Comparison of confidence procedures for type I censored exponential lifetimes

Rolf Sundberg
Stockholm University, Sweden

May 28, 2000

Abstract

In the model of type I censored exponential lifetimes, coverage probabilities are compared for a number of confidence interval constructions proposed in literature. The coverage probabilities are calculated exactly for sample sizes up to 50 and for different degrees of censoring and different degrees of intended confidence. If not only a fair two-sided coverage is desired, but also fair one-sided coverages, only few methods are quite satisfactory. A likelihood-based interval and a third root transformation to normality work almost perfectly, but the χ^2-based method that is perfect under no censoring and under type II censoring can also be advocated.

Keywords: Confidence interval, coverage probability, exponential distribution, failure times, fixed censoring.

1 Introduction

The exponential distribution with its constant hazard rate forms a basic lifetime distribution model, which has an important role in reliability studies. It appears in quite different fields of applications, as well. For example, when splicing optical fibres a Gaussian two-dimensional offset causes a χ^2_2-distributed loss, that is an exponential loss (Tyrcha et al, 2000). For a random sample from an exponential distribution a complete and exact parametric inference methodology (including confidence procedures) is found in statistics textbooks, see for example Cox & Oakes (1984). In practice, however, the pure exponential model is rarely applicable. One of the reasons is that long or very long lifetimes will be censored, that is their existence will be recorded, but not the particular values. The two basic types of censoring are type I and type II, also called time censoring and failure censoring, respectively. In situations of type I, all observations exceeding a fixed value are censored, whereas in type II a specified number of the largest values are censored. The exact theory for complete data has a perfect analogy with type II censored data, whereas no exact confidence theory is available for type I censoring. In this paper we will compare several approximate (asymptotic) confidence procedures for type I censored exponential data, with a common censoring time C, that have been proposed in the literature. Most of them are primarily designed for use with large samples, but we will study their performance on small or medium sized samples.

2 Method descriptions

We will study confidence procedures based on the maximum likelihood (ML) estimator of the mean value parameter θ of the exponential distribution, that is the exponential is parametrized as

$$f(y;\theta) = (1/\theta)\, e^{-y/\theta}\,,\ \text{y} > 0,\, \theta > 0.$$

However, the performance measures to be used depend only on whether the confidence intervals cover the true parameter or not, so they are invariant under reparametrizations. Therefore we could equally well parametrize by the intensity (hazard rate) parameter $\rho = 1/\theta$, and some confidence intervals for θ to be studied below are in fact more easily motivated in terms of ρ.

Before we present the confidence procedures to be compared we give a brief account of the theoretical basics we require for the exponential model without and with censoring.

2.1 Basics for complete data

For a complete sample of size n from the exponential distribution the log-likelihood for θ is

$$\log L(\theta) = -\textstyle\sum_1^n y_i/\theta - n \log\theta\,, \qquad (1)$$

from which the well-known ML estimator follows:

$$\hat{\theta} = \textstyle\sum_1^n y_i/n = \bar{y}\,.$$

From the proportionality of the exponential to the χ^2_2 it follows that the distribution for $\hat{\theta}/\theta$ is $\chi^2_{2n}/2n$, and this distribution also yields exact confidence limits. For example the upper bounded $(1-p)$-interval is

$$\theta \leq 2n\,\hat{\theta}/\chi^2_{2n,p}\,,$$

where $\chi^2_{f,p}$ is the p-quantile of the χ^2 distribution with f degrees of freedom (χ^2_f). A (one-sided) likelihood-based interval takes the same form, $\theta/\hat{\theta} <$ constant.

For large n we can use the asymptotic normal distribution for $\hat{\theta}$, or equivalently the similarity of the χ^2_{2n} to the normal. One way to obtain the asymptotic variance is as the inverse $I(\theta)^{-1}$ of the Fisher information,

$$I(\theta) = n/\theta^2.$$

We obtain the approximation

$$\hat{\theta}/\theta \sim \mathrm{N}(1, 1/n),$$

and from this the asymptotic upper bounded $100(1-p)\%$-interval

$$\theta \leq \hat{\theta}\,(1 + z_{1-p}/\sqrt{n}),$$

if we use $\hat{\theta}/\sqrt{n}$ as the standard error for $\hat{\theta}$, or

$$\theta \leq \hat{\theta}/(1 - z_{1-p}/\sqrt{n}),$$

by solving for θ instead. Here z_p is the (p)-quantile of the standard normal distribution (i.e. $\Phi(z_{1-p}) = 1 - p$).

2.2 Basics under type I censoring.

Under type I censoring, with a binomial $\mathrm{Bin}(n, P_\theta = 1-\exp(-C/\theta))$ number $N \leq n$ of uncensored lifetimes, the log-likelihood is

$$\log L(\theta) = -\textstyle\sum_1^n y_i/\theta - N\log\theta + \text{constant}. \tag{2}$$

In the sum, when taken over all n observations, the $n-N$ censored ones are given the censoring limit value $y_i = C$. Note that the middle term of (2) has the random factor $N \leq n$ where (1) had the full sample size n. Hence the ML estimator is now

$$\hat{\theta} = \textstyle\sum_1^n y_i/N\,. \tag{3}$$

Exact χ^2 proportionality no longer holds, and $\hat{\theta}$ is not unbiased. In type II censoring, however, when N has been fixed instead of C, $\hat{\theta}$ has the same form (3) but $\hat{\theta}/\theta$ is now exactly $\chi^2_{2N}/2N$ distributed.

The observed information is

$$J(\theta) = -\frac{d^2}{d\theta^2}\log L(\theta) = \frac{2\sum_1^n y_i - N\,\theta}{\theta^3},$$

and in particular $J(\hat{\theta}) = N/\hat{\theta}^2$. Its expected value, the Fisher information, is

$$I(\theta) = E\,J(\theta) = n\,\frac{1-\exp(-C/\theta)}{\theta^2},$$

with

$$I(\hat{\theta}) = n\,\frac{1-\exp(-C/\hat{\theta})}{\hat{\theta}^2}.$$

Hence, the expected proportion $1-\exp(-C/\theta)$ of uncensored observations is also the relative reduction in Fisher information due to censoring.

2.3 Confidence procedures under type I censoring.

It is not evident how a confidence interval for θ should be specified. A large number of proposals can be found in the literature. One of them is exact, at least in principle, but of limited interest since it is inefficient. The idea is that since the number N is binomial we can form a confidence interval for the binomial parameter $P_\theta = 1 - \exp(-C/\theta)$ by inverting the binomial distribution function. Next we transform the interval for P_θ into an interval for θ by a monotone transformation. It was discussed by Bartholomew (1963), and Mann, Schafer & Singpurwalla (1974, Sec. 5.1.1) mention no other interval. It might be a reasonable procedure under heavy censoring, when the uncensored observations are approximately uniformly distributed on $(0, C)$, so their particular values carry little information about θ. However, when only little censoring is expected, that is when C/θ is large, this procedure has a low efficiency.

Asymptotic arguments can lead to several possible versions of a normal-based confidence interval, including a likelihood-based interval. An explicit proof that $\hat\theta$ is asymptotically consistent and normal as $n \to \infty$ in this situation was given by Yang & Sirvanci (1977), but in a more general setting such results can be found in Sundberg (1974). Other intervals utilise the χ^2 property from uncensored or type II censored situations. All of the following confidence intervals, C1–C7, have been proposed in the standard literature on life testing. They are here formulated in their two-sided $100(1-p)\%$ versions. As defined in Sec. 2.1, z_p is the standard normal p-quantile, and $\chi^2_{f,p}$ is the p-quantile of the χ^2_f distribution.

(C1) $\hat\theta\,(1 - z_{1-p/2}/\sqrt{N}) \le \theta \le \hat\theta\,(1 + z_{1-p/2}/\sqrt{N})$

This standard type of interval is obtained if we assume that

$$\sqrt{N}(\hat\theta - \theta)/\hat\theta \sim \mathrm{N}(0,1)\,.$$

An argument is that $J(\hat\theta)^{-1} = \hat\theta^2/N$. Note that the lower limit multiplier is negative for very small N. It should then be replaced by zero. In principle this interval is standard, but when mentioned in this particular application it may be accompanied by a warning of need for very large samples (Lawless, 1982). It is the only interval given in Lee (1992).

(C2) $\hat{\theta} \,/\, (1 + z_{1-p/2}/\sqrt{N}) \leq \theta \leq \hat{\theta} \,/\, (1 - z_{1-p/2}/\sqrt{N})$

This is obtained if we assume $\sqrt{N}(\hat{\theta} - \theta)/\theta \sim$ N(0,1), with θ in the denominator instead of $\hat{\theta}$. Note that this is not obtained by using $J(\theta)^{-1}$ instead of $J(\hat{\theta})^{-1}$ for $\text{Var}(\hat{\theta})$. Alternatively, we use the same argument for the intensity parameter $\rho = 1/\theta$ as the one used for θ in interval C1 (see for example Gross & Clark (1975)). We then transform this interval for ρ into an interval for θ. In analogy with interval C1, the right hand side denominator $1 - z_{1-p/2}/\sqrt{N}$ can be negative. The interpretation is then that the interval goes up to $+\infty$.

(C3) $\hat{\theta} \exp(-z_{1-p/2}/\sqrt{N}) \leq \theta \leq \hat{\theta} \exp(z_{1-p/2}/\sqrt{N})$

This is obtained from the normal approximation $\sqrt{N}(\log\hat{\theta} - \log\theta) \sim$ N(0,1). The procedure is mentioned in Kalbfleisch & Prentice (1980, Sec. 3.4.5) and Nelson (1982, formula (2.4), demanding $N > 15$, or so).

(C4) $\hat{\theta} \,/\, (1 + z_{1-p/2}/3\sqrt{N})^3 \leq \theta \leq \hat{\theta} \,/\, (1 - z_{1-p/2}/3\sqrt{N})^3$

This is recommended in the books by Kalbfleisch & Prentice (1980) and Lawless (1982, Sec. 3.2) as based on a refinement of the asymptotic normality property of $\hat{\theta}$. It goes back to Sprott (1973), where it is motivated for the exponential distribution both with complete and censored samples. The argument is that a standard normal distribution fits better to $3\sqrt{N}(\hat{\theta}^{-1/3} - \theta^{-1/3})/\theta^{-1/3}$ than to other functions of $\hat{\theta}$. The transformation is chosen so as to compensate for the skewness of the likelihood function.

(C5) $2N\{\hat{\theta}/\theta - 1 - \log(\hat{\theta}/\theta)\} \leq z_{1-p/2}^2$

This is a likelihood-based (or ML ratio test based) confidence interval. The left hand side is $-2\{\log L(\hat{\theta}) - \log L(\theta)\}$. It is not explicit in θ, which might be the reason it is not proposed in many standard texts, but it is included among the procedures listed in Lawless (1982), and also in Kalbfleisch & Prentice (1980). It is also advocated more generally, of course, without special reference to type I censored exponentials. An example from the reliability literature is Crowder *et al* (1991, Ch. 3), who propose a likelihood-based interval or one of its first order approximations, that is types C1 (expressed in terms of θ) or C2 (expressed in terms of ρ).

(C6) $2N\,\hat{\theta}/\chi^2_{2N,1-p/2} \leq \theta \leq 2N\,\hat{\theta}/\chi^2_{2N,p/2}$
This confidence interval would hold exactly under type II censoring, given N. It should be a natural approximation to try under type I censoring, with a conditional interpretation when conditioning on N. This procedure is proposed by Nelson (1982), and advocated also by Cox & Oakes (1984, Sec. 3.4) as a simpler but reasonable alternative to the likelihood-based interval C5.

(C7) $2N\,\hat{\theta}/\chi^2_{2N+1,1-p/2} \leq \theta \leq 2N\,\hat{\theta}/\chi^2_{2N+1,p/2}$
This modification of the previous interval was suggested by Cox (1953), and taken up by Lawless (1982). It is motivated by analogy with inverse sampling. In comparison with type II censoring we might argue that in type I, when we stop, $\sum_1^N y_i$ is "on the way"to the $(N+1)$st occurrence. A compromise would be to say "halfway", more precisely to treat $2N\hat{\theta}/\theta$ as χ^2_{2N+1}-distributed.

All of the interval types C1–C7 above are of the form

$$c_l\,\hat{\theta} \leq \theta \leq c_u\,\hat{\theta},$$

where the multipliers c_l and c_u are explicit in six of them. How these lower and upper limit multipliers depend on N is illustrated in Figures 1a & 1b respectively, for one-sided intervals aimed at $p = 2.5\%$ (two-sided $p = 5\%$).

The true coverage probability for any one of these intervals may be regarded as a weighted mean over the binomial N of the conditional coverage probabilities, given N. These can be theoretically calculated for each of methods C1–C7 above. For more details about the computations involved, see Section 3 below. Two-sided and one-sided (non-)coverage probabilities are shown in the diagrams of Figures 2–5 for a variety of cases. Figures 2 and 3 represent two-sided intervals, Figures 4 and 5 one-sided intervals. Figures 2, 4 and 5 represent the quite moderate intended confidence level of two-sided 90% and one-sided 95%, whereas Figure 3 has the contrasting high (for small samples extremely high) confidence level of (two-sided) 99%. The conclusions drawn from these diagrams are presented in Section 4.

The intervals above may appear to have a conditional character, since they express $\mathrm{Var}(\hat{\theta}/\theta)$ as approximately $1/N$, thus giving N the character of a precision index. However, this analogy must not be drawn too far. The variate N is not ancillary, since the distribution of N depends on θ, so the conditionality principle does not apply. In fact, under heavy censoring N contains most of the information about θ. More specifically, $\hat{\theta}/\theta$ has a conditional bias that depends linearly on n/N, and from this dependence it also follows that $J(\hat{\theta})^{-1}/\hat{\theta}^2 = 1/N$ is not the conditional variance of $\hat{\theta}/\theta$.

Variants of the intervals C1–C7, which appear less conditional, are obtained if we replace N by an estimate $n\hat{P}_\theta$ of its expected value, where

$$\hat{P}_\theta = 1 - \exp(-C/\hat{\theta}). \tag{4}$$

Some authors regard N/n as a "quick estimate"of P_θ and an inferior substitute for the $\hat{P}_\theta$ in (4) (Bartholomew, 1957, Gross & Clark, 1975, Lee, 1992). This is not the standpoint taken here. A strict use of $\hat{P}_\theta$ requires that the future censoring limit C must be known, even if no observations are censored ($N = n$). To replace n by an estimated expected sample size $E(N) < n$ in such cases appears both counter-intuitive and impractical. This argument has previously been expressed by for example Kalbfleisch & Prentice (1980).

2.4 An illustration with data.

The following data are adapted from Kalbfleisch (1985, Ex. 9.5.1). Ten components were placed on test simultaneously. Suppose first that it was decided to terminate the experiment after 50 days. Eight components failed before that time, yielding life times 4, 5, 8, 11, 20, 29, 35 and 40 days. Hence $N = 8$ and $\hat{\theta} = 252/8 = 31.5$. Alternatively, suppose the experiment should be terminated after 75 days. Then no components had been censored, because the remaining two failed at 66 and 70 days, and we would have had $N = 10$ and $\hat{\theta} = 288/10 = 28.8$. The corresponding log–likelihood functions are shown in Kalbfleisch (1985, Fig. 9.5.2). The two-sided 95% confidence intervals, according to Figures 1 a and b, are shown in Table 1:

Table 1. 95% confidence intervals for different censoring limits and different methods, with data from Kalbfleisch (1985).

Interval	Cens. at 50	Cens. at 75
C1	9.67, 53.3	10.95, 46.7
C2	18.61, 102.6	17.78, 75.7
C3	15.75, 63.0	15.50, 53.5
C4	16.89, 69.3	16.39, 57.7
C5	16.92, 68.9	16.42, 57.5
C6	17.47, 73.0	16.86, 60.1
C7	16.69, 66.6	16.24, 56.0

We see that the intervals vary much between methods, in both position and width. At the same time, some of them are quite similar, in particular C4 and C5. However, this is not evidence that they are more reliable. In fact, our only exact knowledge from theory is that interval C6 would be correct if there were no censoring limit, and this case certainly seems close to the second situation in which the experiment was to be terminated at 75 days.

3 Calculation methods

Given θ, a type I censored exponential sample of size n may be imagined as generated in the following way. A binomial trial decides the random number N of uncensored observations, that is N is generated as

$$\mathrm{Bin}(n; 1 - \exp(-C/\theta)).$$

The $n - N$ censored observations are given the value $y = C$. Given N, the uncensored Y-values are generated as a sample $Z_1, \ldots, Z_N$ from a truncated exponential distribution, truncated at C. Hence the desired coverage probabilities can be found by calculations of type

$$\mathrm{P}_\theta(\hat{\theta} < t) = \mathrm{P}_\theta(\textstyle\sum_1^n Y_i < N\,t) = E_N\,\mathrm{P}_\theta\{\sum_1^N Z_i < Nt - (n-N)C \mid N\}.$$

The distribution of $\sum_1^N Z_i$ has been calculated by Fourier inversion. The characteristic function of a truncated exponential distribution is

$$\varphi(t) = \frac{1}{1 - i\theta t}\,\frac{1 - \exp(-C(1 - i\theta t)/\theta)}{1 - \exp(-C/\theta)},$$

so the distribution function $F(z|N)$ of $\sum_1^N Z_i$, given N, can be computed as

$$F(z|N) = \lim_{T\to\infty} \frac{1}{\pi} \int_{-T}^{T} \frac{\sin zt/2}{t}\, e^{-itz/2}\, \varphi(t)^N\, dt$$

(see for example Cramér (1946, ch. 10)).

In practice T can be chosen large enough to make the integral insensitive to increases in T. Numerically this is not a problem as long as the probability $F(z|N)$ is not very small. Note that the integrand is an oscillating function and that the integral therefore can be numerically relatively unstable. This occurs when the actual infinite integral is very small, but then the contribution from $F(z|N)$ to the coverage probability is very small as well, so this has not caused any great problem. The computations have been carried out in the MATLAB environment.

4 Results

In this section we study the coverage probabilities for the seven confidence procedures C1–C7, as shown in Figures 2-5. These figures show the complementary noncoverage probabilities as functions of sample size. The following levels of confidence are represented:

Figures 2, 4, 5 Moderate confidence: two-sided 90% and one-sided 95%.

Figure 3 High confidence: two-sided 99%.

One-sided high confidence intervals are not represented here, but these diagrams can be obtained from the author on request.

Each of Figures 2-5 has four diagrams, corresponding to the following four censoring probabilities:

a) Heavy censoring, $1 - P_\theta = \mathrm{P}(Y > C) = 0.61$

b) Moderate censoring, $\mathrm{P}(Y > C) = 0.37$

c) Light censoring, $\mathrm{P}(Y > C) = 0.14$

d) No censoring, $C = \infty$

Many curves are seen to be smooth functions of sample size n. Other curves show a rugged form, more pronounced for smaller n. The reason behind the ruggedness is found in the discrete nature of the stochastic N, in combination with a substantial degree of censoring. Of course this ruggedness in not seen in Figure 1.

Figures 2 and 3 show noncoverage probabilities for two-sided intervals. We see that the interval C1 is generally much worse than the others. Except under heavy or moderate censoring and 90% intended confidence (Fig. 2a-b), it has a noncoverage substantially higher than the intended. For $n = 50$ its noncoverage is $> 2\%$ when it should be 1% (Fig. 3). Not only C1 but also C2 has a substantially higher noncoverage than the intended 1% in this case, under heavy or moderate censoring (Fig. 3a-b). On the other hand, with no censoring and 90% intended confidence (Fig. 2d), C2 reacts in the opposite way and has a clear overcoverage. In between, with some little censoring, C2 can appear almost perfect in its coverage (Fig. 2c). However, we cannot control the censoring probability in advance, since we do not know θ.

Also interval C3 shows a tendency to differ more from the intervals C4–C7 than these do between themselves, in most cases showing a higher noncoverage but in Fig. 2b being more perfect than the other intervals. Finally note that when there is no censoring (Figures 2d and 3d), interval C6 fits perfectly, in agreement with theory.

In order to understand the somewhat confusing behaviour of the intervals C1 and C2, in particular, as compared over the eight different cases of Figures 2 and 3, we must go to the corresponding one-sided intervals. Also, more generally, a perfect two-sided coverage is of doubtful value if the corresponding one-sided coverages are very different. Table 1 above yields drastic illustrations of how much the upper and lower ends can vary between methods. Figures 4-5 show one-sided (non-)coverages when the intended value is 5%, that is corresponding to the two-sided intervals in Fig. 2. Note first that the methods spread much more in their one-sided coverages; the scale unit in Figures 4 and 5 is of a higher magnitude than for the two-sided intervals in Fig. 2. Note also the consistency shown over the different degrees of censoring. In particular there is not much difference at all between the curves for weak censoring and no censoring, except at the smallest sample sizes.

Figures 4-5 also include a curve indicating the probability for complete censoring, $\mathrm{P}(N = 0) = (1 - P_\theta)^n$. When $N = 0$, the MLE is $\hat{\theta} = \infty$, and the intervals C1–C7 are taken to include the whole positive real axis. The experimenter will probably not be satisfied with the limited information provided by such a trial, but extend the experiment in some way. Therefore, the coverage probabilities must be taken with a pinch of salt when $\mathrm{P}(N = 0)$ is substantial. However, as soon as $n > 15$ or so, this probability is evidently negligible even under the heavy censoring probability 0.61. Under the moderate censoring probability 0.37, $\mathrm{P}(N = 0)$ is negligible for $n \geq 10$.

Fig. 4 shows that if the lower-bounded interval is constructed by method C1, the noncoverage is $\leq 2\%$ when intended to be 5%. On the other hand Fig. 5 shows a high complementary noncoverage for the upper-bounded interval, $\geq 8\%$ for $n \leq 50$. These two one-sided noncoverages fit reasonably well together to form a 90% two-sided interval. With higher intended confidence, however, the one-sided noncoverages fit less well together. For the two-sided intervals of 99% intended confidence, the noncoverage with method C1 shown in Fig. 3 is not only much too high but also practically one-sided, even for n as large as 50. There are strong reasons to warn against the use of interval C1.

Method C2 behaves in analogy with method C1, but somewhat less extreme, and with reverse roles of the upper- and lower-bounded intervals (as could be expected). Method C2 should also be avoided. The picture for method C3 is similar in character to that for C1, but not as pronounced. The overall impression is that method C3 is inferior to several of the other methods, so there is no reason to use method C3, either.

Finally we compare the remaining four methods, C4–C7. There is no method that dominates the others. Method C6 is best and perfect when there is no censoring ($C = \infty$), whereas the modification C7 shows a small advantage over C6 in the cases of heavy or moderate censoring. Method C7 has a tendency to overcompensate from C6, however. Methods C4 and C5 are often mutually indistinguishable and in all cases as similar as one could expect from their mutual similarity shown in the curves of Figure 1. Method C7 in most cases of lower-bounded intervals (see Fig. 4) differs quite little from C4 and C5, whereas there is a small but clear difference for the upper-bounded intervals (Fig. 5), to the disadvantage of C7. The (lack of) differences are also in accordance with the curves in Figures 1a and 1b, respectively. Overall, methods C4 and C5 fit the intended confidence levels quite well (both for two-sided and one-sided—even for small sample sizes). Methods C6 and C7 do almost as well. The largest deviations from the intended level are due to the fluctuations caused by the discrete character of N and seen for small sample sizes and substantial censoring.

5 Conclusions

Bartholomew (1963) stated that the asymptotic normal distribution for $\hat{\theta}$ "cannot be used unless n is very large". The results above demonstrate that this statement holds true for the simplest constructions, C1 and C2. From a wider point of view, however, Bartholomew was too pessimistic. For example, after a parameter transformation from θ to $\theta^{-1/3}$ (method C4), the confidence levels hold remarkably well for quite small samples, not only for double-sided but also for one-sided intervals.

What interval construction should then be selected? From the overall behaviour, as judged from the study above, C4 or C5 could be advocated. Interval C4 is then the more practical one since it has the important advantage of being explicit, which the likelihood-based C5 is not. Another question is why they are so similar. Much of the explanation lies in that the transformation of θ is selected to make the likelihood function as parabolic as possible around the ML estimator.

However, if heavy censoring is not expected beforehand, the χ^2_{2N}-based method C6 could also be strongly advocated. This method is perfect in the absence of censoring and under type II censoring. It is as good as any other method under light type I censoring, and little worse than the best one even under heavy censoring, albeit not conservative. The advantages of unification makes a strong argument for C6 to become established practice also under type I censoring. It would be practical not having to distinguish type I and type II censoring, or even to decide beforehand if observations will be censored or not, and to allow some unit-to-unit variation in the censoring time without bothering. We would of course also use the same procedure in such less well-specified (hybrid) cases as when we do not realize until the experiment is long begun, that some units live longer than expected and that we must therefore impose some censoring.

References

D.J. Bartholomew, "A problem in life testing," *J. Amer. Statist. Assoc.* vol. 52, pp. 350-355, 1957.

D.J. Bartholomew, "The sampling distribution of an estimate arising in life testing," *Technometrics* vol. 5, pp. 361-374, 1963.

D.R. Cox, "Some simple approximate tests for Poisson variates," *Biometrika*, vol. 40, pp. 354-360, 1953.

D.R. Cox and D. Oakes, *Analysis of Survival Data,* Chapman and Hall: London, 1984.

H. Cramér, *Mathematical Methods of Statistics,* Princeton Univ. Press: Princeton, 1946.

M.J. Crowder, A.C. Kimber, R.L. Smith and T.J. Sweeting, *Statistical Analysis of Reliability Data,* Chapman and Hall: London, 1991.

A.J. Gross and V.A. Clark, *Survival Distributions: Reliability Applications in the Biomedical Sciences,* Wiley: New York, 1975.

J.D. Kalbfleisch and R.L. Prentice, *The Statistical Analysis of Failure Time Data,* Wiley: New York, 1980.

J.G. Kalbfleisch, *Probability and Statistical Inference, Vol. 2, 2nd ed.,* Springer-Verlag: New York, 1985.

J.F. Lawless, *Statistical Models and Methods for Lifetime Data,* Wiley: New York, 1982.

E.T. Lee, *Statistical Methods for Survival Data Analysis, 2nd ed.,* Wiley: New York, 1992.

N.R. Mann, R.E. Schafer and N.D. Singpurwalla, *Methods for Statistical Analysis of Reliability and Life Data,* Wiley: New York, 1974.

W. Nelson, *Applied Life Data Analysis,* Wiley: New York, 1982.

D.A. Sprott, "Normal likelihoods and their relation to large sample theory of estimation," *Biometrika*, vol. 60, pp. 457-465, 1973.

R. Sundberg, "Maximum likelihood theory for incomplete data from an exponential family," *Scand. J. Statist.*, vol. 1, pp. 49-58, 1974.

J. Tyrcha, R. Sundberg, P. Lindskog and B. Sundström, "Statistical modelling and saddle-point approximation of tail probabilities for accumulated splice loss in fibre-optic networks," *J. Applied Statistics*, vol. 27, pp. 245-256, 2000.

G. Yang and M. Sirvanci, "Estimation of a time-truncated exponential parameter used in life testing," *J. Amer. Statist. Assoc.*, vol. 72, pp. 444-447, 1977.

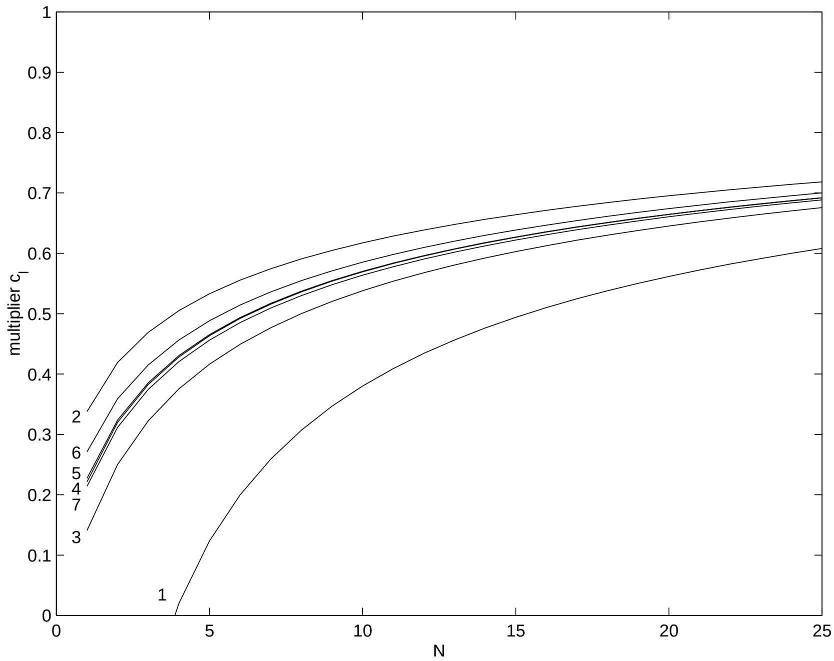

Figure 1a. Multiplier c_l to $\hat{\theta}$ in lower-bounded 97.5% confidence intervals for θ, as a function of the number N of uncensored observations, N=1,...,25. Numbers indicate interval construction methods: C1–C7.

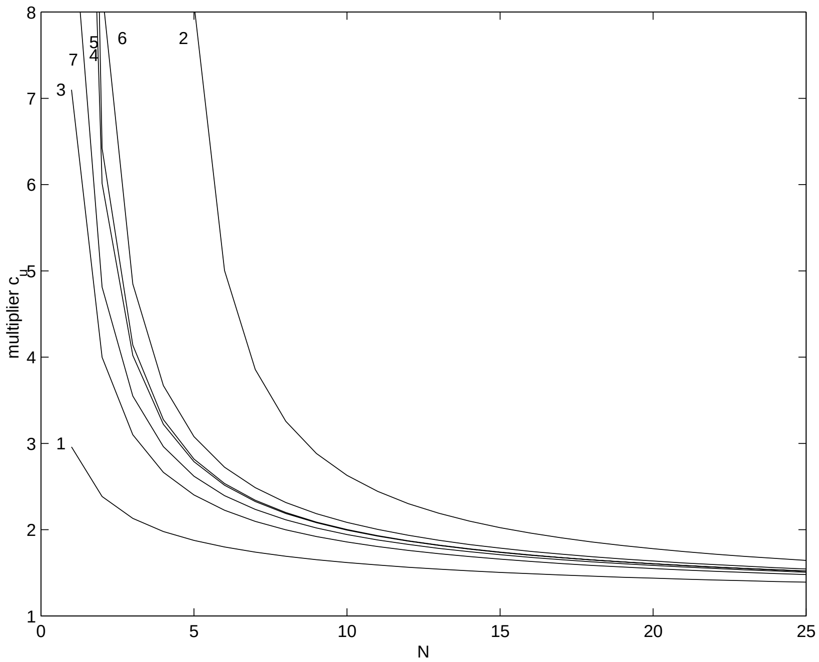

Figure 1b. Multiplier c_u to $\hat{\theta}$ in upper-bounded 97.5% confidence intervals for θ, as a function of the number N of uncensored observations, N=1,...,25. Numbers indicate interval construction methods: C1–C7.

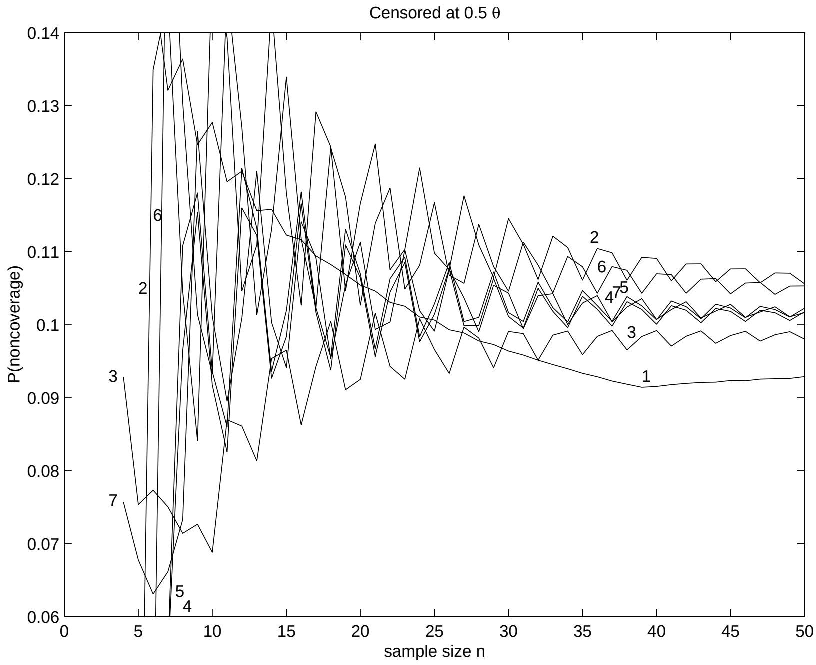

Figure 2a. Noncoverage probabilities for two-sided 90% intervals, with censoring at $0.5\,\theta$, for n=4,...,50. Numbers indicate interval construction methods: C1–C7.

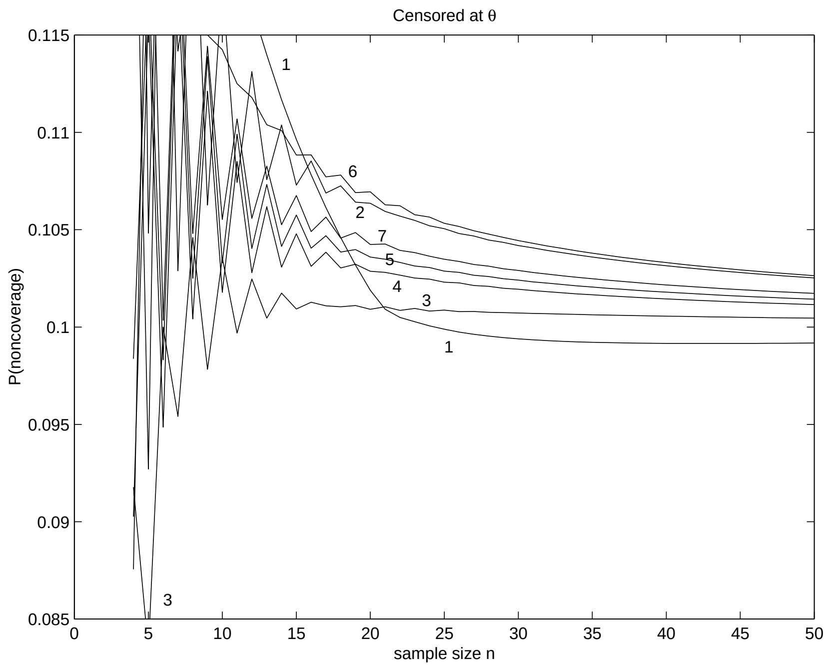

Figure 2b. Noncoverage probabilities for two-sided 90% intervals, with censoring at θ, for n=4,...,50. Numbers indicate the interval construction methods: C1–C7.

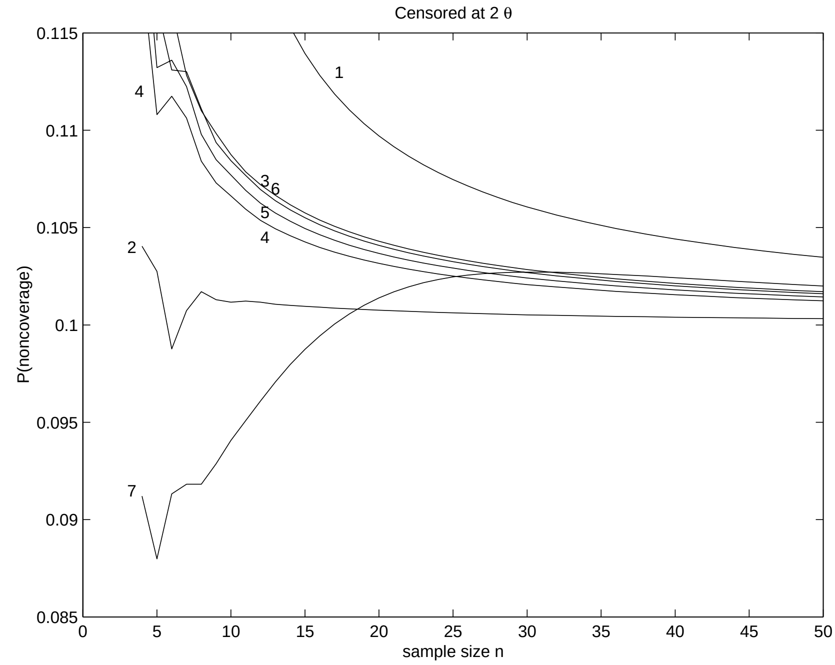

Figure 2c. Noncoverage probabilities for two-sided 90% intervals, with censoring at $2\,\theta$, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

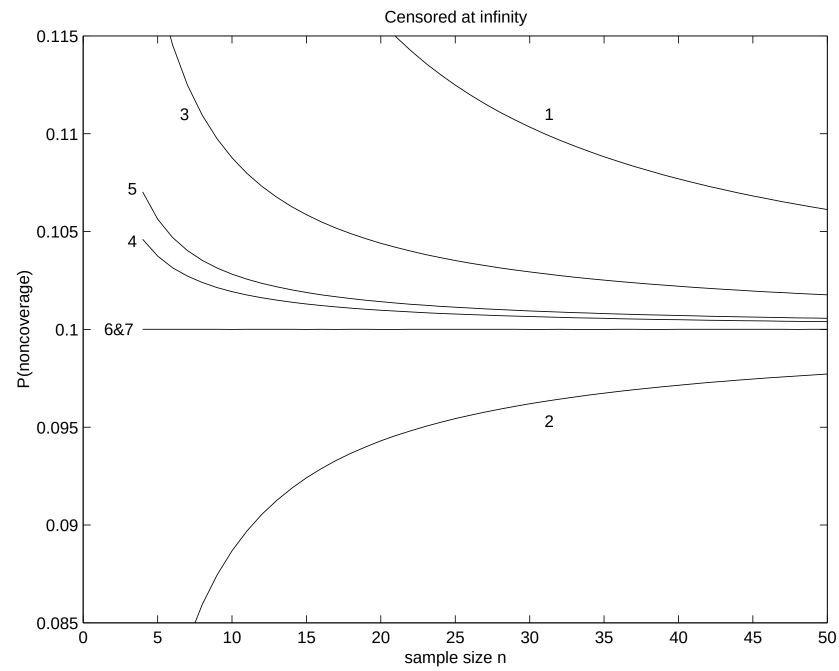

Figure 2d. Noncoverage probabilities for two-sided 90% intervals, no censoring, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

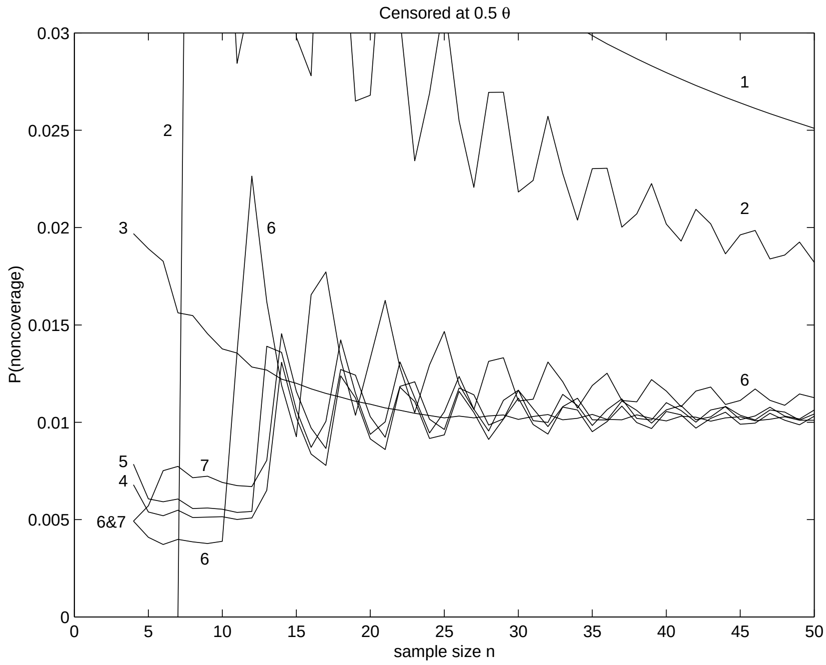

Figure 3a. Noncoverage probabilities for two-sided 99% intervals, with censoring at $0.5\,\theta$, for n=4,...,50. Numbers indicate the interval construction methods: C1–C7.

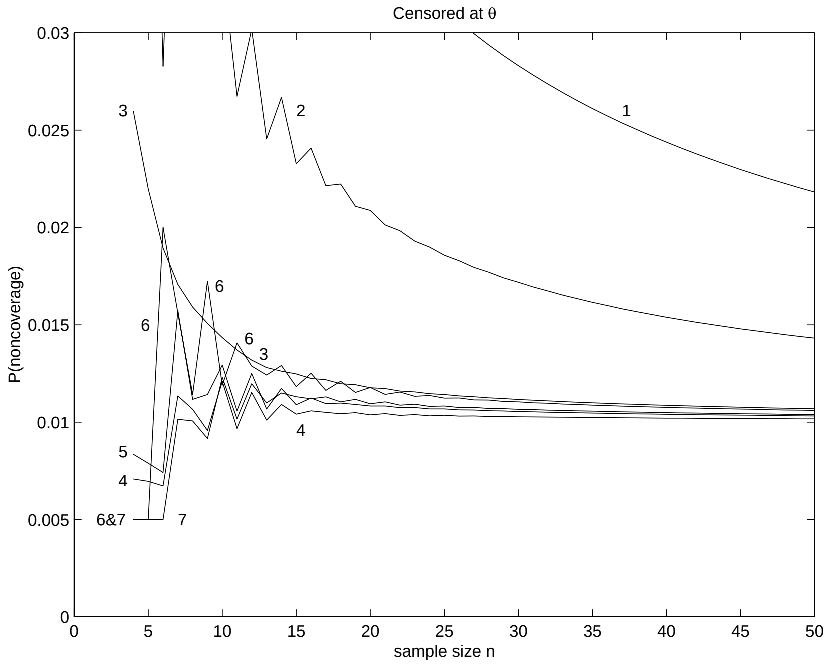

Figure 3b. Noncoverage probabilities for two-sided 99% intervals, with censoring at θ, for n=4,...,50. Numbers indicate the interval construction methods: C1–C7.

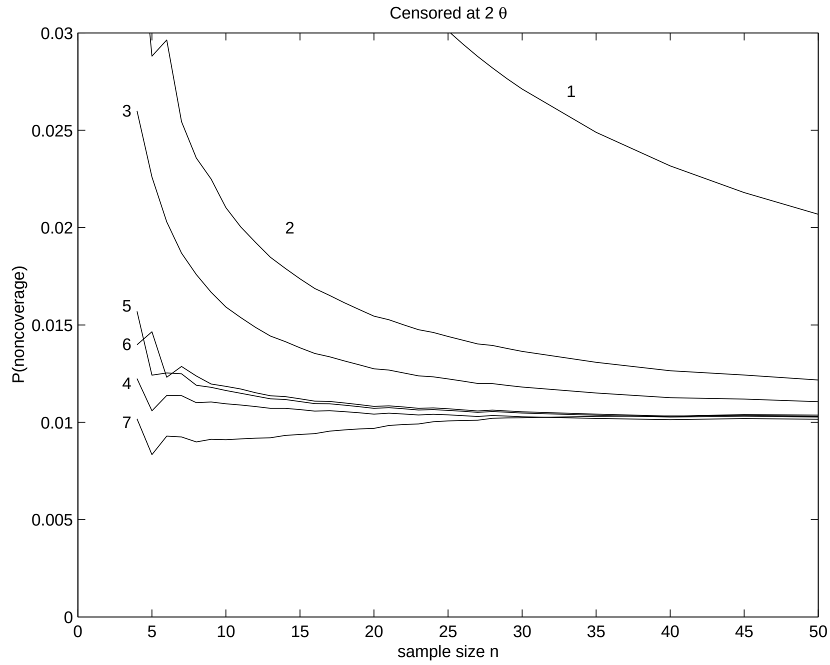

Figure 3c. Noncoverage probabilities for two-sided 99% intervals, with censoring at $2\,\theta$, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

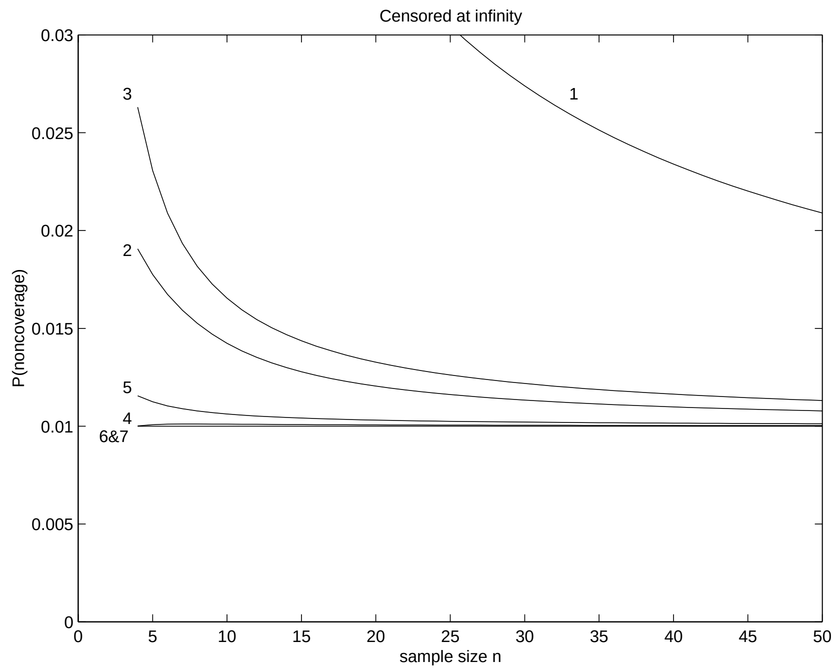

Figure 3d. Noncoverage probabilities for two-sided 99% intervals, no censoring, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

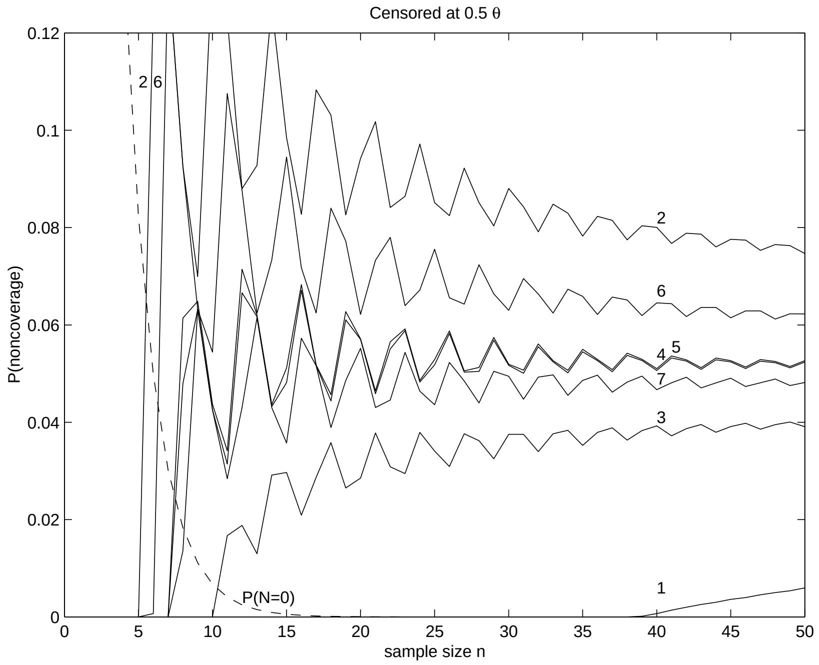

Figure 4a. Noncoverage probabilities for lower-bounded 95% intervals, with censoring at $0.5\,\theta$, and $\mathrm{P}(N = 0)$, for n=4,. . . ,50. Numbers indicate the interval construction methods: C1–C7.

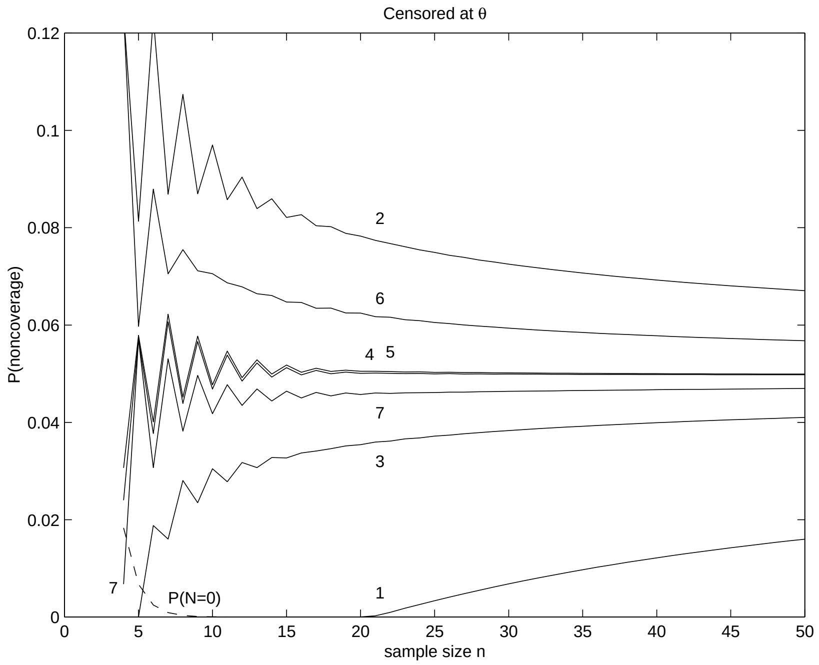

Figure 4b. Noncoverage probabilities for lower-bounded 95% intervals, with censoring at θ, and $\mathrm{P}(N = 0)$, for n=4,. . . ,50. Numbers indicate the interval construction methods: C1–C7.

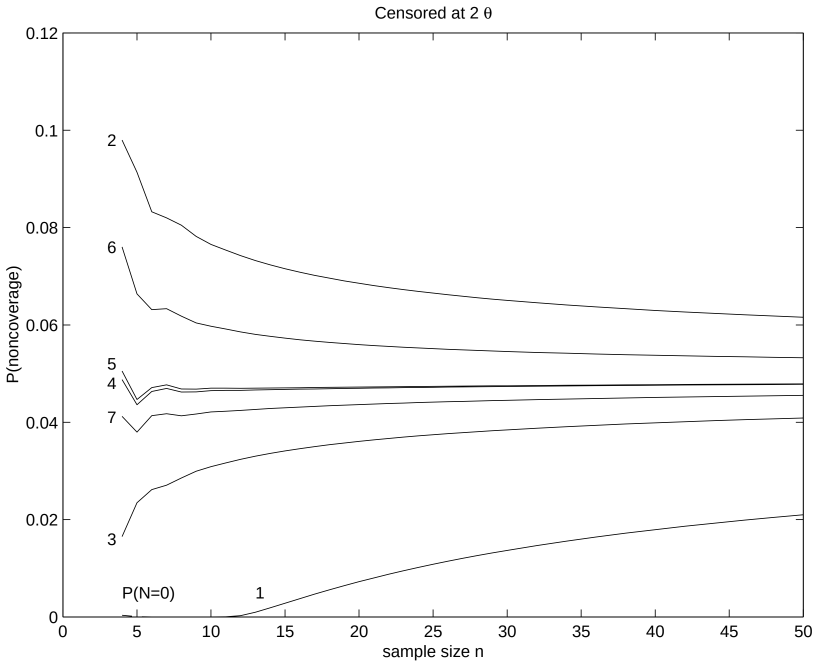

Figure 4c. Noncoverage probabilities for lower-bounded 95% intervals, with censoring at $2\,\theta$, and $\mathrm{P}(N = 0)$, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

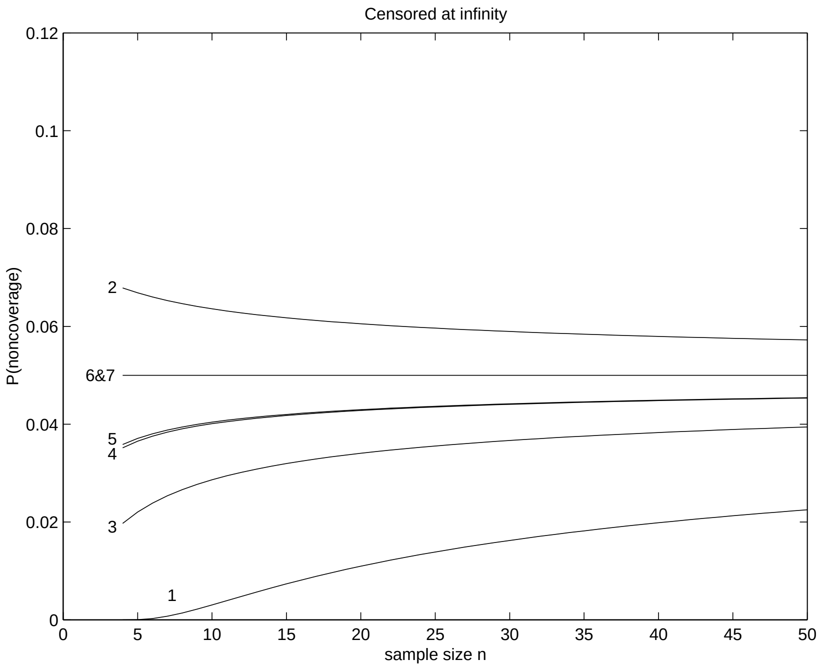

Figure 4d. Noncoverage probabilities for lower-bounded 95% intervals, no censoring, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

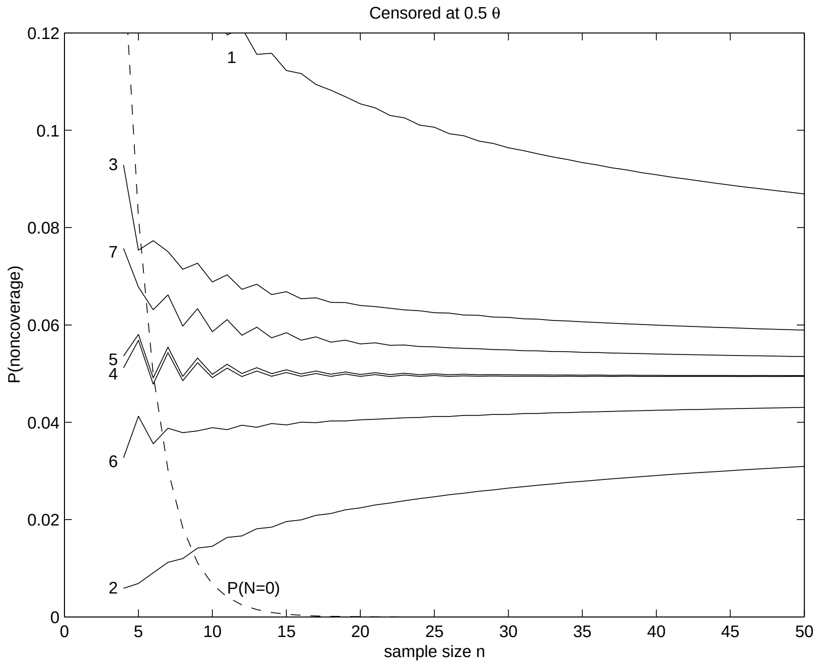

Figure 5a. Noncoverage probabilities for upper-bounded 95% intervals, with censoring at $0.5\,\theta$, and $\mathrm{P}(N = 0)$, for n=4,. . . ,50. Numbers indicate the interval construction methods: C1–C7.

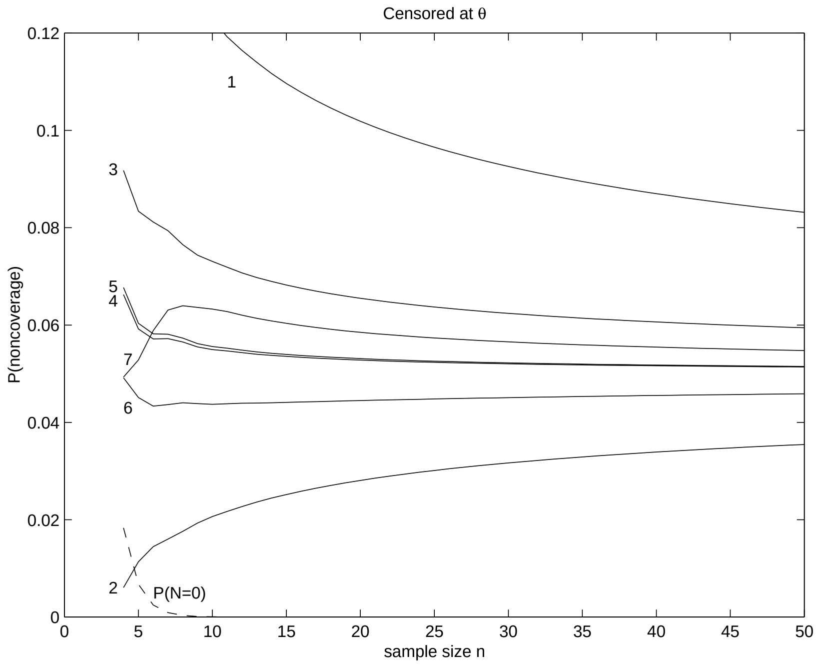

Figure 5b. Noncoverage probabilities for upper-bounded 95% intervals, with censoring at θ, and $\mathrm{P}(N = 0)$, for n=4,. . . ,50. Numbers indicate the interval construction methods: C1–C7.

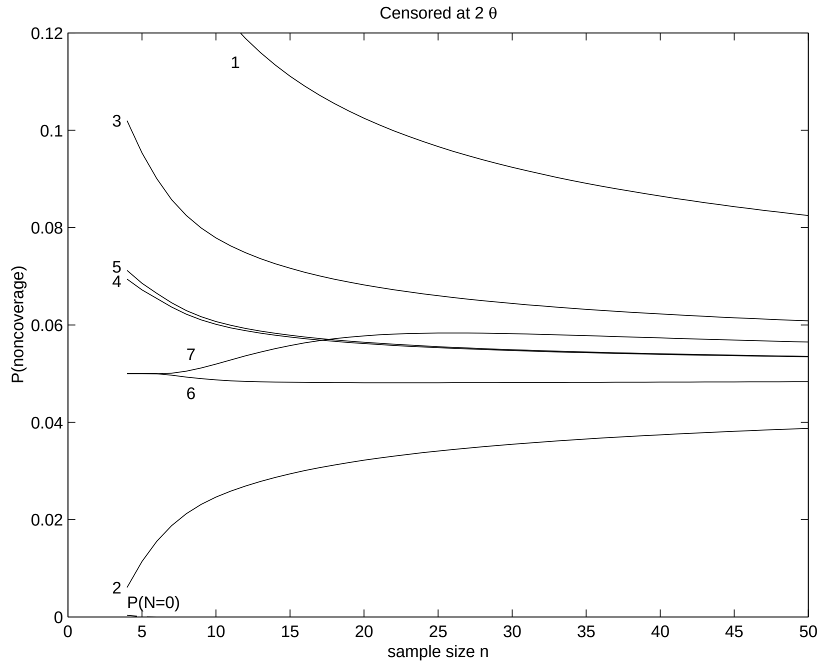

Figure 5c. Noncoverage probabilities for upper-bounded 95% intervals, with censoring at 2θ, and $\mathrm{P}(N = 0)$, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

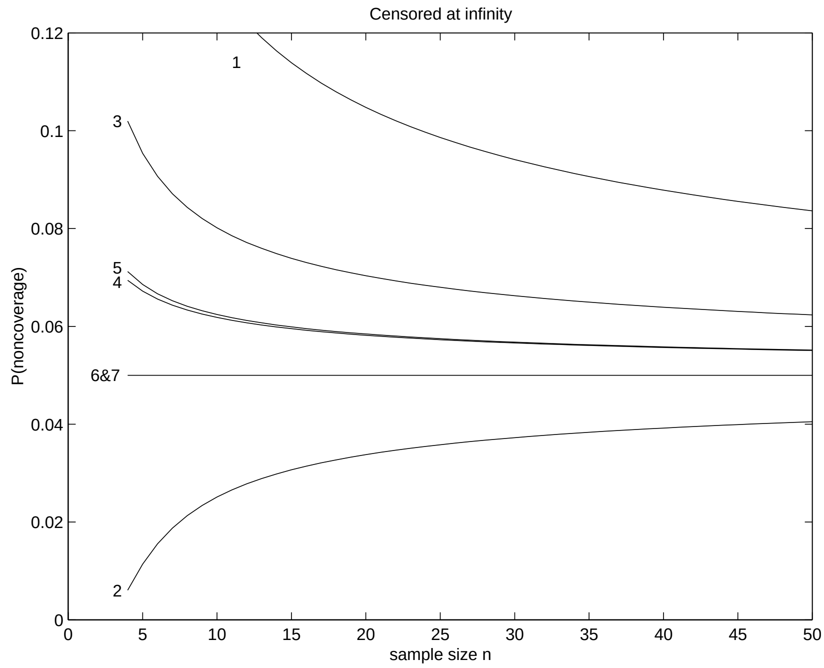

Figure 5d. Noncoverage probabilities for upper-bounded 95% intervals, no censoring, for n=4,…,50. Numbers indicate the interval construction methods: C1–C7.

Author address:

Rolf Sundberg,

Matematical statistics, Stockholm University,

SE–106 91 Stockholm, Sweden

E-mail: rolfs@matematik.su.se

Keywords: Confidence interval, coverage probability, exponential distribution, failure times, fixed censoring.