Bayesian Analysis (2011) **6**, Number 3, pp. 411–428

Recent advances on Bayesian inference for $P(X < Y)$

Laura Ventura* and Walter Racugno†

Abstract. We address the statistical problem of evaluating $R = P(X < Y)$, where X and Y are two independent random variables. Bayesian parametric inference is based on the marginal posterior density of R and has been widely discussed under various distributional assumptions on X and Y. This classical approach requires both elicitation of a prior on the complete parameter and numerical integration in order to derive the marginal distribution of R. In this paper, we discuss and apply recent advances in Bayesian inference based on higher-order asymptotics and on pseudo-likelihoods, and related matching priors, which allow one to perform accurate inference on the parameter of interest R only, even for small sample sizes. The proposed approach has the advantages of avoiding the elicitation on the nuisance parameters and the computation of multidimensional integrals. From a theoretical point of view, we show that the used prior is a strong matching prior. From an applied point of view, the accuracy of the proposed methodology is illustrated both by numerical studies and by real-life data concerning clinical studies.

Keywords: Asymptotic expansions, Frequentist coverage probability, Matching prior, Modified likelihood root, Modified profile likelihood, Nuisance parameter, ROC curve, Stochastic precedence, Stress-strength model, Tail area probability

1 Introduction

This contribution deals with parametric Bayesian inference on $R = P(X < Y)$, where X and Y are two independent random variables. In spite of its apparent simplicity, the topic of inference on $P(X < Y)$ - usually referred to as the stress-strength model - has obtained wide attention in the literature, including quality control, engineering statistics, reliability, medicine, psychology, biostatistics, stochastic precedence, and probabilistic mechanical design (see, e.g., Johnson, 1988, and Kotz, Lumelskii and Pensky, 2003, for a review). For example, in a reliability study, where Y is the strength of a system and X is the stress applied to the system, $(1 - R)$ measures the chance that the system fails. In a clinical study, an example of the application of R is given by treatment comparisons, where X and Y are the responses of a treatment and a control group, respectively, and $1-R$ measures the effectiveness of the treatment. Alternatively, for diagnostic tests used to distinguish between diseased and non-diseased patients, the area under the receiver operating characteristics (ROC) curve, based on the sensitivity and the complement to specificity at different cut-off points of the range of possible test values, is equal to R.

*Department of Statistics, University of Padova, Padova, Italy, mailto:ventura@stat.unipd.it
†Department of Mathematics, University of Cagliari, Cagliari, Italy, mailto:racugno@unica.it

 DOI:10.1214/11-BA616

Bayesian inference about R has been studied under both parametric and nonparametric assumptions; see, among others, Guttman and Papandonatos (1997), Ghosh and Sun (1998), Kotz, Lumelskii and Pensky (2003), Chen and Dunson (2004), Erkanli *et al.* (2006), Dunson and Peddada (2008), Hanson, Kottas and Branscum (2008), and references therein. Here, we focus on Bayesian parametric inference, which has been discussed under various distributional assumptions on $X \sim f_X(x|\theta_x)$ and $Y \sim f_Y(y|\theta_y)$, with $\theta = (\theta_x, \theta_y) \in \Theta \subseteq \mathbb{R}^d$, $d \geq 2$. The classical way to perform Bayesian inference is to derive the marginal posterior probability density function (pdf) of R, using transformation rules. This approach requires both the assumption of a prior pdf on the complete parameter θ and numerical integration.

In this paper we discuss and apply recent advances in Bayesian inference, based on higher-order asymptotics (see, e.g., Reid, 1995, 2003, Brazzale, Davison and Reid, 2007) and on pseudo-likelihoods and related matching priors (Ventura, Cabras and Racugno, 2009), to perform accurate inference on the parameter of interest R, even for small sample sizes. The proposed approach has the advantages of avoiding the elicitation on the nuisance parameters and the computation of multidimensional integrals. From a theoretical point of view, we show that there is a strong agreement between frequentist and Bayesian results, since the corresponding approximate tail areas are proved to be identical. From a computational point of view, the accuracy of the proposed methodology is illustrated both by numerical studies and by real-life data concerning results of two clinical studies.

The outline of the paper is as follows. In Section 2 we briefly review parametric Bayesian inference on R. In Section 3 we discuss recent advances on Bayesian inference on R. Section 4 illustrates numerical studies, when X and Y are both independent exponential or normal random variables. Moreover, we perform two applications to real-life data concerning clinical studies. Concluding remarks are given in Section 5.

2 Background on Bayesian inference

Let X and Y be independent random variables with cumulative distribution functions $F_X(x;\theta_x)$ and $F_Y(y;\theta_y)$, respectively, with $\theta_x \in \Theta_x \subseteq \mathbb{R}^{d_x}$ and $\theta_y \in \Theta_y \subseteq \mathbb{R}^{d_y}$, $d = d_x + d_y$. By definition, R can be evaluated as a function of the entire parameter $\theta = (\theta_x, \theta_y)$, through the relation

$$R = R(\theta) = P(X < Y) = \int F_X(t;\theta_x) f_Y(t;\theta_y)\, dt \ . \tag{1}$$

Theoretical expressions for R are available under several distributional assumptions both for X and Y (see, e.g., Kotz, Lumelskii and Pensky, 2003).

Let $x = (x_1, \ldots, x_{n_x})$ be a random sample of size n_x from X and let $y = (y_1, \ldots, y_{n_y})$ be a random sample of size n_y from Y. Let $\pi(\theta) = \pi(\theta_x, \theta_y)$ be a prior pdf on θ. Let $\pi(\theta|x, y) \propto \pi(\theta)\ L(\theta)$ be the posterior pdf of θ, where $L(\theta) = L(\theta; x, y)$ is the likelihood function for θ based on x and y.

Bayesian inference on R is based on the derivation of the posterior pdf of R, which can be obtained using a suitable one-to-one transformation of θ of the form $F : \theta \rightarrow (R, \lambda)$, with inverse $Q = F^{-1}$. Then, the joint posterior pdf of (R, λ) is given by $\pi(R, \lambda|x, y) = \pi(Q(R, \lambda)|x, y)\, |J_Q(R, \lambda)|$, where $|J_Q(R, \lambda)|$ is the Jacobian of the transformation Q, so that

$$\begin{aligned} \pi_R(R|x, y) &= \int \pi(Q(R, \lambda)|x, y)|J_Q(R, \lambda)|\, d\lambda \\ &= \int \pi(R, \lambda|x, y)\; d\lambda \;. \end{aligned} \tag{2}$$

The most common choices for the Bayes estimator of R are the mode or the expectation over (2). The posterior pdf can also be used for construction of Bayesian credible sets for R. For applications of (2) see, among others, Reiser and Guttman (1986, 1987), Ghosh and Sun (1998), and Guttman and Papandonatos (1997).

The Bayesian approach based on (2) may present some difficulties. First of all, it requires the elicitation of a prior on the complete parameter θ, which may be difficult both in the subjective and objective Bayesian context, in particular when d is large. Second, cumbersome numerical integration may be necessary in order to derive the marginal distribution of R. This latter difficulty can be avoided using higher-order asymptotics (see, e.g., Reid, 1995, 2003, and Brazzale, Davison and Reid, 2007), i.e., accurate approximations of a marginal posterior, which provide very precise inferences on a scalar parameter of interest even when the sample size is small. Frequentist inference on R based on higher-order asymptotics is discussed in Jiang and Wong (2008) and Cortese and Ventura (2009). We now present an asymptotic expansion of (2), and of the corresponding tail area probability, for Bayesian inference on R.

Let us denote by $\ell_p(R) = \log L(R, \hat{\lambda}_R)$ the profile log-likelihood for R, where $\hat{\lambda}_R$ is the constrained maximum likelihood estimate of λ given R. Moreover, let $(\hat{R}, \hat{\lambda})$ be the full maximum likelihood estimate, and let $j_p(R) = -\ell_p''(R)$ be the observed information corresponding to the profile log-likelihood. Standard results for partitioned matrices give $|j_p(R)| = |j(R, \hat{\lambda}_R)|/|j_{\lambda\lambda}(R, \hat{\lambda}_R)|$, where $j(R, \lambda)$ is the observed Fisher information from $\ell(R, \lambda) = \log L(R, \lambda)$ and $j_{\lambda\lambda}(R, \lambda)$ is the (λ, λ)-block of $j(R, \lambda)$.

The marginal posterior pdf (2) can be approximated by expanding $L(R, \lambda)$ as a function of λ about $\hat{\lambda}_R$ at the numerator and by using the Laplace approximation to the denominator (see, e.g., Tierney and Kadane, 1986, Reid, 1995). We have

$$\begin{aligned} \pi_R(R|x, y) &\doteq \pi_R^H(R|x, y) \\ &\propto |j_p(\hat{R})|^{1/2} \exp\{\ell_p(R) - \ell_p(\hat{R})\}\, \frac{|j_{\lambda\lambda}(\hat{R}, \hat{\lambda})|^{1/2}}{|j_{\lambda\lambda}(R, \hat{\lambda}_R)|^{1/2}}\, \frac{\pi(R, \hat{\lambda}_R)}{\pi(\hat{R}, \hat{\lambda})}\;, \end{aligned} \tag{3}$$

where the symbol "$\doteq$" indicates that the approximation to $\pi_R(R|x, y)$ is accurate to $O(n^{-3/2})$ (see Tierney and Kadane, 1986). Note that the approximation (3) depends on simple likelihood quantities evaluated at $(\hat{R}, \hat{\lambda})$ or at $(R, \hat{\lambda}_R)$.

The corresponding $O(n^{-3/2})$ approximation to the marginal posterior tail area probability is (DiCiccio and Martin, 1991)

$$\begin{aligned} \int_{-\infty}^{R_0} \pi_R(R|x,y)\,dR &\doteq \Phi(r_0) + \phi(r_0)\left(\frac{1}{r_0} - \frac{1}{q_0}\right) \\ &= \Phi\left(r_0 + \frac{1}{r_0}\log\frac{q_0}{r_0}\right) \\ &= \Phi\left(r_B^*\right) \ , \end{aligned} \tag{4}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal pdf and the standard normal distribution function, respectively, $r_0 = r_p(R_0)$ with $r_p = r_p(R) = \text{sign}(\hat{R} - R)[2(\ell_p(\hat{R}) - \ell_p(R))]^{1/2}$ is the likelihood root computed from the profile log-likelihood, $q_0 = q(R_0)$ with

$$q(R) = \ell_p'(R)\ |j_p(\hat{R})|^{-1/2}\ \frac{|j_{\lambda\lambda}(R,\hat{\lambda}_R)|^{1/2}}{|j_{\lambda\lambda}(\hat{R},\hat{\lambda})|^{1/2}}\ \frac{\pi(\hat{R},\hat{\lambda})}{\pi(R,\hat{\lambda}_R)}\ ,$$

and $r_B^* = r_B^*(R) = r_p + (1/r_p)\log(q/r_p)$. Formula (4) gives an explicit expression for the quantiles. Moreover, $s(R) = 1 - \Phi(r_B^*)$ gives the Bayesian survivor probability with third-order accuracy.

Following DiCiccio and Stern (1994), asymptotic highest posterior density (HPD) credible sets for R based on (3) can be derived as likelihood ratio type confidence regions based on the adjusted profile log-likelihood

$$\log \pi_R^H(R|x,y) = c + \ell_p(R) + B(R)\ , \tag{5}$$

where $B(R) = -(1/2)\log|j_{\lambda\lambda}(R,\hat{\lambda}_R)| + \log\pi(R,\hat{\lambda}_R)$ and c is a constant. In particular, an asymptotic HPD credible set for R arising from (5) is given by

$$\left\{R : 2\left(\log\pi_R^H(\bar{R}|x,y) - \log\pi_R^H(R|x,y)\right) \leq \chi^2_{1;1-\alpha}\right\}\ , \tag{6}$$

where $\bar{R}$ maximizes $\log\pi_R^H(R|x,y)$ and $\chi^2_{1;1-\alpha}$ is the $(1-\alpha)$-quantile of a chi-squared distribution with one degree of freedom.

The higher-order approximations (3) and (4) allow one to avoid numerical computations in order to obtain the marginal posterior of R, using standard likelihood quantities.

3 Inference based on pseudo-likelihoods

The aim of this section is to discuss some recent advances in Bayesian inference based on pseudo-likelihood functions, and related matching priors, to perform accurate inference on the parameter of interest R only (see Ventura, Cabras and Racugno, 2009). This approach presents the advantages of avoiding the elicitation on λ, as well as the computation of the integral in (2).

As in the frequentist approach, elimination of the nuisance parameter λ may be carried out using appropriate pseudo-likelihoods, functions of the parameter of interest only, with properties similar to those of a likelihood function. Examples of pseudo-likelihoods for a parameter of interest are the marginal, the conditional, the profile, and modifications thereof (see, e.g., Pace and Salvan, 1997, Chapters 4 and 7, Severini, 2000, Chapter 9). Here, we focus on the modified profile likelihood $L_{mp}(R)$ of Barndorff-Nielsen (1983) and on the related matching prior $\pi^*(R)$ (Ventura, Cabras and Racugno, 2009). Treating $L_{mp}(R)$ as a genuine likelihood, the posterior pdf

$$\pi^*(R|x,y) \propto \pi^*(R)\, L_{mp}(R) \tag{7}$$

can be obtained. Although this approach cannot always be considered orthodox in a Bayesian setting, the use of alternative likelihoods is nowadays widely shared, and several papers are devoted to Bayesian interpretation and applications of some well-known pseudo-likelihoods. See, among others, Monahan and Boos (1992), Bertolino and Racugno (1994), Fraser and Reid (1996), Severini (1999), Chang and Mukerjee (2006), Chang, Kim and Mukerjee (2009), Ventura, Cabras and Racugno (2009, 2010), Racugno, Salvan and Ventura (2010), Pauli, Racugno and Ventura (2011) and references therein.

A further advantage in using (7) instead of (2) is that $\pi^*(R|x,y)$ guarantees good frequentist coverages, that is the posterior probability limits are also frequentist limits in the sense that the frequentist coverage of the Bayesian intervals induced by the matching prior $\pi^*(R)$ is equal to the nominal value plus a remainder of order $O(n^{-1})$.

Assume that the minimal sufficient statistic for the model is a one-to-one function of $(\hat{R}, \hat{\lambda}, a)$, where a is an ancillary statistic, so that $\ell(R,\lambda;x,y) = \ell(R,\lambda;\hat{R},\hat{\lambda},a)$. The modified profile likelihood $L_{mp}(R)$ of Barndorff-Nielsen (1983) is

$$L_{mp}(R) = L_{mp}(R;x,y) = L_p(R)\ |j_{\lambda\lambda}(R,\hat{\lambda}_R)|^{-1/2}\ \left|\frac{\partial\hat{\lambda}_R}{\partial\hat{\lambda}}\right|^{-1}, \tag{8}$$

where $|\partial\hat{\lambda}_R/\partial\hat{\lambda}| = |\ell_{\lambda;\hat{\lambda}}(R,\hat{\lambda}_R)|/|j_{\lambda\lambda}(R,\hat{\lambda}_R)|$ involves the sample space derivatives $\ell_{\lambda;\hat{\lambda}}(R,\lambda) = \partial^2\ell(R,\lambda;\hat{R},\hat{\lambda},a)/(\partial\lambda\,\partial\hat{\lambda}^{\intercal})$. Calculation of sample space derivatives is straightforward only in special classes of models, notably exponential and group families. See Severini (2000, Section 9.5) for a review of approximate calculation of sample space derivatives, and Pace and Salvan (2006) for various second-order equivalent versions of $L_{mp}(R)$, which can be used in (7).

When considering the modified profile likelihood (8), the matching prior $\pi^*(R)$ is simply proportional to the square root of the inverse of the asymptotic variance of the maximum likelihood estimator of R (Ventura, Cabras and Racugno, 2009). In particular, the matching prior for R associated with (8) is

$$\pi^*(R) \propto i_{RR.\lambda}(R,\hat{\lambda}_R)^{1/2}\ , \tag{9}$$

where $i_{RR.\lambda}(R,\lambda) = i_{RR}(R,\lambda) - i_{R\lambda}(R,\lambda)i_{\lambda\lambda}(R,\lambda)^{-1}i_{\lambda R}(R,\lambda)$ is the partial information, with $i_{RR}(R,\lambda)$, $i_{R\lambda}(R,\lambda)$, $i_{\lambda\lambda}(R,\lambda)$, and $i_{\lambda R}(R,\lambda)$ blocks of the expected Fisher information $i(R, \lambda)$. The posterior pdf (7) thus becomes

$$
\begin{aligned}
\pi^*(R|x,y) &\propto L_p(R)\, |j_{\lambda\lambda}(R, \hat{\lambda}_R)|^{-1/2} \left| \frac{\partial \hat{\lambda}_R}{\partial \hat{\lambda}} \right|^{-1} i_{RR.\lambda}(R, \hat{\lambda}_R)^{1/2} \\
&\propto L_p(R)\, \frac{|j_{\lambda\lambda}(R, \hat{\lambda}_R)|^{1/2}}{|\ell_{\lambda;\hat{\lambda}}(R, \hat{\lambda}_R)|}\, i_{RR.\lambda}(R, \hat{\lambda}_R)^{1/2} \ . \qquad (10)
\end{aligned}
$$

Accurate tail probabilities are directly computable by direct integration of (10). Let us consider the posterior tail area probability

$$
\int_{-\infty}^{R_0} \pi^*(R|x,y)\, dR = \int_{-\infty}^{R_0} c\ j_p(\hat{R})^{1/2} \exp\left\{ \ell_p(R) - \ell_p(\hat{R}) \right\} \ \frac{\bar{\pi}(R, \hat{\lambda}_R)}{\bar{\pi}(\hat{R}, \hat{\lambda}_R)}\, dR \ , \qquad (11)
$$

with $\bar{\pi}(R, \hat{\lambda}_R) \propto i_{RR.\lambda}(R, \hat{\lambda}_R)^{1/2} |j_{\lambda\lambda}(R, \hat{\lambda}_R)|^{1/2} / |\ell_{\lambda;\hat{\lambda}}(R, \hat{\lambda}_R)|$ and c an arbitrary constant. Considering the change of variable in (11) from R to $r_p = r_p(R)$ (see, e.g., also Reid, 2003), whose Jacobian is $dr_p(R)/dR = \ell'_p(R)/r_p(R)$, we obtain

$$
\begin{aligned}
\int_{-\infty}^{R_0} \pi^*(r|x,y)\, dr &\doteq \int_{-\infty}^{r_0} c\ \exp\left\{ -\frac{1}{2} r_p^2 \right\} \left(\frac{r_p}{q^*} \right) dr_p \\
&= \int_{-\infty}^{r_0} c\ \phi(r_p) \left(\frac{r_p}{q^*} + 1 - 1 \right) dr_p \\
&= \Phi(r_0) + \int_{-\infty}^{r_0} c\ r_p \phi(r_p) \left(\frac{1}{q^*} - \frac{1}{r_p} \right) dr_p \\
&= \Phi(r_0) + \phi(r_0) \left(\frac{1}{r_0} - \frac{1}{q_0^*} \right) \\
&= \Phi \left(r_0 + \frac{1}{r_0} \log \frac{q_0^*}{r_0} \right) \\
&= \Phi \left(r_F^* \right) \ , \qquad (12)
\end{aligned}
$$

where $q_0^* = q^*(R_0)$, with

$$
q^*(R) = \frac{\ell'_p(R)}{j_p(\hat{R})^{1/2}} \frac{i_{RR.\lambda}(\hat{R}, \hat{\lambda})^{1/2}}{i_{RR.\lambda}(R, \hat{\lambda}_R)^{1/2}} \frac{|\ell_{\lambda;\hat{\lambda}}(R, \hat{\lambda}_R)|}{|j_{\lambda\lambda}(\hat{R}, \hat{\lambda})|^{1/2} |j_{\lambda\lambda}(R, \hat{\lambda}_R)|^{1/2}} \ , \qquad (13)
$$

and

$$
r_F^* = r_F^*(R) = r_p + \frac{1}{r_p} \log \frac{q^*}{r_p} \ . \qquad (14)
$$

It is worth noting that the statistic (14), with $q^*(R)$ given by (13), appears also in Barndorff-Nielsen and Chamberlin (1994) and is known as a modified likelihood root (see, e.g., Barndorff-Nielsen and Cox, 1994, and Severini, 2000, Chapter 7). A modified likelihood root, such as r_F^*, is a higher-order pivotal quantity, which allows one to obtain frequentist p-values, confidence limits and accurate point estimators. Since frequentist and Bayesian approximate tail areas are identical when using $\pi^*(R|x, y)$, there is an agreement between frequentist and Bayesian results. Following Fraser and Reid (2002), the prior $\pi^*(R)$ is thus a 'strong' matching prior, in the sense that a frequentist p-value coincides with a Bayesian posterior survivor probability.

Using the posterior distribution $\pi^*(R|x, y)$ and the corresponding (12), the HPD credible set $H(b_\alpha) = \{R : \log \pi^*(R|x, y) \geq b_\alpha\}$ for R is such that $H(b_\alpha) = (1 - \Phi(r_F^*(b_\alpha)))\{1 + O(n^{-3/2})\}$, where b_α is a given constant. In view of this,

$$H(z_{1-\alpha/2}) = \{R : |r_F^*(R)| \leq z_{1-\alpha/2}\} \tag{15}$$

is a HPD credible set for R such that $P_\pi(H(z_{1-\alpha/2})|X, Y) = 1 - \alpha + O(n^{-3/2})$, i.e., it has approximate frequentist validity $(1 - \alpha)$, where $P_\pi(\cdot)$ is the posterior probability measure. Note that (15) is also an accurate likelihood-based confidence interval for R with approximate level $(1-\alpha)$ (see, e.g., Barndorff-Nielsen and Cox, 1994, and Severini, 2000, Chap. 7). The strong matching prior $\pi^*(R)$ is thus also an HPD matching prior (see, e.g., Datta and Mukerjee, 2004) for R based on the modified profile likelihood.

Note also that from (12) the posterior mode of $\pi^*(R|x, y)$ can be computed as the solution in R of the estimating equation $r_F^*(R) = 0$. Then, the posterior mode coincides with the frequentist estimator defined as the zero-level confidence interval based on r_F^*, as explained in Skovgaard (1989). In particular, the solution of $r_F^*(R) = 0$ is a refinement of the maximum likelihood estimator $\hat{R}$, which improves its small sample properties, respecting the requirement of parameterisation equivariance (see Pace and Salvan, 1999, Giummolé and Ventura, 2002).

Note that these theoretical results hold in general when using the modified profile likelihood of Barndorff-Nielsen (1983), and the corresponding matching prior, in (7) for inference about a scalar parameter of interest.

4 Applications and numerical studies

In this section the proposed Bayesian procedures are illustrated for two real-life datasets with small sample sizes, concerning clinical studies. In particular, in the first example about anaplastic large cell lymphoma, it is assumed that X and Y both follow an exponential distribution; in the second example about abdominal aortic aneurysm measurements, X and Y are supposed to be independent normal variables. For discussions on these parametric assumptions in the Bayesian setting see Reiser and Guttman (1986), Ghosh and Sun (1998) and references therein.

The accuracy of the proposed methodology is also illustrated, in both the stress-strength models, by numerical studies which investigate the empirical coverages of Bayesian credible sets from $\pi^*(R|x, y)$, (2) and (3) and the finite-sample properties of their posterior modes. When computing (2) and (3) non-informative priors on θ are considered (Ghosh and Sun, 1998).

4.1 Exponential distribution

Data example. The dataset about anaplastic large cell lymphoma (ALCL), which is a rare cancer disease which affects both children and adults, is part of a retrospective study on the ALCL carried out by the Clinic of Pediatric Hematology Oncology (University of Padova, Italy). The aim of the study was to assess the role of the Hsp70 protein in association with the ALCL. Diseased patients seem to have higher Hsp70 levels than healthy subjects. Moreover, it is known that the presence of the Hsp70 protein can induce the development of pathological states, such as oncogenesis (see Mayer and Bukau, 2005), and that seems to limit the efficacy of the chemotherapy treatment in diseased patients. Thus, Hsp70 protein levels can be studied as a biomarker for detecting early ALCL lymphoma and, therefore, its effectiveness in diagnosing the disease can be evaluated by $R = P(X < Y)$.

The data consist of a small sample: 10 patients with ALCL lymphoma (cases) and 4 healthy subjects (controls). Hsp70 protein level was recorded on a continuous scale for each individual. According to extra-experimental information, two independent exponential random variables were assumed for the protein level in cancer patients and in non-diseased subjects, respectively. Results from a Kolmogorov-Smirnov test supported the choice of an exponential model assumption for these data.

Assume that X and Y are independent and exponentially distributed, with rates α and β, respectively. In this framework, the reliability parameter R can be written as

$$R = R(\alpha, \beta) = \frac{\alpha}{\alpha + \beta} \ .$$

This simple example can be easily extended to the Weibull distribution, which generalizes the exponential distribution by allowing increasing or decreasing failure rates (see Ghosh and Sun, 1998, Kundu and Gupta, 2006).

Both for classical and modern Bayesian inference on R, it is convenient to consider the one-to-one transformation $\theta = (R, \lambda)$, with $R = \alpha/(\alpha + \beta)$ the scalar parameter of interest and $\lambda = \alpha + \beta$ a nuisance parameter. Moreover, we assume the joint Jeffreys' prior $\pi(\alpha, \beta) \propto \alpha^{-1}\beta^{-1}$, or equivalently $\pi(R, \lambda) \propto R^{-1}(1 - R)^{-1}\lambda^{-1}$ (see Ghosh and Sun, 1998). Then, it can be shown that (see, for instance, Kotz, Lumelskii and Pensky, 2003, Chapter 2)

$$\pi_R(R|x, y) \propto R^{n_x - 1}(1 - R)^{n_y - 1}(1 - BR)^{-(n_x + n_y)} \ , \tag{16}$$

with $B = (n_y\bar{y} - n_x\bar{x})/(n_y\bar{y})$, where $\bar{x}$ and $\bar{y}$ denote the sample means.

Let us consider now the Laplace approximation (3). The profile likelihood for R is given by $L_p(R) = \hat{\lambda}_R^{(n_x+n_y)} R^{n_x}(1 - R)^{n_y}$, with $\hat{\lambda}_R = (n_x + n_y)\hat{\lambda}\bar{x}/(n_y(\bar{x} + \bar{y})(1 - BR))$, $\hat{R} = \bar{y}/(\bar{x}+\bar{y})$ and $\hat{\lambda} = (\bar{x}+\bar{y})/(\bar{x}\bar{y})$. Moreover, we have $j_{\lambda\lambda}(R, \lambda) = (n_x+n_y)/\lambda^2$. Then, simple calculations show that the higher-order approximation $\pi_R^H(R|x, y)$ coincides with $\pi_R(R|x, y)$.

Modern Bayesian inference about the parameter of interest R may be based on the modified profile likelihood $L_{mp}(R) = L_p(R)\hat{\lambda}_R^2(n_x + n_y)^{-1/2}/\hat{\lambda}$. Straightforward calculations show that the matching prior $\pi^*(R)$ is given by

$$\pi^*(R) \propto \frac{1}{R(1-R)} \ .$$

The corresponding posterior (7) is thus

$$\pi^*(R|x,y) \propto R^{n_x-1}(1-R)^{n_y-1}(1-BR)^{-(n_x+n_y)} \ , \tag{17}$$

and, in this example, it coincides with (16). To find a credible interval for R the one-to-one transformation $r = (1-R)/(1-BR)$ can be used, since r has a Beta(n_x, n_y) posterior pdf.

For the ALCL data, the two protein level samples have different means (equal to 0.23 and 1.44 in controls and cases, respectively), as observed in Figure 1 (left). Also the posterior pdf $\pi^*(R|x,y)$ and the normalized $L_{mp}(R)$ are reported in Figure 1 (right). Note that the normalized $L_{mp}(R)$ can be interpreted as a posterior distribution for R assuming a uniform prior in (7). The posterior mode from $\pi^*(R|x,y)$, that is the estimated probabilities that a cancer patient has a higher Hsp70 protein level than a healthy patient, is about 0.89, while the posterior mode from the normalized $L_{mp}(\psi)$ is about 0.86. Both these values suggest a sufficiently high effectiveness of the protein level in early detection of ALCL patients. The 0.95% credible sets for R from $\pi^*(R|x,y)$ and from the normalized $L_{mp}(R)$ are, respectively, (0.61,0.95) and (0.55,0.93). Inference based on $\pi^*(R|x,y)$ appears to be more concentrated and thus more accurate in estimating the accuracy of the protein level biomarker.

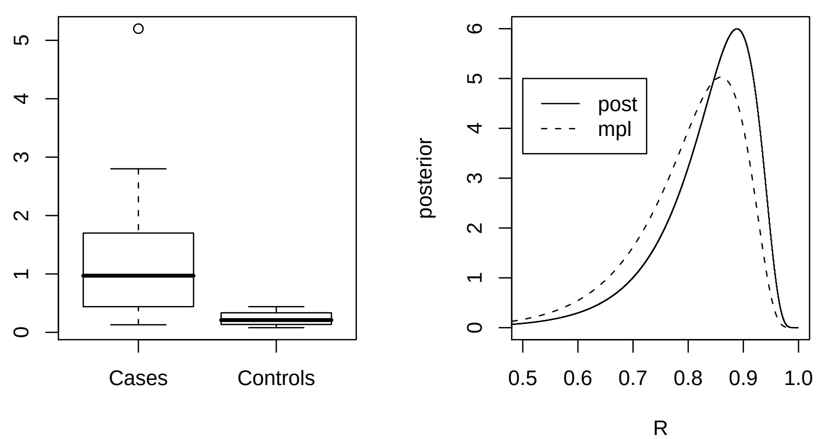

Figure 1: Hsp70 protein levels: (left) boxplot of cases and controls subjects; (right) $\pi^*(R|x,y)$ (solid) and normalized $L_{mp}(R)$ (dashed).

(n_x, n_y)		$R = 0.8$	$R = 0.9$	$R = 0.95$
(5,5)	$\pi^*(R\|x,y)$	0.952	0.949	0.949
		(0.023,0.024)	(0.024,0.025)	(0.026,0.024)
	$L_{mp}(R)$	0.941	0.944	0.943
		(0.028,0.025)	(0.028,0.027)	(0.027,0.029)
(10,10)	$\pi^*(R\|x,y)$	0.948	0.952	0.951
		(0.026,0.025)	(0.022,0.025)	(0.025,0.025)
	$L_{mp}(R)$	0.944	0.946	0.947
		(0.026,0.028)	(0.027,0.023)	(0.026,0.025)
(20,20)	$\pi^*(R\|x,y)$	0.949	0.949	0.950
		(0.023,0.026)	(0.026,0.025)	(0.026,0.025)
	$L_{mp}(R)$	0.949	0.947	0.946
		(0.027,0.024)	(0.026,0.026)	(0.026,0.026)
(30,30)	$\pi^*(R\|x,y)$	0.951	0.951	0.950
		(0.026,0.026)	(0.024,0.025)	(0.026,0.025)
	$L_{mp}(R)$	0.948	0.949	0.949
		(0.027,0.026)	(0.026,0.024)	(0.025,0.026)

Table 1: Frequentist coverage probabilities of approximate 0.95% HPD and of the lower and upper 0.025 quantiles (in brackets), under the exponential model.

Simulation study. The behaviour of (17) under the exponential model is illustrated through simulation studies, based on 10000 Monte Carlo trials. The numerical studies were carried out by fixing the parameter $\alpha = 1$ and determining β values so that $R = 0.8, 0.9, 0.95$, for different combinations of sample sizes (n_x, n_y).

Table 1 gives the empirical frequentist coverages for 95% asymptotic posterior HPD from $\pi^*(R|x,y)$, computed as in (6), and for the lower and upper 0.025 quantiles. For comparison, also the frequentist coverage probabilites from the normalized $L_{mp}(R)$ are given. From Table 1 we observe that, even for small (n_x, n_y), $\pi^*(R|x,y)$ has the correct frequentist coverages. Larger sample sizes $(n_x, n_y > 20)$ show, as one would expect, rather small differences between the results of the two procedures.

In order to compare the behaviour of the posterior pdf $\pi^*(R|x,y)$ with $L_{mp}(R)$, we evaluated the finite-sample properties of their posterior modes. The posterior modes are compared in terms of the usual centering and dispersion measures, i.e., bias and standard deviation. From Table 2 it can be noted that the mode of (17) exhibits a smaller bias than the maximum modified profile estimator. This result is due to the fact that the posterior mode of (17) is an r_F^*-based estimator, as explained in Section 3.

4.2 Normal distribution

Data example. The abdominal aortic aneurysm is a localized blood-filled dilation of the abdominal aorta. Accurate measurements of the diameter of the aneurysm are essential for screening and in assessing the seriousness of the disease. Surgical intervention is planned when the aneurysm diameter exceeds a certain threshold, often fixed at 5 cm, since it is known that the risk of aneurysm rupture increases as the size becomes

		$R = 0.8$		$R = 0.9$		$R = 0.95$	
(n_x, n_y)		bias	sd	bias	sd	bias	sd
(5,5)	$\pi^*(R\|x,y)$	0.012	(0.07)	0.010	(0.04)	0.006	(0.03)
	$L_{mp}(R)$	0.021	(0.07)	0.017	(0.04)	0.010	(0.04)
(10,10)	$\pi^*(R\|x,y)$	0.008	(0.07)	0.006	(0.02)	0.003	(0.02)
	$L_{mp}(R)$	0.010	(0.07)	0.008	(0.02)	0.005	(0.02)
(20,20)	$\pi^*(R\|x,y)$	0.004	(0.05)	0.003	(0.02)	0.001	(0.02)
	$L_{mp}(R)$	0.005	(0.05)	0.004	(0.02)	0.003	(0.02)
(30,30)	$\pi^*(R\|x,y)$	0.002	(0.04)	0.001	(0.02)	0.001	(0.01)
	$L_{mp}(R)$	0.003	(0.04)	0.002	(0.02)	0.001	(0.01)

Table 2: Bias (and standard deviations) of the posterior mode of $\pi^*(R|x,y)$ and of the maximum of $L_{mp}(R)$, under the exponential model.

larger. For decision making about interventions, it is thus important that the available measurement instruments are very accurate and provide the actual diameter values.

The aneurysm study considered two groups of $n_x = n_y = 10$ patients classified with low (L) and high (H) rupture risk, that is with small and large aneurysm diameter. The dataset consists of measurements of the aneurysm diameter on the two groups of patients obtained by a new instrument based on ultrasounds (US). The aim of the study was to evaluate the diagnostic accuracy of this new instrument in discriminating between patients with low and high rupture risk.

According to extra-experimental information, the US measurements can be assumed to be distributed in the two groups as normal variables with different means and equal variances (see, e.g., Nyhsen and Elliott, 2007, Azuma *et al.*, 2010). This is a typical setting commonly used in the literature on two sample comparisons, stress-strength models, and ROC curves. Moreover, this hypothesis was supported by the boxplots in Figure 2 (left) showing a similar variability for the two samples, and the choice of a normal model assumption for these data is supported by the Kolmogorov-Smirnov test.

Let us assume that X and Y are independent normal random variables with equal variances, that is $X \sim N(\mu_x, \sigma^2)$ and $Y \sim N(\mu_y, \sigma^2)$. In this situation, the entire parameter θ is given by $\theta = (\mu_x, \mu_y, \sigma^2)$, and the reliability parameter can be written as

$$R = R(\theta) = \Phi\left(-\frac{\mu_x - \mu_y}{\sigma\sqrt{2}}\right) , \tag{18}$$

which is one-to-one with $\eta = (\mu_x - \mu_y)/\sigma$ (see, for instance, Ghosh and Sun, 1998). This simple example can be easily extended to the situation of unequal variances, for which $R = \Phi\left(-(\mu_x - \mu_y)/\sqrt{\sigma_x^2 + \sigma_y^2}\right)$, or to include linear regression models by assuming that μ_x and μ_y depend on some covariates (Guttman and Papandonatos, 1997).

To perform both classical and modern Bayesian inference on (18), it is convenient to consider the one-to-one transformation $\theta = (\mu_x, \mu_y, \sigma^2) = (\eta, \lambda)$, with $\lambda = (\lambda_1, \lambda_2) = (\mu_y, \sigma)$. To compute the marginal posterior pdf for η we assume the one-at-a-time reference prior recommended by Ghosh and Sun (1998), given by

$$\pi(\eta,\lambda) \propto \frac{1}{\lambda_2}\frac{1}{(2(1+m^{-1})(1+m)+\eta^2)^{1/2}} \ , \tag{19}$$

with $m = n_x/n_y$. As pointed out in Ghosh and Sun (1998), the associated marginal posterior $\pi_\eta(\eta|x,y)$ is analytically intractable, as is the marginal posterior of R. To avoid this drawback, let us consider the higher-order approximation to $\pi_\eta(\eta|x,y)$, based on the profile log-likelihood

$$\ell_p(\eta) = -n\left(\log\tilde{\lambda}_2 + \frac{\hat{\lambda}_2^2}{2\tilde{\lambda}_2^2}\right) - \frac{1}{2\tilde{\lambda}_2^2}\left(n_y(\hat{\lambda}_1-\tilde{\lambda}_1)^2 + n_x(\hat{\lambda}_2\hat{\eta}+\hat{\lambda}_1-\tilde{\lambda}_2\eta-\tilde{\lambda}_1)^2\right) \ , \tag{20}$$

with $n = n_x + n_y$, where the constrained maximum likelihood estimate $\hat{\lambda}_\eta = (\tilde{\lambda}_1, \tilde{\lambda}_2)$ is obtained by numerical procedures. Since

$$|j_{\lambda,\lambda}(\eta,\hat{\lambda}_\eta)| = \frac{2n^2}{\tilde{\lambda}_2^4} - \frac{2n_x n\eta(\hat{\lambda}_2\hat{\eta}+\hat{\lambda}_1-\tilde{\lambda}_2\eta-\tilde{\lambda}_1)}{\tilde{\lambda}_2^5} + \frac{n_x n_y \eta^2}{\tilde{\lambda}_2^4} \ ,$$

we obtain

$$\pi_\eta^H(\eta|x,y) \propto \frac{\exp\{\ell_p(\eta)\}|j_{\lambda,\lambda}(\eta,\hat{\lambda}_\eta)|^{-1/2}}{\tilde{\lambda}_2(2(1+m^{-1})(1+m)+\eta^2)^{1/2}} \ . \tag{21}$$

Finally, since $\eta = \eta(R) = -\sqrt{2}\Phi^{-1}(R)$, the posterior pdf for R is

$$\pi_R^H(R|x,y) \propto \pi_\eta^H(\eta(R)|x,y)\ \frac{\sqrt{2}}{\phi(\Phi^{-1}(R))} \ . \tag{22}$$

Modern Bayesian inference about the parameter η may be based on the modified profile likelihood (8), with (20), (21) and

$$|\ell_{\lambda;\hat{\lambda}}(\eta,\hat{\lambda}_\eta)| = \frac{2n^2\hat{\lambda}_2}{\tilde{\lambda}_2^5} + \frac{2nn_x\hat{\eta}(\hat{\lambda}_2\hat{\eta}+\hat{\lambda}_1-\tilde{\lambda}_2\eta-\tilde{\lambda}_1)}{\tilde{\lambda}_2^5} + \frac{n_x n_y \eta\hat{\eta}}{\tilde{\lambda}_2^4} \ .$$

Straightforward calculations show that the matching prior $\pi^*(\eta)$ is given by

$$\pi^*(\eta) \propto \frac{1}{(2(1+m^{-1})(1+m)+\eta^2)^{1/2}} \ .$$

It may be noted that, as in the previous example, the proposed prior differs from the complete prior (19) because of the absence of the nuisance parameter. The corresponding posterior pdf for η is thus

$$\pi^*(\eta|x,y) \propto \frac{\exp\{\ell_p(\eta)\}|j_{\lambda,\lambda}(\eta,\hat{\lambda}_\eta)|^{1/2}|\ell_{\lambda;\hat{\lambda}}(\eta,\hat{\lambda}_\eta)|^{-1}}{(2(1+m^{-1})(1+m)+\eta^2)^{1/2}} \ . \tag{23}$$

The posterior pdf for R can be obtained from (23), giving

$$\pi^*(R|x,y) \propto \pi^*(\eta(R)|x,y)\ \frac{\sqrt{2}}{\phi(\Phi^{-1}(R))}\ . \tag{24}$$

For the abdominal aortic aneurysm data, the posterior modes from $\pi^*(R|x,y)$ and $\pi_R^H(R|x,y)$ are, respectively, 0.932 and 0.926 (see Figure 2, right). Both these values suggest a high accuracy of the US instrument in discriminating between patients with low and high rupture risk. The 0.95% HPD credible sets for R arising from $\pi^*(R|x,y)$ and $\pi_R^H(R|x,y)$ were found to be similar, i.e., (0.78,0.99) and (0.76,0.99), respectively. Also in this example, $\pi^*(R|x,y)$ gives accurate results.

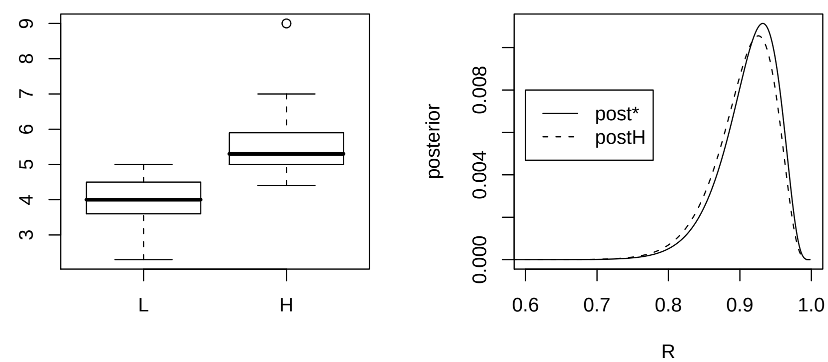

Figure 2: Boxplot of the sample distributions of L and H groups (left); $\pi^*(R|x,y)$ and $\pi_R^H(R|x,y)$ (right).

Simulation study. The accuracy of the proposed posterior $\pi^*(R|x,y)$ under the normal model is illustrated through simulation studies, based on 10000 Monte Carlo trials. The numerical studies were carried out by fixing the parameters (μ_x, μ_y, σ) so that $R = 0.8, 0.9, 0.95$, for different combinations of the sample sizes. The performance of $\pi^*(R|x,y)$ is compared with the higher-order approximation $\pi_R^H(R|x,y)$ and with the normalized modified profile likelihood $L_{mp}(R)$.

Table 3 reports the empirical frequentist coverages for 95% asymptotic posterior HPD credible sets and for the lower and upper 0.025 quantiles from $\pi^*(R|x,y)$, $\pi_R^H(R|x,y)$, and the normalized $L_{mp}(R)$. From Table 3 we observe that inference based on $\pi^*(R|x,y)$ is in general better than that from $\pi_R^H(R|x,y)$ and $L_{mp}(R)$. For $(n_x, n_y) > 20$ the coverages are quite comparable.

In order to compare the behaviour of the the posterior pdf $\pi^*(R|x,y)$ with $\pi_R^H(R|x,y)$ and with $L_{mp}(R)$, we evaluated the finite-sample properties of their posterior modes.

(n_x, n_y)		$R = 0.8$	$R = 0.9$	$R = 0.95$
(5,5)	$\pi_R^H(R\|x,y)$	0.935 (0.022,0.029)	0.934 (0.019,0.030)	0.935 (0.020,0.032)
	$\pi^*(R\|x,y)$	0.936 (0.024,0.027)	0.935 (0.025,0.026)	0.939 (0.025,0.026)
	$L_{mp}(R)$	0.930 (0.025,0.034)	0.934 (0.024,0.035)	0.935 (0.023,0.036)
(10,10)	$\pi_R^H(R\|x,y)$	0.940 (0.024,0.030)	0.940 (0.022,0.033)	0.939 (0.020,0.034)
	$\pi^*(R\|x,y)$	0.944 (0.026,0.029)	0.944 (0.026,0.026)	0.941 (0.024,0.026)
	$L_{mp}(R)$	0.942 (0.026,0.033)	0.943 (0.026,0.031)	0.940 (0.026,0.031)
(20,20)	$\pi_R^H(R\|x,y)$	0.944 (0.023,0.027)	0.943 (0.022,0.028)	0.943 (0.022,0.028)
	$\pi^*(R\|x,y)$	0.948 (0.024,0.025)	0.947 (0.026,0.025)	0.946 (0.027,0.025)
	$L_{mp}(R)$	0.942 (0.024,0.029)	0.945 (0.023,0.029)	0.942 (0.024,0.029)
(30,30)	$\pi_R^H(R\|x,y)$	0.948 (0.026,0.027)	0.948 (0.021,0.024)	0.947 (0.023,0.026)
	$\pi^*(R\|x,y)$	0.948 (0.024,0.024)	0.950 (0.025,0.025)	0.949 (0.025,0.026)
	$L_{mp}(R)$	0.947 (0.023,0.024)	0.949 (0.024,0.024)	0.948 (0.024,0.025)

Table 3: Frequentist coverage probabilities of approximate 0.95% HPD and of lower and upper 0.025 quantiles (in brackets), under the normal model.

As in the previous example, the posterior modes are compared in terms of the usual centering and dispersion measures, i.e., bias and standard deviation (see Table 4). It can be noted that the posterior mode of $\pi^*(R|x,y)$ improves on the mode of $\pi_R^H(R|x,y)$ and on the maximum modified profile estimator.

5 Discussion

We note that, in general, the computation of the proposed posterior $\pi^*(R|x,y)$, and of the associated inferential procedures, is simpler than the computation of $\pi(R|x,y)$. Indeed, to compute $\pi(R|x,y)$, elicitation on the nuisance parameters is required and a Markov-Chain Monte Carlo algorithm may be needed. Moreover, results from $\pi^*(R|x,y)$ appear quite accurate, since it has been shown that the matching prior $\pi^*(R)$ is also a strong matching prior.

On the basis of the simulation results discussed in Section 4, a natural question is what type of adjustment $\pi^*(R)$ makes to the modified profile likelihood $L_{mp}(R)$. This also raises the possibility that the use of matching priors is an alternate way to adjust the profile likelihood.

Finally, we observe that the method we discuss in this paper can be extended to more complex models and that different expressions for the modified profile likelihood,

		$R = 0.8$		$R = 0.9$		$R = 0.95$	
(n_x, n_y)		bias	sd	bias	sd	bias	sd
(5,5)	$\pi_R^H(R\|x,y)$	0.017	(0.13)	0.019	(0.09)	0.025	(0.06)
	$\pi^*(R\|x,y)$	0.006	(0.13)	0.009	(0.09)	0.010	(0.05)
	$L_{mp}(R)$	0.013	(0.13)	0.014	(0.09)	0.013	(0.05)
(10,10)	$\pi_R^H(R\|x,y)$	0.016	(0.10)	0.017	(0.06)	0.012	(0.04)
	$\pi^*(R\|x,y)$	0.003	(0.10)	0.003	(0.06)	0.004	(0.04)
	$L_{mp}(R)$	0.010	(0.10)	0.006	(0.06)	0.008	(0.04)
(20,20)	$\pi_R^H(R\|x,y)$	0.007	(0.06)	0.011	(0.04)	0.008	(0.03)
	$\pi^*(R\|x,y)$	0.000	(0.06)	0.002	(0.04)	0.002	(0.02)
	$L_{mp}(R)$	0.003	(0.06)	0.005	(0.04)	0.004	(0.02)
(30,30)	$\pi_R^H(R\|x,y)$	0.003	(0.05)	0.007	(0.03)	0.006	(0.02)
	$\pi^*(R\|x,y)$	0.001	(0.05)	0.001	(0.03)	0.001	(0.02)
	$L_{mp}(R)$	0.002	(0.05)	0.003	(0.03)	0.002	(0.02)

Table 4: Bias (and standard deviations) of the posterior modes of $\pi^*(R|x,y)$ and $\pi_R^H(R|x,y)$ and of the maximum of $L_{mp}(R)$, under the normal model.

which do not require the sample space derivatives, can be used. Moreover, our proposal might be extended to include linear regression models by assuming that the mean of X and Y depend on some covariates (see, for instance, Guttman *et al.*, 1988), or to truncated or censored data (see, for instance, Jiang and Wong, 2008). A final point concerns the extension of the problem to the partial area under the ROC curve, when only a restricted range of specificity values is of interest.

References

Azuma, J., L. Maegdefessel, T. Kitagawa, R. Lee Dalman and M.V. McConnell 2010. Assessment of elastase-induced murine abdominal aortic aneurysms: comparison of ultrasound imaging with in situ video microscopy. *Journal of Biomedicine and Biotechnology*, in press.

Barndorff-Nielsen, O.E. 1983. On a formula for the distribution of the maximum likelihood estimator. *Biometrika* 70: 343–365.

Barndorff-Nielsen, O.E. and S.R. Chamberlin 1994. Stable and invariant adjusted directed likelihoods. *Biometrika* 81: 485–499.

Barndorff-Nielsen, O.E. and D.R. Cox 1994. *Inference and Asymptotics*. London: Chapman and Hall.

Bertolino, F. and W. Racugno 1994. Robust Bayesian analysis of variance and the χ^2-test by using marginal likelihoods. *The Statistician* 43: 191–201.

Brazzale, A.R., A.C. Davison and N. Reid 2007. *Applied Asymptotics*. Cambridge: Cambridge University Press.

Chang, H. and R. Mukerjee 2006. Probability matching property of adjusted likelihoods. *Statistics and Probability Letters* 76: 838–842.

Chang, H., B.H. Kim and R. Mukerjee 2009. Bayesian and frequentist confidence intervals via adjusted likelihoods under prior specification on the interest parameter. *Statistics* 43: 203–211.

Chen, Z. and D.B. Dunson 2004. Bayesian estimation of survival functions under stochastic precedence. *Lifetime Data Analysis* 10: 159–173.

Cortese, G. and L. Ventura 2009. Accurate likelihood on the area under the ROC curve for small samples. *Working Papers*, 2009.17, Department of Statistics, University of Padova. Submitted.

Datta, G.S. and R. Mukerjee 2004. *Probability Matching Priors: Higher Order Asymptotics*. Berlin: Springer.

DiCiccio, T.J. and M.A. Martin 1991. Approximations of marginal tail probabilities for a class of smooth functions with applications to Bayesian and conditional inference. *Biometrika* 78: 891–902.

DiCiccio, T.J. and S.E. Stern 1994. Frequentist and Bayesian Bartlett corrections to test statistics based on adjusted profile likelihoods. *Journal of the Royal Statistical Society* B 56: 397–408.

Dunson, D.B. and S.D. Peddada 2008. Bayesian nonparametric inference on stochastic ordering. *Biometrika* 95: 859–874.

Erkanli, A., M. Sung, E.J. Costello and A. Angold 2006. Bayesian semi-parametric ROC analysis. *Statistics in Medicine* 25: 3905–3928.

Fraser, D.A.S. and N. Reid 1996. Bayes posteriors for scalar interest parameters. *Bayesian Statistics* 5: 581–585.

—-. 2002. Strong matching of frequentist and Bayesian parametric inference. *Journal of Statistical Planning and Inference* 103: 263–285.

Ghosh, M. and M.C. Yang 1996. Noninformative priors for the two sample normal problem. *Test* 5: 145–157.

Ghosh, M. and D. Sun 1998. Recent development of Bayesian inference for stress-strength models. In: *Frontier of Reliability*, 143–158. New Jersey: World Scientific.

Giummolé, F. and L. Ventura 2002. Practical point estimation from higher-order pivots. *Journal of Statistical Computation and Simulation* 72: 419–430.

Guttman, I., R.A. Johnson, G.K. Bhattacharyya and B. Reiser 1988. Confidence limits for stress-strength models with explanatory variables. *Technometrics* 30: 161–168.

Guttman, I. and G.D. Papandonatos 1997. A Bayesian approach to a reliability problem: theory, analysis and interesting numerics. *Canadian Journal of Statistics* 25: 143–158.

Hanson, T.E., A. Kottas and A.J. Branscum 2008. Modeling stochastic order in the analysis of receiver operating characteristic data: Bayesian non-parametric approaches. *Applied Statistics* 57: 207–225.

Jiang, L. and A.C.M. Wong 2008. A note on inference for $P(X < Y)$ for right truncated exponentially distributed data. *Statistical Papers* 49: 637–651.

Johnson, R.A. 1988. Stress-strength model for reliability. *Handbook of Statistics*, Krishnaiah P.R. and Rao C.R. eds., 7: 27–54. Amsterdam: Elsevier.

Kotz S., Y. Lumelskii and M. Pensky 2003. *The Stress-Strength Model and its Generalizations. Theory and Applications*. Singapore: World Scientific.

Kundu, D. and R.D. Gupta 2006. Estimation of $P(Y < X)$ for Weibull distribution. *IEEE Transactions in Reliability* 55: 270–280.

Mayer, M.P. and B. Bukau 2005. Hsp70 chaperones: cellular functions and molecular mechanism. *Cellular and Molecular Life Sciences* 62: 670–684.

Monahan, J.F. and D.D. Boos 1992. Proper likelihoods for Bayesian analysis. *Biometrika* 79: 271–278.

Nyhsen, C.M. and S.T. Elliott 2007. Rapid assessment of abdominal aortic aneurysms in 3-dimensional ultrasonography. *Journal of Ultrasound in Medicine* 26: 223–226.

Pace, L. and A. Salvan 1997. *Principles of Statistical Inference from a Neo-Fisherian Perspective*. Singapore: World Scientific.

—-. 1999. Point estimation based on confidence intervals: Exponential families. *Journal of Statistical Computation and Simulation* 64: 1–21.

—-. 2006. Adjustments of the profile likelihood from a new perspective. *Journal of Statistical Planning and Inference* 136: 3554–3564.

Pauli, F., W. Racugno and L. Ventura 2011. Bayesian composite marginal likelihoods. *Statistica Sinica*, 21: 149–164.

Racugno, W., A. Salvan and L. Ventura 2010. Bayesian analysis in regression models using pseudo-likelihoods. *Communications in Statistics - Theory and Methods* 39: 3444–3455.

Reid, N. 1995. Likelihood and Bayesian approximation methods. *Bayesian Statistics* 5: 351–368.

—-. 2003. The 2000 Wald memorial lectures: Asymptotics and the theory of inference. *Annals of Statistics* 31: 1695–1731.

Reiser, B. and I. Guttman 1986. Statistical inference for $Pr(Y < X)$: The normal case. *Technometrics* 28: 253–257.

—-. 1987. A comparison of three point estimators for $P(Y < X)$ in the normal case. *Computational Statistics and Data Analysis* 5: 59–66.

Severini, T.A. 1999. On the relationship between Bayesian and non-Bayesian elimination of nuisance parameters. *Statistica Sinica* 9: 713–724.

—-. 2000. *Likelihood Methods in Statistics*. Oxford: Oxford University Press.

Skovgaard, I.M. 1989. A review of higher order likelihood inference. *Bulletin of the International Statistical Institute* 53: 331–351.

Tierney, L.J. and J.B. Kadane 1986. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* 81: 82–86.

Ventura, L., S. Cabras and W. Racugno 2009. Prior distributions from pseudo-likelihoods in the presence of nuisance parameters. *Journal of the American Statistical Association* 104: 768–774.

—-. 2010. Default prior distributions from quasi- and quasi-profile likelihoods. *Journal of Statistical Planning and Inference* 140: 2937–2942.

Acknowledgments

This work was supported in part by grants from Ministero dell'Università e della Ricerca Scientifica e Tecnologica, Italy. The authors acknowledge the associate editor and anonymous referees for the useful comments that greatly improved the paper.