

To appear in the *Journal of Statistical Computation and Simulation*
Vol. 00, No. 00, Month 20XX, 1–16

Modern Likelihood Inference for the Maximum/Minimum of a Bivariate Normal Vector

V. Mamei^{a*} and A.R. Brazzale^b

^a*University of Cagliari, Department of Mathematics and Computer Science, 72 Via Ospedale, Cagliari, Italy;* ^b*University of Padova, Department of Statistical Sciences, 241 Via Cesare Battisti, Padova, Italy*

(Received 00 Month 20XX; final version received 00 Month 20XX)

We consider the use of modern likelihood asymptotics in the construction of confidence intervals for the parameter which determines the skewness of the distribution of the maximum/minimum of an exchangeable bivariate normal random vector. Simulation studies were conducted to investigate the accuracy of the proposed methods and to compare them to available alternatives. Accuracy is evaluated in terms of both coverage probability and expected length of the interval. We furthermore illustrate the suitability of our proposals by means of two data sets, consisting of, respectively, measurements taken on the brains of ten mono-zygotic twins and measurements of mineral content of bones in the dominant and non dominant arms for twenty five elderly women.

Keywords: Bivariate normal distribution; Higher order likelihood inference; Modified likelihood ratio; Skew-Normal distribution.

AMS Subject Classification: *Classification codes:*

1. Introduction

Small sample sizes are rather common to researchers in fields such as biology, genetics, medical sciences and psychology. Inference based on the classical first order Normal and χ^2 approximations may then be unreliable. The last four decades have seen the development of so-called higher order likelihood approximations, which require little more effort than is needed for their first order counterparts while providing highly accurate inferences in small samples. We refer the reader to [1] for a rich collection of realistic examples and case studies, which show how to use the new theory. The aim of this paper is to encourage the use of these modern likelihood-based solutions for the analysis of bivariate normally distributed continuous data when interest focuses on the maximum (or minimum) value of correlated observations. Indeed, a great variety of studies, especially in the medical area, are based on the analysis of extreme observations. For instance, the extremes coming from the left and right sides of the body are commonplace in comparative studies; see [2] for an illustration about the visual acuity of fellow eyes, or [3], [4] and reference therein, for other typical cases.

We set off from a theoretical result — due to [5] and later re-discovered by [6] — which states that the maximum, or minimum, of two random variables, whose joint distribution is bivariate exchangeable normal with correlation coefficient ρ , is skew-normally distributed with skewness parameter γ , or $-\gamma$, where $\gamma = \sqrt{(1-\rho)/(1+\rho)}$. Recently, [7], borrowing from Loperfido (2002) and Fisher's z transform for ρ , obtained an asymp-

*Corresponding author. Email: mamei.valentina@virgilio.it

otic confidence interval for the skewness parameter of the distribution of the maximum/minimum under this framework. Their simulations revealed that the proposed asymptotic confidence interval has similar actual and nominal coverages, though its expected length increases for decreasing sample size and correlation coefficient close to -1 . In this paper we explore the performance of confidence intervals for γ obtained from the small-sample solutions recently proposed in [8], and this in terms of both actual coverage and expected length.

The paper organizes as follows. Section 2 reviews modern likelihood-based inference. The skew-Normal distribution and Loperfido's results will be introduced in Section 3. Inference on γ will be discussed in Section 4. Section 5 analyzes the twin data collected by [9] and the data set presented by [10] using the large- and small-sample solutions of Section 4. The finite-sample properties of confidence intervals will be investigated in Section 6 through simulation. Some concluding remarks are given in Section 7.

2. Likelihood-based inference

2.1. First order theory

Let $y = (y_1, \dots, y_n)$ be a sample of size n with joint log-likelihood function $l(\theta) = l(\theta; y)$, where $\theta = (\psi, \lambda)$ is a k -dimensional parameter, ψ is the scalar parameter of interest, and λ a vector of nuisance parameters of dimension $k - 1$. Under broad regularity conditions, the maximum likelihood estimate of θ , denoted by $\hat{\theta}$, may be obtained by solving the score equation $l_\theta(\hat{\theta}; y) = 0$, with $l_\theta(\theta; y) = \partial l(\theta; y) / \partial \theta$. Let $j(\theta) = \partial^2 l(\theta; y) / \partial \theta \partial \theta^\top$ represent the observed information function for θ and $j(\hat{\theta})$ the observed Fisher information. Here \top denotes matrix transposition. The decomposition of the parameter θ into ψ and λ leads to an analogous decomposition of the score vector $l_\theta(\theta; y)$ and of the observed information function $j(\theta)$.

The recommended likelihood pivot for making inference on ψ is the signed likelihood root

$$r(\psi) = \text{sign}(\hat{\psi} - \psi) \sqrt{2\{l_p(\hat{\psi}) - l_p(\psi)\}}. \quad (1)$$

Here $l_p(\psi) = l(\hat{\theta}_\psi)$, with $\hat{\theta}_\psi = (\psi, \hat{\lambda}_\psi)$, is the profile log-likelihood, while $\hat{\lambda}_\psi$ represents the constrained maximum likelihood estimate obtained by maximizing the log-likelihood $l(\psi, \lambda)$ with respect to λ holding ψ fixed. The signed likelihood root (1) is asymptotically standard Normal up to the order $n^{-1/2}$, which leads to the first order $(1 - \alpha)100\%$ confidence interval for ψ

$$\{\psi : |r(\psi)| \leq z_{1-\alpha/2}\}, \quad (2)$$

where z_p , with $p \in (0, 1)$, is the p th quantile of the standard normal distribution. The standard normal approximation provides a satisfactory approximation for large sample sizes, but can be highly unreliable for small values of n . The value of ψ which satisfies equation (2) can be found numerically by calculating the function $r(\psi)$ on a grid of points ψ , which are then interpolated using a suitable smoothing function. The numerical issues, which may arise in the interpolation step, can be avoided by excluding the values of ψ close to the maximum likelihood estimate $\hat{\psi}$. The details are given in [1, Section 9.3].

2.2. Higher order theory

A nowadays broadly known improvement to the signed likelihood root (1), which was originally introduced by Barndorff-Nielsen (1983), is the modified likelihood ratio

$$r^* = r + \frac{1}{r} \log \left(\frac{q}{r} \right), \quad (3)$$

whose finite-sample distribution may be approximated by the standard Normal up to the order $n^{-3/2}$. The higher order $(1 - \alpha)100\%$ confidence interval for ψ is hence given by

$$\{\psi : |r^*(\psi)| \leq z_{1-\alpha/2}\}. \quad (4)$$

Again, pivot profiling [1, Section 9.3] can be used to identify the upper and lower bounds of the confidence interval. Furthermore, the r^* pivot — like its first order counterpart r — is invariant under interest-respecting re-parametrizations, that is re-parametrizations of the form $\tau(\theta) = \tau(\psi, \lambda) = (\zeta, \eta)$ with $\zeta = \zeta(\psi)$ and $\eta = \eta(\psi, \lambda)$.

Several expressions for the correction term q have been proposed, both from the frequentist and the Bayesian perspective. Here, we will focus on the developments by [12]. Their formula for q is based upon the notion of “tangent exponential model” which, at a fixed value of y , denoted y_0 , approximates the true model by a local exponential model with canonical parameter $\varphi = \varphi(\theta)$, defined as

$$\varphi^\top(\theta) = l_{;V}(\theta; y_0) = \sum_{i=1}^n \frac{\partial l(\theta; y)}{\partial y_i} \Bigg|_{y=y_0} V_i. \quad (5)$$

Here, $l_{;V}$ indicates differentiation of the log-likelihood function in the directions given by the n columns V_1, \dots, V_n of the $n \times k$ matrix V . The matrix V can be constructed using a vector of pivotal quantities $z = \{z_1(y_1, \theta), \dots, z_n(y_n, \theta)\}$, where each component $z_i(y_i, \theta)$ has a fixed distribution under the model. The matrix V is defined through z by

$$V = - \left(\frac{\partial z}{\partial y^\top} \right)^{-1} \left(\frac{\partial z}{\partial \theta^\top} \right) \Bigg|_{(y_0, \hat{\theta}_0)},$$

where $\hat{\theta}_0$ is the maximum likelihood estimate at y_0 . The expression of the correction term q is then

$$q = \frac{|\varphi(\hat{\theta}) - \varphi(\hat{\theta}_\psi)|}{|\varphi_\theta(\hat{\theta})|} \left\{ \frac{|j(\hat{\theta})|}{|j_{\lambda\lambda}(\hat{\theta}_\psi)|} \right\}^{\frac{1}{2}}, \quad (6)$$

where $\varphi_\theta(\theta) = \partial\varphi(\theta)/\partial\theta^\top$ represents the matrix of partial derivatives of $\varphi(\theta)$ with respect to θ , while $\varphi_\lambda(\theta) = \partial\varphi(\theta)/\partial\lambda^\top$ identifies the $k - 1$ columns of this matrix which correspond to the nuisance parameter λ . Analogously, the matrix $j_{\lambda\lambda}(\theta)$ is the $(k - 1) \times (k - 1)$ sub-matrix of the observed information function $j(\theta)$ with respect to the nuisance parameter λ .

The expression of q for the case in which no explicit nuisance parametrization is given can be found in [8].

2.3. Approximations for Bayesian inference

In the Bayesian setting with a prior density $\pi(\theta)$ for θ , the analogue of the first order results of Section 2.1 is the asymptotic normality of the posterior density $\pi(\theta|y)$ for θ . The Bayesian counterpart of the correction term q in (3), which we will denote by q_B , was obtained by [13] under the assumption that the nuisance parametrization is given explicitly, and results to

$$q_B = l'_p(\psi)j_p(\hat{\psi})^{-\frac{1}{2}} \left\{ \frac{|j_{\lambda\lambda}(\hat{\theta}_\psi)|}{|j_{\lambda\lambda}(\hat{\theta})|} \right\}^{\frac{1}{2}} \frac{\pi(\hat{\theta})}{\pi(\hat{\theta}_\psi)}, \quad (7)$$

where $l'_p(\psi) = dl_p(\psi)/d\psi$ is the profile score function and $j_p(\psi) = d^2l_p(\psi)/d\psi^2$ the profile observed information function. Posterior quantiles for the parameter ψ can then be found exploiting the fact that the posterior distribution function

$$\Pi(\psi_0 | y) = \Pr(\psi \leq \psi_0 | y) \doteq 1 - \Phi(r_B^*)$$

may be approximated to the order $n^{-3/2}$ by the standard Normal distribution function $\Phi(r_B^*)$, evaluated at

$$r_B^* = r + \frac{1}{r} \log \left(\frac{q_B}{r} \right). \quad (8)$$

Again, pivot profiling provides the upper and lower bounds of the $(1 - \alpha)100\%$ credible interval for ψ given by

$$\{\psi : |r_B^*(\psi)| \leq z_{1-\alpha/2}\}. \quad (9)$$

Like for q , [8] provide the expression of the correction term q_B for the case in which the nuisance parametrization is not given explicitly.

2.3.1. Matching priors

Given the prior $\pi(\theta)$ for θ , let $\theta_{1-\alpha}^\pi$ denote the $(1 - \alpha)$ th approximate posterior quantile of θ of order n^{-r} , that is, the value of θ for which

$$\Pr_{\theta|y}(\theta \leq \theta_{1-\alpha}^\pi | y) = 1 - \alpha + O_p(n^{-r}), \quad (10)$$

with $r > 0$ and $0 < \alpha < 1$. If we also have that

$$\Pr_{Y|\theta}(\theta_{1-\alpha}^\pi \geq \theta | \theta) = 1 - \alpha + O_p(n^{-r}), \quad (11)$$

with $\theta_{1-\alpha}^\pi$ the upper bound of a frequentist one-sided $(1 - \alpha)100\%$ confidence interval, the prior π is called a probability matching prior to the r th order of approximation. For such priors, Bayesian and frequentist inference for the parameter θ are in perfect agreement up to the order r .

If $r = 1$, $\pi(\theta)$ is called a first order probability matching prior, while for $r = 3/2$ we have a second order probability matching prior. Welch and Peers (1963) showed that the unique first order probability matching prior, when no nuisance parameter is present, is Jeffrey's prior.

The same result does not necessarily hold when θ includes a nuisance component λ . For an orthogonal parametrization, [15] proposed to use the following prior for θ in (7),

$$\pi(\psi, \lambda) \propto i_{\psi\psi}^{1/2}(\psi, \lambda), \quad (12)$$

where $i_{\psi\psi}(\psi, \lambda)$ represents the value of the expected Fisher information function corresponding to ψ . The authors call this prior the “unique prior”, as it leads to an approximation of the marginal posterior distribution of ψ accurate to the order $n^{-3/2}$. When the parametrization $\theta = (\psi, \lambda)$ is not orthogonal, their suggestion is to find an orthogonal parametrization (ψ, η) of the original model for which the prior can be expressed as (12), and then to re-express the prior in the original parametrization (ψ, λ) , leading to

$$\pi(\psi, \lambda) \propto i_{\psi\psi.\lambda}^{1/2}(\psi, \lambda) \left| \frac{\partial \eta}{\partial \lambda} \right|, \quad (13)$$

with $i_{\psi\psi.\lambda}(\psi, \lambda) = i_{\psi\psi}(\psi, \lambda) - i_{\psi\lambda}(\psi, \lambda) i_{\lambda\lambda}^{-1}(\psi, \lambda) i_{\lambda\psi}(\psi, \lambda)$. Here, the indices ψ and λ indicate which sub-blocks of the expected Fisher information function to take. Furthermore, $|\partial \eta / \partial \lambda|$ represents the Jacobian of the transformation from (ψ, η) to (ψ, λ) .

3. The skew-Normal model

The skew-Normal distribution was introduced by [16] to define a class of asymmetric parametric models which includes the standard Normal as a special case. We say that a continuous random variable $Z \sim SN(\gamma)$, distributes as a skew-Normal indexed by the real parameter γ , if it has density function

$$p(z; \gamma) = 2\phi(z)\Phi(\gamma z) \quad \text{with } z \in \mathbb{R}.$$

Here $\phi(\cdot)$ and $\Phi(\cdot)$ denote, respectively, the density and the distribution functions of the standard normal distribution. The class of skew-Normal distributions can be widened by including a location parameter $\mu \in \mathbb{R}$ and a scale parameter $\sigma > 0$. Thus, if $X \sim SN(\gamma)$, then $Y = \mu + \sigma X$ is a skew-Normal random variable with parameters μ, σ, γ , or, $Y \sim SN(\mu, \sigma, \gamma)$ for short. Making inference on the skewness parameter is quite challenging, as the expected Fisher information becomes singular as $\gamma \rightarrow 0$; see the solutions proposed by [17] and [18]. Functions for manipulating the skew-Normal probability distribution and for fitting it to data are given in the R package `sn` [19]. We refer the reader to [20] and [21] for a general treatment of the skew-Normal distribution and its extensions.

The focus of this paper is on the distribution of the maximum (or minimum) of an exchangeable bivariate Normal random vector. The first contribution to this regard is, to our knowledge, [22] who discusses inference for the minimum value with special emphasis on twin studies. We derive our reference model from the below result, as cited in [6], although a first version was already derived by [5] in the special case of the bivariate Normal model with zero means, unit variances and correlation coefficient ρ .

THEOREM 3.1 *Let X_1 and X_2 be two random variables whose joint distribution is bivariate normal with common mean $\mu \in \mathbb{R}$, common variance $\sigma^2 > 0$ and correlation coefficient $\rho \in (-1, 1)$. Then for any two real constants h and $k \neq -h$, the distribution*

of $h \min(X_1, X_2) + k \max(X_1, X_2)$ is

$$SN\left(\mu(h+k), \sigma\sqrt{h^2+k^2+2\rho hk}, \gamma = \frac{k-h}{|k+h|}\sqrt{\frac{1-\rho}{1+\rho}}\right).$$

It follows that the distribution of $\max(X_1, X_2)$ is $SN(\mu, \sigma, \gamma)$ with $\gamma = \sqrt{(1-\rho)/(1+\rho)} \geq 0$, whereas the distribution of $\min(X_1, X_2)$ is $SN(\mu, \sigma, \gamma)$ with $\gamma = -\sqrt{(1-\rho)/(1+\rho)} \leq 0$. The special case of $\rho = 0$ translates into $\gamma = 1$ and $\gamma = -1$, respectively. Theorem 3.1 was subsequently generalized by [3] to the case where X_1 and X_2 are exchangeable, elliptical and continuous random variables.

4. Inference on the skewness parameter γ

This section outlines known exact and large-sample results for making inference on the skewness parameter $\gamma = \sqrt{(1-\rho)/(1+\rho)}$ which characterizes the distribution of the maximum between two correlated normal variables, as well as our newly derived small-sample solutions. In particular, we will consider the following three instances: 1) the no-nuisance parameter case, 2) the complete bivariate Normal model, and 3) the equi-correlated bivariate Normal model. All methods straightforwardly extend to the case $-\gamma = -\sqrt{(1-\rho)/(1+\rho)}$ with the suitable changes.

4.1. The bivariate Normal model without nuisance parameter

Let $Y = (X_1, X_2)$ be a bivariate Normal random vector with common mean 0, common variance 1 and correlation coefficient $\rho \in (-1, 1)$. Given an i.i.d. sample $\{(x_{11}, x_{21}), \dots, (x_{1n}, x_{2n})\}$ of size n from Y , [23] proposed a range-based exact confidence interval for ρ . The construction of the confidence interval makes use of the two random variables $D_+ = \sum_{i=1}^n (X_{1i} + X_{2i})^2$ and $D_- = \sum_{i=1}^n (X_{1i} - X_{2i})^2$. Taking advantage of the independence of $X_{1i} + X_{2i}$ and $X_{1i} - X_{2i}$ along with the fact that $X_{1i} + X_{2i} \sim N(0, 2(1+\rho))$ and $X_{1i} - X_{2i} \sim N(0, 2(1-\rho))$, the authors derive the following pivotal quantity

$$\frac{D_+ (1-\rho)}{D_- (1+\rho)} \sim F_{n,n}, \tag{14}$$

where $F_{n,n}$ is Fisher's F distribution with (n, n) degrees of freedom. This gives an exact $(1-\alpha)100\%$ confidence interval for the parameter γ of the form

$$\left\{ \gamma \in [0, \infty) : \sqrt{\frac{D_-}{D_+} F_{n,n}(\alpha/2)} < \gamma < \sqrt{\frac{D_-}{D_+} F_{n,n}(1-\alpha/2)} \right\}, \tag{15}$$

where $F_{n,n}(p)$, with $p \in (0, 1)$, represents the p th quantile.

4.1.1. Small-sample confidence intervals

[24] provides the expression of the canonical parameter φ when interest relies on $\theta = \rho$, the correlation coefficient of the bivariate Normal vector $Y = (X_1, X_2)$. The reference model in this case is a $(2, 1)$ curved exponential family. Confidence intervals for the parameter γ are derived from the $r(\rho)$, and $r^*(\rho)$ pivots thanks to their invariance under interest-respecting re-parametrizations. Credible intervals for γ are obtained from $r_B^*(\rho)$

by exploiting that $\gamma = \sqrt{(1 - \rho)/(1 + \rho)}$ is a one-to-one transformation of the random variable ρ .

A key quantity for the determination of the canonical parameter (5) of the approximating tangent full exponential model is the vector $V = (1 - \hat{\theta})^{-1}(t - \hat{\theta}s, s - \hat{\theta}t)^\top$, obtained from the two independent pivots $Z_1 = (T + S)/(1 + \theta)$ and $Z_2 = (T - S)/(1 - \theta)$, with $S = n^{-1} \sum_{i=1}^n X_{1i}X_{2i}$ and $T = (2n)^{-1} \sum_{i=1}^n (X_{1i}^2 + X_{2i}^2)$, whose distribution is χ_n^2/n . The canonical parameter takes the form $\varphi(\theta) = n\{(1 - \theta^2)(1 - \hat{\theta}^2)\}^{-1}\{\theta(t - \hat{\theta}s) - (s - \hat{\theta}t)\}$. Later, [25] proposed an alternative formulation, $\bar{\varphi}(\theta) = n\theta/(1 - \theta^2)$, of the canonical parameter. As shown there, both formulations lead to almost the same numerical results as far as the approximation of tail areas is concerned. Turning to the Bayesian world, we may adopt Jeffreys' prior for ρ , given by

$$\pi(\rho) \propto \frac{\sqrt{(1 + \rho^2)}}{(1 - \rho^2)}. \tag{16}$$

As stated in Section 2.3.1, this prior provides a first order probability matching prior for a scalar parameter in the absence of nuisance parameters.

4.2. The complete bivariate Normal model

Let $Y = (X_1, X_2)$ be again a bivariate Normal random vector, this time with parameters $(\mu_1, \mu_2, \sigma_1, \sigma_2, \rho)$ where $(\mu_1, \sigma_1) \in \mathbb{R} \times \mathbb{R}_+$ and $(\mu_2, \sigma_2) \in \mathbb{R} \times \mathbb{R}_+$ are, respectively, the means and variances of X_1 and X_2 , and $\rho \in (-1, 1)$ their correlation. Instructions of how to construct a, yet approximate, confidence interval for ρ given an i.i.d. sample $\{(x_{11}, x_{21}), \dots, (x_{1n}, x_{2n})\}$ of size n from Y , can once more be found in [23]. The first step is to standardize the two components X_{1i} and X_{2i} into X_{1i}^* and X_{2i}^* . An approximate confidence interval for the parameter ρ is obtained by using the fact that $X_{1i}^* - X_{2i}^*$ and $X_{1i}^* + X_{2i}^*$ are nearly independent, $X_{1i}^* + X_{2i}^* \sim N(0, 2(1 + \rho))$ and $X_{1i}^* - X_{2i}^* \sim N(0, 2(1 - \rho))$. The pivot

$$\frac{D_+^* (1 - \rho)}{D_-^* (1 + \rho)} = \left(\frac{1 + R}{1 - R} \right) \frac{(1 - \rho)}{(1 + \rho)}, \tag{17}$$

where $D_+^* = \sum_{i=1}^n (X_{1i}^* + X_{2i}^*)^2$ and $D_-^* = \sum_{i=1}^n (X_{1i}^* - X_{2i}^*)^2$ and R is the sample correlation coefficient defined as

$$R = \frac{\sum_{i=1}^n (X_{1i} - \bar{X}_1)(X_{2i} - \bar{X}_2)}{\sqrt{\sum_{i=1}^n (X_{1i} - \bar{X}_1)^2} \sqrt{\sum_{i=1}^n (X_{2i} - \bar{X}_2)^2}},$$

follows approximately an $F_{n-1, n-1}$ distribution. The corresponding $(1 - \alpha)100\%$ confidence interval for γ becomes

$$\left\{ \gamma \in [0, \infty) : \sqrt{\left(\frac{1 - R}{1 + R} \right) F_{n-1, n-1}(\alpha/2)} < \gamma < \sqrt{\left(\frac{1 - R}{1 + R} \right) F_{n-1, n-1}(1 - \alpha/2)} \right\}. \tag{18}$$

A second approximate solution to the inferential problem we are interested in can be found in [7]. Because of the difficulties of obtaining the finite sample distribution of R , inference for ρ is commonly based on the monotonic transformation

$1/2 \ln((1 + R)/(1 - R))$, called Fisher’s z -transform. In particular, the distribution of

$$Z = \frac{\frac{1}{2} \ln\left(\frac{1+R}{1-R}\right) - \frac{1}{2} \ln\left(\frac{1+\rho}{1-\rho}\right)}{\frac{1}{\sqrt{n-3}}} \tag{19}$$

for $n > 50$ is approximately standard Normal. This turns into an $(1 - \alpha)100\%$ confidence interval for γ , which we will call ACI_5 , of the form

$$\left\{ \gamma \in [0, \infty) : \exp\left(\frac{-z_{1-\frac{\alpha}{2}}}{\sqrt{n-3}}\right) \sqrt{\left(\frac{1-R}{1+R}\right)} < \gamma < \exp\left(\frac{z_{1-\frac{\alpha}{2}}}{\sqrt{n-3}}\right) \sqrt{\left(\frac{1-R}{1+R}\right)} \right\}. \tag{20}$$

Note that the upper and lower bounds of both, the confidence interval (20) proposed by [7] and solution (18) derived from the confidence interval for ρ of [23], include the multiplying factor $\sqrt{(1 - R)/(1 + R)}$.

4.2.1. Small-sample confidence intervals

The computation of small-sample confidence intervals for the parameter ρ from the frequentist point of view is addressed in [26]. Let $\{(x_{11}, x_{21}), \dots, (x_{1n}, x_{2n})\}$ be an i.i.d. sample from the bivariate Normal random vector $Y = (X_1, X_2)$. The log-likelihood function

$$\begin{aligned} l(\theta) = & -n \left(\log(\sigma_1\sigma_2) + \frac{1}{2} \log(1 - \rho^2) + \frac{\mu_1^2}{2(1 - \rho^2)\sigma_1^2} + \frac{\mu_2^2}{2(1 - \rho^2)\sigma_2^2} - \frac{\mu_1\mu_2\rho}{(1 - \rho^2)\sigma_1\sigma_2} \right) + \\ & - \frac{1}{2(1 - \rho^2)\sigma_1^2} \sum_{i=1}^n x_{1i}^2 - \frac{1}{2(1 - \rho^2)\sigma_2^2} \sum_{i=1}^n x_{2i}^2 + \frac{\mu_1\sigma_2 - \mu_2\sigma_1\rho}{(1 - \rho^2)\sigma_1^2\sigma_2} \sum_{i=1}^n x_{1i} + \\ & + \frac{\mu_2\sigma_1 - \mu_1\sigma_2\rho}{(1 - \rho^2)\sigma_1\sigma_2^2} \sum_{i=1}^n x_{2i} + \frac{\rho}{(1 - \rho^2)\sigma_1\sigma_2} \sum_{i=1}^n x_{1i}x_{2i}, \end{aligned}$$

with $\theta = (\rho, \mu_1, \mu_2, \sigma_1, \sigma_2)$, characterizes an exponential family with canonical parameter

$$\varphi(\theta) = \left(-\frac{1}{(1 - \rho^2)\sigma_1^2}, -\frac{1}{(1 - \rho^2)\sigma_2^2}, \frac{\mu_1\sigma_2 - \mu_2\sigma_1\rho}{(1 - \rho^2)\sigma_1^2\sigma_2}, \frac{\mu_2\sigma_1 - \mu_1\sigma_2\rho}{(1 - \rho^2)\sigma_1\sigma_2^2}, \frac{\rho}{(1 - \rho^2)\sigma_1\sigma_2} \right).$$

Setting $\psi = \rho$ and $\lambda = (\mu_1, \mu_2, \sigma_1, \sigma_2)$, q and r^* can readily be obtained from equations (6) and (3), respectively. The computation of the Bayesian credible interval of Section 2.3 requires that we specify a prior for the parameter θ . The “unique prior” defined by [15] may be calculated by referring to the orthogonal re-parametrization

$$p(x_1, x_2 | \psi, \eta) \propto \frac{1}{\eta_4} \exp \left\{ -\frac{1}{2(1 - \psi^2)^{1/2}\eta_4} \left[\frac{(x_1 - \eta_1)^2}{\eta_3} + \eta_3(x_2 - \eta_2)^2 - 2\psi(x_1 - \eta_1)(x_2 - \eta_2) \right] \right\},$$

with $\psi = \rho$ and $\eta_1 = \mu_1$, $\eta_2 = \mu_2$, $\eta_3 = \sigma_1/\sigma_2$ and $\eta_4 = \sigma_1\sigma_2(1 - \rho^2)^{1/2}$. The matching prior for $\theta = (\rho, \lambda)$, with $\lambda = (\mu, \sigma)$, will then be

$$\pi(\theta) \propto i_{\rho\rho,\lambda}^{1/2}(\theta) \left| \frac{\partial \eta}{\partial \lambda} \right|, \tag{21}$$

with $i_{\rho\rho,\lambda}(\theta) = i_{\rho\rho}(\theta) - i_{\rho\lambda}(\theta)i_{\lambda\lambda}^{-1}(\theta)i_{\lambda\rho}(\theta)$. In expression (21) the quantity $i(\theta)$ denotes the expected Fisher information function, while $|\partial \eta / \partial \lambda|$ represents the Jacobian of the transformation from the orthogonal parametrization (ρ, η) to the original parametrization (ρ, λ) . Note that [27, 28] used the same parametrization but with $\mu_1 = \mu_2 = 0$, which

implies that Fisher’s expected information function only includes the parameters ψ , η_3 and η_4 .

As in the previous section, confidence intervals for the parameter γ are readily derived from the $r(\rho)$ and $r^*(\rho)$ pivots thanks to their invariance under interest-respecting re-parametrizations, while credible intervals for γ are obtained from $r_B^*(\rho)$ by using the biunivocal transformation, $\sqrt{(1-\rho)/(1+\rho)}$, of the random variable ρ .

4.3. The equi-correlated bivariate Normal model

Let $(x_{11}, x_{21}), \dots, (x_{1n}, x_{2n})$ be an i.i.d. sample from an equi-correlated bivariate Normal vector $Y = (X_1, X_2)$, where X_1 and X_2 have common mean $\mu \in \mathbb{R}$, common variance $\sigma^2 > 0$ and correlation coefficient $\rho \in (-1, 1)$. Maximum likelihood estimation for $\theta = (\rho, \mu, \sigma)$ yields

$$\hat{\mu} = \frac{1}{2n} \sum_{i=1}^n (x_{1i} + x_{2i}), \quad \hat{\sigma} = \sqrt{\frac{1}{2n} \sum_{i=1}^n [(x_{1i} - \hat{\mu})^2 + (x_{2i} - \hat{\mu})^2]},$$

$$\text{and } \hat{\rho} = \frac{1}{n\hat{\sigma}^2} \sum_{i=1}^n (x_{1i} - \hat{\mu})(x_{2i} - \hat{\mu}).$$

[29] called ρ the intraclass correlation coefficient and its maximum likelihood estimate $\hat{\rho}$ the intraclass correlation.

4.3.1. Small-sample confidence intervals

A first confidence interval for γ , which is accurate up to the second order, can be derived by using an argument equivalent to the one adopted by [7] for the complete bivariate Normal model, and exploits the fact that Fisher’s transformation of $\hat{\rho}$ is again a normalizing and variance stabilizing transformation as it is the case for the well known transformation of the sample correlation. That is, the pivot

$$\sqrt{n} \left[\frac{1}{2} \log \left(\frac{1 + \hat{\rho}}{1 - \hat{\rho}} \right) - \frac{1}{2} \log \left(\frac{1 + \rho}{1 - \rho} \right) + \frac{1}{2n} \right], \tag{22}$$

is asymptotically standard normal up to the order n^{-1} ; see [30]. Expression (22) leads straightforwardly to the following asymptotic confidence interval for γ :

$$\left\{ \gamma \in [0, \infty) : \exp \left(\frac{-z_{1-\frac{\alpha}{2}}}{\sqrt{n}} - \frac{1}{2n} \right) \sqrt{\frac{1-\hat{\rho}}{1+\hat{\rho}}} < \gamma < \exp \left(\frac{z_{1-\frac{\alpha}{2}}}{\sqrt{n}} - \frac{1}{2n} \right) \sqrt{\frac{1-\hat{\rho}}{1+\hat{\rho}}} \right\}, \tag{23}$$

which we will indicate by ACI_3 .

Further small-sample confidence intervals may be derived by setting, as in Section 4.2, $\psi = \rho$ and $\lambda = (\mu, \sigma)$ and using the r^* pivot from equations (6) and (3), respectively. The log-likelihood function

$$l(\theta) = -n \left(2 \log \sigma + \frac{1}{2} \log (1 - \rho^2) + \frac{\mu^2}{(1 + \rho)\sigma^2} \right) + \frac{\sum_{i=1}^n (x_{1i} + x_{2i})}{(1 + \rho)\sigma^2} + \frac{2\rho \sum_{i=1}^n x_{1i}x_{2i} - \sum_{i=1}^n (x_{1i}^2 + x_{2i}^2)}{2(1 - \rho^2)\sigma^2},$$

now characterizes an exponential family with canonical parameter

$$\varphi(\theta) = \left(\frac{\mu}{\sigma^2(1 + \rho)}, \frac{\rho}{\sigma^2(1 - \rho^2)}, -\frac{1}{2(1 - \rho^2)\sigma^2} \right). \tag{24}$$

In order to obtain credible confidence intervals from (8) with frequentist coverage, we assume as matching prior $\pi(\theta)$, the “unique prior” defined in [15]. This requires an orthogonal re-parametrization of the model of the form

$$p(x_1, x_2 | \rho, \eta) \propto \frac{1}{\eta_2} \exp \left\{ -\frac{1}{2(1 - \rho^2)^{1/2}\eta_2} [(x_1 - \eta_1)^2 + (x_2 - \eta_1)^2 - 2\rho(x_1 - \eta_1)(x_2 - \eta_1)] \right\},$$

with $\eta_1 = \mu$ and $\eta_2 = \sigma^2(1 - \rho^2)^{1/2}$. The matching prior for $\theta = (\rho, \lambda)$, with $\lambda = (\mu, \sigma)$, will then be

$$\pi(\theta) \propto i_{\rho\rho,\lambda}^{1/2}(\theta) \left| \frac{\partial \eta}{\partial \lambda} \right|, \tag{25}$$

with $i_{\rho\rho,\lambda}(\theta) = i_{\rho\rho}(\theta) - i_{\rho\lambda}(\theta)i_{\lambda\lambda}^{-1}(\theta)i_{\lambda\rho}(\theta)$. As before, the quantity $i(\theta)$ indicates the expected Fisher information function, while $|\partial\eta/\partial\lambda|$ denotes the Jacobian of the transformation from (ρ, η) to (ρ, λ) .

Once more, confidence and credible intervals for γ are obtained via pivot profiling (see Brazzale *et al.*, 2007) from the pivots $r(\rho)$, $r^*(\rho)$, and $r_B^*(\rho)$.

5. Two real-data examples

5.1. Corpus callosum *surface area in mono-zygotic twins*

Since Sir Francis Galton’s (1876) seminal paper, twin studies have extensively been used for the quantitative ascertainment of genetic and environmental influences. Twin registries worldwide represent nowadays a valuable resource for the investigation of the similarities and dissimilarities between twins. The very large twin studies carried out during the past two decades led to much novel work, especially in genetic research [32]. Classical twin designs remain, nonetheless, a valuable tool in fields such as biomedicine, psychiatry and behavioral sciences, where the number of available observations is far smaller than those typical in modern twin studies. There are several views of how the degree of concordance between twins should be assessed [33, 34]. Here, we promote the use of Azzalini’s (1985) skew-Normal distribution.

The twin study conducted by Tramo *et al.* (1998) focuses on similarities of the brains of mono-zygotic twin pairs such as in forebrain volume, cortical surface area, and callosal area. The data, as available on `StatLib`, include the measurements on the brains of ten pairs of mono-zygotic twins, five male and five female. Our variable of interest is the *corpus callosum* surface area. In particular, we focus on the maximum between the area in the first twin and in the second twin. The order of twins is determined by birth, though this does not affect the method since the two variables are supposed to be exchangeable. The paired measurements exhibit a strong positive linear relationship with correlation coefficient amounting to $R = 0.90$. To assure that all conditions of Theorem 3.1 hold, we first standardize the pairs of observations $\{(x_{11}, x_{21}), \dots, (x_{1n}, x_{2n})\}$ as on page 7. The bivariate Shapiro-Wilk normality test ($W = 0.97$, p -value = 0.86) supports the hypothesis of bivariate normality of the standardized data. Maximum likelihood yields $\hat{\gamma} = 0.229$.

The five 95% confidence intervals for γ , computed using the methods outlined in Section 4.2, are given in Table 1. The interval based on the third order Bayesian solution r_B^* is wider than the confidence intervals obtained from the first order pivot r , the higher order frequentist pivot r^* , the large sample (*HP*) confidence interval by [23] and the ACI_5 confidence interval by [7]. Figure 1 shows how to compute the lower and upper bounds numerically. The intervals based on r (1st), r^* (3rd), r_B^* (*Bayes*) and ACI_5 can be read off from the intersections of the corresponding pivots with the horizontal black lines, which represent the 2.5% and the 97.5% quantiles of the standard Normal distribution. The lower and upper bounds of the *HP* confidence interval are computed similarly, but this time by referring to the horizontal grey lines, which represent the 2.5%, and the 97.5% quantiles of the F distribution with (9, 9) degrees of freedom.

Figure 1 and Table 1 about here.

5.2. Mineral content measurements in the ulna of elderly women

A number of situations, especially in clinical comparative studies, foresee that the data are collected on the left and right sides of the body of each individual. In this framework, it can be relevant to estimate the distribution of the extremes (maximum or minimum) of these contra-lateral observations; see for example [3] and [4] and reference therein for details.

Here, we consider the analysis presented in [4] on data by [10]. The observations consist of mineral content measurements of three bones, among these the ulna, in the dominant and non dominant arm for twenty five elderly women. We focus our attention on the minimum between the contra-lateral measurements of the ulna. To take account of the bilateral symmetries in the human body, we constrain the model parameters to $\mu_1 = \mu_2 = \mu$ and $\sigma_1 = \sigma_2 = \sigma$. The bivariate Shapiro-Wilk normality test confirms the hypothesis of bivariate normality ($W = 0.9708$, p -value = 0.6645). Furthermore, as expected, the two contra-lateral measurements present a strong positive linear relationship with intraclass correlation $\hat{\rho} = 0.724$. The maximum likelihood estimate of $-\gamma$ is $-\hat{\gamma} = -0.4004$.

Table 2 shows the four 95% confidence intervals for the parameter $-\gamma$ obtained by using the methods outlined in Section 4.3. The ACI_3 confidence interval of equation (23) is smaller than the confidence intervals obtained from the first order pivot r and the higher order frequentist and Bayesian pivots r^* and r_B^* , respectively. Figure 2 plots the pivot functions r (1st), r^* (3rd), r_B^* (*Bayes*) and ACI_3 for the ulna data. As before, intervals can be read off from the intersections of the corresponding pivots with the horizontal lines which represent the 2.5% and the 97.5% quantiles of the standard Normal distribution.

Figure 2 and Table 2 about here.

6. Numerical assessment

We designed three simulations studies to assess and compare the finite-sample properties of the methods discussed in this paper. The summary statistics used are: empirical coverage (*CP*), upper error probability (*UE*), that is, the percentage of the true parameter values falling above the upper bound, lower error probability (*LE*), that is, the percentage of the true parameter values falling below the lower bound, and average length (*AL*)

of the confidence intervals considered in Section 4. All simulations were run using the numerical computing environment R [35]. For every occurrence, 10,000 replicates were generated for the four sample sizes $n = 5, 10, 15, 20$. The simulation error amounts to ± 0.004 . We only report the summary statistics for the two samples sizes $n = 5, 10$; the omitted results are available as supplementary material.

6.1. Considered scenarios

6.1.0.1. Simulation 1. Considers a bivariate Normal model with zero means, unit variances and unknown correlation ρ , which takes values from -0.9 to $+0.9$, with step size 0.1 . The purpose is to compare the behavior of the higher order pivot r^* with its first order counterpart r , the Bayesian small-sample solution r_B^* and the exact method (14), which apply when no nuisance parameter is present, and this for small sample sizes.

6.1.0.2. Simulation 2. Wants to investigate the finite-sample performance of the confidence intervals obtained when nuisance parameters are present, again while emphasizing small sample sizes. The pivots used are the higher order solution r^* , its first order counterpart r , the Bayesian competitor r_B^* , and the large-sample solutions (18) and (20). We used $\mu_1 = \mu_2 = 7$, $\sigma_1 = \sigma_2 = 0.9$, while again $\rho \in \{-0.9, -0.8, \dots, 0.8, 0.9\}$. Note that the simulation set-up borrows from the twin data example of Section 5.1, for which the maximum likelihood estimate is $\hat{\theta} = (0.900, 7.061, 6.924, 0.905, 0.872)$.

6.1.0.3. Simulation 3. Explores the finite-sample performance of the different confidence intervals which are available for the equi-correlated bivariate Normal model, again while emphasizing on small sample sizes. The pivots used are the higher order solution r^* , its first order counterpart r , the Bayesian competitor r_B^* , and the approximate solution (22). We set $\mu = 0.7$, $\sigma = 0.1$, while again $\rho \in \{-0.9, -0.8, \dots, 0.8, 0.9\}$. Note that the simulation set-up borrows from the real data example of Section 5.2, for which the maximum likelihood estimate is $\hat{\theta} = (0.724, 0.699, 0.103)$.

6.2. Discussion

6.2.1. Simulation 1

Figure 3 shows the actual coverage of the nominal 95% confidence intervals for γ derived from equation (15) (*exact*), and by using the pivots r (*1st*), r^* (*3rd*), and r_B^* (*Bayes*) for the no-nuisance parameter case. The higher order likelihood pivots r^* and, to a somewhat lesser extent, r_B^* outperform their first order counterpart r , even for the very limited sample sizes considered in the four scenarios of Simulation 1. The differences among the four pivots fade out as the sample size increases.

Figure 3 and Tables 3 and 4 about here.

Tables 3 and 4 summarize the performance of the nominal 95% confidence intervals for γ derived from the four pivots considered in Simulation 1. Being there no nuisance parameter present, the corresponding true coverage probabilities are very close to the nominal level, as we may have expected. Inspection of the upper and lower error probabilities reveals that the r^* and r_B^* pivots also improve in terms of symmetry over r . The exact method produces confidence intervals for γ which are, on average, wider than the confidence intervals obtained from the first order solution r and the higher order pivots

r^* and r_B^* . The average length of all four confidence intervals is larger for negative values of ρ , and increases when the correlation tends to -1 .

6.2.2. Simulation 2

Figure 4 reports the actual coverage of the nominal 95% confidence intervals for γ obtained from expressions (17) (*HP*) and (19) (*ACI₅*), and by using the pivots r (*1st*), r^* (*3rd*) and r_B^* (*Bayes*). In terms of real coverage, r^* again outperforms its first order counterpart r . It also outperforms the large-sample (*HP*) proposal by [23] and, surprisingly, the Bayesian solution r_B^* . The most accurate method is the large-sample confidence interval developed by [7], although the differences fade out for increasing sample size.

Figure 4 and Tables 5 and 6 about here.

Tables 5 and 6 summarize the performance of the nominal 95% confidence intervals for γ derived from the five methods considered in Simulation 2. The results reveal that r^* is more accurate than r , especially when the sample size is small, because of both an, on average, larger width and its capability of correctly centering the confidence intervals. The r_B^* pivot consistently over-estimates the real coverage, while guaranteeing symmetry on the tails, because of the, on average, longer confidence intervals it produces. The *ACI₅* and *HP* methods lead to confidence intervals for γ which are remarkably asymmetric. Their better performance with respect to, respectively, r^* and r may be explained by the, on average, larger widths of the corresponding confidence intervals. For all five methods considered, the expected length becomes larger for negative values of ρ , especially when ρ is close to -1 . This is in agreement with [7], who noted the same behavior for their *ACI₅* confidence interval.

6.2.3. Simulation 3

Figure 5 shows the actual coverage of the nominal 95% confidence intervals for γ by using the pivots r (*1st*), r^* (*3rd*), r_B^* (*Bayes*) and (22) (*ACI₃*) in the equi-correlated bivariate Normal model. In terms of real coverage, r^* again outperforms its first order counterpart r , the Bayesian solution r_B^* and the approximate solution (23).

Figure 5 and Tables 7 and 8 about here.

Tables 7 and 8 reveal that the small-sample pivots r^* and r_B^* exhibit more reliable coverage than the confidence intervals obtained from their large-sample counterpart r and the pivot (22), although the Bayesian solution r_B^* somewhat overestimates the nominal level. Moreover, r and, to a somewhat lesser extent (22), exhibit, for small sample sizes, an unsatisfactory behavior as far as the symmetry of errors is concerned, although the *ACI₃* confidence intervals seem to improve as soon as $n \geq 10$. In addition, r_B^* and (22) produce confidence intervals with, respectively, larger and shorter expected widths as compared with the intervals obtained from r^* and r . The finite-sample differences among the four pivots vanish as the sample size increases. For all four methods, the expected length becomes wider for negative values of ρ , especially when ρ approaches -1 .

7. Concluding remarks

In this paper we investigate the behavior of modern likelihood-based small-sample procedures to compute confidence intervals for the parameter of skewness which characterizes the distribution of the maximum/minimum of a bivariate Normal and exchangeable random vector. This distribution may, for example, be used for assessing the degree of concordance of a continuous mono-zygotic twin trait when interest focuses on the pairwise maximum or minimum, as in Section 5.1. Or, it may represent the reference model for making inference on contra-lateral measurements, as in Section 5.2.

Extensive numerical investigation revealed that the higher order frequentist pivot r^* is highly accurate, especially for the rather small sample sizes which may be encountered, and for the challenging situation where ρ is close to -1 . This is in agreement with the findings by [26], though their contribution focuses on ρ and does not consider the custom-tailored statistics of Section 4.2. When no nuisance parameter is present, r^* yields confidence intervals which, for practical purposes, may be considered exact. Among the four alternatives available in the presence of nuisance parameters, the only real competitor to r^* , in terms of both real coverage and required computational efforts, is the ACI_5 confidence interval, though it leads to, on average, longer confidence intervals which counterbalance the lack of symmetry on the tails. The potential applicability of the ACI_5 method to studies on twins was already put forward in [7]. The construction of confidence sets for γ in the three-parameter bivariate Normal model shows that the small-sample frequentist confidence interval based on r^* outperforms its three alternatives with respect to both average length and symmetry of errors. The ACI_3 confidence interval requires a minor computational effort than the likelihood-based. However, despite the fact that its construction borrows from [7], it performs poorly in terms of both coverage probability and symmetry of error rates. This is especially true for small sample sizes.

The straightforward extension of our proposal is to the skew-Normal model. For instance, Loperfido's Theorem provides the reference models for mono-zygotic twin studies for which information on the pair (X_1, X_2) is missing, and only their maximum (or minimum) value is recorded. This may happen because of practical reasons; see [22] for a rather early treatment. As pointed out there, because healthy mono-zygotic twins share an identical genetic mark-up, time of onset for a particular event in the first twin — such as getting a cold or developing leukaemia — is likely to closely follow in the second twin, so that only the smaller or larger record may be kept. The same holds true for clinical comparative studies. Furthermore, working with the maximum (or minimum) of two correlated measurements can be, at times, more reliable than the study of the original values, especially if the measurements of the smaller (or larger) values are more accurate.

Acknowledgements

The data on twins used in Section 5.1 were obtained from the electronic software repository *StatLib*. We would like to express our gratitude to [9] who have submitted the dataset on twins in this archive and have given us the authorization to use them. This research was partially supported by the University of Padova grant no. CPDA101912.

References

- [1] Brazzale AR, Davison AC, Reid N. Applied Asymptotics: Case Studies in Small-Sample Statistics. Cambridge (UK): Cambridge University Press; 2007.

- [2] Viana MAG, Lee HM. Correlation analysis of ordered symmetrically dependent observations and their concomitants of order statistics. *The Canadian Journal of Statistics*. 2006; 34:327–340.
- [3] Loperfido N. A note on skew-elliptical distributions and linear functions of order statistics. *Statistics & Probability Letters*. 2008; 78:3184–3186.
- [4] Sheikhi A, Tata M. The exact joint distribution of concomitants of order statistics and their order statistics under normality. *REVSTAT Statistical Journal*. 2013; 11:121–134.
- [5] Nagaraja HN. A note on linear functions of ordered correlated normal random variables. *Biometrika*. 1982; 69:284–285.
- [6] Loperfido N. Statistical implications of selectively reported inferential results. *Statistics & Probability Letters*. 2002; 56:13–22.
- [7] Mameli V, Musio M, Saleau E, Biggeri A. Large sample confidence intervals for the skewness parameter of the skew-normal distribution based on Fisher’s transformation. *Journal of Applied Statistics*. 2012; 39:1693–1702.
- [8] Fraser DAS, Reid N, Wu J. A simple general formula for tail probabilities for frequentist and Bayesian inference. *Biometrika*. 1999; 86: 249–264.
- [9] Tramo MJ, Loftus WC, Green RL, Stukel TA, Weaver JB, Gazzaniga, M. S. Brain size, head size, and IQ in monozygotic twins. *Neurology*. 1998; 50:1246–1252.
- [10] Johnson RA, Wichern DW. *Applied Multivariate Statistical Analysis* (5th Ed.), Englewood Cliffs, New Jersey: Prentice-Hall; 2002.
- [11] Barndorff-Nielsen O. On a formula for the distribution of the maximum likelihood estimator. *Biometrika*. 1983; 70:343–365.
- [12] Fraser DAS, Reid N. Ancillary and third order significance. *Utilitas Mathematica*. 1995; 47:33–53.
- [13] DiCiccio TJ, Martin A. Approximations of tail probabilities for a class of smooth functions with applications to Bayesian and conditional inference. *Biometrika*. 1991; 78:891–902.
- [14] Welch BL, Peers BL. On formulae for confidence points based on integrals of weighted likelihoods. *Journal of the Royal Statistical Society Series B*. 1963; 25:318–329.
- [15] Staicu AM, Reid N. On probability matching priors. *The Canadian Journal of Statistics*. 2008; 36:613–622.
- [16] Azzalini A. A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*. 1985; 12:171–178.
- [17] Azzalini A, Capitanio A. Statistical applications of the multivariate skew-normal distributions. *Journal of the Royal Statistical Society Series B*. 1999; 61:579–602.
- [18] Sartori N. Bias prevention of maximum likelihood estimates for scalar skew normal and skew t distributions. *Journal of Statistical Planning and Inference*. 2006; 136:4259–4275.
- [19] Azzalini A. R package `sn`: The skew-normal and skew- t distributions (version 0.4-18). Available from: <http://azzalini.stat.unipd.it/SN>.
- [20] Genton MG. *Skew-Elliptical Distributions and Their Applications: A Journey Beyond Normality*. Boca Raton (FL): Chapman & Hall / CRC; 2004.
- [21] Azzalini A, Capitanio A. *The Skew-Normal and Related Families*. Cambridge (UK): Cambridge University Press; 2014.
- [22] Roberts C. A correlation model useful in the study of twins. *Journal of the American Statistical Association*. 1966; 61:1184–1190.
- [23] Haddad JH, Provost SB. Approximations to the distribution of the sample correlation coefficient. *World Academy of Science, Engineering and Technology*. 2011; 52:910–915.
- [24] Reid N. Asymptotics and the theory of inference. *Annals of Statistics*. (2003); 31:1695–1731.
- [25] Reid N, Fraser DAS. Mean log-likelihood and higher-order approximations. *Biometrika*. 2010; 97:159–170.
- [26] Sun Y, Wong ACM. Interval estimation for the normal correlation coefficient. *Statistics & Probability Letters*. 2007; 77: 1652–1661.
- [27] Gosh M, Mukherjee B, Santra U. Objective priors: an introduction for frequentists. *International Journal of Statistical Sciences*. 2009; 9:255–271.
- [28] Gosh M, Mukherjee B, Santra U, Kim D. Bayesian and likelihood-based inference for the bivariate normal correlation coefficient. *Journal of Statistical Planning and Inference*. 2010; 140:1410–1416.

- [29] Fisher RA. On the “probable error” of a coefficient of correlation deduced from a small sample. *Metron*. 1921; 1:3–32.
- [30] Konishi S. Normalizing and variance stabilizing transformation for intraclass correlations. *Annals of the Institute of Statistical Mathematics*. 1985; 37:87–94.
- [31] Galton F. The history of twins, as a criterion of the relative powers of nature and nurture. *The Journal of the Anthropological Institute of Great Britain and Ireland*. 1876; 5:391–406.
- [32] van Dongen J, Slagboom PE, Draaisam HHM, Martin NG, Boomsma DI. The continuing value of twin studies in the omics era. *Nature Reviews Genetics*. 2012; 13:640–653.
- [33] Kraemer HC. What is the ‘right’ statistical measure of twin concordance (or diagnostic reliability and validity)? *Archives of Genetic Psychiatry*. 1997; 54:1121–1124.
- [34] Lyons MJ, Faraone SV, Tsuang MT, Goldberg J, Eaves LJ, Meyer JM, True WR, Eisen SA. Another view on the ‘right’ statistical measure of twin concordance. *Archives of Genetic Psychiatry*. 1997; 54:1126–1128.
- [35] R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. (2013). Available from: <http://www.R-project.org>.




Figure 1. [9] *corpus callosum* surface area data. Pivot functions for the parameter γ obtained from: likelihood root r (solid), modified likelihood root r^* (bold); Bayesian modified likelihood root r_B^* (long-dashes); expression (17) by [23] (dotted); approximate pivot used in (19) by [7] (dashed). Black horizontal lines: 2.5% and 97.5% Normal quantiles; grey horizontal lines: 2.5% and 97.5% quantiles of the $F(9, 9)$ distribution.

Table 1. [9] *corpus callosum* surface area data. Lower (LB) and upper (UB) bounds of 95% confidence intervals for the parameter γ . Pivots used, with corresponding confidence intervals: 1st – likelihood root r (2); 3rd – modified likelihood root r^* (4); Bayes – Bayesian modified likelihood root r_B^* (9); HP – [23] (18); ACI_5 – [7] (20).

Method	LB	UB	Length
1st	0.121	0.435	0.314
3rd	0.119	0.493	0.374
Bayes	0.114	0.518	0.404
HP	0.114	0.460	0.346
ACI_5	0.109	0.481	0.372






Figure 2. [10] mineral content measurements of the ulna data. Pivot functions for the parameter γ obtained from: likelihood root r (solid), modified likelihood root r^* (bold); Bayesian modified likelihood root r_B^* (long-dashes); approximate pivot (23) (dashed). Horizontal lines: 2.5% and 97.5% Normal quantiles.

Table 2. [10] mineral content measurements of the ulna data. Lower (LB) and upper (UB) bounds of 95% confidence intervals for the parameter γ . Pivots used, with corresponding confidence intervals: 1st – likelihood root r (2); 3rd – modified likelihood root r^* (4); Bayes – Bayesian modified likelihood root r_B^* (9); ACI_3 – approximate pivot (22).


Method	LB	UB	Length
1st	-0.596	-0.269	0.327
3rd	-0.588	-0.261	0.327
Bayes	-0.573	-0.253	0.320
ACI_3	-0.581	-0.265	0.316




(a) $n = 5$



(b) $n = 10$




(c) $n = 15$




(d) $n = 20$


Figure 3. Simulation 1: bivariate Normal with means 0 and variances 1. Empirical coverage of nominal two-sided 95% confidence intervals for γ for varying values of ρ and sample sizes $n = 5, 10, 15, 20$. Pivots used: likelihood root r (*1st*), modified likelihood root r^* (*3rd*); Bayesian modified likelihood root r_B^* (*Bayes*); expression (14) by [23] (*exact*). Based on 10,000 replicates.




(a) $n = 5$



(b) $n = 10$




(c) $n = 15$




(d) $n = 20$


Figure 4. Simulation 2: bivariate Normal with means 7 and variances 0.9. Empirical coverage of nominal equi-tailed 95% confidence intervals for γ for varying values of ρ and sample sizes $n = 5, 10, 15, 20$. Pivots used: likelihood root r (1st), modified likelihood root r^* (3rd); Bayesian modified likelihood root r_B^* (Bayes); expression (17) by [23] (HP) and expression (19) by [7] (ACI_5). Based on 10,000 replicates.




(a) $n = 5$



(b) $n = 10$



(c) $n = 15$



(d) $n = 20$

Figure 5. Simulation 3: bivariate Normal with means 0.7 and variances 0.1. Empirical coverage of nominal equi-tailed 95% confidence intervals for γ for varying values of ρ and sample sizes $n = 5, 10, 15, 20$. Pivots used: likelihood root r (1st), modified likelihood root r^* (3rd); Bayesian modified likelihood root r_B^* (Bayes); and expression (23) (ACI_3). Based on 10,000 replicates.

Table 3. Summary statistics for Simulation 1: bivariate Normal with means 0 and variances 1. Empirical coverage (CP), upper (UE) and lower (LE) error probability and average length (AL) of nominal two-sided 95% confidence intervals for γ , for varying values of ρ and sample size $n = 5$. Pivots used: likelihood root r ($1st$), modified likelihood root r^* ($3rd$); Bayesian modified likelihood root r_B^* ($Bayes$); expression (15) by [23] ($exact$). Based on 10,000 replicates; simulation error: ± 0.004 .

(a)						(b)					
ρ	Method	CP	UE	LE	AL	ρ	Method	CP	UE	LE	AL
-0.9	$1st$	0.924	0.039	0.037	6.714	0.1	$Bayes$	0.937	0.031	0.032	1.758
	$3rd$	0.954	0.024	0.022	6.556		$exact$	0.952	0.023	0.025	2.579
	$Bayes$	0.953	0.024	0.022	6.604		$1st$	0.927	0.038	0.035	1.806
	$exact$	0.950	0.024	0.026	11.329		$3rd$	0.951	0.025	0.024	1.732
-0.8	$1st$	0.938	0.038	0.024	4.753	$Bayes$	0.941	0.031	0.029	1.647	
	$3rd$	0.952	0.023	0.026	4.600	$exact$	0.953	0.024	0.022	2.346	
	$Bayes$	0.947	0.024	0.029	4.587	0.2	$1st$	0.925	0.033	0.042	1.695
	$exact$	0.951	0.023	0.026	7.803		$3rd$	0.948	0.025	0.026	1.636
-0.7	$1st$	0.939	0.038	0.023	3.861		$Bayes$	0.936	0.030	0.034	1.550
	$3rd$	0.948	0.024	0.029	3.702		$exact$	0.949	0.024	0.028	2.110
	$Bayes$	0.940	0.024	0.036	3.648	0.3	$1st$	0.925	0.036	0.040	1.577
	$exact$	0.949	0.025	0.026	6.144		$3rd$	0.948	0.027	0.024	1.526
-0.6	$1st$	0.934	0.040	0.026	3.345		$Bayes$	0.938	0.031	0.030	1.439
	$3rd$	0.948	0.024	0.028	3.198		$exact$	0.950	0.026	0.024	1.903
	$Bayes$	0.943	0.023	0.034	3.130	0.4	$1st$	0.929	0.032	0.039	1.460
	$exact$	0.952	0.025	0.024	5.244		$3rd$	0.948	0.028	0.024	1.421
-0.5	$1st$	0.930	0.038	0.033	2.946		$Bayes$	0.942	0.032	0.026	1.347
	$3rd$	0.945	0.024	0.031	2.807		$exact$	0.952	0.025	0.024	1.701
	$Bayes$	0.942	0.025	0.033	2.727	0.5	$1st$	0.927	0.032	0.040	1.350
	$exact$	0.947	0.025	0.028	4.483		$3rd$	0.943	0.031	0.026	1.324
-0.4	$1st$	0.932	0.037	0.031	2.659		$Bayes$	0.940	0.033	0.027	1.257
	$3rd$	0.950	0.022	0.027	2.531		$exact$	0.947	0.026	0.026	1.499
	$Bayes$	0.945	0.024	0.030	2.441	0.6	$1st$	0.928	0.031	0.041	1.236
	$exact$	0.953	0.024	0.023	3.948		$3rd$	0.940	0.035	0.025	1.209
-0.3	$1st$	0.922	0.042	0.035	2.426		$Bayes$	0.937	0.037	0.025	1.157
	$3rd$	0.942	0.027	0.031	2.314		$exact$	0.948	0.027	0.025	1.313
	$Bayes$	0.934	0.032	0.033	2.234	0.7	$1st$	0.935	0.023	0.042	1.061
	$exact$	0.948	0.026	0.027	3.552		$3rd$	0.945	0.029	0.026	1.054
-0.2	$1st$	0.926	0.039	0.034	2.241		$Bayes$	0.935	0.038	0.026	1.039
	$3rd$	0.949	0.026	0.025	2.141		$exact$	0.950	0.026	0.024	1.097
	$Bayes$	0.940	0.030	0.030	2.055	0.8	$1st$	0.936	0.025	0.038	0.865
	$exact$	0.952	0.025	0.023	3.189		$3rd$	0.948	0.027	0.024	0.864
-0.1	$1st$	0.927	0.037	0.036	2.087		$Bayes$	0.946	0.029	0.024	0.905
	$3rd$	0.951	0.023	0.026	1.995		$exact$	0.951	0.026	0.023	0.874
	$Bayes$	0.937	0.031	0.032	1.911	0.9	$1st$	0.925	0.035	0.040	0.554
	$exact$	0.953	0.023	0.024	2.854		$3rd$	0.954	0.022	0.024	0.563
0	$1st$	0.925	0.037	0.038	1.927		$Bayes$	0.953	0.023	0.024	0.652
	$3rd$	0.949	0.024	0.027	1.842		$exact$	0.952	0.026	0.023	0.598

Table 4. Summary statistics for Simulation 1: bivariate Normal with means 0 and variances 1. Empirical coverage (CP), upper (UE) and lower (LE) error probability and average length (AL) of nominal two-sided 95% confidence intervals for γ , for varying values of ρ and sample size $n = 10$. Pivots used: likelihood root r ($1st$), modified likelihood root r^* ($3rd$); Bayesian modified likelihood root r_B^* ($Bayes$); expression (15) by [23] ($exact$). Based on 10,000 replicates; simulation error: ± 0.004 .

(a)						(b)					
ρ	Method	CP	UE	LE	AL	ρ	Method	CP	UE	LE	AL
-0.9	$1st$	0.942	0.037	0.021	4.907	0	$Bayes$	0.945	0.029	0.026	1.227
	$3rd$	0.951	0.024	0.025	4.826		$exact$	0.951	0.025	0.024	1.485
	$Bayes$	0.951	0.024	0.025	4.834	0.1	$1st$	0.936	0.031	0.034	1.192
	$exact$	0.952	0.022	0.025	6.496		$3rd$	0.949	0.026	0.025	1.177
					$Bayes$		0.945	0.027	0.028	1.138	
-0.8	$1st$	0.945	0.034	0.020	3.174	$exact$	0.951	0.024	0.025	1.345	
	$3rd$	0.950	0.025	0.025	3.144	0.2	$1st$	0.935	0.033	0.032	1.106
	$Bayes$	0.949	0.025	0.026	3.154		$3rd$	0.948	0.029	0.023	1.098
	$exact$	0.949	0.024	0.027	4.484		$Bayes$	0.946	0.030	0.024	1.060
-0.7	$1st$	0.946	0.035	0.019	2.588	$exact$	0.952	0.027	0.021	1.223	
	$3rd$	0.953	0.025	0.023	2.556	0.3	$1st$	0.937	0.028	0.036	1.011
	$Bayes$	0.949	0.025	0.026	2.551		$3rd$	0.947	0.026	0.027	1.008
	$exact$	0.953	0.025	0.022	3.545		$Bayes$	0.945	0.026	0.029	0.974
-0.6	$1st$	0.943	0.033	0.024	2.242	$exact$	0.949	0.024	0.027	1.091	
	$3rd$	0.948	0.023	0.029	2.205	0.4	$1st$	0.938	0.027	0.035	0.920
	$Bayes$	0.944	0.023	0.034	2.185		$3rd$	0.948	0.028	0.025	0.922
	$exact$	0.952	0.024	0.024	2.968		$Bayes$	0.948	0.027	0.025	0.896
-0.5	$1st$	0.944	0.034	0.022	1.999	$exact$	0.951	0.024	0.025	0.974	
	$3rd$	0.951	0.023	0.026	1.961	0.5	$1st$	0.941	0.023	0.035	0.817
	$Bayes$	0.947	0.025	0.028	1.931		$3rd$	0.947	0.026	0.027	0.827
	$exact$	0.954	0.025	0.021	2.583		$Bayes$	0.945	0.027	0.027	0.807
-0.4	$1st$	0.937	0.036	0.027	1.801	$exact$	0.950	0.024	0.026	0.860	
	$3rd$	0.946	0.026	0.028	1.766	0.6	$1st$	0.942	0.022	0.036	0.704
	$Bayes$	0.943	0.028	0.030	1.730		$3rd$	0.948	0.027	0.025	0.719
	$exact$	0.949	0.025	0.026	2.275		$Bayes$	0.943	0.031	0.026	0.716
-0.3	$1st$	0.937	0.034	0.029	1.649	$exact$	0.954	0.023	0.023	0.743	
	$3rd$	0.947	0.025	0.028	1.615	0.7	$1st$	0.945	0.020	0.036	0.574
	$Bayes$	0.944	0.027	0.029	1.577		$3rd$	0.950	0.023	0.026	0.594
	$exact$	0.952	0.023	0.024	2.034		$Bayes$	0.946	0.028	0.026	0.606
-0.2	$1st$	0.933	0.036	0.031	1.503	$exact$	0.949	0.025	0.026	0.627	
	$3rd$	0.945	0.027	0.028	1.472	0.8	$1st$	0.943	0.020	0.037	0.415
	$Bayes$	0.940	0.029	0.031	1.433		$3rd$	0.954	0.022	0.024	0.438
	$exact$	0.947	0.027	0.026	1.825		$Bayes$	0.952	0.023	0.025	0.459
-0.1	$1st$	0.936	0.034	0.030	1.395	$exact$	0.952	0.023	0.026	0.496	
	$3rd$	0.948	0.027	0.025	1.368	0.9	$1st$	0.943	0.022	0.034	0.259
	$Bayes$	0.947	0.028	0.025	1.328		$3rd$	0.949	0.029	0.023	0.275
	$exact$	0.951	0.026	0.024	1.648		$Bayes$	0.949	0.028	0.023	0.283
0	$1st$	0.937	0.032	0.031	1.290	$exact$	0.949	0.026	0.025	0.343	
	$3rd$	0.949	0.026	0.025	1.268						

Table 5. Summary statistics for Simulation 2: bivariate Normal with means 7 and variances 0.9. Empirical coverage (*CP*), upper (*UE*) and lower (*LE*) error probability and average length (*AL*) of nominal two-sided 95% confidence intervals for γ , for varying values of ρ and sample size $n = 5$. Pivots used: likelihood root r (*1st*), modified likelihood root r^* (*3rd*); Bayesian modified likelihood root r_B^* (*Bayes*); expression (18) by [23] (*HP*); expression (20) by [7] (*ACI₅*). Based on 10,000 replicates; simulation error: ± 0.004 .

(a)						(b)					
ρ	Method	CP	UE	LE	AL	ρ	Method	CP	UE	LE	AL
-0.9	<i>1st</i>	0.842	0.118	0.041	11.313	0	<i>HP</i>	0.901	0.049	0.050	3.561
	<i>3rd</i>	0.938	0.035	0.026	12.490		<i>ACI₅</i>	0.953	0.023	0.024	4.808
	<i>Bayes</i>	0.970	0.016	0.014	14.316	0.1	<i>1st</i>	0.839	0.081	0.080	2.473
	<i>HP</i>	0.904	0.072	0.024	18.031		<i>3rd</i>	0.938	0.034	0.028	3.073
	<i>ACI₅</i>	0.956	0.034	0.010	24.343		<i>Bayes</i>	0.969	0.016	0.015	3.784
-0.8	<i>1st</i>	0.849	0.107	0.044	8.565	<i>HP</i>	0.905	0.048	0.047	3.197	
	<i>3rd</i>	0.940	0.035	0.025	9.071	<i>ACI₅</i>	0.953	0.024	0.023	4.316	
	<i>Bayes</i>	0.968	0.017	0.015	10.617	0.2	<i>1st</i>	0.843	0.070	0.087	2.149
	<i>HP</i>	0.909	0.066	0.025	12.216		<i>3rd</i>	0.941	0.028	0.031	2.721
	<i>ACI₅</i>	0.953	0.034	0.013	16.492		<i>Bayes</i>	0.970	0.014	0.016	3.360
-0.7	<i>1st</i>	0.847	0.102	0.051	7.001	<i>HP</i>	0.906	0.041	0.053	2.766	
	<i>3rd</i>	0.938	0.034	0.028	7.743	<i>ACI₅</i>	0.954	0.020	0.026	3.735	
	<i>Bayes</i>	0.968	0.018	0.014	9.189	0.3	<i>1st</i>	0.843	0.068	0.089	1.898
	<i>HP</i>	0.906	0.064	0.030	9.559		<i>3rd</i>	0.940	0.029	0.030	2.433
	<i>ACI₅</i>	0.955	0.031	0.014	12.906		<i>Bayes</i>	0.969	0.016	0.016	3.010
-0.6	<i>1st</i>	0.843	0.101	0.056	5.865	<i>HP</i>	0.906	0.041	0.053	2.444	
	<i>3rd</i>	0.939	0.034	0.027	6.644	<i>ACI₅</i>	0.953	0.021	0.026	3.300	
	<i>Bayes</i>	0.969	0.017	0.014	7.967	0.4	<i>1st</i>	0.845	0.062	0.094	1.657
	<i>HP</i>	0.906	0.062	0.031	7.785		<i>3rd</i>	0.944	0.027	0.029	2.156
	<i>ACI₅</i>	0.954	0.031	0.015	10.510		<i>Bayes</i>	0.972	0.014	0.014	2.674
-0.5	<i>1st</i>	0.839	0.100	0.061	5.078	<i>HP</i>	0.908	0.037	0.055	2.133	
	<i>3rd</i>	0.940	0.033	0.028	5.883	<i>ACI₅</i>	0.956	0.019	0.026	2.880	
	<i>Bayes</i>	0.969	0.017	0.014	7.096	0.5	<i>1st</i>	0.838	0.058	0.104	1.412
	<i>HP</i>	0.906	0.059	0.035	6.647		<i>3rd</i>	0.936	0.026	0.038	1.879
	<i>ACI₅</i>	0.954	0.029	0.016	8.974		<i>Bayes</i>	0.970	0.012	0.018	2.337
-0.4	<i>1st</i>	0.840	0.096	0.064	4.439	<i>HP</i>	0.905	0.033	0.062	1.813	
	<i>3rd</i>	0.937	0.034	0.029	5.206	<i>ACI₅</i>	0.953	0.015	0.033	2.448	
	<i>Bayes</i>	0.967	0.017	0.015	6.317	0.6	<i>1st</i>	0.841	0.056	0.104	1.217
	<i>HP</i>	0.905	0.058	0.038	5.807		<i>3rd</i>	0.942	0.026	0.033	1.652
	<i>ACI₅</i>	0.950	0.031	0.020	7.840		<i>Bayes</i>	0.971	0.013	0.016	2.061
-0.3	<i>1st</i>	0.845	0.089	0.067	3.892	<i>HP</i>	0.909	0.031	0.060	1.559	
	<i>3rd</i>	0.939	0.032	0.030	4.637	<i>ACI₅</i>	0.956	0.014	0.030	2.105	
	<i>Bayes</i>	0.969	0.016	0.015	5.650	0.7	<i>1st</i>	0.850	0.052	0.098	1.012
	<i>HP</i>	0.906	0.053	0.040	5.044		<i>3rd</i>	0.939	0.029	0.032	1.415
	<i>ACI₅</i>	0.954	0.027	0.020	6.810		<i>Bayes</i>	0.968	0.016	0.016	1.772
-0.2	<i>1st</i>	0.835	0.088	0.077	3.426	<i>HP</i>	0.909	0.031	0.060	1.290	
	<i>3rd</i>	0.938	0.030	0.032	4.143	<i>ACI₅</i>	0.955	0.015	0.030	1.741	
	<i>Bayes</i>	0.969	0.013	0.018	5.066	0.8	<i>1st</i>	0.849	0.045	0.105	0.794
	<i>HP</i>	0.903	0.050	0.047	4.417		<i>3rd</i>	0.939	0.027	0.034	1.160
	<i>ACI₅</i>	0.951	0.024	0.025	5.963		<i>Bayes</i>	0.968	0.014	0.018	1.459
-0.1	<i>1st</i>	0.842	0.082	0.076	3.055	<i>HP</i>	0.911	0.026	0.063	0.995	
	<i>3rd</i>	0.937	0.032	0.031	3.726	<i>ACI₅</i>	0.955	0.012	0.033	1.344	
	<i>Bayes</i>	0.969	0.014	0.017	4.569	0.9	<i>1st</i>	0.852	0.041	0.106	0.569
	<i>HP</i>	0.905	0.048	0.047	3.956		<i>3rd</i>	0.939	0.029	0.032	0.894
	<i>ACI₅</i>	0.951	0.025	0.024	5.341		<i>Bayes</i>	0.968	0.017	0.016	1.125
0	<i>1st</i>	0.833	0.083	0.084	2.762	<i>HP</i>	0.913	0.024	0.064	0.675	
	<i>3rd</i>	0.937	0.031	0.032	3.405	<i>ACI₅</i>	0.956	0.012	0.032	0.912	
	<i>Bayes</i>	0.969	0.015	0.015	4.186						

Table 6. Summary statistics for Simulation 2: bivariate Normal with means 7 and variances 0.9. Empirical coverage (*CP*), upper (*UE*) and lower (*LE*) error probability and average length (*AL*) of nominal two-sided 95% confidence intervals for γ , for varying values of ρ and sample size $n = 10$. Pivots used: likelihood root r (*1st*), modified likelihood root r^* (*3rd*); Bayesian modified likelihood root r_B^* (*Bayes*); expression (18) by [23] (*HP*); expression (20) by [7] (*ACI₅*). Based on 10,000 replicates; simulation error: ± 0.004 .

(a)						(b)					
ρ	Method	CP	UE	LE	AL	ρ	Method	CP	UE	LE	AL
-0.9	<i>1st</i>	0.914	0.055	0.031	7.005	0	<i>HP</i>	0.935	0.033	0.033	1.615
	<i>3rd</i>	0.948	0.026	0.026	7.427		<i>ACI₅</i>	0.947	0.027	0.026	1.736
	<i>Bayes</i>	0.960	0.020	0.020	7.903		0.1	<i>1st</i>	0.909	0.046	0.045
	<i>HP</i>	0.938	0.041	0.021	7.477	<i>3rd</i>		0.949	0.026	0.026	1.463
	<i>ACI₅</i>	0.952	0.032	0.016	8.037	<i>Bayes</i>		0.963	0.018	0.018	1.569
-0.8	<i>1st</i>	0.911	0.059	0.031	4.830	<i>HP</i>	0.934	0.033	0.033	1.465	
	<i>3rd</i>	0.949	0.027	0.025	5.128	<i>ACI₅</i>	0.950	0.024	0.026	1.575	
	<i>Bayes</i>	0.962	0.020	0.018	5.487	0.2	<i>1st</i>	0.912	0.040	0.048	1.178
	<i>HP</i>	0.934	0.044	0.022	5.137		<i>3rd</i>	0.950	0.025	0.025	1.304
	<i>ACI₅</i>	0.951	0.033	0.016	5.521		<i>Bayes</i>	0.960	0.020	0.020	1.398
-0.7	<i>1st</i>	0.917	0.053	0.031	3.749	<i>HP</i>	0.935	0.029	0.036	1.298	
	<i>3rd</i>	0.948	0.029	0.023	4.015	<i>ACI₅</i>	0.950	0.023	0.027	1.395	
	<i>Bayes</i>	0.959	0.022	0.019	4.304	0.3	<i>1st</i>	0.913	0.040	0.047	1.060
	<i>HP</i>	0.939	0.040	0.021	4.049		<i>3rd</i>	0.951	0.024	0.025	1.180
	<i>ACI₅</i>	0.949	0.033	0.018	4.352		<i>Bayes</i>	0.962	0.019	0.019	1.266
-0.6	<i>1st</i>	0.911	0.053	0.035	3.091	<i>HP</i>	0.936	0.029	0.036	1.168	
	<i>3rd</i>	0.948	0.026	0.026	3.318	<i>ACI₅</i>	0.952	0.021	0.027	1.255	
	<i>Bayes</i>	0.961	0.020	0.019	3.557	0.4	<i>1st</i>	0.914	0.042	0.044	0.941
	<i>HP</i>	0.936	0.038	0.026	3.371		<i>3rd</i>	0.947	0.030	0.024	1.055
	<i>ACI₅</i>	0.951	0.030	0.019	3.623		<i>Bayes</i>	0.961	0.022	0.016	1.132
-0.5	<i>1st</i>	0.910	0.051	0.038	2.647	<i>HP</i>	0.936	0.031	0.033	1.037	
	<i>3rd</i>	0.949	0.026	0.026	2.846	<i>ACI₅</i>	0.949	0.025	0.026	1.114	
	<i>Bayes</i>	0.963	0.018	0.019	3.050	0.5	<i>1st</i>	0.914	0.037	0.050	0.821
	<i>HP</i>	0.934	0.039	0.027	2.903		<i>3rd</i>	0.951	0.025	0.024	0.928
	<i>ACI₅</i>	0.951	0.029	0.020	3.121		<i>Bayes</i>	0.961	0.020	0.019	0.997
-0.4	<i>1st</i>	0.910	0.050	0.040	2.305	<i>HP</i>	0.937	0.026	0.037	0.905	
	<i>3rd</i>	0.948	0.026	0.026	2.487	<i>ACI₅</i>	0.952	0.021	0.028	0.972	
	<i>Bayes</i>	0.961	0.019	0.020	2.664	0.6	<i>1st</i>	0.915	0.032	0.053	0.703
	<i>HP</i>	0.935	0.036	0.029	2.534		<i>3rd</i>	0.950	0.024	0.026	0.803
	<i>ACI₅</i>	0.951	0.028	0.021	2.723		<i>Bayes</i>	0.963	0.018	0.020	0.863
-0.3	<i>1st</i>	0.917	0.049	0.034	2.059	<i>HP</i>	0.938	0.023	0.038	0.774	
	<i>3rd</i>	0.952	0.026	0.022	2.227	<i>ACI₅</i>	0.952	0.018	0.031	0.832	
	<i>Bayes</i>	0.962	0.020	0.018	2.386	0.7	<i>1st</i>	0.912	0.034	0.054	0.590
	<i>HP</i>	0.939	0.037	0.024	2.266		<i>3rd</i>	0.948	0.027	0.025	0.682
	<i>ACI₅</i>	0.951	0.028	0.020	2.435		<i>Bayes</i>	0.962	0.019	0.018	0.734
-0.2	<i>1st</i>	0.910	0.046	0.043	1.816	<i>HP</i>	0.933	0.025	0.042	0.649	
	<i>3rd</i>	0.947	0.025	0.028	1.972	<i>ACI₅</i>	0.951	0.018	0.031	0.698	
	<i>Bayes</i>	0.960	0.019	0.021	2.113	0.8	<i>1st</i>	0.918	0.029	0.054	0.465
	<i>HP</i>	0.933	0.034	0.033	2.000		<i>3rd</i>	0.952	0.023	0.025	0.551
	<i>ACI₅</i>	0.948	0.026	0.026	2.150		<i>Bayes</i>	0.962	0.019	0.019	0.593
-0.1	<i>1st</i>	0.912	0.047	0.041	1.637	<i>HP</i>	0.940	0.021	0.039	0.510	
	<i>3rd</i>	0.946	0.027	0.026	1.785	<i>ACI₅</i>	0.954	0.016	0.030	0.548	
	<i>Bayes</i>	0.958	0.021	0.021	1.913	0.9	<i>1st</i>	0.914	0.027	0.059	0.328
	<i>HP</i>	0.934	0.035	0.031	1.804		<i>3rd</i>	0.946	0.024	0.030	0.420
	<i>ACI₅</i>	0.947	0.028	0.025	1.938		<i>Bayes</i>	0.960	0.018	0.021	0.453
0	<i>1st</i>	0.914	0.044	0.042	1.466	<i>HP</i>	0.935	0.019	0.046	0.347	
	<i>3rd</i>	0.946	0.027	0.026	1.605	<i>ACI₅</i>	0.948	0.015	0.037	0.373	
	<i>Bayes</i>	0.959	0.021	0.020	1.720						

Table 7. Summary statistics for Simulation 3: bivariate Normal with means 0.7 and variances 0.1. Empirical coverage (CP), upper (UE) and lower (LE) error probability and average length (AL) of nominal two-sided 95% confidence intervals for γ , for varying values of ρ and sample size $n = 5$. Pivots used: likelihood root r ($1st$), modified likelihood root r^* ($3rd$); Bayesian modified likelihood root r_B^* ($Bayes$); expression (23) (ACI_3). Based on 10,000 replicates; simulation error: ± 0.004 .

(a)						(b)					
ρ	Method	CP	UE	LE	AL	ρ	Method	CP	UE	LE	AL
-0.9	1st	0.916	0.066	0.018	10.378	0	Bayes	0.971	0.015	0.014	3.409
	3rd	0.951	0.028	0.021	10.671		ACI ₃	0.900	0.061	0.039	2.381
	Bayes	0.972	0.015	0.013	12.069		0.1	1st	0.909	0.069	0.021
	ACI ₃	0.904	0.058	0.038	10.431	3rd		0.945	0.030	0.025	2.640
-0.8	1st	0.913	0.067	0.020	8.075	Bayes		0.971	0.014	0.014	3.106
	3rd	0.951	0.026	0.023	8.208	ACI ₃		0.900	0.060	0.040	2.175
	Bayes	0.972	0.014	0.014	9.445	0.2	1st	0.910	0.067	0.023	2.330
	ACI ₃	0.904	0.056	0.040	7.105		3rd	0.947	0.028	0.025	2.353
-0.7	1st	0.911	0.066	0.023	6.641		Bayes	0.972	0.013	0.016	2.769
	3rd	0.949	0.026	0.025	6.722		ACI ₃	0.900	0.059	0.041	1.939
	Bayes	0.972	0.013	0.015	7.809	0.3	1st	0.912	0.066	0.021	2.097
	ACI ₃	0.901	0.058	0.041	5.620		3rd	0.949	0.026	0.024	2.118
-0.6	1st	0.912	0.068	0.020	5.717		Bayes	0.973	0.013	0.014	2.493
	3rd	0.949	0.028	0.023	5.776		ACI ₃	0.902	0.058	0.040	1.746
	Bayes	0.973	0.014	0.013	6.748	0.4	1st	0.914	0.066	0.020	1.884
	ACI ₃	0.903	0.058	0.039	4.794		3rd	0.951	0.027	0.022	1.904
-0.5	1st	0.913	0.066	0.021	4.996		Bayes	0.974	0.014	0.012	2.240
	3rd	0.952	0.025	0.024	5.041		ACI ₃	0.906	0.058	0.036	1.571
	Bayes	0.973	0.013	0.014	5.912	0.5	1st	0.907	0.073	0.020	1.684
	ACI ₃	0.906	0.056	0.038	4.142		3rd	0.944	0.033	0.023	1.702
-0.4	1st	0.905	0.074	0.021	4.485		Bayes	0.970	0.017	0.013	2.003
	3rd	0.944	0.031	0.024	4.528		ACI ₃	0.895	0.065	0.040	1.405
	Bayes	0.970	0.017	0.013	5.314	0.6	1st	0.909	0.069	0.022	1.436
	ACI ₃	0.895	0.064	0.040	3.731		3rd	0.947	0.027	0.026	1.453
-0.3	1st	0.909	0.070	0.021	3.927		Bayes	0.972	0.014	0.014	1.710
	3rd	0.947	0.028	0.024	3.962		ACI ₃	0.897	0.061	0.042	1.198
	Bayes	0.973	0.014	0.012	4.656	0.7	1st	0.909	0.070	0.021	1.212
	ACI ₃	0.899	0.060	0.041	3.264		3rd	0.945	0.031	0.024	1.228
-0.2	1st	0.913	0.066	0.021	3.532		Bayes	0.973	0.014	0.013	1.446
	3rd	0.950	0.028	0.022	3.566		ACI ₃	0.898	0.063	0.039	1.008
	Bayes	0.971	0.015	0.014	4.194	0.8	1st	0.908	0.073	0.019	0.969
	ACI ₃	0.904	0.057	0.038	2.942		3rd	0.947	0.030	0.022	0.986
-0.1	1st	0.912	0.068	0.020	3.216		Bayes	0.971	0.016	0.013	1.162
	3rd	0.950	0.028	0.022	3.247		ACI ₃	0.897	0.064	0.040	0.801
	Bayes	0.974	0.014	0.012	3.819	0.9	1st	0.914	0.066	0.020	0.686
	ACI ₃	0.903	0.060	0.037	2.675		3rd	0.949	0.028	0.023	0.705
0	1st	0.910	0.070	0.021	2.870		Bayes	0.973	0.014	0.013	0.831
	3rd	0.949	0.028	0.022	2.898		ACI ₃	0.905	0.057	0.038	0.548

Table 8. Summary statistics for Simulation 3: bivariate Normal with means 0.7 and variances 0.1. Empirical coverage (CP), upper (UE) and lower (LE) error probability and average length (AL) of nominal two-sided 95% confidence intervals for γ , for varying values of ρ and sample size $n = 10$. Pivots used: likelihood root r ($1st$), modified likelihood root r^* ($3rd$); Bayesian modified likelihood root r_B^* ($Bayes$); expression (23) (ACI_3). Based on 10,000 replicates; simulation error: ± 0.004 .

(a)						(b)					
ρ	Method	CP	UE	LE	AL	ρ	Method	CP	UE	LE	AL
-0.9	1st	0.935	0.045	0.020	6.952	0	Bayes	0.959	0.021	0.021	1.644
	3rd	0.953	0.024	0.023	6.971		ACI ₃	0.927	0.039	0.034	1.413
	Bayes	0.963	0.019	0.018	7.391		0.1	1st	0.932	0.046	0.022
	ACI ₃	0.930	0.038	0.032	6.154	3rd		0.948	0.026	0.026	1.395
-0.8	1st	0.935	0.043	0.022	4.772	Bayes		0.959	0.020	0.021	1.486
	3rd	0.951	0.024	0.025	4.768	ACI ₃		0.926	0.039	0.035	1.278
	Bayes	0.962	0.018	0.020	5.084	0.2	1st	0.934	0.045	0.021	1.259
	ACI ₃	0.931	0.036	0.033	4.215		3rd	0.949	0.027	0.024	1.262
-0.7	1st	0.933	0.046	0.020	3.766		Bayes	0.960	0.021	0.019	1.344
	3rd	0.950	0.026	0.024	3.766		ACI ₃	0.928	0.038	0.034	1.155
	Bayes	0.962	0.020	0.019	4.018	0.3	1st	0.935	0.046	0.019	1.131
	ACI ₃	0.929	0.039	0.032	3.387		3rd	0.951	0.027	0.022	1.133
-0.6	1st	0.931	0.047	0.021	3.101		Bayes	0.963	0.020	0.018	1.207
	3rd	0.947	0.028	0.025	3.104		ACI ₃	0.930	0.038	0.031	1.038
	Bayes	0.960	0.020	0.020	3.310	0.4	1st	0.930	0.049	0.021	1.011
	ACI ₃	0.926	0.040	0.034	2.820		3rd	0.947	0.029	0.024	1.013
-0.5	1st	0.934	0.048	0.018	2.693		Bayes	0.958	0.022	0.020	1.079
	3rd	0.950	0.028	0.022	2.697		ACI ₃	0.927	0.042	0.031	0.928
	Bayes	0.963	0.020	0.017	2.874	0.5	1st	0.935	0.047	0.018	0.895
	ACI ₃	0.930	0.040	0.030	2.460		3rd	0.952	0.026	0.022	0.897
-0.4	1st	0.932	0.046	0.022	2.347		Bayes	0.963	0.020	0.017	0.955
	3rd	0.950	0.025	0.026	2.352		ACI ₃	0.929	0.041	0.030	0.821
	Bayes	0.961	0.018	0.020	2.505	0.6	1st	0.937	0.046	0.018	0.772
	ACI ₃	0.928	0.038	0.033	2.150		3rd	0.954	0.025	0.021	0.774
-0.3	1st	0.931	0.047	0.022	2.100		Bayes	0.964	0.019	0.017	0.824
	3rd	0.948	0.027	0.026	2.104		ACI ₃	0.933	0.039	0.028	0.709
	Bayes	0.959	0.021	0.020	2.241	0.7	1st	0.933	0.044	0.023	0.646
	ACI ₃	0.927	0.040	0.033	1.925		3rd	0.946	0.027	0.026	0.648
-0.2	1st	0.933	0.047	0.020	1.894		Bayes	0.959	0.020	0.020	0.690
	3rd	0.950	0.026	0.025	1.898		ACI ₃	0.927	0.038	0.035	0.593
	Bayes	0.962	0.020	0.019	2.021	0.8	1st	0.934	0.046	0.020	0.516
	ACI ₃	0.927	0.039	0.034	1.737		3rd	0.949	0.027	0.024	0.518
-0.1	1st	0.938	0.044	0.018	1.699		Bayes	0.961	0.020	0.019	0.552
	3rd	0.954	0.024	0.022	1.703		ACI ₃	0.930	0.039	0.030	0.472
	Bayes	0.966	0.017	0.017	1.814	0.9	1st	0.932	0.045	0.022	0.368
	ACI ₃	0.933	0.038	0.030	1.559		3rd	0.947	0.027	0.026	0.372
0	1st	0.930	0.048	0.022	1.540		Bayes	0.959	0.021	0.020	0.396
	3rd	0.948	0.028	0.025	1.544		ACI ₃	0.929	0.038	0.033	0.324