] also proposed the alternative *Type-II hybrid sampling* scheme, where the experiment is terminated at the time $\max\left(X_{(r)}, c\right)$ for r and c fixed in advance. This scheme has the advantage of guaranteeing that at least r failures are observed, and the problem of no observed failures disappears.

A data generating function can now be obtained as follows. Let R be the number of observed failures before time c. Then [7] the maximum likelihood estimator of θ is given by

$$\theta^* = \begin{cases} \frac{r}{\sum_{i=1}^{r} X_{(i)} + (n-r)X_{(r)}} & \text{if } R = 0, 1, \ldots, r-1 \\ \\ \frac{R}{\sum_{i=1}^{R} X_{(i)} + (n-R)c} & \text{if } R = r, r+1, \ldots, n \end{cases} \tag{7}$$

Let $\mathbf{U} = (U_1, \ldots, U_n)$ be iid exponentially distributed variables with expected value 1, so that U_i/θ is exponentially distributed with hazard rate θ. Hence θ^*, when the true parameter is θ, is distributed as $\tau^*(\mathbf{U}, \theta)$, where

$$\tau^*(\mathbf{u}, \theta) = \begin{cases} \frac{r}{\sum_{i=1}^{r}(u_{(i)}/\theta)+(n-r)(u_{(r)}/\theta)} & \text{if } R = 0, 1, \ldots, r-1 \\ \\ \frac{R}{\sum_{i=1}^{R} u_{(i)}/\theta+(n-R)c} & \text{if } R = r, r+1, \ldots, n \end{cases} \tag{8}$$

Here we now represent R as

$$R = \sum_{i=1}^{n} I\left(u_{(i)}/\theta < c\right).$$

For fixed $\mathbf{u}$, R makes jumps, as a function of θ, at each consecutive $u_{(i)}/c$. Thus, until immediately before $u_{(r)}/c$, the upper expression in (8) is in charge, and is in fact linear and strictly increasing starting from 0 for $\theta = 0$. In the point where $\theta = u_{(r)}/c$, the upper and lower expressions of (8) are equal, while for $\theta > u_{(r)}/c$ the lower expression is the same as (3) and is hence strictly increasing, with jumps at $u_{(i)}/c$ for $i = r+1, r+2, \ldots, n$. The conditions of Corollary 4 are, therefore, satisfied, and exact confidence intervals can, therefore, be found in the same manner as was done in Section 5.

For illustration, in the case $n = 2, \quad r = 1$, the curve $\tau^*(\mathbf{u}, \theta)$ would follow the solid line in Figure 1 from $\theta = u_{(1)}/c$ and upwards, but be a strictly increasing straight line for $\theta \leq u_{(1)}/c$, starting at 0 for $\theta = 0$ and equal $1/(2c)$ at $\theta = u_{(1)}/c$.

Childs et al. [7] derived the probability distribution of θ^* and assumed without proof that the survival function of θ^* is increasing in θ. With this assumption, they were able to derive exact confidence intervals. The claimed stochastic monotonicity follows, on the other hand, directly from the monotonicity of $\tau^*(\mathbf{u}, \theta)$ as a function of θ for fixed $\mathbf{u}$ as described above.

As noted earlier, Balakrishnan and Iliopoulos [9] have already settled the question of monotonicity for this and several related censoring schemes.

8 | Concluding Remarks

Starting from a general approach of obtaining simulated confidence intervals for single parameters, we have shown how the approach may be utilized for various types of censoring schemes for samples from the exponential distribution. Similar confidence intervals have already been derived by several authors, as cited. In these papers, the common method to obtain confidence intervals has been through the method of "pivoting the cdf" after calculating the distribution of the maximum likelihood estimator and inverting the corresponding cumulative distribution function. As already noted, such a procedure has often been complicated and an apparent challenge.

In the present paper, we have instead based the confidence intervals on data generating functions, which are closer to the intrinsic mechanisms of the involved failure and censoring processes. As was seen in both the ordinary type I censoring (Section 5) and the hybrid censoring (Section 7), stochastic monotonicity with respect to the parameter was a natural consequence of the modeling, while exactness of the confidence bounds was a consequence of the theoretical results of Section 3.

While the given censoring time c is fixed in the examples we have considered, it is interesting to see that the data generating functions for type I censoring easily can be generalized to allow the more general censoring called administrative censoring [24]. Administrative censoring is common in clinical and epidemiological studies, and means that potential censoring times are known even for subjects who fail. (It is perhaps less common in life testing of components as in reliability engineering.)

Consider equation (3). By replacing c by c_i in the sums of the numerator and denominator, it is seen that for fixed $\mathbf{u}$, the resulting data generating function jumps at $\theta = u_{(i)}/c_i$, while the denominator is still a decreasing function of θ. Monotonicity of the modified data generating function, therefore, follows, which in turn gives rise to exact confidence limits as seen in Section 5.

Regarding the hybrid sampling of Section 7, it is however not clear how to extend it to a case where the censoring time c varies among the items. The point there was essentially that if too few items failed by time c, then the censoring times for the unfailed ones were extended until the required number of items had ultimately failed. In some sense, a decision is then taken at time c whether to continue testing the remaining unfailed items. But with different c_i, if, say, one or more of the items are censored early (low c_i), then it is not clear how to define the "type II" part of the censoring scheme.

As discussed in Section 7, the conditioning on at least one failure in type I censoring, does not induce a straightforward modification of the data generating functions considered in the present paper. Of course, if the cdf of the maximum likelihood estimator for θ is known, then the inverse cdf is a valid data generating function being formulated in terms of a single uniform variable U. However, in view of (3), it would be nice to have a data generating function depending more on the process itself. We leave it as an open problem to find a reasonable solution to this issue.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

References

1. W. Q. Meeker and L. A. Escobar, *Statistical Methods for Reliability Data* (New York: Wiley, 1998).

2. N. D. Singpurwalla, *Reliability and Risk: A Bayesian Perspective* (Hoboken , NJ: John Wiley & Sons, 2006).

3. R. A. Fisher, "Inverse Probability," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 26 (Cambridge, UK: Cambridge University Press, 1930), 528–535.

4. B. Efron, "R. A. Fisher in the 21st Century," *Statistical Science* 13, no. 2 (1998): 95–114.

5. G. Casella and R. L. Berger, *Statistical Inference*, 2nd ed. (Pacific Grove, CA: Duxbury, 2002).

6. S. M. Chen and G. K. Bhattacharyya, "Exact Confidence Bounds for an Exponential Parameter Under Hybrid Censoring," *Communications in Statistics—Theory and Methods* 16, no. 8 (1987): 2429–2442.

7. A. Childs, B. Chandrasekar, N. Balakrishnan, and D. Kundu, "Exact Likelihood Inference Based on Type-I and Type-II Hybrid Censored Samples From the Exponential Distribution," *Annals of the Institute of Statistical Mathematics* 55 (2003): 319–330.

8. B. Chandrasekar, A. Childs, and N. Balakrishnan, "Exact Likelihood Inference for the Exponential Distribution Under Generalized Type-I and Type-II Hybrid Censoring," *Naval Research Logistics* 51, no. 7 (2004): 994–1004.

9. N. Balakrishnan and G. Iliopoulos, "Stochastic Monotonicity of the MLE of Exponential Mean Under Different Censoring Schemes," *Annals of the Institute of Statistical Mathematics* 61 (2009): 753–772.

10. E. Bølviken and E. Skovlund, "Confidence Intervals From Monte Carlo Tests," *Journal of the American Statistical Association* 91, no. 435 (1996): 1071–1078.

11. M. Lillegård and S. Engen, "Exact Confidence Intervals Generated by Conditional Parametric Bootstrapping," *Journal of Applied Statistics* 26, no. 4 (1999): 447–459.

12. B. H. Lindqvist, "Exact Confidence Intervals by Simulation," in *Proceedings of 53rd Session of International Statistical Institute (ISI)* (Seoul, 2001).

13. J. Hannig, H. Iyer, R. C. Lai, and T. C. Lee, "Generalized Fiducial Inference: A Review and New Results," *Journal of the American Statistical Association* 111, no. 515 (2016): 1346–1361.

14. D. V. Lindley, "Fiducial Distributions and Bayes' Theorem," *Journal of the Royal Statistical Society: Series B: Methodological* 20, no. 1 (1958): 102–107.

15. L. J. Savage, "The Foundations of Statistics Reconsidered," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics* (University of California Press, 1961), 575–587.

16. J. Hannig, "On Generalized Fiducial Inference," *Statistica Sinica* 19, no. 2 (2009): 491–544.

17. G. Taraldsen and B. H. Lindqvist, "Fiducial Theory and Optimal Inference," *Annals of Statistics* 41, no. 1 (2013): 323–341.

18. C. J. Clopper and E. S. Pearson, "The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial," *Biometrika* 26, no. 4 (1934): 404–413.

19. T. Schweder and N. L. Hjort, "Confidence and Likelihood," *Scandinavian Journal of Statistics* 29, no. 2 (2002): 309–332.

20. T. Schweder and N. L. Hjort, *Confidence, Likelihood, Probability* (Cambridge, UK: Cambridge University Press, 2016).

21. J. E. Gentle, *Computational Statistics* (New York: Springer, 2010).

22. M. Thulin, "The Cost of Using Exact Confidence Intervals for a Binomial Proportion," *Electronic Journal of Statistics* 8 (2014): 817–840.

23. D. Bartholomew, "The Sampling Distribution of an Estimate Arising in Life Testing," *Technometrics* 5, no. 3 (1963): 361–374.

24. M. M. Joffe, "Administrative and Artificial Censoring in Censored Regression Models," *Statistics in Medicine* 20, no. 15 (2001): 2287–2304.

Appendix A

The General Solution for Type-I Censoring

As a function of θ, the data generating function $\tau(\mathbf{u}, \theta)$ in (3) makes upward jumps at $\theta = u_{(i)}/c$ for $i = 1, 2, \ldots, n$, respectively

$$\text{from} \quad \frac{(i-1)\ u_{(i)}}{\left(\sum_{k=1}^{i} u_{(k)} + (n-i)u_{(i)}\right)c} \quad \text{to} \quad \frac{i\ u_{(i)}}{\left(\sum_{k=1}^{i} u_{(k)} + (n-i)u_{(i)}\right)c}$$

The value of the function $\tau(\mathbf{u}, \theta)$, for fixed $\mathbf{u}$, in the interval for θ from $u_{(i-1)}/c$ to $u_{(i)}/c$ is

$$\frac{(i-1)\theta}{\sum_{k=1}^{i-1} u_{(k)} + (n-i+1)c\theta}$$

while for $\theta > u_{(n)}/c$ it is

$$\frac{n\theta}{\sum_{i=1}^{n} u_{(i)}}$$

Thus, we can express $\overline{\theta}(\mathbf{u}, t)$ as

$$\overline{\theta}(\mathbf{u}, t) = \begin{cases} u_{(1)}/c & \text{for } 0 \le t \le \frac{1}{nc} \\ u_{(i)}/c & \text{for } \frac{(i-1)\ u_{(i)}}{(\sum_{k=1}^{i} u_{(k)} + (n-i)u_{(i)})c} < t \le \frac{i\ u_{(i)}}{(\sum_{k=1}^{i} u_{(k)} + (n-i)u_{(i)})c}; \quad i = 2, \ldots, n \\ \frac{t\sum_{k=1}^{i} u_{(k)}}{i-(n-i)ct} & \text{for } \frac{i\ u_{(i)}}{(\sum_{k=1}^{i} u_{(k)} + (n-i)u_{(i)})c} < t \\ & \le \frac{i\ u_{(i+1)}}{(\sum_{k=1}^{i+1} u_{(k)} + (n-i-1)u_{(i+1)})c}; \; i = 1, \ldots, n-1 \\ \frac{t\sum_{k=1}^{n} u_{(k)}}{n} & \text{for } t > \frac{nu_{(n)}}{(\sum_{k=1}^{n} u_{(k)})c} \end{cases} \tag{A1}$$