AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Confidence distributions from likelihoods by median bias correction

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1691287> since 2019-04-16T17:10:12Z

Published version:

DOI:10.1016/j.jspi.2017.09.010

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Confidence distributions from likelihoods by median bias correction

Pierpaolo De Blasi*

University of Torino and Collegio Carlo Alberto

Tore Schweder

Department of Economics, University of Oslo, Norway

September 18, 2018

Abstract

By the modified directed likelihood, higher order accurate confidence limits for a scalar parameter are obtained from the likelihood. They are conveniently described in terms of a confidence distribution, that is a sample dependent distribution function on the parameter space. In this paper we explore a different route to accurate confidence limits via tail-symmetric confidence curves, that is curves that describe equal tailed intervals at any level. Instead of modifying the directed likelihood, we consider inversion of the log-likelihood ratio when evaluated at the median of the maximum likelihood estimator. This is shown to provide equal tailed intervals, and thus an exact confidence distribution, to the third-order of approximation in regular one-dimensional models. Median bias correction also provides an alternative approximation to the modified directed likelihood which holds up to the second order in exponential families.

Keywords: Asymptotic expansion; Confidence curve; Confidence distribution; Exponential family; Modified directed likelihood; Normal transformation family.

1 Introduction

The level of reported confidence intervals are most often 95%, with equal probability of missing the target at both sides. Sometimes other levels are used, but rarely are several intervals at their different levels reported in applied work. Instead of only reporting one confidence interval we suggest to report a family of nested confidence intervals for parameters of primary interest. The family is indexed by the confidence level $1 - \alpha$ for $\alpha \in (0, 1)$ and is conveniently represented by what is called a *confidence curve*, a quantity introduced by Birnbaum (1961) to give a complete picture of the estimation uncertainty. As an example, take $\hat{\theta} \sim N(\theta, \sigma^2)$ for σ known. It yields the curve $cc(\theta) = |1 - 2\Phi((\theta - \hat{\theta})/\sigma)|$ for $\Phi(z)$ the cumulative distribution function of a $N(0, 1)$. This is a confidence curve since, for all $\alpha \in (0, 1)$, $\{\theta : cc(\theta) \leq 1 - \alpha\} =$

*Supported by the European Research Council (ERC) through St G "N-BNP" 306406.

$(\hat{\theta} + \sigma\Phi^{-1}(\alpha/2), \hat{\theta} + \sigma\Phi^{-1}(1 - \alpha/2))$ is the respective confidence interval of level $1 - \alpha$. In the example the confidence curve has its minimum at $\hat{\theta}$ which is a point estimate of θ . The normal confidence curve is tail-symmetric, i.e. the probability of missing the parameter to the left equals that to the right and is $\alpha/2$ at level $1 - \alpha$. A tail-symmetric confidence curve represents uniquely a confidence distribution, that is confidence curves that describe upper confidence limits. Confidence distribution is a term coined by Cox (1958) and formally defined in Schweder and Hjort (2002). For scalar parameters the fiducial distributions developed by Fisher (1930) are confidence distributions. Neyman (1934) saw that the fiducial distribution leads to confidence intervals. Cox (2013) sees confidence distributions as “simple and interpretable summaries of what can reasonably be learned from the data (and an assumed model)”. Confidence distributions are reviewed by Xie and Singh (2013), and more broadly and with more emphasis on confidence curves by Schweder and Hjort (2016). In location models and other simple models the confidence distribution is obtained from pivots, e.g. the normal pivot $(\hat{\theta} - \theta)/\sigma$ in the above example. A canonical pivot is $C(\theta) = C(\theta; \hat{\theta}) = 1 - G(\hat{\theta}, \theta)$ where $G(y, \theta) = P(\hat{\theta} \leq y; \theta)$ is the distribution function of the maximum likelihood estimator $\hat{\theta}$, assumed to be absolutely continuous with respect to the Lebesgue measure and non-increasing in θ . See Section 2 for precise definitions and notation. The confidence distribution $C(\theta)$ is a canonical pivot in the sense of being uniformly distributed on the unit interval when $\hat{\theta}$ is distributed according to θ . When $\hat{\theta}$ is a sufficient statistic with monotone likelihood ratio, $C(\theta)$ is also optimal in the Neyman-Pearson sense, that is it describes smaller confidence intervals at a given level when compared to any other confidence distribution for the parameter θ (Schweder and Hjort, 2016, Section 5.4). An equal tailed confidence curve is readily obtained from $C(\theta)$ by $cc(\theta) = |1 - 2C(\theta)|$.

In this paper we shall be concerned with confidence curves obtained from the log-likelihood ratio $w(\theta)$, and we shall study the properties of median bias correction. Median bias correction of a confidence curve, proposed by Schweder (2007), is a method to make the resulting confidence curve approximately tail-symmetric. In the normal example $w(\theta) = ((\hat{\theta} - \theta)/\sigma)^2$ and the confidence curve mentioned above is also given by $cc(\theta) = Q(w(\theta))$ where Q is the cumulative chi-square distribution function with one degree of freedom. This confidence curve is tail-symmetric, as mentioned, and the confidence interval of level 0 is the single point $\hat{\theta}$ which thus has median θ and is said to be median unbiased. In general $w(\theta)$ hits zero at the maximum likelihood estimator, which might not be median unbiased. Let $\hat{\theta}$ have median $b(\theta)$. The median bias corrected confidence curve is the confidence curve of the parameter $b(\theta)$. The idea is to probability transform the bias corrected log likelihood ratio $w^*(\theta) = w(b(\theta))$ rather than $w(\theta)$. With $F^*(y; \theta)$ denoting the sampling distribution of $w^*(\theta)$ when the data is distributed according to θ , the bias corrected confidence curve is $cc^*(\theta) = F^*(w^*(\theta); \theta)$. Since $\arg \min(cc^*(\theta)) = b^{-1}(\hat{\theta})$ is median unbiased, the level set at $\alpha = 0$ is typically the single point $b^{-1}(\hat{\theta})$ and, by continuity, $cc^*(\theta)$ is close to be equal tailed at low levels. We undertake a theoretical study of the asymptotic properties of $cc^*(\theta)$ by showing that $cc^*(\theta)$ is third-order tail-symmetric for $(\theta, \hat{\theta})$ in the normal deviation range in two important classes of parametric models with parameter dimension one. First, we consider parametric models that belong to the Efron’s normal transformation family (Efron, 1982). Then, we extend the result to regular one dimensional exponen-

tial families, where we also discuss the relation between median bias corrected and modified directed likelihood of Barndorff-Nielsen (1986), thus providing an alternative approximation to the latter. Since median bias correction works so well in these cases, it is reasonable to expect the method to work well quite generally. However, when a canonical confidence distribution is available, as in the exponential family models, we do of course not advocate to use median bias correction rather than using the canonical confidence distribution.

The rest of the paper is organized as follows. In Section 2, we recast confidence estimation in terms of confidence curves and introduce the notation we use in the sequel. We also define the confidence curve based on inverting the median bias corrected version of the log-likelihood ratio. In Section 3 and 4, we investigate its asymptotic properties in terms of tail symmetry in the Efron's normal transformation family and in one dimensional exponential families, respectively. Finally, in Section 5 some concluding remarks and lines of future research are presented, together with an example that provides a preliminary illustration of the use of median bias correction in the presence of nuisance parameters. Some proofs and a technical lemma are deferred to the Appendix.

2 Likelihood-based confidence curves

Let $X = (X_1, \dots, X_n)$ be a continuous random sample with density $f(x; \theta)$ depending on a real parameter $\theta \in \Theta \subset \mathbb{R}$ and let $P(\cdot; \theta)$ indicate probabilities calculated under $f(x; \theta)$. The log-likelihood is $\ell(\theta) = \ell(\theta; x) = \log f(x; \theta)$, and the log-likelihood ratio is $w(\theta) = w(\theta; x) = 2(\ell(\hat{\theta}; x) - \ell(\theta; x))$, where $\hat{\theta}$ is the maximum likelihood estimate. We drop the second argument in sample-dependent functions like w and ℓ whenever it is clear from the context whether we refer to a random quantity or to its observed value. Unless otherwise specified, all asymptotic approximations are for $n \rightarrow \infty$ and stochastic term $O_p(\cdot)$ refers to convergence in probability with respect to $f(x; \theta)$. We assume that the model is sufficiently regular for the validity of first order asymptotic theory, cfr. Barndorff-Nielsen and Cox (1994, Chapter 3). In particular, $w(\theta)$ converges in distribution to a chi-squared random variable, hence, by contouring $w(\theta)$ with respect to this distribution we obtain intervals of θ values given by the level sets for the curve $Q(w(\theta))$ where Q is the distribution function of the chi-squared distribution with 1 degree of freedom. This curve depends on the sample x and has its minimum at $\hat{\theta}$. However its level sets are not in general exact confidence intervals since the chi-squared approximation for the distribution of $w(\theta)$ is valid only for n large and the coverage probabilities equal the nominal levels only in the limit. As a consequence, $Q(w(\theta))$ is not uniformly distributed on the unit interval under $P(\cdot; \theta)$, a property we require for a *regular* confidence curve as spelled in the following definition.

DEFINITION 1 *A function $cc : \Theta \times \mathbb{R}^n \rightarrow [0, 1)$ is a regular confidence curve when $\min_{\theta} cc(\theta; x) = 0$, the level sets $\{\theta : cc(\theta; x) \leq 1 - \alpha\}$ are finite intervals for all $\alpha \in (0, 1)$, and $cc(\theta; X) \sim \text{Unif}(0, 1)$ under $P(\cdot; \theta)$.*

Confidence curves might be defined for parameters of higher dimension and also for irregular curves that even might have more than one local minimum or might have

infinite level sets for $\alpha < 1$, see Schweder and Hjort (2016, Section 4.6). Note that, under Definition 1, $I = \{\theta : cc(\theta; x) \leq 1 - \alpha\}$ is an exact confidence region of level $1 - \alpha$ since $P(I \ni \theta; \theta) = P(cc(\theta; X) \leq 1 - \alpha; \theta) = 1 - \alpha$. Among confidence curves, of special importance are confidence distributions, which are confidence curves that describe upper confidence limits. The definition is as follows.

DEFINITION 2 *A function $C : \Theta \times \mathbb{R}^n \rightarrow [0, 1]$ is a confidence distribution when $C(\cdot; x)$ is a cumulative distribution function in θ for all x and $C(\theta; X) \sim \text{Unif}(0, 1)$ under $P(\cdot; \theta)$.*

Keep in mind that the realized confidence curve and confidence distribution depend on the data, and prior to observation they are random variables (with distribution depending on the parameter value from which the data are generated). To keep the notation simple, we drop the second argument x in $cc(\theta; x)$ and $C(\theta; x)$. Moreover, we will confine ourselves to regular confidence curves $cc(\theta)$ with only one local minimum. In this setting $cc(\theta)$ can be transformed into a distribution via

$$H(\theta) = \frac{1}{2}\{1 - \text{sign}(\tilde{\theta} - \theta)cc(\theta)\}, \quad \tilde{\theta} = \arg \min_{\theta} cc(\theta) \quad (1)$$

so that the left and right endpoints of the interval $I = \{\theta : cc(\theta) \leq 1 - \alpha\} = (\underline{\theta}, \bar{\theta})$, are given by

$$\underline{\theta} = H^{-1}\left(\frac{\alpha}{2}\right), \quad \bar{\theta} = H^{-1}\left(1 - \frac{\alpha}{2}\right), \quad (2)$$

respectively. We refer to $\tilde{\theta} = H^{-1}(\frac{1}{2})$ as the median confidence estimator for θ . By construction, $\underline{\theta}$ and $\bar{\theta}$ satisfy $P(\underline{\theta} > \theta; \theta) + P(\bar{\theta} < \theta; \theta) = \alpha$. We then say that $cc(\theta)$ is tail-symmetric when the interval $(\underline{\theta}, \bar{\theta})$ is equal tailed, that is

$$P(\underline{\theta} > \theta; \theta) = P(\bar{\theta} < \theta; \theta) = \alpha/2, \quad \forall \alpha \in (0, 1).$$

This is equivalent to $H(\theta)$ defining a confidence distribution according to Definition 2.

DEFINITION 3 *A confidence curve $cc(\theta)$ is tail-symmetric if $H(\theta)$ in (1) is a confidence distribution according to Definition 2.*

The relation obviously works in the other direction: given a confidence distribution $C(\theta)$, $cc(\theta) = 1 - 2\min\{C(\theta), 1 - C(\theta)\} = |1 - 2C(\theta)|$ defines a tail-symmetric confidence curve, see Birnbaum (1961). Note that the median confidence estimator $\tilde{\theta}$ of a tail-symmetric $cc(\theta)$ is median-unbiased, i.e. $P(\tilde{\theta} > \theta; \theta) = 0.5$. See Lehmann (1986, Section 5.6). The relation between median-unbiased estimators and equal tailed intervals have been noted by Skovgaard (1989) in connection with the maximum likelihood estimator.

We now focus on confidence distributions derived from the likelihood. It is convenient to set as the exact confidence distribution the one obtained from the sampling distribution of the maximum likelihood estimator, namely

$$C(\theta) = 1 - G(\hat{\theta}; \theta), \quad (3)$$

where we assume that G , the distribution function of $\hat{\theta}$, is continuous and non-increasing in θ . In order to have $C(\theta)$ being a proper cumulative distribution function,

it is also required that $\lim_{\theta \downarrow a} G(\hat{\theta}; \theta) = 0$ and $\lim_{\theta \uparrow b} G(\hat{\theta}; \theta) = 1$, where a and b are the infimum and supremum of the parameter space Θ , respectively. The α -quantile is denoted by $\hat{\theta}(\alpha) = C^{-1}(\alpha)$. In particular, $\hat{\theta}(\frac{1}{2})$ corresponds to the median-unbiased estimator of θ . The exact distribution $G(y; \theta)$ is generally unknown and the asymptotic approximation of the confidence limit $\hat{\theta}(\alpha)$ has been object of an extensive research which goes beyond first order accuracy. See DiCiccio and Efron (1996) for a review. Third-order approximations to $G(y; \theta)$, and thus to $C(\theta)$, can be obtained from the modified directed likelihood of Barndorff-Nielsen (1986), see Section 4.2 for a discussion. We will instead look for a route to such good approximations by transforming the scale at which the log-likelihood ratio is presented. To this aim, let $F(y; \theta)$ be the sampling distribution function of $w(\theta)$ under $P(\cdot; \theta)$, and define

$$cc(\theta) = F(w(\theta); \theta). \quad (4)$$

According to (1)–(2), $\underline{\theta}$ and $\bar{\theta}$ are the endpoints of a confidence interval of level $1 - \alpha$. It is clear that, in general, $cc(\theta)$ is not tail-symmetric according to Definition 3, in particular when $\hat{\theta}$ is not median unbiased. More generally, the distribution estimator

$$H(\theta) = \frac{1}{2} \{1 - \text{sign}(\hat{\theta} - \theta) F(w(\theta); \theta)\} \quad (5)$$

is not uniformly distributed on the unit interval under $P(\cdot; \theta)$. According to first-order asymptotics, $cc(\theta)$ is tail-symmetric up to the first order of approximation, that is

$$H(\theta) = C(\theta) + O_p(n^{-1/2}). \quad (6)$$

Consequently $\underline{\theta} = \hat{\theta}(\frac{\alpha}{2}) + O_p(n^{-1})$ and $\bar{\theta} = \hat{\theta}(1 - \frac{\alpha}{2}) + O_p(n^{-1})$, and $P(\underline{\theta} > \theta; \theta) = P(\bar{\theta} < \theta; \theta) + O(n^{-1/2})$.

In order to improve on (6), we consider the median bias correction to $w(\theta)$. Let $b(\theta)$ be the median of $\hat{\theta}$ as function of θ , that is

$$G(b(\theta); \theta) = 0.5.$$

By assumption, $b(\theta)$ is continuously increasing in θ and $b^{-1}(\hat{\theta}) = \hat{\theta}(\frac{1}{2})$, as a simple calculation reveals. The *median bias corrected log-likelihood ratio* is defined as

$$w^*(\theta) = w^*(\theta; x) = w(b(\theta); x), \quad (7)$$

By construction, $w^*(\theta)$ attains its minimum at $\hat{\theta}(\frac{1}{2})$, the median unbiased estimator of θ . Since both the likelihood function $\ell(\theta)$ and the median function $b(\theta)$ are invariant to monotone parameter transformations, invariance is preserved for $w^*(\theta)$. See Firth (1993) for a different type of likelihood correction, aimed at reducing the bias of the maximum likelihood estimator. The median bias corrected confidence curve is defined as

$$cc^*(\theta) = F^*(w^*(\theta); \theta), \quad (8)$$

where $F^*(y; \theta)$ stands for the sampling distribution of $w^*(\theta)$ under $P(\cdot; \theta)$. According to (1), it yields the distribution estimator

$$H^*(\theta) = \frac{1}{2} \left\{ 1 - \text{sign} \left(\hat{\theta}(\frac{1}{2}) - \theta \right) F^*(w^*(\theta); \theta) \right\}. \quad (9)$$

For illustration, we consider confidence distributions for the variance parameter θ in the normal model $N(0, \theta)$. For $\hat{\theta} = n^{-1} \sum_{i=1}^n x_i^2$, the log-likelihood ratio is $w(\theta) = n(\hat{\theta}/\theta - \log(\hat{\theta}/\theta) - 1)$. Based on $n\hat{\theta}/\theta \sim \chi_n^2$, one finds $G(y; \theta) = P(\chi_n^2 \leq ny/\theta)$ and $b(\theta) = \chi_{n,.5}^2 \theta/n$ (in obvious notation). Using the pivotal distribution of $n\hat{\theta}/\theta$, $F(y; \theta)$ and $F^*(y; \theta)$ can be computed via Monte Carlo. Based on a simulated sample of size $n = 10$ with $\theta = 4$, the left panel of Figure 1 displays $cc^*(\theta)$ according to (8) while the right panel reports $C(\theta)$, $H(\theta)$ and $H^*(\theta)$, according to (3), (5) and (9), respectively. Note that $H^*(\theta)$ and $C(\theta)$ are on top of each other and are almost indistinguishable.

Figure 1: *Confidence distributions for θ in the normal model $N(0, \theta)$ for $n = 10$ observations generated according to $\theta = 4$. Left panel: $cc^*(\theta)$ together with some of its confidence intervals. Right panel: $H(\theta)$ (dashed line), $H^*(\theta)$ (solid line), and nearly on top $C(\theta)$ (dotted line). $cc^*(\theta)$, $H(\theta)$ and $H^*(\theta)$ are based on 50000 Monte Carlo simulations.*

Hence, the median correction in (8), by making the median confidence estimator of $cc^*(\theta)$ coincide with $\hat{\theta}(\frac{1}{2})$, shifts the whole curve $H^*(\theta)$ towards $C(\theta)$, thus inducing nearly exact tail symmetry. We return to this example in Section 4 where we give a theoretical justification to the fact $H^*(\theta)$ and $C(\theta)$ coincide to the third order of approximation in Theorem 2.

We conclude this section by noting that, while w^* can be interpreted as the log-likelihood ratio for the parameter $\psi = b^{-1}(\theta)$, that is $w^{(\psi)}(\psi) = w(b(\psi)) = w^*(\psi)$, $cc^*(\psi)$ does not correspond to the confidence curve in the ψ -parametrization, that is $cc^{(\psi)}(\psi) := F(w^{(\psi)}(\psi; x); \psi)$ where $F(y; \psi)$ stands now for the sampling distribution of the log-likelihood ratio in terms of the ψ parameter. As an example, consider the exponential model, $f(x; \theta) = \theta e^{-\theta x}$. By standard calculation one finds that $w(\theta) = 2n(\theta/\hat{\theta} - 1 - \log(\theta/\hat{\theta}))$, $2n\theta/\hat{\theta} \sim \chi_{2n}^2$ and $b(\theta) = 2n\theta/\chi_{2n,.5}^2$. Hence, for $\psi = \chi_{2n,.5}^2 \theta/2n$, we get $cc^*(\psi) = P(2n h(X/\chi_{2n,.5}^2) \leq w(b(\psi)))$ while $cc^{(\psi)}(\psi) = P(2n h(X/2n) \leq w(b(\psi)))$, where $X \sim \chi_{2n}^2$ and $h(x) = x - 1 - \log(x)$. On the other hand, $cc^*(\theta)$ shares with $cc(\theta)$ the property of invariance with respect to monotone transformation of the parameter: if $\psi = g(\theta)$ for g invertible, then, it is easy to see

that bias corrected confidence curve in the ψ -parametrization, say $cc^{*(\psi)}(\psi)$, corresponds to $cc^*(g^{-1}(\psi))$. This can be easily verified in the exponential model above by taking, e.g., $g(\theta) = 1/\theta$ so that ψ represents the mean parameter. In the sequel, for ease of notation, we avoid superscripts as in $w^{(\psi)}$ and $cc^{(\psi)}$ whenever the parametrization the likelihood is referring to will be clear from the context.

3 Normal transformation family

In this section we establish third-order tail symmetry of the bias corrected confidence curve $cc^*(\theta)$ when $\hat{\theta}$ is a sufficient statistic and belongs to the normal transformation family of Efron (1982). This family of distributions was used by Efron (1987) to introduce bias and acceleration corrected bootstrapped confidence intervals that achieve second order accuracy. The idea is that standard intervals are based on assuming that the normal approximation of $(\hat{\theta} - \theta)/\hat{\sigma}$ is exact, with $\hat{\sigma}$ a fixed constant and, hence, convergence to normality can be improved by considering a monotone transformation of $\hat{\theta}$ and θ which is exactly normalizing and variance stabilizing. Second order accuracy was later extended to regular statistical models such as the exponential family, see DiCiccio and Efron (1992). We follow a similar path here, as we first prove, in Theorem 1, tail symmetry in the normal transformation family as this case provides a simple illustration of the generalized inverse mapping argument reported in Lemma 1 of the Appendix. Theorem 2 of Section 4 addresses tail symmetry in the exponential family, where an additional Cornish–Fisher expansion of the distribution of the maximum likelihood estimator is needed. Theorem 2 is indeed a more general result than Theorem 1 since, by Pitman–Koopman–Darmois Theorem, cfr. Lehmann and Casella (1999, Theorem 6.18), if the data are independent and identically distributed and the dimension of the sufficient statistic does not depend on n , as we are assuming here, then the model is an exponential family.

Let $\hat{\theta}$ be a sufficient estimator for θ , not necessarily maximizing the likelihood, but behaving asymptotically like the maximum likelihood estimator in terms of order of magnitude of its bias, standard deviation, skewness, and kurtosis:

$$\hat{\theta} - \theta \sim (B_\theta/n, C_\theta/n^{1/2}, D_\theta/n^{1/2}, E_\theta/n), \quad (10)$$

where $B_\theta, C_\theta, D_\theta$, and E_θ are functions of θ and n (the latter suppressed in the notation) bounded in n . See equations (5.1)–(5.3) in Efron (1987). Next, suppose there exists a monotone increasing transformation g and constants z_0 (bias constant) and a (acceleration constant) such that $\hat{\phi} = g(\hat{\theta})$ and $\phi = g(\theta)$ satisfy

$$\hat{\phi} = \phi + (1 + a\phi)(Z - z_0), \quad Z \sim N(0, 1), \quad (11)$$

where $\phi > -1/a$ when $a > 0$ and $\phi < -1/a$ when $a < 0$. Model (11) has standard deviation linear in ϕ on the transformed scale. It provides a pivot with accompanying confidence distribution $C^{(\phi)}(\phi)$. The latter is directly transformed back to a confidence distribution for θ , that is $C^{(\theta)}(\theta) = C^{(\phi)}(g(\theta))$. Theorem 1 states that $cc^{*(\phi)}(\phi)$ as well as $cc^{*(\theta)}(\theta)$, are third order tail-symmetric according to Definition 3, an improvement up to $O_p(n^{-3/2})$ in the asymptotic order displayed in (6). The proof relies on the asymptotic inversion of convex functions reported in Lemma 1 in the Appendix.

THEOREM 1 *Let $\hat{\theta}$ be a sufficient estimator of θ based on a sample of size n satisfying (10), and assume there exists a monotone increasing function g such that (11) holds. Then, for $C(\theta)$ and $H^*(\theta)$ defined in (3) and (9), respectively,*

$$H^*(\theta) = C(\theta) + O_p(n^{-3/2}), \quad \text{for } n^{1/2}(\theta - \hat{\theta})/C_\theta = O_p(1). \quad (12)$$

PROOF. Since a confidence curve for $\phi = g(\theta)$ translates into one for θ for the invertible transformation g , it is sufficient to prove (12) in the transformed normal model. Under (10), the normalizing transformation g is locally linear in its argument with a scale factor of order $n^{1/2}$. In particular, from (11), the normal deviation range $n^{1/2}(\theta - \hat{\theta})/C_\theta = O_p(1)$ in (12) corresponds to

$$(\phi - \hat{\phi})/(1 + a\phi) = O_p(1). \quad (13)$$

According to Efron (1987, Theorem 2), $z_0 = a[1 + O(n^{-1})]$ and both z_0 and a are $O(n^{-1/2})$ as long as the $\hat{\theta}$ satisfies (10). We will make repeatedly use of these asymptotic behaviors throughout the proof, even though we suppress the dependence of a and z_0 on n in the notation. The log-likelihood $\ell(\phi; \hat{\phi}) = -[(\hat{\phi} - \phi)/(1 + a\phi) + z_0]^2/2 - \log(1 + a\phi)$ is not maximized at $\phi = \hat{\phi}$, unless $z_0 = a$, rather at

$$\hat{\phi}^c = \hat{\phi} - c(1 + a\hat{\phi}), \quad c = \frac{1}{a} - \frac{(1 - az_0)}{2a^3} \left[\left(1 + \frac{4a^2}{(1 - az_0)^2} \right)^{1/2} - 1 \right] \quad (14)$$

as a simple calculation reveals. One finds that $c = a - z_0 + O(n^{-3/2})$ and, consequently, $c = O(n^{-3/2})$. Actually, $\hat{\phi}^c$ belongs to the normal transformation family (11) since it can be written as $\hat{\phi}^c = \phi + (1 - ac)(1 + a\phi)(Z - z_0^c)$ for $z_0^c = z_0 + c/(1 - ac)$, see Efron (1987, Section 11), with distribution

$$G(y; \phi) = \Phi \left(\frac{y - \phi}{(1 - ac)(1 + a\phi)} + z_0^c \right), \quad (15)$$

and median function $b(\phi) = \phi - z_0^c(1 - ac)(1 + a\phi)$. Note that $b(\phi)$ is increasing in ϕ when $z_0^c a(1 - ac) < 1$, which we assume without loss of generality since it certainly is for large n . Since $\hat{\phi}^c$ is a sufficient statistic, the log-likelihood ratio for ϕ is

$$w(\phi; \hat{\phi}^c) = -(z_0^c)^2 + \left(\frac{\hat{\phi}^c - \phi}{(1 - ac)(1 + a\phi)} + z_0^c \right)^2 - 2 \log \frac{1 + a\hat{\phi}^c}{1 + a\phi}.$$

It is easy to check that $w(\phi; \hat{\phi}^c)$ is convex in both arguments, and so is its bias corrected version $w^*(\phi; \hat{\phi}^c) = w(b(\phi); \hat{\phi}^c)$. Let $H^*(\phi)$ be defined according to (9). We are interested in expressing $H^*(\phi)$ in terms of tail probabilities associated to $\hat{\phi}^c$ for comparison with the confidence distribution $C(\phi) = 1 - G(\hat{\phi}^c; \phi)$. To this aim, let $\hat{\phi}^*$ be implicitly defined in function of $\hat{\phi}^c$ and ϕ by $w^*(\phi; \hat{\phi}^*) = w^*(\phi; \hat{\phi}^c)$. Then, $w^*(\phi; x) \leq w^*(\phi; \hat{\phi}^c)$ for $\hat{\phi}^* \leq x \leq \hat{\phi}^c$ when $\hat{\phi}^c > b(\phi)$, for $\hat{\phi}^c \leq x \leq \hat{\phi}^*$ when $\hat{\phi}^c < b(\phi)$. We only consider the first case, where the equality of interest is $H^*(\phi) = \frac{1}{2} \{1 - cc^*(\phi)\} = \frac{1}{2} \{1 - G(\hat{\phi}^c; \phi) + G(\hat{\phi}^*; \phi)\}$. Hence, for $\hat{\phi}^c > b(\phi)$, the normal deviation range $n^{1/2}(\theta - \hat{\theta})/C_\theta = O_p(1)$ in (12) corresponds to

$$G(\hat{\phi}^*; \phi) = 1 - G(\hat{\phi}^c; \phi) + O_p(n^{-3/2}) \quad (16)$$

for $(\phi, \hat{\phi})$ in (13). As for the right hand side of (16), from (15) it follows that, when $\hat{\phi}^c > b(\phi)$, $1 - G(\hat{\phi}^c; \phi) = G(2b(\phi) - \hat{\phi}^c; \phi)$, so that (16) is implied by

$$\frac{\hat{\phi}^* - \phi}{(1 - ac)(1 + a\phi)} = \frac{(2b(\phi) - \hat{\phi}^c) - \phi}{(1 - ac)(1 + a\phi)} + O_p(n^{-3/2}), \quad (17)$$

for $(\phi, \hat{\phi})$ in (13). In order to establish (17), we derive an asymptotic expansion of $\hat{\phi}^* := \hat{\phi}^*(\hat{\phi}^c, \phi)$ locally at $\hat{\phi}^c = b(\phi)$ by an application of the generalized inverse mapping argument of Lemma 1. Let $f_n(x) = w^*(\phi; b(\phi) + [1 + ab(\phi)]x)$ so that $w^*(\phi; \hat{\phi}^c) = f_n(x_0)$ for $x_0 = [\hat{\phi}^c - b(\phi)]/[1 + ab(\phi)]$. Also, let $g_n(x)$ be implicitly defined by $f_n(x) = f_n(g_n(x))$ so that $g_n(x_0) = [\hat{\phi}^* - b(\phi)]/[1 + ab(\phi)]$. One finds $f_n^{(2)}(0) = 2[(1 - ac)^{-2} + a^2]$ and $f_n^{(k)}(0) = (-1)^k 2(k - 1)!a^k$ for $k \geq 3$, so that the hypotheses of Lemma 1 are satisfied. Hence, $g_n(x_0) = -x_0 + O(n^{-3/2})$ for $x_0 = O(1)$, that is

$$\frac{\hat{\phi}^* - b(\phi)}{1 + ab(\phi)} = -\frac{\hat{\phi}^c - b(\phi)}{1 + ab(\phi)} + O_p(n^{-3/2}),$$

for $[\hat{\phi}^c - b(\phi)]/[1 + ab(\phi)] = O_p(1)$. Since $1 + ab(\phi) = [1 - az_0^c(1 - ac)](1 + a\phi) = (1 - az_0 - c)(1 + a\phi)$ and both $1 - az_0 - c$ and $1 - ac$ are $O(1)$, we get (17) for $[\hat{\phi}^c - b(\phi)]/(1 + a\phi) = O_p(1)$. The latter corresponds to $(\phi, \hat{\phi})$ in the normal deviation range (13) upon substitution for $\hat{\phi}^c$ and $b(\phi)$. The proof is then complete. $\square$

As an illustration of (12), assume the coefficients in (11) are in agreement so that $\hat{\phi}$ is the maximum likelihood estimator, that is $z_0 = a$, see equation (14). Therefore, let $\hat{\phi} = \phi + (1 + a\phi)(Z - a)$, which yields $b(\phi) = \phi - a(1 + a\phi)$. The distribution function of $\hat{\phi}$ is $G(y; \phi) = \Phi[(y - \phi)/(1 + a\phi) + a]$, so that $C(\phi) = 1 - \Phi[(\hat{\phi} - \phi)/(1 + a\phi) + a]$, with median-unbiased estimator $\hat{\phi}(\frac{1}{2}) = (\hat{\phi} + a)/(1 - a^2)$. The log-likelihood ratio and its bias corrected version $w^*(\phi; \hat{\phi})$ are easily derived. The exact distribution F^* of w^* is recovered from the standard normal distribution, $w^*(\phi; \hat{\phi}) \sim -a^2 + [Z/(1 - a^2) + a]^2 - 2\log[1 + aZ/(1 - a^2)]$, and $cc^*(\phi)$ can be calculated by Monte-Carlo. In the left panel of Figure 2 we plot $C(\phi)$ and $H^*(\phi)$ for $\hat{\phi} = 10$ and $a = 0.3$. Even for a non-negligible acceleration a (later we argue that $a = O(n^{-1/2})$, so it roughly corresponds to $n = 10$), the median corrected confidence curve $cc^*(\phi)$ nearly exactly recovers, through $H^*(\phi)$, the confidence distribution $C(\phi)$. The right panel shows that the difference between the two confidence distributions is very small, approximately of order n^{-3} , suggesting that the order of magnitude in (12) might be conservative.

4 Exponential family

4.1 Tail symmetry. In this section we establish third-order tail symmetry for the mean value parameter of regular one-parameter exponential families. Following DiCiccio and Efron (1992, Section 5), let $X \sim f(x; \bar{\eta}) = \exp[\bar{\eta}t(x) - \bar{\psi}(\bar{\eta}) - d(x)]$, where $\bar{\eta}$ is an unknown real parameter. Given a random sample of size n , the log-likelihood for $\bar{\eta}$ based on $y = n^{-1} \sum_{i=1}^n t(x_i)$ has form $\ell(\bar{\eta}; y) = n[\bar{\eta}y - \bar{\psi}(\bar{\eta})]$. Upon defining $\eta = n\bar{\eta}$ and $\psi(\eta) = n\bar{\psi}(\bar{\eta})$, the log-likelihood for η is $\ell(\eta; y) = \eta y - \psi(\eta)$. Since the cumulant

Figure 2: Normal transformation model (11) with $a = z_0 = .3$, $\hat{\phi} = 10$. Left panel: confidence distributions $H^*(\phi)$ (solid line), and nearly on top $C(\phi)$ (dotted line). Right panel: difference $H^*(\phi) - C(\phi)$. H^* is based on 100000 Monte Carlo simulations.

generating function for y is $\log E(e^{\xi y}) = \psi(\eta + \xi) - \psi(\eta)$, the r -th order cumulant of y is $\psi^{(r)}(\eta)$, the r -th order derivative of $\psi(\eta)$. We set $\theta = \psi'(\eta)$, so that $\theta = E(y)$ and $\hat{\theta} = y$. Consequently, $\sigma_\theta = \psi^{(2)}(\eta)^{1/2}$ is the standard error of $\hat{\theta}$, where we use the subscript θ in σ_θ to highlight the dependence on θ . Note that $\sigma_\theta = O(n^{-1/2})$ since $\psi^{(k)} = O(n^{1-k})$. The following result can be stated.

THEOREM 2 *Let $\hat{\theta}$ and $w(\theta)$ be the maximum likelihood estimator and the log-likelihood ratio for the mean value parameter in a continuous one-dimensional exponential model based on a random sample of size n . Also, let σ_θ be the standard error of $\hat{\theta}$ and $C(\theta)$ and $H^*(\theta)$ be defined in (3) and (9), respectively. Then, as $n \rightarrow \infty$,*

$$H^*(\theta) = C(\theta) + O_p(n^{-3/2}), \quad \text{for } (\theta - \hat{\theta})/\sigma_\theta = O_p(1). \quad (18)$$

The proof is deferred to the Appendix and we only provide here in this paragraph a sketch. Reasoning as in the proof of Theorem 1, take $\hat{\theta} > b(\theta)$ so that $H^*(\theta) = [1 - G(\hat{\theta}; \theta) + G(\hat{\theta}^*; \theta)]/2$, where $\hat{\theta}^*$ is implicitly defined by $w^*(\theta; \hat{\theta}^*) = w^*(\theta; \hat{\theta})$. When for $\hat{\theta} > b(\theta)$, (18) corresponds to $G(\hat{\theta}^*; \theta) = 1 - G(\hat{\theta}; \theta) + O_p(n^{-3/2})$, cfr. (16). An asymptotic expansion of $\hat{\theta}^*$ as function of $\hat{\theta}$ and θ is obtained via the generalized inverse mapping result of Lemma 1. In order to work with left tail probabilities, we further define $\hat{\theta}^{**}$ as function of $\hat{\theta}$ and θ to satisfy $1 - G(\hat{\theta}; \theta) = G(\hat{\theta}^{**}; \theta)$. While for the normal transformation family $\hat{\theta}^{**}$ can be derived in explicit form, in the present setting an additional asymptotic expansion is required. This is achieved by an Edgeworth expansion and the allied Cornish–Fisher inversion for the standardized distribution of $(\hat{\theta} - \theta)/\sigma_\theta$. The proof is completed by showing that $\hat{\theta}^*$ and $\hat{\theta}^{**}$ coincide up to the required asymptotic order for $(\hat{\theta} - \theta)/\sigma_\theta = O_p(1)$.

Note that tail symmetry of $cc^*(\theta)$ in the $N(0, \theta)$ example of Section 2 holds by Theorem 2 for $\sigma_\theta = \sqrt{2\theta}/n^{1/2}$, and, hence, a theoretical justification of what we

observed in Figure 1 is obtained. It is worth noting that the chi-squared distribution of the maximum likelihood estimator is the running example in Efron (1987) where it is shown that the transformation g leading to (11) nearly exists, see remark E in Section 11 therein (actually, Efron (1987) considers the sampling distribution of the unbiased estimate of the variance when the mean is unknown). Hence, this example also provides an illustration of tail symmetry in the normal transformation family as stated in Theorem 1.

4.2 Comparison with the modified directed likelihood. We adopt here the notation in Barndorff-Nielsen and Cox (1994, Section 5) for the partial derivatives of ℓ with respect to θ and $\hat{\theta}$, namely

$$\ell_{k;s}(\theta; \hat{\theta}) = -\frac{\partial}{\partial \theta^k} \frac{\partial}{\partial \hat{\theta}^s} \ell(\theta; \hat{\theta})$$

for nonnegative integers k and s . We also adopt the convention of a slash through ℓ indicating the substitution of θ for $\hat{\theta}$ and a hat sign indicating the substitution of $\hat{\theta}$ for θ after any differentiation. The observed information is defined either as $\hat{J} = -\ell_2$ or as $\hat{j} = -\hat{\ell}_2$ according to whether it is considered as a quantity depending on the parameter or the data only. Recall the definition $w(\theta; x) = 2(\ell(\hat{\theta}; x) - \ell(\theta; x))$ the log-likelihood ratio and

$$r(\theta) = r(\theta; x) = \text{sign}(\hat{\theta} - \theta) w(\theta; x)^{1/2} \quad (19)$$

for the directed likelihood. Since $r(\theta)$ is increasing in $\hat{\theta}$, $1 - P(r(\theta; X) \leq r(\theta; x); \theta) = C(\theta)$, where $C(\theta)$ has been defined as $C(\theta) = 1 - G(\hat{\theta}; \theta)$ assuming that the mle $\hat{\theta}$ has distribution $G(y, \theta) = P(\hat{\theta} \leq y; \theta)$ non-increasing in θ . The modified directed likelihood is defined as

$$r^*(\theta) = r(\theta) - \frac{1}{r(\theta)} \log \frac{r(\theta)}{u(\theta)}, \quad u(\theta) = \hat{j} \{\hat{\ell}_{;1} - \ell_{;1}(\theta)\}^{1/2}, \quad (20)$$

see Barndorff-Nielsen and Cox (1994, Section 6.6). It is a higher order pivot, that is it has normal distribution with error $O(n^{-3/2})$ in the normal deviation range $\sqrt{n}(\hat{\theta} - \theta) = O_p(1)$, so that $1 - \Phi(r^*(\theta)) = C(\theta) + O(n^{-3/2})$. Consider now the median bias corrected directed likelihood,

$$r(b(\theta)) = \text{sign}(\hat{\theta} - b(\theta)) w(b(\theta))^{1/2} \quad (21)$$

where we recall that $b(\theta)$ has been defined as the median of the $\hat{\theta}$, i.e. the function of θ that satisfies $G(b(\theta); \theta) = 0.5$. Notice that, since $r(b(\theta))$ is increasing in $\hat{\theta}$, we also have that $1 - P(r(b(\theta); X) \leq r(b(\theta); x); \theta) = C(\theta)$. In the next theorem we establish that, in regular one parameter exponential families, $r(b(\theta))$ and the modified directed likelihood $r^*(\theta)$ are second order equivalent in the normal deviation range.

THEOREM 3 *Let $\hat{\theta}$ and $w(\theta)$ be the maximum likelihood estimator and the log-likelihood ratio for the mean value parameter in a continuous one-dimensional exponential model based on a random sample of size n . Also, let σ_θ be the standard error of $\hat{\theta}$ and $r^*(\theta)$ and $r(b(\theta))$ be defined according to (19)–(21). Then, as $n \rightarrow \infty$,*

$$r(b(\theta)) = r^*(\theta) + O_p(n^{-1}), \quad \text{for } (\theta - \hat{\theta})/\sigma_\theta = O_p(1). \quad (22)$$

The proof is deferred to the Appendix. Note that, because of the higher order pivotal property of $r^*(\theta)$, (22) implies that $1 - \Phi(r(b(\theta))) = C(\theta) + O_p(n^{-1})$ in the normal deviation range, that is $r(b(\theta))$ has sampling distribution closer to normality than $r(\theta)$.

5 Discussion

There has been a renewed interest in confidence distributions in recent years, see Xie and Singh (2013) and Schweder and Hjort (2016). In this paper we have undertaken an asymptotic investigation of the merits of median bias correction in deriving higher order accurate confidence limits. We found that, in regular one-dimensional models, the confidence distribution obtained from the bias corrected log-likelihood ratio is third-order equivalent to the unique exact confidence distribution based on the maximum likelihood estimator. Moreover, the bias corrected directed likelihood provides a second order approximation to the modified directed likelihood of Barndorff-Nielsen (1986), thus consisting in a high order pivot. It shows, from a different perspective, the close connection between the log-likelihood ratio and the distribution of the maximum likelihood estimator so extensively studied in the literature, a key example being the p^* approximation of Barndorff-Nielsen (1983). We are not aware of similar results in the literature on higher order asymptotics.

We have not discussed the effect of the bias correction on the sampling distribution of the log-likelihood ratio $w^*(\theta) = w(b(\theta))$. With the median function $b(\theta)$ at hand, the chi-squared transformation of w^* will typically provide more equal tailed intervals than the usual chi-squared calibration of w . We found however that median bias correction is second order equivalent to what is found via the modified directed likelihood. The convergence to chi-squared distribution of w^* is thus at least to the second order in regular one-parameter exponential families.

An important direction for future research is the extension of the results of Theorem 2 to models with nuisance parameters. In full p -dimensional exponential models when the interest parameter θ is a linear function of the canonical parameters, or a ratio of two canonical parameters, a reparametrization from the canonical parameter vector η to (θ, λ) , where λ is a $(p - 1)$ -dimensional nuisance parameter, can be made and the canonical statistic y can be re-expressed as (y_1, y_2) having density $f(y_1, y_2; \theta, \lambda) = \exp[\theta y_1 + \lambda y_2 - \psi(\theta, \lambda) - d(y_1, y_2)]$. Exact inference on θ can be based on the conditional distribution of y_1 given y_2 , which depends on η only through θ . See Pierce and Peters (1992), and Schweder and Hjort (2016) who find the conditional confidence distribution to be uniformly most powerful. The definition of $C(\theta)$ and $b(\theta)$ are to be interpreted conditionally on y_2 as well. We expect the median bias corrected confidence curve based on the profile likelihood to be tail-symmetric to the third order, and to the second order to be chi-square distributed. The investigation of the relation of the bias corrected profile likelihood with other versions of adjusted profile likelihoods that have been proposed in the literature would also be of interest. Outside the exponential family, the evaluation of sample space derivatives of the likelihood requires the identification of an ancillary statistic. Moreover, the distribution of the maximum likelihood estimator has to be evaluated conditionally upon this statistic.

The asymptotic approximations used in Theorem 2 can be adapted to this setting, a natural extension being for transformation families. Next is a preliminary illustration of the use of median bias correction to confidence curves in a multidimensional statistical model. The model in the example below is not in the exponential family, nor an ancillary statistic is available, and we there use brute force to handle the nuisance parameter.

Example. We consider the “Bolt from heaven” data example from Section 7.4 in Schweder and Hjort (2016). Data consists of $n = 195$ winning times in the fastest 100-m races from 2000 to 2007, that is races that clocked at 10.00 seconds or better. Schweder and Hjort (2016) translate these races results r_i as $x_i = 10.005 - r_i$ in order to apply extreme value statistics. Specifically, the data is modeled using the Generalized Pareto Distribution (GPD) which has density

$$f(x; a, \sigma) = \frac{1}{\sigma} \left(1 - a \frac{x}{\sigma} \right)^{\frac{1}{a} - 1}, \quad 0 \leq x \leq \sigma/a,$$

for $a, \sigma > 0$. Cfr. Sections 3.4 and 6.5 in Embrechts et al. (1997). Interest is in estimating

$$p = p(a, \sigma) = 1 - \exp\{-\lambda(1 - aw/\sigma)\},$$

for $\lambda = 195/8$ and $w = 10.005 - 9.72 = 0.285$. It takes on the interpretation of the probability, as seen at the start of 2008, that in the $N \sim \text{Pois}(\lambda)$ fastest races of 2008 one should experience a race of 9.72 or better, where 9.72 is the world record time scored by Usain Bolt on 31 May 2008. See Schweder and Hjort (2016) for details. The authors compute a confidence curve for the parameter p by profiling the log-likelihood, $\ell_P(p_0) = \max\{\ell(a, \sigma) : p(a, \sigma) = p_0\}$ and by inverting the profile log-likelihood ratio $w(p_0) = 2(\ell_P(\hat{p}) - \ell_P(p_0))$ with respect to the chi-squared distribution after Bartlett correction,

$$cc(p) = Q(w(p)/(1 + \epsilon)),$$

where $(1 + \epsilon) = E(w(p)) \approx 1.07$ (found through simulations) and $Q(\cdot)$ is the chi-squared distribution function with 1 degree of freedom. By construction, $cc(p)$ points at $\hat{p} = p(\hat{a}, \hat{\sigma}) = 0.0345$ according to maximum likelihood estimates $\hat{a} = 0.1821$ (0.0702) and $\hat{\sigma} = 0.0745$ (0.0074) (with approximate standard errors in parentheses) and has 90% confidence interval $[0.0002, 0.1965]$. We proceed next with median bias correction of $w(p)$ so to produce the bias corrected confidence curve $cc^*(p)$. To this aim, the median function of $\hat{p}$ needs to be estimated. The problem here is that, since $\hat{p}$ is not a sufficient statistics for p , its sampling distribution is not uniquely determined by the value of p . If an ancillary statistic for p was available, say $s = s(x)$, then the median function to be used would be

$$b(p) : P(\hat{p} \leq b(p) | s; p) = 0.5$$

where conditioning is intended with respect to the realized value of s in the data. Note that $b(p)$ does not depend on the parameter a since the conditional distribution of $\hat{p}$ given s does not depend on a by definition. The median bias corrected log-likelihood ratio would then be $w^*(p) = w(b(p))$ with sampling distribution $F^*(y|s; p) =$

$P(w(b(p)) \leq y|s; p)$ so that the bias corrected confidence curve would be $cc^*(p) = F^*(w^*(p)|s; p)$.

For the sake of illustration we proceed here without relying on such ancillary statistic. It is convenient to introduce the reparameterization $(a, \sigma) \mapsto (a, p)$ so to define the median function of $\hat{p}$ to depend on both p and a , say

$$b(p|a) : P(\hat{p} \leq b(p|a); a, p) = 0.5.$$

In the left panel of Figure 3 we plot $b(p|a)$ for $a = (0.05, \hat{a}, 0.3)$ (this range has about 95% confidence for a). Each curve is obtained by spline interpolation (constrained to be 0 at $p = 0$) of the median of $\hat{p}$ for a fine grid of p values. For this we used Monte Carlo simulations: for each combination of a and p , 15000 samples of size $n = 195$ from the GPD were drawn and $b(p|a)$ is estimated via the 15000 realizations of the sample median. Hence we replace $b(p)$ with $\hat{b}(p) = b(p|\hat{a})$, i.e. the solid line plotted in the left panel of Figure 3. The median bias corrected log-likelihood ratio is then defined as $w(\hat{b}(p))$ and we estimate its sampling distribution $F^*(y; \hat{a}, p) = P(w(\hat{b}(p)) \leq y; \hat{a}, p)$ for a grid of p values through simulations. Finally, we compute the bias corrected confidence curve

$$cc^*(p) = F^*(w(\hat{b}(p)); \hat{a}, p)$$

and we plot it together with $cc(p)$ (based on the chi-squared approximation with Bartlett correction) in the right panel of Figure 3. Median bias correction moves the confidence curves slightly to the right to the effect that the upper 5% confidence quantile is 0.2278 instead of 0.1965.

Figure 3: *Bolt from heaven data example with $\hat{a} = 0.1821$, $\hat{\sigma} = 0.0745$ and $\hat{p} = 0.0345$. Left panel: median function $b(p|a)$ for $a = (0.05, \hat{a}, 0.3)$. Right panel: confidence curve $cc(p)$ (dashed line) together with bias corrected confidence curve $cc^*(p)$ (solid line).*

Appendix

PROOF OF THEOREM 2. We adopt the notation $\eta = \eta(\theta)$ for the inverse of $\psi'(\eta)$ with the convention that, whenever we suppress the dependence of η on θ , we mean that it is evaluated at θ . Moreover, the “hat” sign denotes evaluation at $\theta = \hat{\theta}$ as in $\hat{\psi}^{(k)} = \psi^{(k)}(\eta(\hat{\theta}))$ or in $\hat{\eta}' = \eta'(\hat{\theta})$. Since $\psi^{(k)}(\eta) = O(n^{1-k})$, the standardized skewness and kurtosis of $\hat{\theta}$ are $\rho_3 = n^{1/2}\psi^{(3)}(\eta)/\sigma_\theta^3$ and $\rho_4 = n\psi^{(4)}(\eta)/\sigma_\theta^4$, where $\sigma_\theta = \psi^{(2)}(\eta)^{1/2}$ is the standard error of $\hat{\theta}$. For $U = (\hat{\theta} - \theta)/\sigma_\theta$ and $K_\theta(x) = P(U \leq x; \theta)$, a Cornish-Fisher expansion of $K_\theta^{-1}(\alpha)$ yields

$$K_\theta^{-1}(\alpha) = z_\alpha + \frac{\rho_3}{6n^{1/2}}(z_\alpha^2 - 1) + \frac{\rho_4}{24n}(z_\alpha^3 - 3z_\alpha) + \frac{\rho_3^2}{36n}(2z_\alpha^3 - 5z_\alpha) + O(n^{-3/2}), \quad (23)$$

see Barndorff-Nielsen and Cox (1989, Section 4.4). The following asymptotic expansion of the median of $\hat{\theta}$ is readily derived:

$$b(\theta) = \theta - \frac{1}{6\sqrt{n}}\rho_3\sigma_\theta + O(n^{-2}). \quad (24)$$

The log-likelihood ratio for θ is $w(\theta; \hat{\theta}) = 2[\hat{\theta}(\hat{\eta} - \eta) - [\hat{\psi} - \psi(\eta)]]$. Let $w^*(\theta; \hat{\theta}) = w(b(\theta); \hat{\theta})$ according to (7). It is easy to check that the first three sample derivatives of w^* are $\partial w^*/\partial \hat{\theta} = 2[\hat{\eta} - \eta(b(\theta))]$, $\partial^2 w^*/\partial \hat{\theta}^2 = 2\hat{\eta}'$ and $\partial^3 w^*/\partial \hat{\theta}^3 = 2\hat{\eta}''$. Using the formulae for the derivative of the inverse of a function, one obtains

$$\eta'(x) = 1/\psi^{(2)}(\eta(x)), \quad \eta''(x) = -\psi^{(3)}(\eta(x))/\psi^{(2)}(\eta(x))^3 \quad (25)$$

so that $\eta'(\theta) = 1/\sigma_\theta^2$ and $\eta''(\theta) = -\rho_3/(n^{1/2}\sigma_\theta^3)$. Let θ^* be implicitly defined as a function of $\hat{\theta}$ and θ by $w^*(\theta; \theta^*) = w^*(\theta; \hat{\theta})$ and assume that $\hat{\theta} > b(\theta)$. Then (18) corresponds to

$$K_\theta((\theta^* - \theta)/\sigma_\theta) = 1 - K_\theta(U) + O(n^{-3/2}) \quad (26)$$

for $U = O_p(1)$, cfr. (16). Next, let θ^{**} be implicitly defined in function of $\hat{\theta}$ and θ by $1 - K_\theta(U) = K_\theta((\theta^{**} - \theta)/\sigma_\theta)$. Hence, (26) is implied by

$$(\theta^* - \theta)/\sigma_\theta = (\theta^{**} - \theta)/\sigma_\theta + O(n^{-3/2}) \quad (27)$$

cfr. (17). We prove (27) by matching the asymptotic expansions of θ^* and θ^{**} via an application of Lemma 1 and Edgeworth expansion of K_θ , respectively.

As for the former, we define, according to Lemma 1, $f_n(x) = w^*(\theta; b(\theta) + \sigma_\theta x)$ and $g_n(x)$ by $f_n(x) = f_n(g_n(x))$, so that $g_n((\hat{\theta} - b(\theta))/\sigma_\theta) = (\theta^* - b(\theta))/\sigma_\theta$. One finds $f_n^{(2)}(0) = 2\eta'(b(\theta))\sigma_\theta^2$ and $f_n^{(k)}(0) = 2\eta^{(k-1)}(b(\theta))\sigma_\theta^k$ for $k \geq 3$. It is easy to show that $\eta^{(k)}(x) = O(n)$ for any integer k , cfr. (25), so the hypothesis of Lemma 1 are satisfied. Hence an application of Lemma 1 yields

$$\begin{aligned} \frac{\theta^* - b(\theta)}{\sigma_\theta} &= -\frac{\hat{\theta} - b(\theta)}{\sigma_\theta} - \frac{1}{3} \frac{\eta'(b(\theta))\sigma_\theta}{\eta''(b(\theta))} \left(\frac{\hat{\theta} - b(\theta)}{\sigma_\theta} \right)^2 \\ &\quad - \left(\frac{1}{3} \frac{\eta'(b(\theta))\sigma_\theta}{\eta''(b(\theta))} \right)^2 \left(\frac{\hat{\theta} - b(\theta)}{\sigma_\theta} \right)^3 + O(n^{-3/2}) \end{aligned} \quad (28)$$

for $[\hat{\theta} - b(\theta)]/\sigma_\theta = O_p(1)$. Based on (24), we have $[\theta^{**} - b(\theta)]/\sigma_\theta = (\theta^{**} - \theta)/\sigma_\theta + \rho_3/6n^{1/2} + O(n^{-3/2})$ and $(\hat{\theta} - b(\theta))/\sigma_\theta = U + \rho_3/6n^{1/2} + O(n^{-3/2})$. Moreover, using a simple Taylor expansion, (25) and $\eta^{(k)}(x) = O(n)$ for any k , it can be shown that $\eta'(b(\theta))/\eta''(b(\theta)) = -\rho_3/n^{1/2}\sigma_\theta + O(n^{-1})$. Hence, we can reduce (28) to

$$\begin{aligned} \frac{\theta^* - \theta}{\sigma_\theta} &= -U - \frac{2\rho_3}{6n^{1/2}} + \frac{1}{3} \frac{\rho_3}{n^{1/2}} U^2 + \frac{1}{3} \frac{\rho_3}{n^{1/2}} \left(2U \frac{\rho_3}{6n^{1/2}} \right) - \left(\frac{1}{3} \frac{\rho_3}{n^{1/2}} \right)^2 U^3 + O(n^{-3/2}) \\ &= -U + \frac{\rho_3}{3n^{1/2}} (U^2 - 1) - \left(\frac{1}{3} \frac{\rho_3}{n^{1/2}} \right)^2 (U^3 - U) + O(n^{-3/2}) \end{aligned} \quad (29)$$

for $U = O_p(1)$. As for the asymptotic expansion of θ^{**} in (27), let $h_n(x)$ satisfy $K_\theta(h_n(x)) = 1 - K_\theta(x)$ so that $(\theta^{**} - \theta)/\sigma_\theta = h_n(U)$. Using Cornish-Fisher expansion (23) one finds that, for any $\alpha \in (0, 1)$,

$$K^{-1}(\alpha) = -K^{-1}(1 - \alpha) + \rho_3(z_{1-\alpha}^2 - 1)/3n^{1/2} + O(n^{-3/2})$$

so that $h_n(x) = -x + \rho_3[\Phi^{-1}(1 - K_\theta(x))^2 - 1]/3n^{1/2} + O(n^{-3/2})$ for $x = O(1)$. Hence,

$$\frac{\theta^{**} - \theta}{\sigma_\theta} = -U + \frac{\rho_3}{3n^{1/2}} [\Phi^{-1}(1 - K_\theta(U))^2 - 1] + O(n^{-3/2}) \quad (30)$$

for $U = O_p(1)$. Next, use the Edgeworth expansion for $K_\theta(U)$ up to the first term, i.e. $1 - K_\theta(U) = \Phi(-U) + \phi(-U)\rho_3(U^2 - 1)/6n^{1/2} + O(n^{-1})$ and a Taylor expansion of $\Phi^{-1}(x + \Delta x)$ at $x = \Phi(-U)$ for $\Delta x = 1 - K_\theta(U) - \Phi(-U)$ to get

$$\Phi^{-1}(1 - K_\theta(U)) = -U + \rho_3(U^2 - 1)/6n^{1/2} + O(n^{-1}).$$

Substitution into (30) leads to an asymptotic expansion of $(\theta^{**} - \theta)/\sigma_\theta$ which corresponds to (29). Hence, (27) follows and the proof is complete. $\square$

PROOF OF THEOREM 2. In order to prove (22), we proceed by deriving two asymptotic expansions for $r^*(\theta)$ and $r(b(\theta))$ and by showing that they coincide up to the required order. As for $r^*(\theta)$, we resort to equation (2.4)–(2.6) in Barndorff-Nielsen (1990). After some algebra and further expansion,

$$r(\theta) = u(\theta) \left(1 + \frac{1}{6} (\ell_3 + 3\ell_{2;1})(\hat{\theta} - \theta)j^{-1} + O(n^{-1}) \right)$$

so that

$$r^*(\theta) = r(\theta) - \frac{1}{6} \frac{1}{r(\theta)} (\ell_3 + 3\ell_{2;1})(\hat{\theta} - \theta)j^{-1} + O(n^{-1}) \quad (31)$$

where we have also used $\log(1 + x) = x + O(x^2)$ for $|x|$ small. As for $r(b(\theta))$, a Taylor expansion around θ gives

$$r(b(\theta)) = r(\theta) - \frac{1}{r(\theta)} \ell_1(\theta)(b(\theta) - \theta) + R_n, \quad (32)$$

with R_n denoting the remainder. In the one-parameter exponential family, borrowing the notation from the proof of Theorem 2, we have

$$\ell_1(\theta) = \eta'(\theta)(\hat{\theta} - \theta) = j(\hat{\theta} - \theta) \quad (33)$$

since $\eta'(\theta) = 1/\psi^{(2)}(\eta) = \sigma_\theta^{-2} = \mathcal{J}$. Moreover, (24) in the proof of Theorem 2 implies that

$$b(\theta) - \theta = -\frac{1}{6\sqrt{n}}\rho_3\sigma_\theta + O(n^{-2}) = \frac{1}{6}\mathcal{J}^{-2}(\ell_3 + 3\ell_{2;1}) + O(n^{-2}) \quad (34)$$

since $\ell_3 = -2\eta''(\theta)$, $\ell_{2;1} = \eta''(\theta)$ and $\eta''(\theta) = -\psi^{(3)}(\eta)/\psi^{(2)}(\eta)^3 = \sigma_\theta^{-2} = \mathcal{J} = -\rho_3/(\sigma_\theta^3\sqrt{n})$. Inserting (33) and (34) into (32) we obtain the same expansion in (31) provided that the remainder R_n is $O(n^{-1})$. This can be shown by using $\ell_1(\theta) = O(n^{1/2})$, $\ell_k(\theta) = O(n)$, $k \geq 2$ and $r(\theta) = O(1)$ in the normal deviation range, together with $b(\theta) - \theta = O(n^{-1})$. Hence (22) follows. $\square$

LEMMA 1 *Let $\{f_n(x)\}_{n \geq 1}$ be a sequence of infinitely differentiable convex functions with minimum at $x = 0$ and $f_n(0) = 0$, and let $g_n(x)$ be defined by $f_n(x) = f_n(g_n(x))$. For $b_{n,k} = 2f_n^{(k)}(0)/k!f_n^{(2)}(0)$, assume that, as $n \rightarrow \infty$, $b_{n,k} = O(b_{n,k-1}n^{-1/2})$ for any $k \geq 3$. Then, $g_n(x)$ admits asymptotic expansion*

$$g_n(x) = -x - \sum_{k \geq 2} a_{n,k}x^k,$$

where $a_{n,2} = b_{n,3}$, $a_{n,3} = b_{n,3}^2$ and

$$a_{n,k} = \begin{cases} O(b_{n,k+1}) & (k \text{ even}), \\ O(b_{n,3}b_{n,k}) & (k \text{ odd}). \end{cases} \quad (35)$$

PROOF. We omit the subscript n for ease of notation. Taylor expansion of f at $x = 0$ gives $f(x) = (1/2)f^{(2)}(0)(x^2 + b_3x^3 + \dots + b_kx^k + \dots)$. Substitute $g(x) = -x - a_2x^2 - \dots$, into $f(x) = f(g(x))$ and equate coefficients of successive order to obtain

$$\begin{cases} b_3 = (a_1a_2 + a_2a_1) - b_3 \\ b_4 = (a_1a_3 + a_2a_2 + a_3a_1) - b_3(a_1a_1a_2 + a_1a_2a_1 + a_2a_1a_1) + b_4 \\ \dots & \dots \\ b_k = \sum_{i_1+i_2=k} a_{i_1}a_{i_2} - b_3 \sum_{i_1+i_2+i_3=k} a_{i_1}a_{i_2}a_{i_3} + b_4 \sum_{i_1+\dots+i_4=k} a_{i_1} \dots a_{i_4} \\ \quad + \dots + (-1)^{k-1}b_{k-1} \sum_{i_1+\dots+i_{k-1}=k} a_{i_1} \dots a_{i_{k-1}} + (-1)^k b_k \end{cases}$$

where the i_j 's are positive integers and we set $a_1 = 1$ for notational convenience. Rearranging terms, the first 4 equations are

$$\begin{cases} b_3 = -b_3 + 2a_1a_2 \\ b_4 = +b_4 - b_3(3a_1^2a_2) + 2a_1a_3 + a_2^2 \\ b_5 = -b_4 + b_4(4a_1^3a_2) - b_3(3a_1^2a_3 + 3a_1a_2^2) + 2a_1a_4 + 2a_2a_3 \\ b_6 = +b_6 - b_5(5a_1^4a_2) + b_4(4a_1^3a_3 + 6a_1^2a_2^2) \\ \quad - b_3(3a_1^2a_4 + 6a_1a_2a_3 + a_2^3) + 2a_1a_5 + 2a_2a_4 + a_3^2 \end{cases}$$

A similar expression for b_k can be given by means of multinomial coefficients. Now substitute back $a_1 = 1$, and solve for a_2, a_3, a_4, a_5 to get $a_2 = b_3$, $a_3 = b_3^2$ and

$$\begin{cases} a_4 = b_5 - 2b_3b_4 + 2b_3^3 & = O(b_5) \\ a_5 = 3b_3b_5 - 6b_3^2b_4 + 4b_3^4 & = O(b_3b_5) \end{cases}$$

where the order of asymptotics of a_4 and a_5 are determined by the hypothesis $b_k = O(b_{k-1}n^{-1/2})$. An argument by induction leads to (35). $\square$

Acknowledgements

The authors are grateful to two reviewers for comments that have helped to improve the paper substantially. Special thanks are also due to Igor Prünster and to Mattia Ciollaro for comments on an earlier version of this work. P. De Blasi was supported by the European Research Council (ERC) through StG “N-BNP” 306406.

References

- Barndorff-Nielsen, O.E. (1983). On a formula for the distribution of the maximum likelihood estimator. *Biometrika* **70**, 343–365.
- Barndorff-Nielsen, O.E. (1986). Inference on full and partial parameters based on the standardized signed log likelihood ratio. *Biometrika* **73**, 307–322.
- Barndorff-Nielsen, O.E. (1990). Approximate interval probabilities. *J. R. Stat. Soc. Ser. B* **52**, 485–496.
- Barndorff-Nielsen, O.E. and Cox, D.R. (1989). *Asymptotic Techniques for Use in Statistics*. Chapman & Hall, London.
- Barndorff-Nielsen, O.E. and Cox, D.R. (1994). *Inference and Asymptotics*. Chapman & Hall, London.
- Birnbaum, A. (1961). Confidence curves: an omnibus technique for estimation and testing statistical hypothesis. *J. Amer. Statist. Assoc.* **56**, 246–249.
- Cox, D. R. (1958). Some problems with statistical inference. *The Annals of Mathematical Statistics*, **29**, 357–372.
- Cox, D. R. (2013). Discussion of M. Xie and K. Singh’s paper. *International Statistical Review*, **81**, 40–41.
- DiCiccio, T.J. and Efron, B. (1992). More accurate confidence intervals in exponential families. *Biometrika* **79**, 231–245.
- DiCiccio, T.J. and Efron, B. (1996). Bootstrap confidence intervals. *Statist. Sci.* **11**, 189–228.
- Efron, B. (1982). Transformation theory: how normal is a family of distributions? *Ann. Statist.* **10**, 323–339.
- Efron, B. (1987). Better bootstrap confidence intervals. *J. Amer. Statist. Assoc.* **82**, 171–185.
- Embrechts, P., Klüppelberg, C. and Mikosch T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer-Verlag Berlin Heidelberg.
- Firth, D. (1993). Bias reduction of maximum likelihood estimates. *Biometrika* **80**, 27–38.
- Fisher, R.A. (1930). Inverse probability. *Proc. Cambridge Philos. Soc.* **26**, 528–35.
- Lehmann, E.L. (1986). *Testing Statistical Hypothesis*, 2ed. Springer-Verlag, New York.
- Lehmann, E.L. and Casella, G. (1998). *Theory of Point Estimation*, 2ed. Springer-Verlag, New York.

- Neyman, J. (1934). On the two different aspects of the representative method: The method of stratified sampling and the method of purposive selection *J. R. Stat. Soc. Ser. A* **97**, 558–625.
- Pierce, D.A. and Peters, D. (1992). Practical use of higher order asymptotics for multiparameter exponential families. *J. R. Stat. Soc. Ser. B* **54**, 701–725.
- Schweder, T. (2007). Confidence nets for curves. In *Advances in Statistical Modeling and Inference, Essay in Honour of Kjell A. Doksum*, V. Nair ed. World Scientific, 593–609.
- Schweder, T. and Hjort, N.L. (2002). Confidence and likelihood. *Scand. J. Statist.* **29**, 309–322.
- Schweder, T. and Hjort, N. L. (2016). *Confidence, Likelihood, Probability: Statistical inference with confidence distributions*. Cambridge University Press.
- Skovgaard, I.M. (1989). A review of higher order likelihood inference. *Bull. Int. Statist. Inst.* **53**, 331–351.
- Xie, M. and Singh, K. (2013). Confidence Distribution, the Frequentist distribution estimator of a parameter – a review. *International Statistical Review* **81**, 3-39.