An algorithm to construct Monte Carlo confidence intervals for an arbitrary function of probability distribution parameters

Hristos Tyralis*, Demetris Koutsoyiannis and Stefanos Kozanis

Department of Water Resources, Faculty of Civil Engineering, National Technical University, Athens

Heroon Polytechneiou 5, GR-157 80 Zographou, Greece

* Corresponding author, montchrister@gmail.com

Abstract We derive a new algorithm for calculating an exact confidence interval for a parameter of location or scale family, based on a two-sided hypothesis test on the parameter of interest, using some pivotal quantities. We use this algorithm to calculate approximate confidence intervals for the parameter or a function of the parameter of one-parameter continuous distributions. After appropriate heuristic modifications of the algorithm we use it to obtain approximate confidence intervals for a parameter or a function of parameters for multi-parameter continuous distributions. The advantage of the algorithm is that it is general and gives a fast approximation of an exact confidence interval. Some asymptotic (analytical) results are shown which validate the use of the method under certain regularity conditions. In addition, numerical results of the method compare well with those obtained by other known methods of the literature on the exponential, the normal, the gamma and the Weibull distribution.

Key words Confidence interval; Monte Carlo; maximum likelihood; location-scale family; hydrological statistics.

1 Introduction

Various general methods for the calculation of a confidence interval for a parameter of interest exist. Casella & Berger (2002, p.496-497) suggest the use of the asymptotic distribution of the maximum likelihood estimator (MLE) to construct a confidence interval for a function of the parameter of a one-parameter distribution. Wilks (1938) constructs a confidence interval based on the score statistic (see also Casella & Berger, 2002, p. 498). Kite (1988) gives approximate confidence intervals for the parameters of various distributions, by performing separate analyses for each distribution and each parameter estimation method. Garthwaite & Buckland (1992) make a new use of the Robbins-Monro search process to generate Monte Carlo confidence intervals for a one-parameter probability distribution. The Jacknife method is another general technique to obtain confidence intervals (see e.g. Román-Montoya, 2008). Ripley (1987, p.176-178) constructs simple Monte Carlo confidence intervals which depend on the type of local properties (location or scale) of the parameter of interest.

In this paper we generalize the method proposed by Ripley (1987) retaining its simplicity. The method we study here incorporates Ripley's two suggested equations into one new equation. The algorithm of the method was derived by Koutsoyiannis & Kozanis (2005), has already been used (as an intuitive tool without mathematical proofs) by Koutsoyiannis et al. (2007) and is a main tool of the statistical software Hydrognomon (2009-2011). The method has a general character and does not make a distinction for location or scale family, while other methods make such distinction. It provides single results without requiring user choices. These are strong advantages which make the proposed method a useful statistical computation tool.

Initially, we show that our algorithm is asymptotically equivalent to a Wald-type interval (Casella & Berger, 2002, p.499) of a parameter or a function of a parameter of any one-parameter probability distribution. We also show how this algorithm works for certain distributions. Then we generalize this new algorithm to construct confidence intervals for the parameters or functions of parameters for multi-parameter probability distributions. We show that these intervals are asymptotically equivalent to Wald-type intervals. We also show analytically how this algorithm works for the normal distribution. We compare the results of the algorithm with those obtained by other exact and approximate methods for the exponential, normal, gamma and Weibull distributions, and it turns out that the algorithm works well even for small samples. The approximate methods described here include Wald-type intervals given in the literature or derived using the formula in Casella & Berger, 2002, p. 497, Ripley's two equations, and bias-corrected and accelerated (BCa) bootstrap non-parametric intervals (see also DiCiccio 1984; Di Ciccio & Efron 1996; Di Ciccio & Romano 1995; Efron 1987; Efron & Tibshirani 1993; Hall 1988; Kisielinska 2012).

The proposed algorithm is partly heuristic and simultaneously so general that needs no assumptions about the statistical behaviour of the statistics under study, i.e. it can perform for any continuous distribution with any number of parameters, and for any distributional or derivative parameter. Only the theoretical probabilistic model is needed and all other calculations are done by a number of Monte Carlo simulations without additional assumptions.

2 Terminology and notation

We use the terminology of Casella & Berger (2002) as well as the Dutch convention for notation, according to which random variables are underlined (Hemelrijk, 1966). We recall that an interval estimate of a parameter $\theta \in R$ is any pair of functions, $l(\boldsymbol{x})$ and $u(\boldsymbol{x})$, of a sample $\boldsymbol{x} = (x_1,\ldots,x_n)$ that satisfy $l(\boldsymbol{x}) \leq u(\boldsymbol{x})$ for all $\boldsymbol{x}$. If $\underline{\boldsymbol{x}}$ is the random variable whose realization is $\boldsymbol{x}$, the inference $l(\boldsymbol{x}) \leq \theta \leq u(\boldsymbol{x})$ is made. The random interval $[l(\underline{\boldsymbol{x}}),u(\underline{\boldsymbol{x}})]$ is called an interval estimator.

The following result from Casella & Berger (2002, p. 421,422) is necessary for the proofs of the next section and shows how we can construct a confidence interval from a hypothesis testing procedure:

For each $\theta_0 \in \Theta \subseteq R$, let $A(\theta_0)$ be the acceptance region of a level α test of H_0: $\theta \in \Theta_0$. For each $\boldsymbol{x}$, we define an interval $C(\boldsymbol{x})$ in the parameter space by

$$C(\boldsymbol{x}) = \{\theta_0 : \boldsymbol{x} \in A(\theta_0)\} \tag{1}$$

Then the random set $C(\underline{\boldsymbol{x}})$ is a $1 - \alpha$ confidence interval. Conversely, let $C(\underline{\boldsymbol{x}})$ be a $1 - \alpha$ confidence interval. For any $\theta_0 \in \Theta$, we define

$$A(\theta_0) = \{\boldsymbol{x} : \theta_0 \in C(\boldsymbol{x})\} \tag{2}$$

Then $A(\theta_0)$ is the acceptance region of a level α test of H_0: $\theta = \theta_0$. Note that the above terminology is not precise when the test is randomized (Shao, 2003, p.477).

3 Construction of confidence intervals for one-parameter distributions

Now we proceed to the construction of a confidence interval for one-parameter continuous probability distributions. The following result which is a consequence of (1) and (2) is necessary for the construction of the confidence interval.

We suppose that $\underline{b} := b(\underline{\boldsymbol{x}})$ is a MLE of the parameter θ of a one-parameter distribution with density $f(x|\theta)$. Then $\hat{\theta} = b(\boldsymbol{x})$ is the estimate of the parameter. We suppose now that the probability density of the statistic $b(\underline{\boldsymbol{x}})$ is $g(b|\theta)$. Then we seek two functions $\lambda(\theta)$, $\upsilon(\theta)$ such that $P\{\lambda(\theta) \leq \underline{b}(\underline{\boldsymbol{x}}) \leq \upsilon(\theta)\} = 1 - \alpha$. We define $\lambda(\theta)$, $\upsilon(\theta)$ as those functions that satisfy:

$$P\{b(\underline{\boldsymbol{x}}) < \lambda(\theta)\} = P\{b(\underline{\boldsymbol{x}}) > \upsilon(\theta)\} = \alpha/2 \tag{3}$$

The above equation implies that:

$$\lambda(\theta) = G^{-1}(\alpha/2|\theta) \text{ and } \upsilon(\theta) = G^{-1}(1 - \alpha/2|\theta) \tag{4}$$

where $G^{-1}(\cdot|\theta)$ denotes the inverse of the distribution function G.

Now we construct a test H_0: $\theta = \hat{\theta}$ vs H_1: $\theta \neq \hat{\theta}$ with acceptance region:

$$A(\hat{\theta}) = \{\boldsymbol{x}: G^{-1}(\alpha/2|\hat{\theta}) \leq b(\boldsymbol{x}) \leq G^{-1}(1 - \alpha/2|\hat{\theta})\} \tag{5}$$

which is a size α test because $\beta(\hat{\theta}) = 1 - P(G^{-1}(\alpha/2|\hat{\theta}) \leq b(\underline{\boldsymbol{x}}) \leq G^{-1}(1 - \alpha/2|\hat{\theta})|\theta = \hat{\theta}) = 1 - [\mathrm{G}(G^{-1}(1 - \alpha/2|\hat{\theta})|\hat{\theta}) - \mathrm{G}(G^{-1}(\alpha/2|\hat{\theta})|\hat{\theta})] = 1 - (1 - \alpha/2 - \alpha/2) = \alpha$. From this test and according to (1) and (2) we obtain the following $1 - \alpha$ confidence interval for θ:

$$C(\boldsymbol{x}) = \{\hat{\theta}: G^{-1}(\alpha/2|\hat{\theta}) \leq b(\boldsymbol{x}) \leq G^{-1}(1 - \alpha/2|\hat{\theta})\} \tag{6}$$

In our case we assume that we have an observation $\boldsymbol{y} = (y_1,\ldots,y_n)$. We obtain the following $1 - \alpha$ confidence interval for θ:

$$C(\boldsymbol{y}) = \{\theta: G^{-1}(\alpha/2|\theta) \leq b(\boldsymbol{y}) \leq G^{-1}(1 - \alpha/2|\theta)\} \tag{7}$$

Now we define l and u as the solutions of the equations:

$$\upsilon(l) = b(\boldsymbol{y}) \text{ and } \lambda(u) = b(\boldsymbol{y}) \tag{8}$$

From the above equation we obtain that:

$$G^{-1}(\alpha/2|u) = b(\boldsymbol{y}) \text{ and } G^{-1}(1 - \alpha/2|l) = b(\boldsymbol{y}) \tag{9}$$

We assume that $C(\boldsymbol{y}) = [\theta_1,\theta_2]$ where θ_1,θ_2 are the solutions of the equations

$$G^{-1}(\alpha/2|\theta_2) = b(\boldsymbol{y}) \text{ and } G^{-1}(1 - \alpha/2|\theta_1) = b(\boldsymbol{y}) \tag{10}$$

Now it is obvious that $[l,u]$ is a $1 - \alpha$ confidence interval estimate for θ.

3.1 Construction of the confidence interval

Having proved that $[l,u]$ is a $1 - \alpha$ confidence interval estimate for θ, we can use it to construct an approximate confidence interval that can be easily computed numerically. From Figure 1 we observe that

$$\frac{\upsilon(\hat{\theta}) - \hat{\theta}}{\hat{\theta} - l} = \frac{CA}{CB} \approx \left(\frac{d\upsilon}{d\theta}\right)_{|\theta = \hat{\theta}} \tag{11}$$

Solving for l we find

$$l \approx \hat{\theta} + \frac{\hat{\theta} - \upsilon(\hat{\theta})}{(d\upsilon/d\theta)_{|\theta = \hat{\theta}}} \tag{12}$$

and in a similar way we find that

$$u \approx \hat{\theta} + \frac{\hat{\theta} - \lambda(\hat{\theta})}{(d\lambda/d\theta)_{|\theta = \hat{\theta}}} \tag{13}$$

We can thus claim that

$$[l(\underline{\mathbf{x}}), u(\underline{\mathbf{x}})] = [b(\underline{\mathbf{x}}) + \frac{b(\underline{\mathbf{x}}) - \upsilon(b(\underline{\mathbf{x}}))}{(d\upsilon/d\theta)_{|\theta = b(\underline{\mathbf{x}})}}, b(\underline{\mathbf{x}}) + \frac{b(\underline{\mathbf{x}}) - \lambda(b(\underline{\mathbf{x}}))}{(d\lambda/d\theta)_{|\theta = b(\underline{\mathbf{x}})}}] \tag{14}$$

is an approximate $1 - \alpha$ confidence interval for θ.

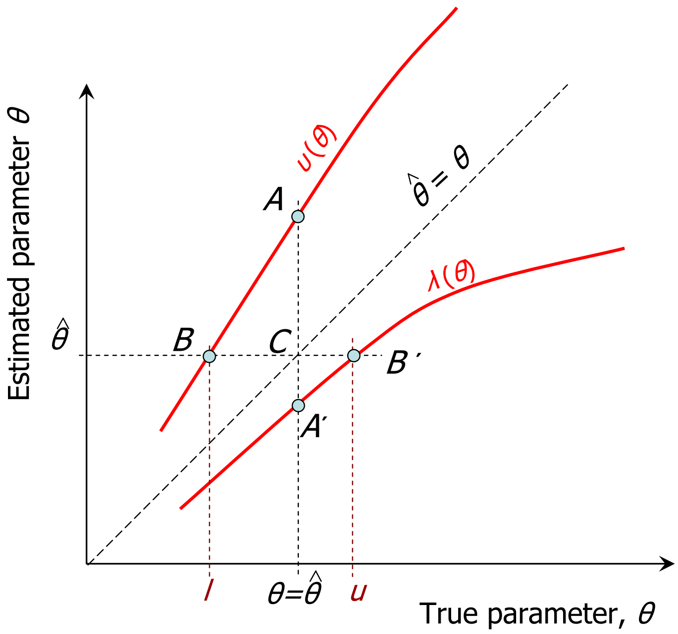

Figure 1. Sketch explaining the determination of confidence limits l and u from an inversion of a hypothesis test.

Under suitable regularity conditions (i.e. Casella & Berger, 2002, p.516) the density of the MLE is given by Hillier & Armstrong (1999). The necessary conditions for the equations (12) and (13) to hold are that λ and υ are continuous and differentiable at a region of $\hat{\theta}$. The validity of these assumptions for certain cases could be investigated, using Hillier & Armstrong (1999) formula, but at such situations this is not always possible.

3.2 Some theoretical results

It is useful to find cases where the above derived confidence interval is exact (i.e. (12) and (13) are exact). We can easily prove that this happens in the case where $\upsilon(\theta) = c_1\theta + c_2$, where c_1 and c_2 are any real numbers:

$$\left(\frac{d\upsilon}{d\theta}\right)_{|\theta=\hat{\theta}} = c_1, \text{ and } \frac{\upsilon(\hat{\theta}) - \hat{\theta}}{\hat{\theta} - l} = \frac{c_1\hat{\theta} + c_2 - \hat{\theta}}{\hat{\theta} - [(\hat{\theta} - c_2)/c_1]} = c_1$$

(The proof for u can be conducted in a similar way and is omitted). Special cases of this are (i) when $\upsilon(\theta) = \theta + c$, and (ii) when $\upsilon(\theta) = c\theta$. These two correspond to the first and second method described by Ripley (1987) respectively (p.176, eq.3 and p.177, eq.6, after substitution of $\left(\frac{d\upsilon}{d\theta}\right)_{|\theta=\hat{\theta}} = c = \upsilon(\theta)/\theta$ in (14)). We can also easily prove that location families correspond to the first case and scale families correspond to the second case. The proof is given below:

(a) For location families the quantity $\underline{\mu} - \mu$ (where $\underline{\mu}$ is a MLE of the location parameter μ) is a pivotal quantity (see Lawless, 2003, p.562). Then from (3) we have that $P\{\underline{\mu} < \lambda(\mu)\} = \alpha/2$, which implies that $P\{\underline{\mu} - \mu < \lambda(\mu) - \mu\} = \alpha/2$ and we obtain that $\lambda(\mu) = \mu + G^{-1}(\alpha/2)$, where G is the distribution function of $\underline{\mu} - \mu$ that does not depend on μ. In a similar way we obtain that $\upsilon(\mu) = \mu + G^{-1}(1 - \alpha/2)$. Now it is obvious from (i) above that the confidence interval obtained by (14) is an exact confidence interval.

(b) For scale families the quantity $\underline{\sigma}/\sigma$ (where $\underline{\sigma}$ is a MLE of the location parameter σ) is a pivotal quantity (see Lawless, 2003, p. 562). Then from (3) we have that $P\{\underline{\sigma} < \lambda(\sigma)\} = \alpha/2$, which implies that $P\{\underline{\sigma}/\sigma < \lambda(\sigma)/\sigma\} = \alpha/2$ and we obtain that $\lambda(\sigma) = \sigma\, G^{-1}(\alpha/2)$, where G is the distribution function of $\underline{\sigma}/\sigma$ and is independent of σ. In a similar way we obtain that $\upsilon(\sigma) = \sigma G^{-1}(1 - \alpha/2)$. Now it is obvious from (ii) above that the confidence interval obtained by (14) is an exact confidence interval.

While in the above cases our method provides exact confidence intervals, when the equation $\upsilon(\theta) = c_1\theta + c_2$ does not hold, it can only provide approximate confidence intervals, where the level of approximation depends on the form of λ and υ and for certain cases will be examined in the next sections. It is also easy to prove that the confidence interval given by (14) is asymptotically equivalent to a Wald-type interval for any function of the parameter θ (and hence for the parameter itself) under certain regularity conditions. The proof is given below.

We want to find a confidence interval for a function $h(\theta)$ of θ. We assume that $\underline{\theta}$ is a MLE of θ. Then according to Casella & Berger (2002, p.497) and Efron & Hinkley (1978), the variance of the function $h(\underline{\theta})$ can be approximated by

$$\hat{Var}(h(\underline{\theta})|\theta) \approx \frac{[h'(\theta)]^2|_{\theta=\hat{\theta}}}{-\frac{\partial^2}{\partial\theta^2}\ln l(\theta|\boldsymbol{x})|_{\theta=\hat{\theta}}} \tag{15}$$

where $\hat{\theta}$ is the maximum likelihood estimate of θ and $l(\theta|\boldsymbol{x})$ the likelihood function of θ. Now according to Casella & Berger 2002, p.497) we have the following result:

$$\frac{h(\underline{\theta}) - h(\theta)}{\sqrt{\hat{Var}(h(\underline{\theta})|\theta)}} \to N(0,1) \tag{16}$$

Then from (3) we get that $P\{h(\underline{\theta}) < \lambda(\theta)\} = \alpha/2$ and $P\{h(\underline{\theta}) > \upsilon(\theta)\} = \alpha/2$ which imply that

$$\lambda(\theta) = h(\theta) + \sqrt{\hat{Var}(h(\underline{\theta})|\theta)}\,\Phi^{-1}(\alpha/2) \text{ and } \upsilon(\theta) = h(\theta) + \sqrt{\hat{Var}(h(\underline{\theta})|\theta)}\,\Phi^{-1}(1 - \alpha/2) \tag{17}$$

where Φ^{-1} denotes the inverse of the standard normal distribution function. If we substitute θ for $h(\theta)$ then (17) becomes identical to case (i) above.

3.3 Construction of the algorithm

Having found an expression for the confidence interval, we can construct a Monte Carlo algorithm to calculate it when there do not exist analytical expressions for the functions of interest. The algorithm has the following steps:

Step 1. We find the MLE of θ, and its maximum likelihood estimate say $\hat{\theta}$.

Step 2. We produce k samples of size n, from $f(x|\hat{\theta})$.

Step 3. We use these k samples to compute $\lambda(\hat{\theta})$ and $\upsilon(\hat{\theta})$.

Step 4. We produce additional k samples of size n, from $f(x|\hat{\theta}+\delta\theta)$, where $\delta\theta$ is a small increment.

Step 5. We use these additional k samples to compute $\lambda(\hat{\theta}+\delta\theta)$ and $\upsilon(\hat{\theta}+\delta\theta)$.

Step 6. We substitute $(\frac{d\upsilon}{d\theta})_{|\theta=\hat{\theta}}$ of (12) with $[\upsilon(\hat{\theta}+\delta\theta) - \upsilon(\hat{\theta})]/d\theta$, and $(\frac{d\lambda}{d\theta})_{|\theta=\hat{\theta}}$ of (13) with $[\lambda(\hat{\theta}$

$+\delta\theta) - \lambda(\hat{\theta})]/\delta\theta$.

Step 7. We compute l and u from (12) and (13).

We conclude based on the construction of the algorithm that it could be applied to cases where θ is estimated by a different estimator. Below we give an application of the algorithm on the normal distribution where we used the unbiased estimator of θ and obtained good results.

4 Construction of confidence intervals for multi-parameter probability distributions

We assume now that we have a multi-parameter probability distribution with density $f(x|\boldsymbol{\theta})$ and parameter $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_k)$, whose estimator is $\underline{\boldsymbol{T}} = (\underline{T}_1, \underline{T}_2, \ldots, \underline{T}_k)$. We wish to calculate a $1-\alpha$ confidence interval for a scalar function $\beta := h(\boldsymbol{\theta})$ of $\boldsymbol{\theta}$. If we assume that $\underline{\boldsymbol{T}}$ is a MLE then $b(\underline{\boldsymbol{x}}) := h(\underline{\boldsymbol{T}})$ is a MLE for $h(\boldsymbol{\theta})$ and $b(\boldsymbol{x}) = h(\boldsymbol{t})$ is its estimate. To extend the method, described by (12) and (13) in the multiple parameter case, the derivatives $d\lambda/d\theta$ and $d\upsilon/d\theta$ should be evaluated at appropriate directions $\boldsymbol{d}_\lambda$ and $\boldsymbol{d}_\upsilon$.

Let $\boldsymbol{\gamma} := (\lambda, \beta, \upsilon)^{\mathrm{T}}$ where λ, υ have been defined by (3) and let $\mathrm{Var}(\underline{\boldsymbol{T}}) = \mathrm{diag}(\mathrm{Var}(\underline{T}_1), ..., \mathrm{Var}(\underline{T}_k))$. The latter can be easily computed during the same Monte Carlo simulation that is performed to compute $\boldsymbol{\gamma}$. It is reasonable to assume that $\boldsymbol{d}_\lambda$ and $\boldsymbol{d}_\upsilon$ will depend on $\mathrm{Var}(\underline{\boldsymbol{T}})$ as well as of the matrix of derivatives of $\boldsymbol{\gamma}$,

$$\frac{d\boldsymbol{\gamma}}{d\boldsymbol{\theta}} = \begin{bmatrix} \dfrac{d\lambda}{d\boldsymbol{\theta}} \\ \dfrac{d\beta}{d\boldsymbol{\theta}} \\ \dfrac{d\upsilon}{d\boldsymbol{\theta}} \end{bmatrix} = \begin{bmatrix} \dfrac{\partial\lambda}{\partial\theta_1} & \dfrac{\partial\lambda}{\partial\theta_2} & \cdots & \dfrac{\partial\lambda}{\partial\theta_k} \\ \dfrac{\partial\beta}{\partial\theta_1} & \dfrac{\partial\beta}{\partial\theta_2} & \cdots & \dfrac{\partial\beta}{\partial\theta_k} \\ \dfrac{\partial\upsilon}{\partial\theta_1} & \dfrac{\partial\upsilon}{\partial\theta_2} & \cdots & \dfrac{\partial\upsilon}{\partial\theta_k} \end{bmatrix} \tag{18}$$

Heuristically, we can assume a simple relation of the form

$$\boldsymbol{d}_\lambda = \mathrm{Var}(\underline{\boldsymbol{T}}) \left(\frac{d\boldsymbol{\gamma}}{d\boldsymbol{\theta}}\right)^{\mathrm{T}} \boldsymbol{e}_\lambda \tag{19}$$

where $\mathbf{e}_\lambda$ is a size 3 vector of constants needed to transform the matrix product of the first two terms of the right hand side into a vector. The elements of this vector could be thought of as weights corresponding to each of the derivatives of the three elements of $\boldsymbol{\gamma}$. The simplest choice is to assume equal weights, i.e. $\boldsymbol{e}_\lambda = (1, 1, 1)^{\mathrm{T}}$. However, numerical investigations showed that the choice $\boldsymbol{e}_\lambda = (0, 1, 1)^{\mathrm{T}}$ yields better approximations and the theoretical analysis below showed that it yields asymptotically good results under certain regularity conditions.

The derivatives of λ and β with respect to $\boldsymbol{\theta}$ on direction $\boldsymbol{d}_\lambda$ will then be

$$\left(\frac{d\lambda}{d\boldsymbol{\theta}}\right) \boldsymbol{d}_\lambda = \left(\frac{d\lambda}{d\boldsymbol{\theta}}\right) \mathrm{Var}(\underline{\boldsymbol{T}}) \left(\frac{d\boldsymbol{\gamma}}{d\boldsymbol{\theta}}\right)^{\mathrm{T}} \boldsymbol{e}_\lambda, \quad \left(\frac{d\beta}{d\boldsymbol{\theta}}\right) \boldsymbol{d}_\lambda = \left(\frac{d\beta}{d\boldsymbol{\theta}}\right) \mathrm{Var}(\underline{\boldsymbol{T}}) \left(\frac{d\boldsymbol{\gamma}}{d\boldsymbol{\theta}}\right)^{\mathrm{T}} \boldsymbol{e}_\lambda \tag{20}$$

and are both scalars, so by taking their ratio we can calculate $d\lambda/d\beta$. By symmetry, similar relationships can be written for υ and $\boldsymbol{d}_\upsilon$ with $\boldsymbol{e}_\upsilon = (1, 1, 0)^{\mathrm{T}}$. The two groups of relationships can be unified in terms of the 3×3 matrix $\boldsymbol{q}$ defined as

$$\boldsymbol{q} := \frac{d\gamma}{d\boldsymbol{\theta}} \mathrm{Var}(\underline{\boldsymbol{T}}) \left(\frac{d\gamma}{d\boldsymbol{\theta}}\right)^{\mathrm{T}} \tag{21}$$

It can then be easily shown that on the directions $\boldsymbol{d}_\lambda$ and $\boldsymbol{d}_\upsilon$,

$$\frac{d\lambda}{d\beta} = \frac{q_{12} + q_{13}}{q_{22} + q_{23}}, \quad \frac{d\upsilon}{d\beta} = \frac{q_{31} + q_{32}}{q_{21} + q_{22}} \tag{22}$$

In Appendix A we show that the confidence interval for the parameter μ of a normal distribution $N(\mu,\sigma^2)$ is asymptotically equivalent to a Wald-type interval. We also show that the confidence interval obtained by our method is asymptotically equivalent to a Wald-type interval for two-parameter regular distributions and hence for any multi-parameter distribution.

4.1 Construction of the algorithm

Now the algorithm for the calculation of the intervals follows:

Step 1. We find the MLE of $\boldsymbol{\theta}$ namely $\underline{\boldsymbol{\theta}}$, and its maximum likelihood estimate say $\hat{\boldsymbol{\theta}}$.

Step 2. The MLE of β is $h(\underline{\boldsymbol{\theta}})$, and its maximum likelihood estimate is $h(\hat{\boldsymbol{\theta}})$.

Step 3. We produce m samples of size n, from $f(x|\hat{\boldsymbol{\theta}})$.

Step 4. We use these m samples to compute $\lambda(\hat{\boldsymbol{\theta}})$, $\upsilon(\hat{\boldsymbol{\theta}})$, $h(\hat{\boldsymbol{\theta}})$ and Var($\boldsymbol{T}$).

Step 5. We produce additional m samples of size n, from $f(x|\hat{\boldsymbol{\theta}}+\delta\boldsymbol{\theta}_i)$, where $\delta\boldsymbol{\theta}_i$ is a vector with all elements zero except the ith element, which is a small quantity $\delta\theta_i$

Step 6. We use these additional m samples to compute $\lambda(\hat{\boldsymbol{\theta}}+\delta\boldsymbol{\theta}_i)$, $\upsilon(\hat{\boldsymbol{\theta}}+\delta\boldsymbol{\theta}_i)$ and $h(\hat{\boldsymbol{\theta}}+\delta\boldsymbol{\theta}_i)$.

Step 7. We repeat steps 4 and 5 for $i = 1,2,\ldots,k$.

Step 8. We substitute in (18) $[\lambda(\hat{\boldsymbol{\theta}}+\delta\boldsymbol{\theta}_i) - \lambda(\hat{\boldsymbol{\theta}})]/\delta\theta_i$ for $\frac{\partial\lambda}{\partial\theta_i}$, $[\upsilon(\hat{\boldsymbol{\theta}}+\delta\boldsymbol{\theta}_i) - \upsilon(\hat{\boldsymbol{\theta}})]/\delta\theta_i$ for $\frac{\partial\upsilon}{\partial\theta_i}$ and $[h(\hat{\boldsymbol{\theta}}+\delta\boldsymbol{\theta}_i) - h(\hat{\boldsymbol{\theta}})]/\delta\theta_i$ for $\frac{\partial h}{\partial\theta_i}$.

Step 9. We calculate $\boldsymbol{q}$ from (21).

Step 10. We compute l and u from (12) and (13).

5 Simulation results

To test the algorithm in specific cases, we construct confidence intervals for the scale parameter of the exponential distribution, the location parameter and the pth percentile of the normal distribution, the scale and shape parameter of the gamma distribution and the scale parameter and the pth percentile of the Weibull distribution. Then we compare the numerical results with known, mostly analytical, results from the literature.

5.1 Confidence interval for the scale parameter of the exponential distribution

The density of the exponential distribution is $f(x|\sigma) = (1/\sigma)\exp(-x/\sigma)$, $x \geq 0$, $\sigma > 0$. The MLE of σ is $\underline{\sigma} = \underline{\overline{x}}$. A $1-\alpha$ Wald-type confidence interval (Papoulis & Pillai, 2002, p.310), is

$$[l(\underline{\mathbf{x}}),u(\underline{\mathbf{x}})] = \left[\frac{\overline{x}}{1 + \Phi^{-1}(1-\alpha/2)/\sqrt{n}}, \frac{\overline{x}}{1 - \Phi^{-1}(1-\alpha/2)/\sqrt{n}}\right] \tag{23}$$

We find a $1-\alpha$ exact confidence interval, using the pivotal quantity $\underline{\sigma}/\sigma$. The distribution of $\underline{\sigma}$ is gamma, with shape parameter n and scale parameter σ/n and a $1-\alpha$ exact confidence interval is obtained by the following equations.

$$F(\overline{x}|n,l/n) = 1-\alpha/2, \; F(\overline{x}|n,u/n) = \alpha/2 \tag{24}$$

where $F(x|k,\theta)$ is the gamma distribution whose density is $f(x|\theta,k) = \dfrac{e^{-x/\theta}x^{k-1}}{\Gamma(k)\theta^k}$, $x \geq 0$ where $\theta > 0$ is the scale parameter and $k > 0$ is the shape parameter.

The confidence interval obtained by (24) is exact and the confidence interval obtained by (23) is Wald-type. These two are intercompared also with the BCa bootstrap non-parametric interval, designated as "bootstrap", the two confidence intervals obtained by the two Ripley's methods, designated as "Ripley location" and "Ripley scale", respectively, and the confidence interval obtained by our algorithm, designated as MCCI (Monte Carlo Confidence Interval). Figure 2 compares the confidence intervals obtained by all six methods for a simulated sample with 50 elements from an exponential distribution with σ = 1. For this sample $\hat{\sigma}$ = 1.002. As we see, MCCI is close to the exact and the "Ripley scale" and gives a better approximation than the Wald-type, the "bootstrap" and the "Ripley location".

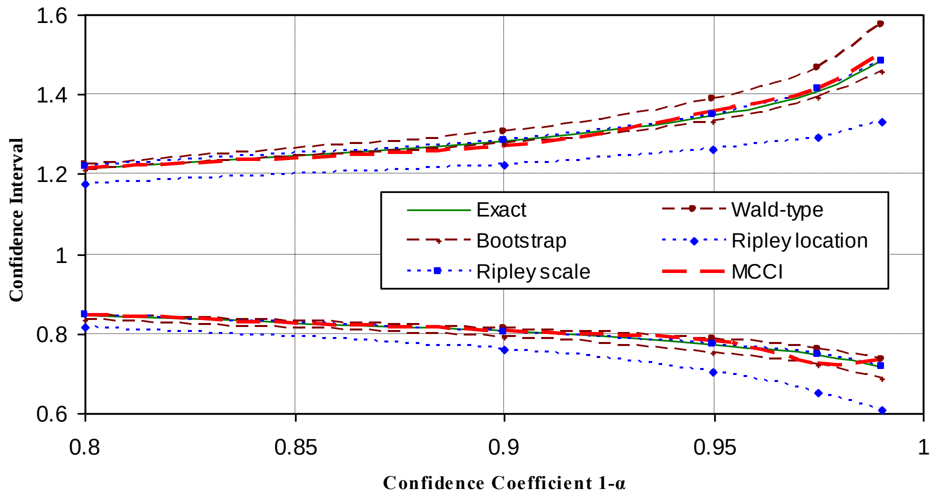

Figure 2. Confidence intervals for the scale of an exponential distribution with n = 50, $\hat{\sigma}$ = 1.002. Here the number of samples k = 50 000 for MCCI, "Ripley location" and "Ripley scale" cases and $\delta\sigma$ = 0.05.

5.2 Confidence interval for the location parameter of the normal distribution

The density of the normal distribution is $f(x|\mu,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}}\exp(\frac{(x-\mu)^2}{2\sigma^2})$, where μ is the location parameter, and $\sigma > 0$ is the scale parameter. A $1-\alpha$ exact confidence interval (Papoulis & Pillai, 2002, p.309) is

$$[l(\mathbf{x}),u(\mathbf{x})] = [\bar{x} - t_{n-1}(1-\alpha/2)\frac{s}{\sqrt{n}}, \bar{x} + t_{n-1}(1-\alpha/2)\frac{s}{\sqrt{n}}] \quad (25)$$

A $1-\alpha$ Wald-type confidence interval (Papoulis & Pillai, 2002, p.309) is

$$[l(\underline{\mathbf{x}}),u(\underline{\mathbf{x}})] = [\bar{x} - \Phi^{-1}(1-\alpha/2)\frac{s}{\sqrt{n}}, \bar{x} + \Phi^{-1}(1-\alpha/2)\frac{s}{\sqrt{n}}] \quad (26)$$

We compare the MCCI with the exact interval obtained by (25), as well as with the Wald-type interval, the BCa interval and the intervals obtained by Ripley's two methods. Figure 3 compares the confidence intervals obtained by the six methods for a simulated sample with 10 elements from a normal distribution with μ = 0 and σ = 1. For this sample $\hat{\mu}$ = 0.026 and $\hat{\sigma}$ = 1.023. In this case for the calculation of the confidence interval we use the unbiased estimators of μ and σ^2 (instead of the MLE). As we see, MCCI gives a better approximation than the other four approximate methods.

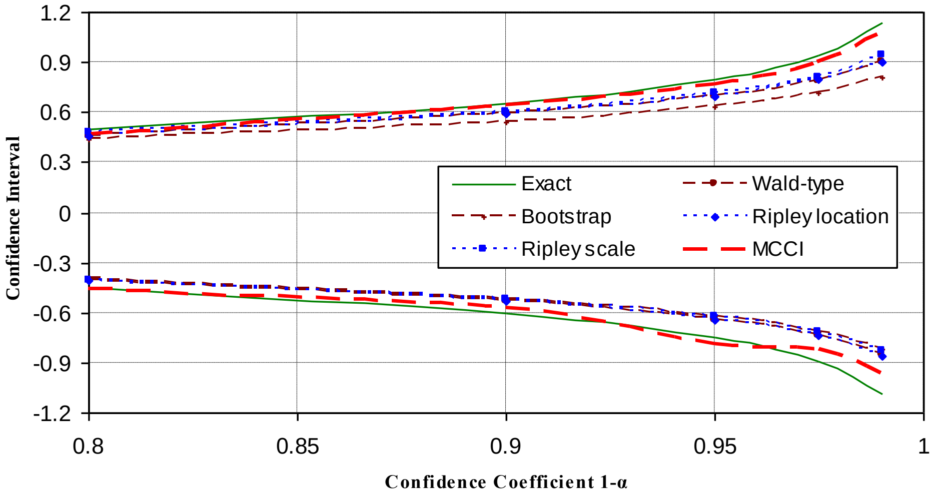

Figure 3. Confidence intervals for the location parameter of a normal distribution with n = 10, $\hat{\mu}$ = 0.026 and $\hat{\sigma}$ = 1.023. Here the number of samples k = 100 000 for MCCI, "Ripley location" and "Ripley scale" cases, $\delta\mu$ = 0.1 and $\delta\sigma$ = 0.1.

5.3 Confidence interval for the percentile of the normal distribution

The pth percentile is $t_p := \mu + z_p\sigma$, where z_p is pth percentile of the standard normal distribution. A $1-\alpha$ Wald-type confidence interval estimate is given by the following equation (e.g. Koutsoyiannis, 1997, p.69).

$$[l(\mathbf{x}),u(\mathbf{x})] = [\bar{x} + z_p s - \Phi^{-1}(1-\alpha/2)\,\frac{s}{\sqrt{n}}\sqrt{1+z_p^2/2},\ \bar{x} + z_p s + \Phi^{-1}(1-\alpha/2)\,\frac{s}{\sqrt{n}}\sqrt{1+z_p^2/2}] \quad (27)$$

Another way to obtain a confidence interval is by using Bayesian analysis (see Gelman et al., 2004, p.75,76). Then if we chose a prior $P(\mu,\sigma) \propto 1/\sigma^2$, we can construct a sampler based on the following mixture.

$$\underline{\sigma}^2|\mathbf{x} \sim \text{inv-}\chi^2(n-1,s^2) \text{ and } \underline{\mu}|\sigma^2,\mathbf{x} \sim \mathrm{N}(\bar{x},\sigma^2/n) \tag{28}$$

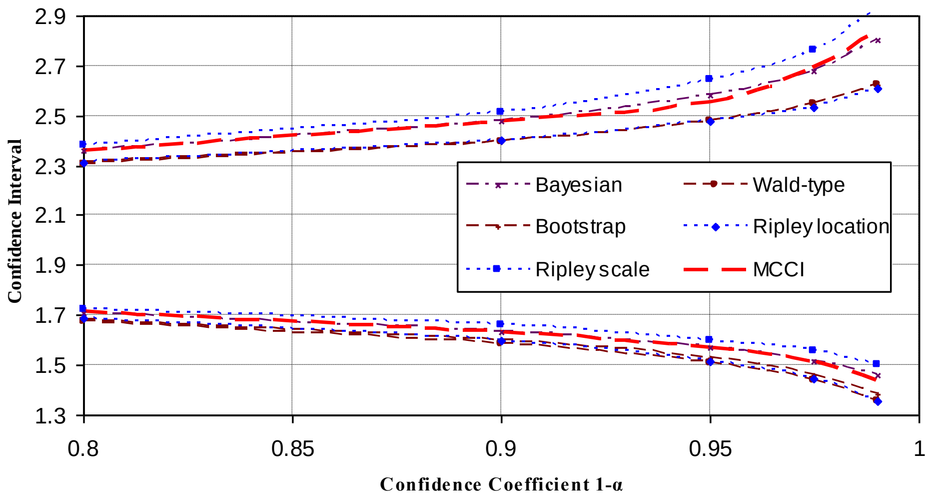

Figure 4. Confidence intervals for μ + 2σ of a normal distribution with n = 50, $\hat{\mu}$ = −0.027 and $\hat{\sigma}$ = 0.998. Here the number of samples m = 50 000 for MCCI, "Ripley location" and "Ripley scale" cases, δμ = 0.1 and δσ = 0.1.

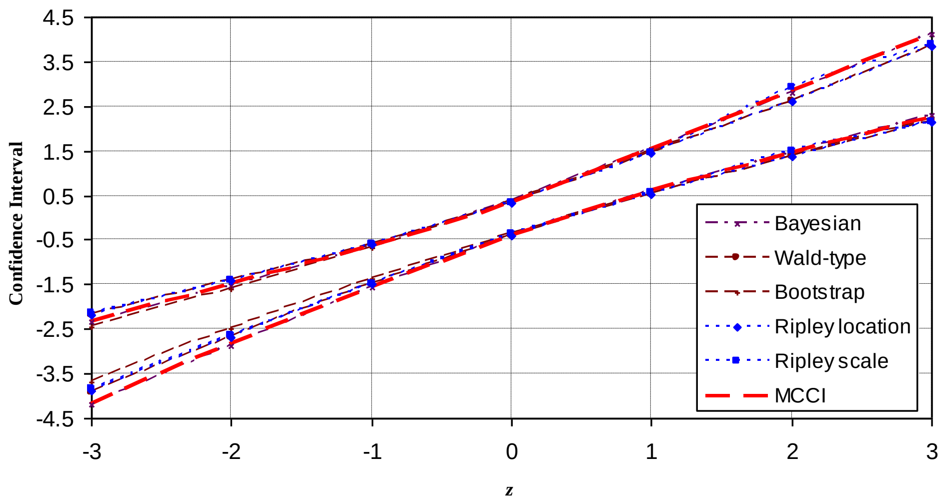

Figure 5. Confidence intervals with confidence coefficient 1−0.01 for μ + $z\sigma$ of a normal distribution with n = 50, $\hat{\mu}$ = −0.027 and $\hat{\sigma}$ = 0.998. Here the number of samples m = 50 000 for MCCI, "Ripley location" and "Ripley scale" cases, δμ = 0.1 and δσ = 0.1.

Thus, here we compare six confidence intervals, the Bayesian confidence region, the Wald-type of equation (27), the BCa interval, the intervals obtained by Ripley's two methods, and the MCCI. Figure 4 compares the confidence intervals for μ + 2σ obtained by the six methods for a simulated sample with 50 elements from a normal distribution with $\mu = 0$ and $\sigma = 1$. For this sample $\hat{\mu} = -0.027$ and $\hat{\sigma} = 0.998$. As we see, Bayesian and MCCI are almost indistinguishable, and MCCI is better when compared to "Ripley location" and "Ripley scale". "Ripley location" is close to the Wald-type and the "bootstrap". The same holds for Figure 5 which compares all methods on the calculation of a 1−0.01 confidence interval, when z_p varies from −3 to 3.

5.4 Confidence interval for the scale parameter of the gamma distribution

First we show how we can calculate an approximate confidence interval for the scale parameter of the gamma distribution. We define $\underline{R}_n := \ln(\overline{\underline{x}}/\tilde{\underline{x}})$, where $\overline{\underline{x}}$ and $\tilde{\underline{x}}$ are the arithmetic and the geometric mean of a size-n sample, which, according to Bhaumik et al. (2009) and Bain & Engelhardt (1975), has a distribution independent of the scale parameter. The maximum likelihood estimates of θ and k according to Bhaumik et al. (2009) and Choi & Wette (1969), denoted by $\hat{k}$ and $\hat{\theta}$ are the solutions of the following equations:

$$R_n = \ln(k) - \psi(k) \text{ and } k\theta = \overline{x} \tag{29}$$

where ψ denotes the digamma function.

We have that

$$\mathrm{E}(\underline{R}_n) = -\ln(n) - \psi(k) + \psi(nk) \text{ and } \mathrm{Var}(R_n) = (1/n)\psi'(k) - \psi'(nk) \tag{30}$$

We also define as c and v, functions of k and n, the solutions of the following system of equations

$$2nk\,\mathrm{E}(\underline{R}_n) = cv \text{ and } (2nk)^2\mathrm{Var}(\underline{R}_n) = 2c^2v \tag{31}$$

From (31) we obtain

$$c = \frac{nk\mathrm{Var}(\underline{R}_n)}{\mathrm{E}(\underline{R}_n)} \text{ and } v = \frac{2\mathrm{E}^2(\underline{R}_n)}{\mathrm{Var}(\underline{R}_n)} \tag{32}$$

For the construction of the confidence interval see Bhaumik et al. (2009) and Engelhardt & Bain (1977). The statistic $\underline{z} = 2\underline{x}/\theta$ has approximately a chi-square distribution with $2n\hat{k}$ degrees of freedom, specifically $\underline{z} \sim \chi^2_{2n\hat{k}}$. The statistic $\underline{T}_1 = 2n\hat{k}\underline{R}_n/c + \underline{z} \sim \chi^2_{v+2n\hat{k}}$. Now using the $\underline{T}_1$ statistic we obtain the following $1 - \alpha$ confidence interval for the scale parameter θ.

$$[l(\underline{\boldsymbol{x}}),u(\underline{\boldsymbol{x}})] = [\frac{2n\bar{\underline{x}}}{\chi^2_{v+2n\hat{k}}(1-\alpha/2) - 2n\hat{k}\,R_n/c},\frac{2n\bar{\underline{x}}}{\chi^2_{v+2n\hat{k}}(\alpha/2) - 2n\hat{k}\,R_n/c}] \tag{33}$$

We will designate the confidence interval obtained by (33) as "approximate". Another way to obtain a confidence interval is by using Bayesian analysis (See Robert, 2007). According to Son & Oh (2006), if we chose a prior $P(k,\theta) \propto 1/\theta$, we construct a Gibbs sampler using the following equations

$$\underline{\theta}|k,\boldsymbol{x} \sim IG(nk,1/\sum_{i=1}^{n} x_i) \tag{34}$$

where $IG(a,b)$ denotes the inverse gamma distribution with parameters a and b, with density function $f(x|a,b) = [\Gamma(a)b^a]^{-1}x^{-(a+1)}\exp(-1/bx)$, x, a, $b > 0$. Also

$$P(k|\theta,\boldsymbol{x}) \propto [\Gamma(k)]^{-n}\theta^{-nk}\prod_{i=1}^{n} x_i^{k-1} \tag{35}$$

A Wald-type interval is calculated, using the formula in Casella & Berger (2002, p. 497)

$$[l(\boldsymbol{x}),u(\boldsymbol{x})] = [\hat{\theta} - \Phi^{-1}(1-\alpha/2)\sqrt{-l''(\hat{\theta})},\ \hat{\theta} + \Phi^{-1}(1-\alpha/2)\sqrt{-l''(\hat{\theta})}] \tag{36}$$

where $-\,l''(\hat{\theta})$ is an estimate of the Hessian at $(\hat{k}, \hat{\theta})$, when optimizing the log-likelihood function.

We designate the confidence region obtained by (34) and (35) as Bayesian, the BCa interval as "bootstrap", the confidence interval obtained by Ripley's two methods as "Ripley location" and "Ripley scale" and the confidence interval obtained by our algorithm as MCCI.

Figure 6 compares the confidence intervals obtained by all seven methods for a simulated sample with 50 elements from a gamma distribution with k = 2 and θ = 3. For this sample $\hat{k}$ = 1.979 and $\hat{\theta}$ = 3.007. As we can see, the MCCI, "Ripley scale" and "bootstrap" limits are close to the Bayesian ones, but the approximate, "Wald-type and "Ripley location" limits lie far apart, which shows that they do not provide a satisfactory approximation (perhaps owing to too many assumptions involved in their derivation).

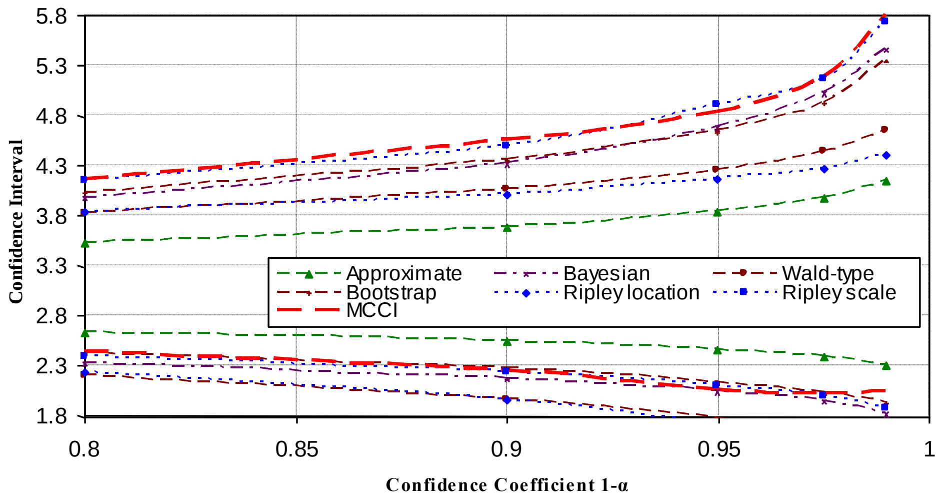

Figure 6. Confidence intervals for the scale parameter of a gamma distribution with n = 50, $\hat{k}$ = 1.979 and $\hat{\theta}$ = 3.007. Here the number of samples m = 20 000 for MCCI, "Ripley location" and "Ripley scale" cases, δk = 0.3 and $\delta\theta$ = 0.3.

5.5 Confidence interval for the shape parameter of the gamma distribution

To obtain a 1 − α confidence interval for the shape parameter k, according to Bhaumik et al. (2009, see also Engelhardt & Bain, 1978), we use the statistic $\underline{T}_1 = 2nkR_n \sim c\chi^2_v$, approximately. Then a 1 − α confidence interval corresponds to the following inequality

$$\frac{\mathrm{Var}(R_n)}{\mathrm{E}(R_n)}\chi^2_v(\alpha/2) < 2R_n < \frac{\mathrm{Var}(R_n)}{\mathrm{E}(R_n)}\chi^2_v(1-\alpha/2) \qquad (37)$$

where we solve for k.

A Wald-type interval is calculated, using the formula in Casella & Berger (2002, p. 497)

$$[l(\mathbf{x}),u(\mathbf{x})]=[\,\hat{k}-\Phi^{-1}(1-\alpha/2)\sqrt{-l''(\hat{k})},\,\hat{k}+\Phi^{-1}(1-\alpha/2)\sqrt{-l''(\hat{k})}] \qquad (38)$$

where $-\,l''(\hat{k})$ is an estimate of the Hessian at $(\hat{k},\,\hat{\theta})$, when optimizing the log-likelihood function.

We designate the confidence interval obtained by (37) as "approximate", the confidence region obtained by (34), (35) as Bayesian, the confidence interval obtained by (38) as Wald-type, the BCa confidence interval as "bootstrap", the confidence intervals obtained by the two Ripley's methods as "Ripley location" and "Ripley scale" and the confidence interval obtained by our algorithm as MCCI.

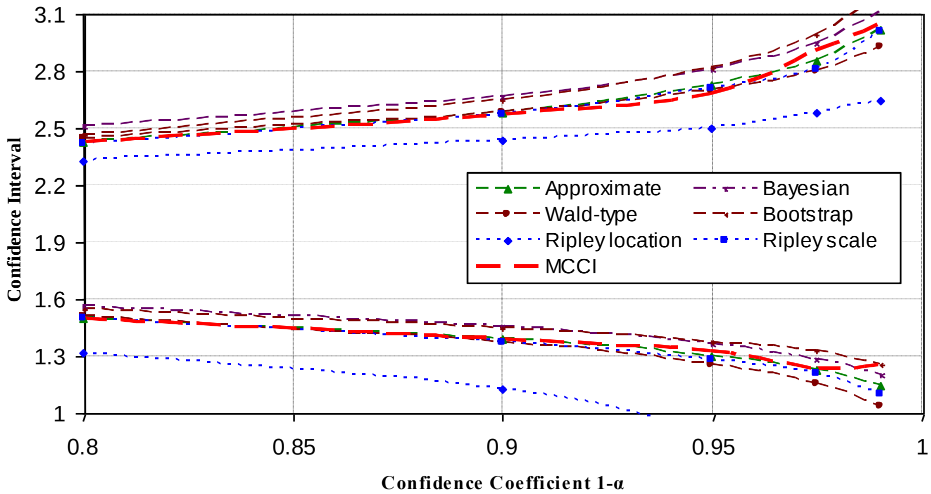

Figure 7. Confidence intervals for the shape parameter of a gamma distribution with n = 50, $\hat{k}$ = 1.979 and $\hat{\theta}$ = 3.007. Here the number of samples m = 20 000 for MCCI, "Ripley location" and "Ripley scale" cases, δk = 0.3 and $\delta\theta$ = 0.3.

Figure 7 compares the confidence intervals obtained by all seven methods for a simulated sample with 50 elements from a gamma distribution with k = 2 and θ = 3. For this sample $\hat{k}$ = 1.979 and $\hat{\theta}$ = 3.007. As we can see, the "approximate", Wald-type, "Ripley location" and MCCI confidence intervals are close. The Bayesian confidence region is close to the "approximate" which, in our opinion, gives a good approximation of the exact confidence interval. "Ripley location" is far from the other intervals.

5.6 Confidence interval for the scale parameter of the Weibull distribution

The density of the Weibull distribution is $f(x|a,b) = \frac{b}{a}\left(\frac{x}{a}\right)^{b-1} \exp[-\left(\frac{x}{a}\right)^{b}]$, $x > 0$ where $a > 0$ is the scale parameter and $b > 0$ is the shape parameter. According to Yang et al. (2007) first we must find a modified MLE of b, according to the following equation, which is a modification of the equations discussed in Cohen (1965).

$$l(b) := \frac{n-2}{b} - (n\sum_{i=1}^{n} x_i^{b}\ln x_i)(\sum_{i=1}^{n} x_i^{b})^{-1} + \sum_{i=1}^{n}\ln x_i = 0 \tag{39}$$

We denote $\hat{b}$ the modified maximum likelihood estimate given by (39) and $\hat{a}$ the modified maximum likelihood estimate given by the following equation.

$$\hat{a} = [(1/n)\sum_{i=1}^{n} x_i^{\hat{b}}]^{1/\hat{b}} \tag{40}$$

We define $S(b) := \sum_{i=1}^{n} x_i^{b}$ and $c_1 := \sqrt{1 + 0.607927\cdot 0.422642^2}$. Then a $1 - \alpha$ confidence interval estimate is given by the following equation.

$$[l(\boldsymbol{x}),u(\boldsymbol{x})] = [(\frac{2S(\hat{b})}{c_1\chi^2_{2n}(1-\alpha/2) - 2n(c_1-1)})^{1/\hat{b}},(\frac{2S(\hat{b})}{c_1\chi^2_{2n}(\alpha/2) - 2n(c_1-1)})^{1/\hat{b}}] \tag{41}$$

A Wald-type interval is calculated, using the formula in Casella & Berger (2002, p. 497)

$$[l(\boldsymbol{x}),u(\boldsymbol{x})]=[\hat{a} - \Phi^{-1}(1-\alpha/2)\sqrt{-l''(\hat{a})}, \hat{a} + \Phi^{-1}(1-\alpha/2)\sqrt{-l''(\hat{a})}] \tag{42}$$

where $-l''(\hat{a})$ is an estimate of the Hessian at $(\hat{a}, \hat{b})$, when optimizing the log-likelihood function.

We designate the interval obtained by (41) as "approximate", the interval obtained by (42) as Wald-type, the BCa interval as "bootstrap", the confidence interval obtained by Ripley’s two method as "Ripley location" and "Ripley scale" and the confidence interval obtained by our algorithm as MCCI.

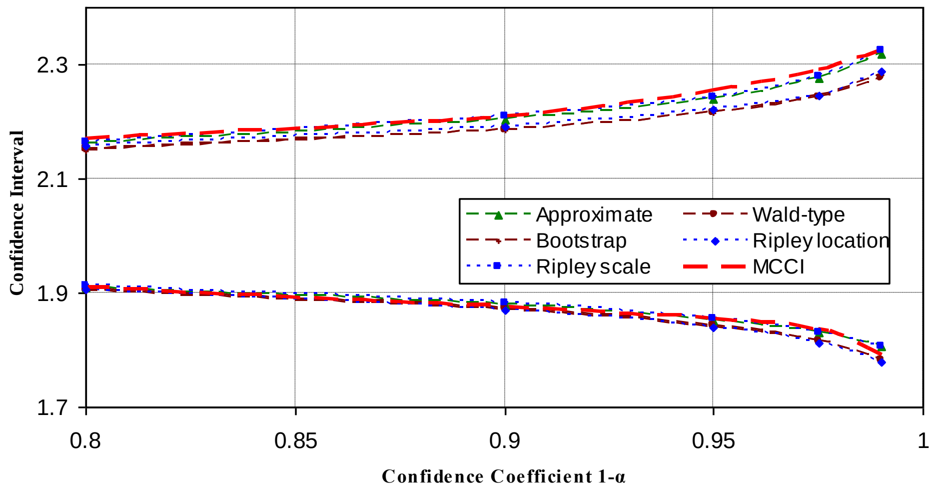

Figure 8. Confidence intervals for the scale parameter of a Weibull distribution with n = 50, $\hat{a}$ = 2.022 and $\hat{b}$ = 3.097. Here the number of samples m = 20 000 for MCCI, "Ripley location" and "Ripley scale" cases, δa = 0.1 and δb = 0.1.

Figure 8 compares the confidence intervals obtained by the six methods for a simulated sample with 50 elements from a Weibull distribution with a = 2 and b = 3. For this sample $\hat{a}$ = 2.022 and $\hat{b}$ = 3.097. As we can see, the "approximate", "Ripley scale" and MCCI confidence intervals are almost indistinguishable and the Wald-type, "bootstrap" and "Ripley location" are far from the previous intervals.

5.7 Confidence interval for the *p*th percentile of the Weibull distribution

According to Yang et al. (2007) the pth percentile of the Weibull distribution is $t_p = a[-\ln(1-p)]^{1/b}$. Then a $1 - \alpha$ approximate confidence interval estimate is given by the following equation.

$$[l(\mathbf{x}),u(\mathbf{x})] = [(-\frac{2S(\hat{b})\ln(1-p)}{c_2\chi^2_{2n}(1-\alpha/2)-2n(c_2-1)})^{1/\hat{b}},(-\frac{2S(\hat{b})\ln(1-p)}{c_2\chi^2_{2n}(\alpha/2)-2n(c_2-1)})^{1/\hat{b}}] \tag{43}$$

where $c_2 := \sqrt{1+0.607927\cdot\{0.422642-\ln[-\ln(1-p)]\}^2}$.

Figure 9 compares the confidence intervals obtained by the five methods, the "approximate", the "bootstrap", the "Ripley location", the "Ripley scale" and the MCCI for a simulated sample with 50 elements from a Weibull distribution with a = 2 and b = 3. For this sample $\hat{a}$ = 2.022 and $\hat{b}$ = 3.097. "Bootstrap" and "Ripley location" are close to each other but far from the other three confidence intervals.

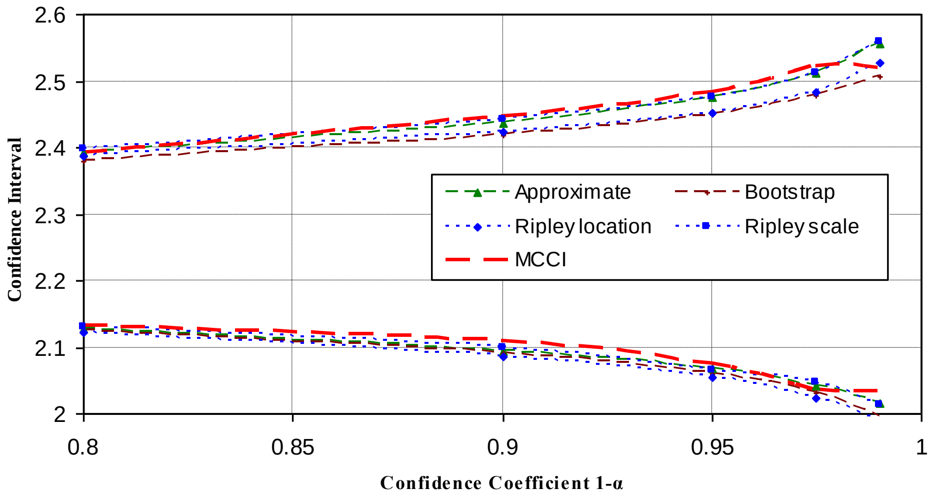

Figure 9. Confidence intervals for the 75th percentile of a Weibull distribution with n = 50, $\hat{a}$ = 2.022 and $\hat{b}$ = 3.097. Here the number of samples m = 20 000 for MCCI, "Ripley location" and "Ripley scale" cases, δa = 0.1 and δb = 0.1.

5.8 Summary results

Table 1 shows the results of all previous methods summarized. MCCI is similar to "exact" (when "exact" can be calculated analytically, cases 1,2). MCCI is also similar to "approximate", in cases 5,6,7. In these cases "approximate" seems to be a good approximation of an exact confidence interval. This implies that MCCI is a good approximation of an exact confidence interval.

On the other hand in case 3 MCCI is almost identical to "Bayesian" which we think is a good property. In case 4, MCCI is closer to "Bayesian" than the "approximate". We believe that the "approximate" is not a good approximation of an exact confidence interval, because it involves a lot of assumptions and transformations. MCCI was also better in our opinion than Wald-type and bootstrap intervals in all cases. We should also keep in mind that confidence intervals and Bayesian "confidence regions" are not directly comparable (see also the chapter dedicated to matching priors in Robert, 2007, p.137).

As an additional means of intercomparison, coverage probabilities using Monte Carlo methods were calculated for all methods except for the Bayesian confidence regions and the algorithm behaved relatively well in all cases (Table 2). MCCI was better when estimating the confidence intervals for the normal and the gamma distribution parameters, and had the best mean rank for all the examined cases.

An application of the algorithm, using historical river flow data is given in appendix B.

Table 1. Summary results of the case studies examined. Smaller numbers mean that the corresponding result is better. Equal numbers mean that there is a similarity between the different results. For example, in the case of the percentile of the normal distribution, MCCI, "Ripley scale" and "Bayesian" methods (marked as 1) gave similar results, whereas Wald-type, "bootstrap" and "Ripley location" methods (marked as 2, 3 and 3 correspondingly) gave results worse than the former methods.

				Methods							
Case	Figure No	Distribution	Parameter	Exact	Bayesian	Approximate	Ripley location	Ripley scale	Wald-type	Bootstrap	MCCI
1	2	Exponential	Scale	1			4	1	3	2	1
2	3	Normal	Location	1			2	2	2	3	1
3	4	Normal	Percentile		1		3	1	2	3	1
4	6	Gamma	Scale		1	3	2	1	2	1	1
5	7	Gamma	Shape		2	1	3	1	1	2	1
6	8	Weibull	Scale			1	2	1	2	2	1
7	9	Weibull	Percentile			1	2	1		2	1

Table 2. Monte Carlo coverage probabilities and rank of each method when calculating 0.975 confidence intervals after 10 000 iterations (rank 1 is assigned to the method of best performance).

						Coverage probabilities (with ranks in parentheses) for all methods					
Case	Distribution	Parameter	Sample size	Parameter value	Parameter value	Approximate	Ripley location	Ripley scale	Wald-type	Bootstrap	MCCI
1	Exponential	Scale	10	$\sigma = 2$			0.889 (5)	0.977 (2)	0.975 (1)	0.916 (4)	0.966 (3)
2	Normal	Location	10	$\mu = 0$	$\sigma = 1$		0.946 (3)	0.946 (3)	0.947 (2)	0.931 (5)	0.968 (1)
3	Normal	Percentile	10	$\mu = 0$	$\sigma = 1$		0.919 (4)	0.929 (2)	0.929 (2)	0.867 (5)	0.973 (1)
4	Gamma	Scale	50	$k = 2$	$\theta = 3$	0.753	0.923 (5)	0.976 (1)	0.940 (4)	0.957 (3)	0.974 (1)
5	Gamma	Shape	50	$k = 2$	$\theta = 3$	0.976	0.948 (5)	0.972 (2)	0.978 (2)	0.956 (4)	0.974 (1)
6	Weibull	Scale	50	$a = 2$	$b = 3$	0.971	0.969 (3)	0.970 (2)	0.966 (4)	0.965 (5)	0.973 (1)
7	Weibull	Percentile	50	$a = 2$	$b = 3$	0.971	0.968 (3)	0.970 (1)		0.961 (4)	0.969 (2)
	mean rank						4.000	1.857	2.500	4.286	1.429

6 Sensitivity of the algorithm on the choice of the increment and the simulation sample size

In this chapter we test the sensitivity of the algorithm on the choice of the increments $\delta\mu$ and $\delta\sigma$ and the simulated sample size in the case of the location parameter and the percentile of the normal distribution.

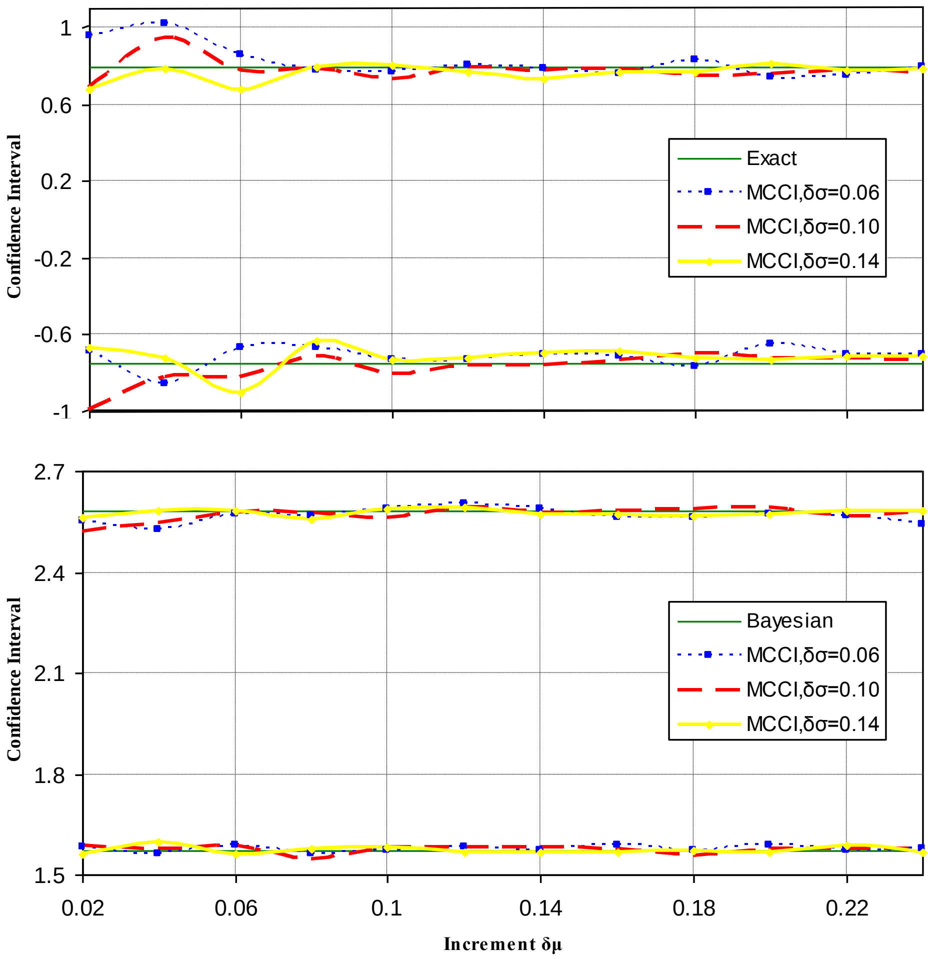

Figure 10. 0.95 confidence intervals for a normal distribution estimated for different $\delta\mu$ and $\delta\sigma$ (the parameter increments denoted in text as $\delta\theta_i$): (upper) confidence interval for the location parameter μ from a sample with $n = 10$, $\hat{\mu} = 0.026$ and $\hat{\sigma} = 1.023$ and number of samples drawn m = 100 000; (lower) confidence interval for the quantity $\mu + 2\sigma$ from a sample with $n = 50$, $\hat{\mu} = -0.027$ and $\hat{\sigma} = 0.998$ and number of samples drawn m = 50 000.

Figure 10 tests the sensitivity of the algorithm on the choice of the increments $\delta\mu$ and $\delta\sigma$ in the cases of the location and the percentile parameters of the normal distribution, for $n = 10$

(upper panel) and n = 50 (lower panel), where for the calculation of the confidence interval the unbiased estimators of μ and σ^2 were used. As we see, the algorithm gives good approximations, regardless of the choice of $\delta\mu$ and $\delta\sigma$. For small n, a slight problem appears if $\delta\mu$ is too small (< 0.5). Figure 11 describes the convergence of the algorithm for the same cases. The speed of convergence is low since ~50 000 iterations are needed for its stabilization, although reasonable results are obtained even for ~10 000 iterations.

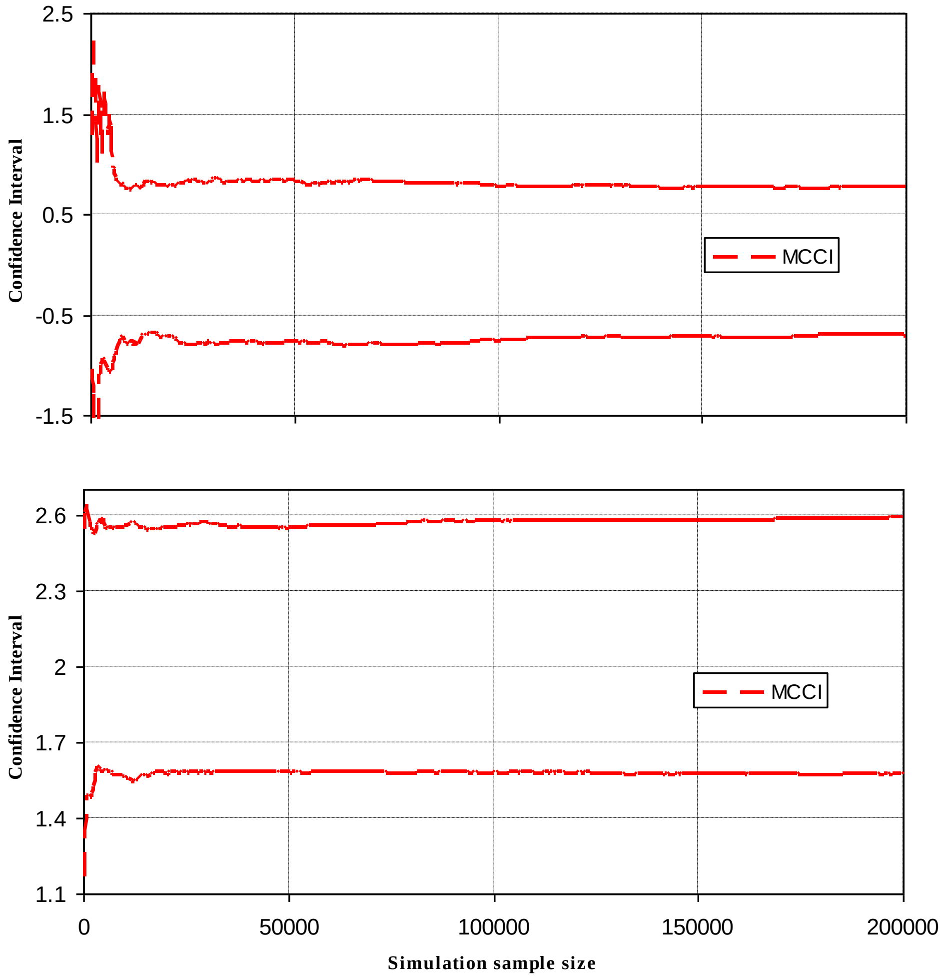

Figure 11. 0.95 confidence intervals for a normal distribution estimated for varying simulation sample size: (upper) confidence interval for the location parameter μ from a sample with n = 10, $\hat{\mu}$ = 0.026 and $\hat{\sigma}$ = 1.023; (lower) confidence interval for the quantity $\mu + 2\sigma$ from a sample with n = 50, $\hat{\mu}$ = −0.027 and $\hat{\sigma}$ = 0.998.

7 Conclusions

By modifying two Monte Carlo methods used by Ripley (1987), associated with the computation of a confidence interval for a parameter of a probability distribution, we derive a new equation and a general algorithm which gives a single solution for a confidence interval, which combines the advantages of these two methods without requiring discrimination for the type of parameter. We show that this algorithm is exact for a single parameter of distribution of either location or scale family. It is also asymptotically equivalent to a Wald-type interval for parameters of regular continuous distributions.

After appropriate modification of the algorithm we make it appropriate for calculating confidence intervals for a parameter of multi-parameter distributions. We show that this algorithm is asymptotically equivalent to a Wald-type interval for regular distributions.

We tested the algorithm in seven cases, namely the construction of a confidence interval for the scale parameter of the exponential distribution, the location parameter and the *p*th percentile of the normal distribution, the scale and shape parameter of a gamma distribution, and the scale parameter and the *p*th percentile of the Weibull distribution. We found that in general this algorithm works well and results in correct coverage probabilities.

We propose the use of the algorithm for an approximation of a confidence interval of any parameter for any continuous distribution because it is easily applicable in every case and gives better approximations than other known algorithms as shown in specific cases above. An additional advantage compared to Ripley's two methods is that it is not needed to select one of the methods. Our algorithm worked equally well or better from the best of Ripley's methods in all the examined cases. Thanks to its generality, the algorithm has been implemented in a hydrometeorological software package (Hydrognomon, 2009-2012), which fits various distributions in data records and calculates point and interval estimates for parameters and distribution quantiles, which are then used for hydrological design.

The confidence intervals obtained by the algorithm are approximate and the algorithm was not developed with the intention to replace the exact confidence intervals, when their calculation is possible. Further research is needed to evaluate the influence of the choice of the numerical parameters (increments $\delta\boldsymbol{\theta}_i$ and the simulation sample size) to the results of the algorithm. A disadvantage of the algorithm is that a lot of repetitions are needed to converge.

Acknowledgements The authors are grateful to two anonymous reviewers, for their positive comments which were very helpful for substantial improvement of the manuscript.

Appendix A: Some theoretical results

First we show that the confidence interval for the parameter μ of a normal distribution $N(\mu,\sigma^2)$ is asymptotically equivalent to a Wald-type interval. For the normal distribution we define $\boldsymbol{\theta} = (\mu, \sigma)$, $\boldsymbol{T}(\underline{\boldsymbol{x}}) = (T_1(\underline{\boldsymbol{x}}), T_2(\underline{\boldsymbol{x}}))$, where $T_1(\underline{\boldsymbol{x}}) = \underline{\mu}$, and $T_2(\underline{\boldsymbol{x}}) = \underline{\sigma}$ are the MLE of μ and σ respectively.

Then, following the notation of the preceding sections we have $\beta := h(\mu,\sigma) = \mu$, $h(\boldsymbol{T}) = T_1$ and $P(b(\underline{\boldsymbol{x}}) < \lambda(\boldsymbol{\theta})) = \alpha/2$, $P(b(\underline{\boldsymbol{x}}) > \upsilon(\boldsymbol{\theta})) = \alpha/2$ which imply that $\lambda = \mu + \Phi^{-1}(\alpha/2)\sigma/\sqrt{n}$ and $\upsilon = \mu + \Phi^{-1}(1-\alpha/2)\sigma/\sqrt{n}$. Now from (18) we obtain

$$\frac{d\boldsymbol{\gamma}}{d\boldsymbol{\theta}} = \begin{bmatrix} \dfrac{d\lambda}{d\boldsymbol{\theta}} \\ \dfrac{d\beta}{d\boldsymbol{\theta}} \\ \dfrac{d\upsilon}{d\boldsymbol{\theta}} \end{bmatrix} = \begin{bmatrix} 1 & \Phi^{-1}(\alpha/2)/\sqrt{n} \\ 1 & 0 \\ 1 & \Phi^{-1}(1-\alpha/2)/\sqrt{n} \end{bmatrix} \tag{44}$$

It is also easy to prove that asymptotically $\begin{bmatrix} \underline{\mu}-\mu \\ \underline{\sigma}-\sigma \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \frac{\sigma^2}{n}\begin{bmatrix} 1 & 0 \\ 0 & 1/2 \end{bmatrix}\right)$, thus $\underline{\mu} \sim N(\mu, \sigma^2/n)$ and $\underline{\sigma} \sim N(\sigma, \sigma^2/2n)$. We also have that $\Phi^{-1}(1-\alpha/2) = -\Phi^{-1}(\alpha/2)$.

From (12), (13) we derive $l = \underline{\mu} - \dfrac{\underline{\sigma}\Phi^{-1}(1-\alpha/2)}{\sqrt{n}\,\dfrac{d\upsilon}{d\mu}}$ and $u = \underline{\mu} + \dfrac{\underline{\sigma}\Phi^{-1}(1-\alpha/2)}{\sqrt{n}\,\dfrac{d\lambda}{d\mu}}$. From (22) we have that $\dfrac{d\lambda}{d\mu} = \dfrac{d\upsilon}{d\mu} = 1 - \dfrac{(\Phi^{-1}(1-\alpha/2))^2}{4n}$. A $1-\alpha$ confidence interval for μ is $(\underline{\mu} - t_{n-1}(1-\alpha/2)$

$\frac{\sigma}{\sqrt{n}}, \underline{\mu} + t_{n-1}(1 - \alpha/2)\frac{\sigma}{\sqrt{n}})$ (e.g. Papoulis & Pillai, 2002, p. 309). Now we have that $\lim\limits_{n\to\infty} \frac{\Phi^{-1}(1-\alpha/2)/(\frac{d\upsilon}{d\mu})}{\mathrm{t}_{n-1}(1-\alpha/2)} = 1$, which proves that the confidence interval obtained by (14) is asymptotically exact.

We will also show that the confidence interval obtained by our method is asymptotically equivalent to a Wald-type interval for two-parameter regular distributions. According to Casella & Berger (2002, p. 472) $\sqrt{n}(\underline{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{\mathrm{d}} N(\mathbf{0},\boldsymbol{I}^{-1})$, where $\underline{\boldsymbol{\theta}}$ is the MLE of $\boldsymbol{\theta}$, and $\boldsymbol{I}$ is the Fisher Information Matrix with elements $\boldsymbol{I}_{jk} = \mathrm{E}(-\frac{\partial^2 \ln f(x|\boldsymbol{\theta})}{\partial\theta_j\theta_k})$. This means that $\sqrt{n}(\underline{\theta}_1 - \theta_1) \xrightarrow{\mathrm{d}} N(0,\boldsymbol{I}^{-1}_{11})$ and $\sqrt{n}(\underline{\theta}_2 - \theta_2) \xrightarrow{\mathrm{d}} N(0,\boldsymbol{I}^{-1}_{22})$. We conclude that $\sqrt{n}(\underline{\beta} - \beta) \xrightarrow{\mathrm{d}} \mathrm{N}(0,\sigma^2_\beta)$, where σ^2_β depends only on θ_1 and θ_2. Suppose that we seek a $1 - \alpha$ confidence interval for β. Then it is easy to show that asymptotically $\lambda(\beta) = \beta - \Phi^{-1}(1-\alpha/2)\sigma_\beta/\sqrt{n}$, $\upsilon(\beta) = \beta + \Phi^{-1}(1-\alpha/2)\sigma_\beta/\sqrt{n}$. Now we have $\mathrm{Var}(\underline{\theta}_1) = \boldsymbol{I}^{-1}_{11}/n$, $\mathrm{Var}(\underline{\theta}_2) = \boldsymbol{I}^{-1}_{22}/n$ and

$$\frac{d\boldsymbol{\gamma}}{d\boldsymbol{\theta}} = \begin{bmatrix} \frac{d\lambda}{d\boldsymbol{\theta}} \\ \frac{d\beta}{d\boldsymbol{\theta}} \\ \frac{d\upsilon}{d\boldsymbol{\theta}} \end{bmatrix} = \begin{bmatrix} \frac{\partial\beta}{\partial\theta_1} - \Phi^{-1}(1-\alpha/2)\frac{\partial\sigma_\beta}{\partial\theta_1}/\sqrt{n} & \frac{\partial\beta}{\partial\theta_2} - \Phi^{-1}(1-\alpha/2)\frac{\partial\sigma_\beta}{\partial\theta_2}/\sqrt{n} \\ \frac{\partial\beta}{\partial\theta_1} & \frac{\partial\beta}{\partial\theta_2} \\ \frac{\partial\beta}{\partial\theta_1} + \Phi^{-1}(1-\alpha/2)\frac{\partial\sigma_\beta}{\partial\theta_1}/\sqrt{n} & \frac{\partial\beta}{\partial\theta_2} + \Phi^{-1}(1-\alpha/2)\frac{\partial\sigma_\beta}{\partial\theta_2}/\sqrt{n} \end{bmatrix} \tag{45}$$

$$\frac{q_{31} + q_{32}}{q_{21} + q_{22}} = \frac{[\Phi^{-1}(1-\alpha/2)]^2[(\frac{\partial\sigma_\beta}{\partial\theta_1})^2\boldsymbol{I}^{-1}_{11} + (\frac{\partial\sigma_\beta}{\partial\theta_2})^2\boldsymbol{I}^{-1}_{22}] - \Phi^{-1}(1-\alpha/2)\sqrt{n}[\frac{\partial\sigma_\beta\,\partial\beta}{\partial\theta_1\partial\theta_1}\boldsymbol{I}^{-1}_{11} + \frac{\partial\sigma_\beta\,\partial\beta}{\partial\theta_2\partial\theta_2}\boldsymbol{I}^{-1}_{22}]}{\sqrt{n}[\Phi^{-1}(1-\alpha/2)(\frac{\partial\sigma_\beta\,\partial\beta}{\partial\theta_1\partial\theta_1} + \frac{\partial\sigma_\beta\,\partial\beta}{\partial\theta_2\partial\theta_2}) - 2\sqrt{n}((\frac{\partial\beta}{\partial\theta_1})^2\boldsymbol{I}^{-1}_{11} + (\frac{\partial\beta}{\partial\theta_2})^2\boldsymbol{I}^{-1}_{22})]} -$$

$$- \frac{2n((\frac{\partial\beta}{\partial\theta_1})^2\boldsymbol{I}^{-1}_{11} + (\frac{\partial\beta}{\partial\theta_2})^2\boldsymbol{I}^{-1}_{22})}{\sqrt{n}[\Phi^{-1}(1-\alpha/2)(\frac{\partial\sigma_\beta\,\partial\beta}{\partial\theta_1\partial\theta_1} + \frac{\partial\sigma_\beta\,\partial\beta}{\partial\theta_2\partial\theta_2}) - 2\sqrt{n}((\frac{\partial\beta}{\partial\theta_1})^2\boldsymbol{I}^{-1}_{11} + (\frac{\partial\beta}{\partial\theta_2})^2\boldsymbol{I}^{-1}_{22})]} \tag{46}$$

It is obvious that $\lim\limits_{n\to\infty}\dfrac{d\upsilon}{d\beta} = \lim\limits_{n\to\infty}\dfrac{q_{31} + q_{32}}{q_{21} + q_{22}} = 1$. In a similar way we can find that $\lim\limits_{n\to\infty}\dfrac{d\lambda}{d\beta} = \lim\limits_{n\to\infty}\dfrac{q_{12} + q_{13}}{q_{22} + q_{23}} = 1$. Now substituting to (12), (13) we obtain $l = \underline{\beta} - \Phi^{-1}(1 - \alpha/2)\sigma_\beta/\sqrt{n}$, $u = \underline{\beta} + \Phi^{-1}(1 - \alpha/2)\sigma_\beta/\sqrt{n}$ which is an asymptotically equivalent to a Wald-type interval according to Casella & Berger (2002, p.497).

Repeating the same procedure for three-parameter distributions, we obtain the same results.

Appendix B: Application of the algorithm on a historical river flows dataset

In this Appendix we apply the algorithm on a historical river flow data set using the hydrological statistical software Hydrognomon (2009-2012), suitable for the processing and the analysis of hydrological time series, which has already incorporated the proposed method. The case study is performed on an important basin in Greece, which is currently part of the water supply system of Athens and has a history, as regards hydraulic infrastructure and management, that goes back to at least 3500 years ago. Modelling attempts with good performance have already been done on the hydrosystem (Rozos et al. 2004). A long-term dataset of the catchment runoff, extending from 1906 to 2008, is available. The example presented in Figure 12 is for the January monthly flow record at the Boeoticos Kephisos river outlet at the Karditsa station measured in hm^3. The gamma distribution is often used to model monthly river flows. Confidence limits of quantiles of distributions are of interest to hydrologists. Here we derived confidence intervals for the scale and the shape parameters of the gamma distribution. Comparison of the results of the different methods used show that the MCCI and "Ripley scale" limits are close to the Bayesian ones. In addition, Figure 13 gives confidence limits of the distribution percentiles using the same dataset, this time constructed using Hydrognomon (2009-2012).

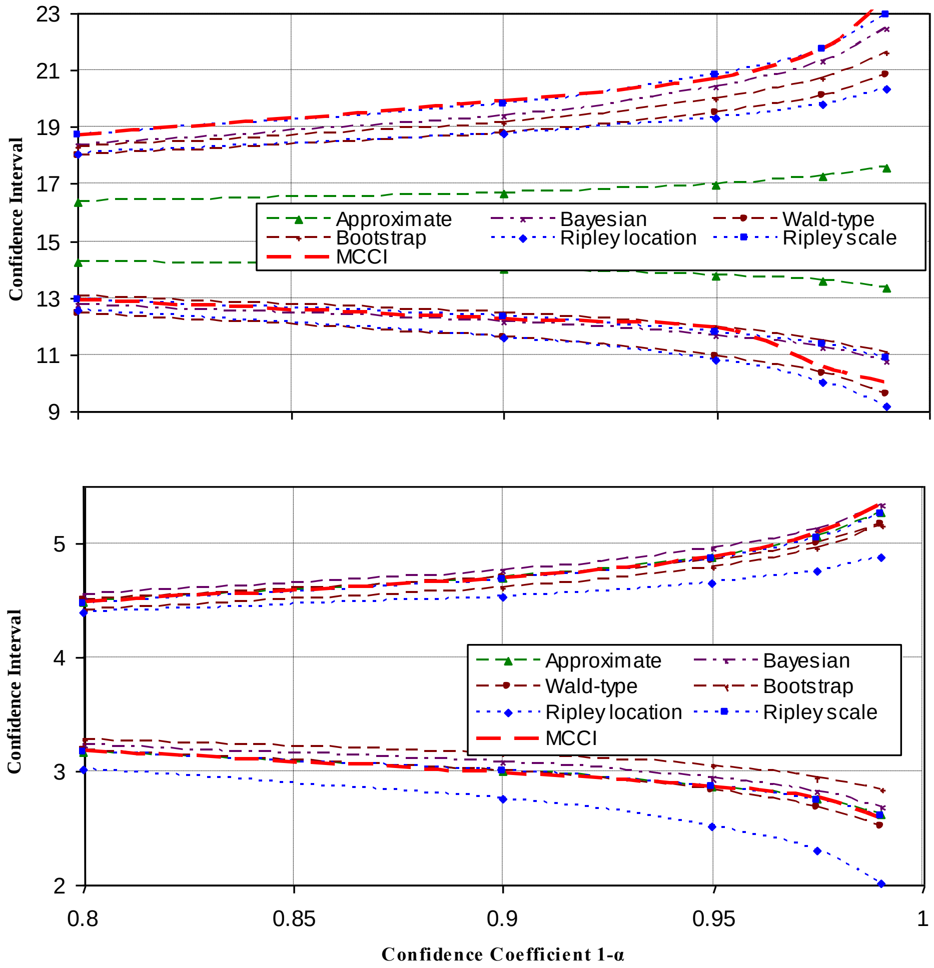

Figure 12. Confidence intervals for the scale (upper) and shape (lower) parameter of a gamma distribution, used to model the Boeoticos Kephisos river January monthly flows with n = 102, $\hat{k}$ = 3.842 and $\hat{\theta}$ = 15.218. Here the number of samples m = 120 000 for MCCI and m = 60 000 for the "Ripley location" and "Ripley scale" cases, δk = 0.3 and $\delta\theta$ = 0.3.

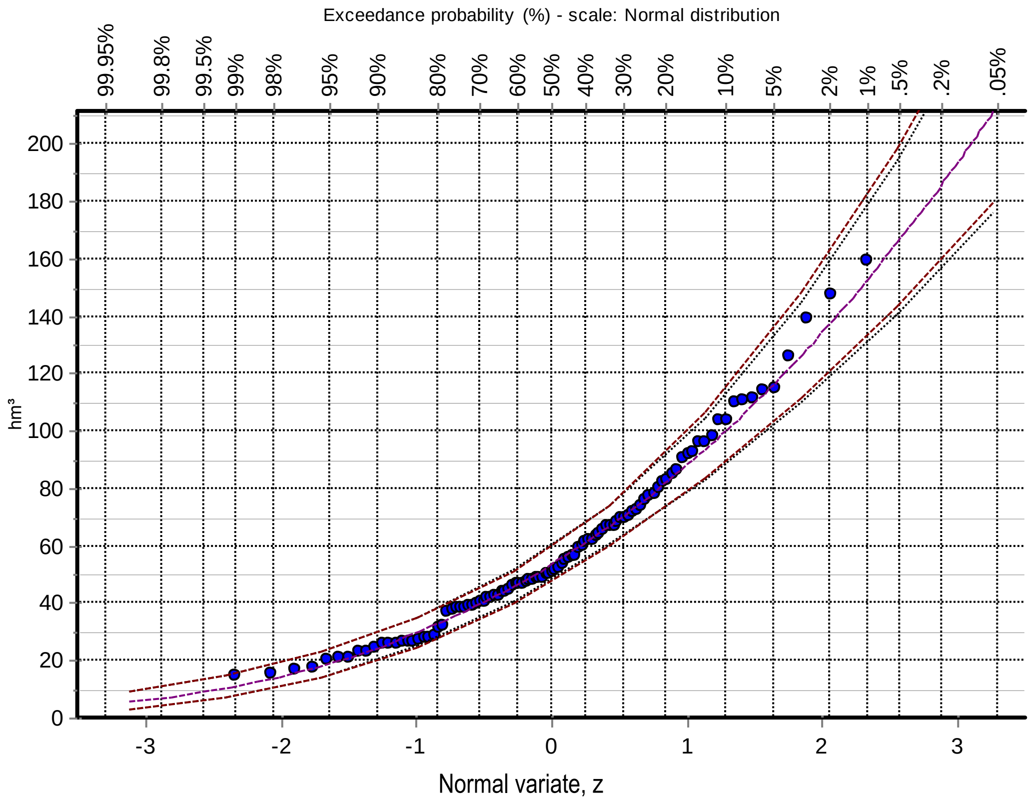

Figure 13. A graph (normal probability plot) produced by the Hydrognomon software referring to the monthly flow of Boeoticos Kephisos river for the month of January (1993-2006). The sample (dots plotted using Weibul plotting positions) was modelled by a gamma distribution (central line) with $\hat{k}$ = 3.842 and $\hat{\theta}$ = 15.218. Dotted lines represent 95% prediction intervals for these parameter values (denoted as λ and υ in the text) and dashed lines represent 95% confidence intervals (MCCI denoted as *l* and *u* in the text) for the distribution percentiles.

References

Bain L, Engelhardt M (1975) A Two-Moment Chi-Square Approximation for the Statistic Log($\bar{x}/\tilde{x}$). Journal of the American Statistical Association 70(352):948-950

Bhaumik D, Kapur K, Gibbons R (2009) Testing Parameters of a Gamma Distribution for Small Samples. Technometrics 51(3):326-334

Casella G, Berger R (2002) Statistical Inference. Duxbury press, Belmont

Choi SC, Wette R (1969) Maximum Likelihood Estimation of the Parameters of the Gamma Distribution and Their Bias. Technometrics 11(4):683-690

Cohen C (1965) Maximum Likelihood Estimation in the Weibull Distribution Based on

Complete and on Censored Samples. Technometrics 7(4):579-588

DiCiccio T (1984) On Parameter Transformations And Interval Estimation. Biometrica 71(3):477-485

DiCiccio T, Efron B (1996) On Bootstrap Procedures For Second-Order Accurate Confidence Limits In Parametric Models. Statistica Sinica 5:141-160

DiCiccio T, Romano J (1995) Bootstrap Confidence Intervals. Statistical Science 11(3):189-228

Efron B (1987) Better Bootstrap Confidence Intervals. Journal of the American Statistical Association 82(397):171-185

Efron B, Hinkley D (1978) Assessing the Accuracy of the Maximum Likelihood Estimator: Observed Versus Expected Fisher Information. Biometrica 65(3):457-482

Efron B, Tibshirani R (1993) An introduction to the bootstrap. Chapman & Hall, New Yourk.

Engelhardt M, Bain L (1978) Construction of Optimal Inference Procedures for the Parameters of the Gamma Distribution. Technometrics 20(4):485-489

Engelhardt M, Bain L (1977) Uniformly Most Powerful Unbiased Tests on the Scale Parameter of a Gamma Distribution with a Nuisance Shape Parameter. Technometrics 19(1):77-81

Garthwaite P, Buckland S (1992) Generating Monte Carlo Confidence Intervals by the Robbins-Monro Process. Journal of the Royal Statistical Society, Series C (Applied Statistics) 41(1):159-171

Gelman A, Carlin J, Stern H, Rubin D (2004) Bayesian Data Analysis, second edition. Chapman & Hall, London

Hall P (1988) Theoretical Comparison of Bootstrap Confidence Intervals. The Annals of Statistics 16(3):927-953

Hemelrijk J (1966) Underlining random variables. Statistica Neerlandica 20:1-7. doi:10.1111/j.1467-9574.1966.tb00488.x

Hillier G, Armstrong M (1999) The Density of the Maximum Likelihood Estimator. Econometrica 67(6):1459-1470

Hydrognomon (2009-2012) Hydrological time series processing software, http://hydrognomon.org/

Kisielinska J (2012) The exact bootstrap method shown on the example of the mean and variance estimation. Comput Stat doi:10.1007/s00180-012-0350-0

Kite GW (1988) Frequency and Risk Analyses in Hydrology. Water Recources Publications

Koutsoyiannis D (1997) Statistical Hydrology, Edition 4. National Technical University of Athens, Athens

Koutsoyiannis D, Kozanis S (2005) A simple Monte Carlo methodology to calculate generalized approximate confidence intervals. Research report. Hydrologic Research Center

Koutsoyiannis D, Efstratiadis A, Georgakakos KP (2007) Uncertainty Assessment of Future Hydroclimatic Predictions: A Comparison of Probabilistic and Scenario-Based Approaches. Journal of Hydrometeorology 8(3):261-281

Lawless J (2003) Statistical Models and Methods for Lifetime Data. John Wiley & Sons

Papoulis A, Pillai U (2002) Probability, Random Variables and Stochastic Processes. McGraw-Hill

Ripley B (1987) Stochastic Simulation. John Wiley & Sons

Robert C (2007) The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation. Springer, New York

Román-Montoya Y, Rueda M, Arcos M (2008) Confidence intervals for quantile estimation using Jackknife techniques. Comput Stat 23:573-585

Rozos E, Efstratiadis A, Nalbantis I, Koutsoyiannis D (2004) Calibration of a semi-distributed model for conjuctinve simulation of surface and groundwater flows. Hydrological Sciences Journal 49 (5):819-842

Shao J (2003) Mathematical Statistics, Second edition, Springer

Son YS, Oh M (2006) Bayesian Estimation of the Two-Parameter Gamma Distribution. Communications in Statistics: Simulation and Computation 35(2):285-293

Wilks SS (1938) Shortest Average Confidence Intervals from Large Samples. The Annals of Mathematical Statistics 9(3):166-175

Yang Z, Xie M, Wong A (2007) A unified confidence interval for reliability-related quantities of two-parameter Weibull distribution. Journal of Statistical Computation and Simulation 77(5):365-378