The Canadian Journal of Statistics
Vol. 25, No. 1, 1997, Pages 91–99
La Revue Canadienne de Statistique

Simple and accurate inference for the mean of the gamma model*

D.A.S. FRASER, N. REID, and A. WONG

University of Toronto (D.A.S.F., N.R.) and *York University* (A.W.)

Key words and phrases: Asymptotic approximation, averaging, canonical parameter, conditional likelihood, exponential model, location model, saddlepoint approximation, significance function, third-order approximation.
AMS 1991 subject classifications: 62F03, 62E20.

ABSTRACT

The two-parameter gamma model is widely used in reliability, environmental, medical and other areas of statistics. It has a two-dimensional sufficient statistic, and a two-dimensional parameter which can be taken to describe shape and mean. This makes it closely comparable to the normal model, but it differs substantially in that the exact distribution for the minimal sufficient statistic is not available. Some recently developed asymptotic theory is used to derive an approximation for observed levels of significance and confidence intervals for the mean parameter of the model. The approximation is as easy to apply as first-order methods, and substantially more accurate.

RÉSUMÉ

Le modèle gamma à deux paramètres est communément utilisé dans les domaines de la fiabilité, des statistiques environnementales et médicales et dans d'autres domaines des statistiques. Ce modèle comporte une statistique bi-dimensionnelle suffisante et un paramètre bi-dimensionnel qui peut être pris pour décrire la forme et la moyenne. Cela le rend étroitement comparable au modèle normal, mais celui-là diffère substantiellement de celui-ci en cela que la distribution exacte de la statistique minimale suffisante n'est pas disponible. Nous utilisons de la théorie asymptotique récente pour dériver une approximation des niveaux de signification observés et des intervalles de confiance pour le paramètre de moyenne du modèle. L'approximation est aussi facile à appliquer que la méthode du premier ordre, et nettement plus précise.

1. INTRODUCTION

In this paper we consider application of an approximation derived in Fraser and Reid (1995) to the problem of constructing inference for the mean of a gamma distribution with unknown shape. The development in Fraser and Reid (1995) is substantially more general than is required here, and this obscures to some extent the simplicity of the method. Section 2 describes the general method and its application to the gamma mean problem, and some numerical comparisons with other approaches are provided in Section 3. The remainder of this section describes the gamma mean model and reviews several approaches suggested in the literature.

Consider a sample $(y_1, \ldots, y_n)$ from the two-parameter gamma model with shape β and mean μ. The joint density is

$$f(y_1, \ldots, y_n; \beta, \mu) = \Gamma^{-n}(\beta) \left(\frac{\beta}{\mu}\right)^{n\beta} \exp\left(\beta t_1 - \frac{\beta t_2}{\mu}\right) \prod \frac{1}{y_i}, \tag{1}$$

*This research was partially supported by the Natural Sciences and Engineering Research Council.

where $(t_1, t_2) = (\sum \log y_i, \sum y_i)$ is the minimal sufficient statistic for (β, μ). An alternative version of the minimal sufficient statistic for (β, μ) is given by (d, t_2), where d is the log offset of the arithmetic mean from the geometric mean,

$$d = \log \frac{\sum y_i}{n} - \log \left(\prod y_i\right)^{1/n} = \log \frac{t_2}{n} - \frac{t_1}{n}.$$

For fixed β, the density of $\log y_i$ has location model form, and it follows that the conditional density for t_2 given the residuals $(\log y_1 - \sum \log y_i/n, \ldots, \log y_n - \sum \log y_i/n)$ and thus given d is readily available. We then have that the joint density factors as $f(d, t_2; \beta, \mu) = f(d; \beta) f(t_2|d; \beta, \mu)$ into marginal and conditional components,

$$f(d;\, \beta) = \Gamma(n\beta)\Gamma^{-n}(\beta) n^{-(n\beta - \frac{1}{2})} \exp(-n\beta d)\, h_n(d), \tag{2a}$$

$$f(t_2|d;\, \beta,\, \mu) = \Gamma^{-1}(n\beta) \left(\frac{\beta}{\mu}\right)^{n\beta} \exp\left(n\beta \log t_2 - \frac{\beta t_2}{\mu}\right) t_2^{-1}, \tag{2b}$$

where $h_n(d)$ requires $(n-2)$-dimensional integration and is available exactly only for very small values of n. Approximating $h_n(d)$ is the focus of Jensen (1986), although he works with a function f_n, where $\sqrt{n}\, h_n(d) = f_n(d^{-1}) d^{-2}$. Note that the distribution of t_2 given d is free of d, so that t_2 and d are statistically independent. Also note that the joint distribution of (t_1, t_2) then has the form

$$f(t_1,\, t_2;\, \beta,\, \mu) = \Gamma^{-n}(\beta) \left(\frac{\beta}{\mu}\right)^{n\beta} \exp\left(\beta t_1 - \frac{\beta t_2}{\mu}\right) \frac{1}{t_2} \frac{1}{\sqrt{n}} h_n \left(\log \frac{t_2}{n} - \frac{t_1}{n}\right). \tag{3}$$

which is an exponential family model with canonical variable (t_1, t_2) and canonical parameter $(\beta, -\beta/\mu)$.

The log-likelihood function in terms of $\boldsymbol{\theta} = (\beta, \mu)$ is

$$l(\boldsymbol{\theta}) = l(\beta,\, \mu) = -n \log \Gamma(\beta) + n\beta \log \beta - n\beta \log \mu + \beta t_1 - \frac{\beta t_2}{\mu}. \tag{4}$$

For later reference, the first and negative second derivatives of $l(\boldsymbol{\theta})$ are

$$l_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \begin{pmatrix} l_\beta(\boldsymbol{\theta}) \\ l_\mu(\boldsymbol{\theta}) \end{pmatrix} = \begin{pmatrix} -ng(\beta) + n + n \log \beta - n \log \mu + t_1 - \dfrac{t_2}{\mu} \\ -\dfrac{n\beta}{\mu} + \dfrac{\beta t_2}{\mu^2} \end{pmatrix}, \tag{5}$$

$$\mathbf{j}_{\boldsymbol{\theta\theta}}(\boldsymbol{\theta}) = \begin{pmatrix} -l_{\beta\beta}(\boldsymbol{\theta}) & -l_{\beta\mu}(\boldsymbol{\theta}) \\ -l_{\mu\beta}(\boldsymbol{\theta}) & -l_{\mu\mu}(\boldsymbol{\theta}) \end{pmatrix} = \begin{pmatrix} ng'(\beta) - \dfrac{n}{\beta} & \dfrac{n}{\mu} - \dfrac{t_2}{\mu^2} \\ \dfrac{n}{\mu} - \dfrac{t_2}{\mu^2} & -\dfrac{n\beta}{\mu^2} + \dfrac{2\beta t_2}{\mu^3} \end{pmatrix}, \tag{6}$$

where $g(t) = d \log \Gamma(t)/dt$ is the digamma function; the subscripts $\boldsymbol{\theta}$ and $\boldsymbol{\theta\theta}$ denote first- and second-order differentiation with respect to $\boldsymbol{\theta}$.

Perhaps the most straightforward method for testing μ is based on the first-order normal approximation for the maximum-likelihood estimate $\hat{\boldsymbol{\theta}} = (\hat{\beta}, \hat{\mu})$ obtained from the equation $l_{\boldsymbol{\theta}}(\hat{\boldsymbol{\theta}}) = \mathbf{0}$: this is discussed in Gross and Clark (1975). The full maximum-likelihood estimate $\hat{\boldsymbol{\theta}}$ is asymptotically normal with mean $\boldsymbol{\theta}$ and variance $\mathbf{j}^{\boldsymbol{\theta\theta}}(\hat{\boldsymbol{\theta}}) = \mathbf{j}_{\boldsymbol{\theta\theta}}^{-1}(\hat{\boldsymbol{\theta}})$, so $\hat{\mu}$ is asymptotically normal with mean μ and variance estimated by $\hat{\mu}^2/n\hat{\beta}$. This method has first-order accuracy, meaning the error in the approximation is $O(n^{-\frac{1}{2}})$.

Several methods of inference can be obtained from the location model factorizaton (2) in which the simple gamma distribution $f(t_2|d;\beta,\mu) = f(t_2;\beta,\mu)$ can be viewed as containing all the information concerning μ; however, to reliably test a value μ, the effect of the nuisance parameter β needs to be accommodated. Grice and Bain (1980) approached this by first substituting the maximum-likelihood estimate $\hat{\beta}$ into (2b), thus treating $\hat{\beta}t_2/(n\mu) = \hat{\beta}\bar{y}/\mu$ as gamma$(n,\hat{\beta})$. They then used Monte Carlo simulations to adjust for the effect of this substitution. Shiue, Bain and Engelhardt (1988) reviewed this approximation and extended the method to obtain confidence intervals for the difference in means for two independent gamma models. Shiue and Bain (1990) discussed further the method of Grice and Bain (1980) and fine-tuned it using a Weibull approximation for the negative log of the probability integral transform of $\hat{\beta}\bar{y}/\mu$.

In addition to its use in applications, the two-parameter gamma is often taken as a test case for higher-order approximations. It is a relatively simple example of an exponential family model in which the parameter of interest is a ratio of canonical parameters, and for which exact calculations are not possible.

Wong (1993) used a parameter-averaging method described in Fraser and Wong (1993) to eliminate β from $f(t_2|d;\beta,\mu) = f(t_2;\beta,\mu)$: it is shown in Fraser and Wong (1996) that the averaging method yields second-order accuracy, i.e., with relative error $O(n^{-1})$. The averaging method has some similarities to Grice and Bain's (1980) method and also to the partially Bayes method of Cox (1975). Another second-order method based on Bartlett adjustment of the likelihood-ratio statistic is given in Jensen and Kristensen (1991).

Jensen (1986) obtained a third-order approximation, with relative error $O(n^{-\frac{3}{2}})$, by examining the conditional density of t_2 given $t_1 - t_2/\mu_0$, which is an exponential family with canonical parameter $\gamma = \beta(1/\mu_0 - 1/\mu)$: the nuisance parameter β is eliminated under the null hypothesis $\mu = \mu_0$. The explicit form of this conditional density is typically unavailable except for small n. By using saddlepoint methods for h_n and numerical integration, Jensen derived tables for constructing 95% and 98% confidence intervals for μ with sample sizes $n = 10$, 20, 40 and ∞. A different evaluation of the conditional density is needed for each value of μ_0, and confidence intervals are constructed iteratively. The method is of third-order accuracy, but the calculations for other than tabulated values are substantial. Jensen (1991) discussed another saddlepoint approximation which simplifies the calculation of the observed level of significance. This method is accurate only to second order, but has the advantage that the error is uniformly bounded in a large derivation region. Jensen (1992) examined large-deviation properties of some third-order tests.

Fraser and Reid (1995) used tangent exponential models to derive a fairly simple third-order approximation to the significance level for testing an interest parameter. An advantage of the method is that it needs only maximum-likelihood estimates and observed information for a recalibrated parameter. As we will see below, for inference about a ratio of canonical parameters in an exponential family model the approximation is the same as that derived by Barndorff-Nielsen (1986). The relationship between the two approaches in more general settings is the focus of as yet unpublished work by Wu, Reid and Fraser.

2. THIRD-ORDER INFERENCE FOR A COMPONENT PARAMETER

Consider a model with parameter $\boldsymbol{\theta} = (\lambda,\psi)$ (7) of dimension p, where ψ is a scalar parameter of interest. Many recently developed third-order methods for inference on ψ can be presented in terms of one or the other of the formulae

$$\Phi^1(R,Q) = \Phi(R) + \phi(R)(R^{-1} - Q^{-1}), \qquad \Phi^2(R,Q) = \Phi\left(R - R^{-1}\log\frac{R}{Q}\right), \quad (8)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal density and distribution functions, and R and Q are statistics that are equivalent to first order and whose limiting distribution is standard normal: their exact form depends on the type of problem. The two formulae were introduced respectively by Lugannani and Rice (1980) for a one-parameter exponential family model and by Barndorff-Nielsen (1986) for a reduced model after a sufficiency or ancillary reduction. As $\Phi(R,R)$, $\Phi_Q(R,R)$, $\Phi_{QQ}(R,R)$ are the same for the two definitions in (8), it follows easily that the approximations are third-order equivalent. Various expressions have been developed for R and Q appropriate to problems of varying generality; in most cases R is the signed square root of a log-likelihood ratio statistic. Use of the approximations has generally been hampered by the difficulty in calculating Q, which is typically some version of a standardized maximum likelihood departure.

A reduction by approximate ancillarity is usually needed to compute Q, although it is not needed for the gamma mean problem, since the sufficient statistic and the full parameter are of the same dimension. The third-order procedure developed by Fraser and Reid (1995, Section 5) gives a method for obtaining the conditionality reduction without explicit computation of the approximate ancillary statistic. A further reduction of dimension to a scalar pivotal quantity is necessary for testing the interest parameter ψ; this is outlined in Fraser and Reid (1995, Section 6). The resulting expression for Q is relatively simple in requiring only the observed likelihood function $l^0(\boldsymbol{\theta})$ and the likelihood gradient $l^0_{\mathbf{V}}(\boldsymbol{\theta})$ at the observed data point. In general problems the directions $\mathbf{V}$ of differentiation are determined using a first-derivative ancillary, but in the gamma mean problem we simply need to compute the likelihood gradient relative to the minimal sufficient statistic.

The computation of Q is obtained from a reparametrization that is determined by the gradient of the log-likelihood in the direction of the minimal sufficient statistic, evaluated at the observed value:

$$\boldsymbol{\varphi} = l_{;\mathbf{t}}(\boldsymbol{\theta};\mathbf{t}^0) = \frac{\partial}{\partial \mathbf{t}} l(\boldsymbol{\theta};\mathbf{t})\bigg|_{\mathbf{t}^0}. \tag{9}$$

We define the related Jacobians

$$\mathbf{J}_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \frac{\partial \boldsymbol{\varphi}}{\partial \boldsymbol{\theta}'} = l_{\boldsymbol{\theta};t}(\boldsymbol{\theta};\mathbf{t}^0), \qquad \mathbf{J}^{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\varphi}'} = \mathbf{J}_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\theta}), \tag{10}$$

so that the observed information can be recalibrated in the new parametrization

$$|\mathbf{j}_{(\boldsymbol{\theta\theta})}(\hat{\boldsymbol{\theta}})| = |\mathbf{j}_{\boldsymbol{\theta\theta}}(\hat{\boldsymbol{\theta}})||\mathbf{J}_{\boldsymbol{\theta}}(\hat{\boldsymbol{\theta}})|^2, \tag{11}$$

$$|\mathbf{j}_{(\boldsymbol{\lambda\lambda})}(\hat{\boldsymbol{\theta}}_\psi)| = |\mathbf{j}_{\boldsymbol{\lambda\lambda}}(\hat{\boldsymbol{\theta}}_\psi)||\mathbf{J}'_{\boldsymbol{\lambda}}(\hat{\boldsymbol{\theta}}_\psi)\mathbf{J}_{\boldsymbol{\lambda}}(\hat{\boldsymbol{\theta}}_\psi)|^{-2}, \tag{12}$$

where $\hat{\boldsymbol{\theta}}_\psi = (\psi, \hat{\boldsymbol{\lambda}}_\psi)$ is the restricted maximum-likelihood estimate, and $\mathbf{J}_{\boldsymbol{\lambda}}$ records the columns of $\mathbf{J}_{\boldsymbol{\theta}}$ associated with $\boldsymbol{\lambda}$. We also need the signed square root of the profile log-likelihood ratio statistic

$$R = \operatorname{sgn}(\hat{\psi} - \psi_0)\,[2\{l^0(\hat{\boldsymbol{\theta}}) - l^0(\hat{\boldsymbol{\theta}}_{\psi_0})\}]^{\frac{1}{2}} \tag{13}$$

and the standardized maximum likelihood departure

$$Q = (\hat{\chi} - \hat{\chi}_{\psi_0})\frac{|\mathbf{j}_{(\boldsymbol{\theta\theta})}(\hat{\boldsymbol{\theta}})|^{\frac{1}{2}}}{|\mathbf{j}_{(\boldsymbol{\lambda\lambda})}(\hat{\boldsymbol{\theta}}_{\psi_0})|^{\frac{1}{2}}} \tag{14}$$

for a scalar parameter χ that measures departure from the value ψ_0. The parameter χ is a scaled linear version of $\boldsymbol{\varphi}(\boldsymbol{\theta})$,

$$\chi = \frac{\mathbf{J}^{\psi}(\hat{\boldsymbol{\theta}}_{\psi_0})\boldsymbol{\varphi}(\boldsymbol{\theta})}{|\mathbf{J}^{\psi}(\hat{\boldsymbol{\theta}}_{\psi_0})|}, \tag{15}$$

that corresponds to $d\psi$ at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}_{\psi_0}$. The vector $\mathbf{J}^{\psi}$ is the row of $\mathbf{J}^{\boldsymbol{\theta}}$ that corresponds to ψ. The significance obtained from (8) is viewed as probability left of the data point, appropriately defined; for some discussion, see Fraser (1991). Significance corresponds to values near 0 or 1.

For testing $\mu = \mu_0$ in the gamma model (1), (2) or (3), the parameter of interest ψ is replaced by μ, and the nuisance parameter $\boldsymbol{\lambda}$ is replaced by β, in the formulae given above. The maximum-likelihood estimates $\hat{\boldsymbol{\theta}} = (\hat{\beta}, \hat{\mu})$, $\hat{\boldsymbol{\theta}}_{\mu_0} = (\hat{\beta}_{\mu_0}, \mu_0)$ are routinely obtained by solving $l_{\boldsymbol{\theta}}(\hat{\boldsymbol{\theta}}) = \mathbf{0}$ and $l_{\beta}(\hat{\beta}_{\mu_0}, \mu_0) = 0$ from (5). The observed information matrices $\mathbf{j}_{\boldsymbol{\theta\theta}}(\hat{\boldsymbol{\theta}})$ and $j_{\beta\beta}(\hat{\boldsymbol{\theta}}_{\mu_0})$ are obtained by substituting $\hat{\boldsymbol{\theta}}$ in (6) and $\hat{\boldsymbol{\theta}}_{\mu_0}$ in $j_{\beta\beta}(\boldsymbol{\theta})$. The new parametrization is the canonical parametrization $\boldsymbol{\varphi} = (\beta, -\beta/\mu)^{\mathrm{T}}$ with Jacobians

$$\mathbf{J}_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \begin{pmatrix} 1 & 0 \\ -1/\mu & \beta/\mu^2 \end{pmatrix}, \qquad \mathbf{J}^{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \begin{pmatrix} 1 & 0 \\ \mu/\beta & \mu^2/\beta \end{pmatrix}, \tag{16}$$

and the scalar version of $\boldsymbol{\varphi}$ specific to $\mu = \mu_0$ is

$$\chi = \frac{\left(\dfrac{\mu_0}{\hat{\beta}_{\mu_0}}, \dfrac{\mu_0^2}{\hat{\beta}_{\mu_0}}\right)\begin{pmatrix} \beta \\ -\beta/\mu \end{pmatrix}}{\left\{\left(\dfrac{\mu_0}{\hat{\beta}_{\mu_0}}\right)^2 + \left(\dfrac{\mu_0^2}{\hat{\beta}_{\mu_0}}\right)^2\right\}^{\frac{1}{2}}} = \frac{\beta - \beta\mu_0/\mu}{(1+\mu_0^2)^{\frac{1}{2}}}. \tag{17}$$

After some simplification the expression for Q is given by

$$Q = \sqrt{n}\hat{\beta}\left(\frac{\hat{\mu}}{\mu_0} - 1\right)\{g'(\hat{\beta}) - \hat{\beta}^{-1}\}^{\frac{1}{2}}\{g'(\hat{\beta}_{\mu_0}) - \hat{\beta}_{\mu_0}^{-1}\}^{-\frac{1}{2}}, \tag{18}$$

which is the same as the Q used in the approximation derived by Barndorff-Nielsen [1986, Equation (3.28)]. This equivalence provides an alternative derivation of the third order accuracy of (8), as well as an alternative route to Barndorff-Nielsen's r^*-approximation in (k, k) exponential families. The p-value or significance for testing μ_0 is then given by (8) using the signed log profile likelihood ratio (13) and the standardized maximum likelihood departure (18).

The extension to the comparison of several gamma means with the same shape parameter is straightforward: in the case of two means, for example, we could write $\psi = \mu_1 - \mu_2$, $\boldsymbol{\theta} = (\beta, \mu_2, \psi)^{\mathrm{T}}$, giving $\varphi = \left(\beta, -\beta/(\psi+\lambda), \beta/\lambda\right)^{\mathrm{T}}$.

The general case of inference for a ratio of canonical parameters in an exponential family is similarly easy to solve using Equations (9) through (15). Writing the model as

$$l(\boldsymbol{\theta}) = \frac{\lambda}{\psi}t_1 + \lambda t_2 - nc(\lambda, \psi),$$

we have

$$Q = \sqrt{n}\left(\frac{\hat{\psi}}{\psi_0} - 1\right)\hat{\psi}\{c_{\lambda\lambda}(\hat{\psi}, \hat{\lambda})c_{\psi\psi}(\hat{\psi}, \hat{\lambda}) - c_{\psi\lambda}^2(\hat{\psi}, \hat{\lambda})\}^{\frac{1}{2}}\{c_{\lambda\lambda}(\psi_0, \hat{\lambda}_{\psi_0})\}^{-\frac{1}{2}},$$

which again is the same expression as is obtained by following the prescription outlined in Barndorff-Nielsen (1986).

3. EXAMPLES

EXAMPLE 1. We first examine four of the approximations in the small-sample case with $n = 2$, where exact tail probabilities are readily computable. As an observed data point we use $(y_1^0, y_2^0) = (1, 4)$.

For the present sample size, the function $h_2(d)$ is available by marginalizing from (y_1, y_2) to the scalar variable $d = \log\{(y_1 + y_2)/2\} - \log(y_1 y_2)^{\frac{1}{2}}$. This gives

$$h_2(d) = 2\sqrt{2}(1 - e^{-2d})^{-\frac{1}{2}}.$$

The conditional relative density of t_2 given $t_1 - t_2/\mu_0$ is available exactly:

$$f_c(t_2 | t_1 - t/\mu_0;\, \beta,\, \gamma) = c\, \exp(\gamma t_2)\, (t_2^2 - 4e^{t_1})^{-\frac{1}{2}}.$$

The sample space for (t_2, t_1) is the region beneath and to the right of the curve $d = 0$ or $t_2^2 = 4e^{t_1}$ in Figure 1; the observed data point $(t_2^0, t_1^0) = (5, \log 4)$ is also plotted. The conditional distributions for testing $\mu = 1$, 3, 5, 7, and 9 are distributions along the lines correspondingly marked through the observed data point. These conditional distributions are recorded on the right side, in each case with the location of the data marked $*$; the left tail probabilities $p(1)$, $p(3)$, $p(5)$, $p(7)$, and $p(9)$ were obtained by numerical integration. It is clear from Figure 1 that the conditional distributions are very far from normal.

Table 1 records the significance $p(\mu)$ for $\mu = 1$, 3, 5, 7, and 9 obtained from the exact distribution and from six approximations: the first-order method, Shiue and Bain's (1990) Weibull approximation, the parameter-averaging method, Jensen's third-order method, and the new third-order methods using (13) and (14) with Φ^1 or Φ^2 in (8). The third-order methods are all of comparably good accuracy, with that based on Φ^2 perhaps slightly worse than Jensen's method and Φ^1. The averaging method also gives satisfactory results, while the first-order method and Shiue and Bain's (1990) Weibull approximation are not very good.

EXAMPLE 2. In order to examine smaller significance levels than those attained in Example 1, it seems appropriate to increase the sample size. Table 2 summarizes the results for 10 observations from a gamma with mean 1 and shape parameter $\beta = 2$. For 10,000 simulations, the observed p-values for testing $\mu = 1$ were computed by the various approximate methods, and the percentages of two-sided p-values less than the nominal 1%, 2.5%, and 5.0% values are recorded. The averaging method is satisfactory; the difference between the third-order approximations and the exact values are within the simulation error. Again the first-order method and Shiue and Bain's (1990) Weibull approximation are rather poor. These comparisons are unconditional: the first-order methods would likely be even worse if compared with the corresponding exact test. The p-value for the exact test can only be approximated, though: Barndorff-Nielsen (1986) uses Jensen's (1986) approximation as the nominally exact benchmark.

EXAMPLE 3. Grice and Bain (1980) discuss the Gross and Clark (1975) data on survival times on 20 mice exposed to 240 rads of gamma radiation:

152	152	115	109	137	88	94	77	160	165
125	40	128	123	136	101	62	153	83	69.

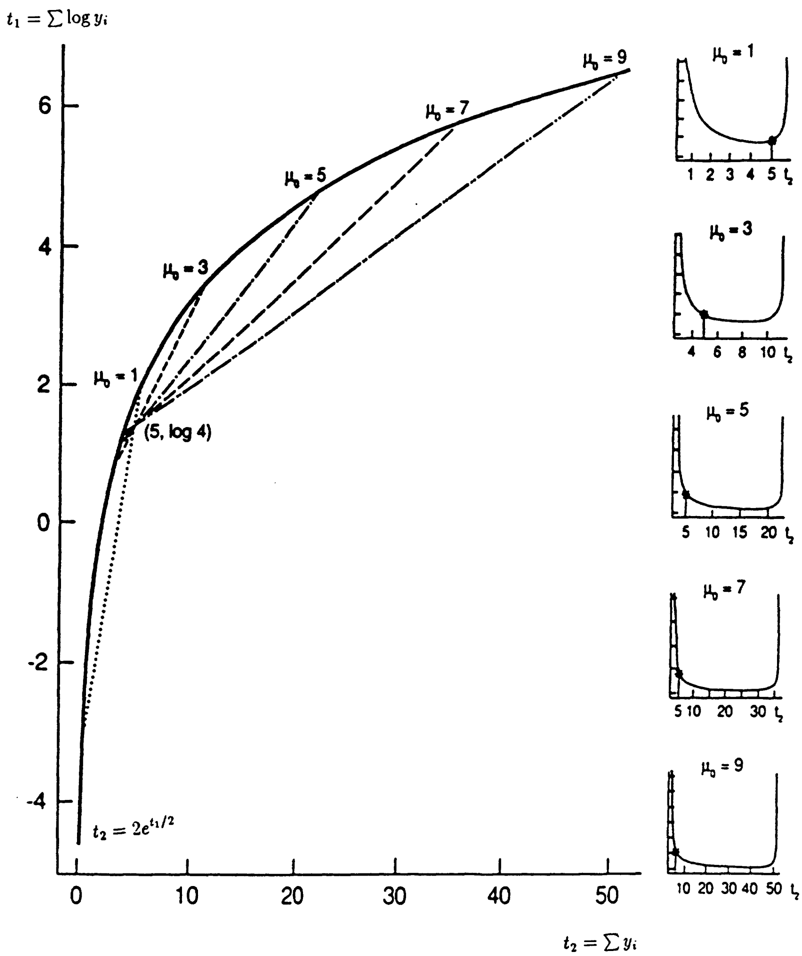

FIGURE 1: Sample space when $n = 2$; conditional distributions of t_2 for various μ_0 values.

TABLE 1: Approximate and exact p-values for a sample of size 2.

Method	$p(\mu)$				
	$\mu = 1$	3	5	7	9
First-order mle	0.905	0.331	0.014	4.09×10^{-5}	6.37×10^{-9}
Shiue and Bain	0.944	0.429	0.257	0.184	0.145
Averaging	0.892	0.512	0.330	0.280	0.257
Jensen	0.908	0.464	0.292	0.231	0.200
From (13),(14) with Φ^1	0.910	0.466	0.291	0.230	0.200
From (13),(14) with Φ^2	0.911	0.464	0.280	0.215	0.182
Exact	0.901	0.464	0.318	0.256	0.225

The 95% confidence intervals are recorded in Table 3 for six different approximation methods: the first-order method, Shiue and Bain's (1990) Weibull approximation, the method of parameter averaging, Jensen's third-order method, and the third-order methods from Equations (13) and (14) with Φ^1 and Φ^2 in (8). The three third-order methods are

TABLE 2: Results of 10,000 simulations with sample size 10.

Method	Percentage of p-values					
	< 0.5%	< 1.25%	< 2.5%	> 97.75%	> 98.75%	> 99.5%
First-order mle	5.00	5.73	8.53	1.52	0.74	0.35
Shiue and Bain	0.28	0.70	1.19	1.20	0.67	0.30
Averaging	0.29	0.69	2.18	2.05	0.95	0.41
From (13),(14) with Φ^1	0.37	0.90	2.30	2.41	1.13	0.47
From (13),(14) with Φ^2	0.37	0.91	2.30	2.41	1.13	0.47

TABLE 3: Confidence intervals for survival-time data.

Method	95% confidence interval for μ
First-order mle	(96.7, 130.2)
Shiue and Bain	(96.7, 130.8)
Averaging	(97.0, 134.7)
Jensen	(96.8, 133.5)
From (13),(14) with Φ^1	(97.2, 134.2)
From (13),(14) with Φ^2	(97.2, 134.2)

very close. In terms of ease of calculation, the two proposed methods using maximum-likelihood estimates and information are computationally simple.

ACKNOWLEDGEMENTS

We thank the referees and associate editors for helpful comments on an earlier version.

REFERENCES

Barndorff-Nielsen, O.E. (1986). Inference on full or partial parameters based on the standardized signed log-likelihood ratio. *Biometrika*, 73, 307–322.

Cox, D.R. (1975). A note on partially Bayes inference and linear models. *Biometrika*, 62, 651–654.

Fraser, D.A.S. (1991). Statistical inference: Likelihood to significance. *J. Amer. Statist. Assoc.*, 86, 258–265.

Fraser, D.A.S., and Reid, N. (1995). Ancillaries and third order significance. *Utilitas Math.*, 47, 33–53.

Fraser, D.A.S., and Wong, A. (1993). Approximate studentization with marginal and conditional inference. *Canad. J. Statist.*, 21, 212–320.

Fraser, D.A.S., and Wong, A. (1996). On the accuracy of approximation studentization. *Statist. Papers*. To appear.

Grice, J.V., and Bain, L.J. (1980). Inference concerning the mean of the gamma distribution. *J. Amer. Statist. Assoc.*, 75, 929–933.

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

Jensen, J.L. (1986). Inference for the mean of a gamma distribution with unknown shape parameter. *Scand. J. Statist.*, 13, 135–151.

Jensen, J.L. (1991). Uniform saddlepoint approximations and log-concave densities. *J. Roy. Statist. Soc. Ser. B*, 53, 157–172.

Jensen, J.L. (1992). The modified signed likelihood statistic and saddlepoint approximation. *Biometrika*, 79, 693–703.

Jensen, J.L., and Kristensen, L.B. (1991). Saddlepoint approximation to exact test and improved likelihood ratio tests for the gamma distribution. *Comm. Statist. Theory Methods*, 20, 1515–1532.

Lugannani, R., and Rice, S. (1980). Saddlepoint approximation for the distribution of the sum of independent random variables. *Adv. Appl. Probab.*, 12, 479–490.

Shiue, W.K., and Bain, L.J. (1990). Simple approximate inference procedures for the mean of the gamma model. *J. Statist. Comput. Simulation*, 34, 67–73.

Shiue, W.K., Bain, L.J., and Engelhardt, M. (1988). Test of equal gamma-distribution means with unknown and unequal shape parameters. *Technometrics*, 30, 169–174.

Wong, A. (1993). A note on inference for the mean parameter of the gamma distribution. *Statist. Probab. Lett.*, 17, 61–66.

Received 11 September 1994
Revised 9 September 1995
Accepted 5 July 1996

Department of Statistics
University of Toronto
Toronto, Ontario
Canada M5S 3G3

Department of Mathematics and Statistics
York University
North York, Ontario
Canada