

Likelihood Inference in the Presence of Nuisance Parameters

N. Reid, D.A.S. Fraser

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

We describe some recent approaches to likelihood based inference in the presence of nuisance parameters. Our approach is based on plotting the likelihood function and the p -value function, using recently developed third order approximations. Orthogonal parameters and adjustments to profile likelihood are also discussed. Connections to classical approaches of conditional and marginal inference are outlined.

1. INTRODUCTION

We take the view that the most effective form of inference is provided by the observed likelihood function along with the associated p -value function. In the case of a scalar parameter the likelihood function is simply proportional to the density function. The p -value function can be obtained exactly if there is a one-dimensional statistic that measures the parameter. If not, the p -value can be obtained to a high order of approximation using recently developed methods of likelihood asymptotics. In the presence of nuisance parameters, the likelihood function for a (one-dimensional) parameter of interest is obtained via an adjustment to the profile likelihood function. The p -value function is obtained from quantities computed from the likelihood function using a canonical parametrization $\varphi = \varphi(\theta)$, which is computed locally at the data point. This generalizes the method of eliminating nuisance parameters by conditioning or marginalizing to more general contexts. In Section 2 we give some background notation and introduce the notion of orthogonal parameters. In Section 3 we illustrate the p -value function approach in a simple model with no nuisance parameters. Profile likelihood and adjustments to profile likelihood are described in Section 4. Third order p -values for problems with nuisance parameters are described in Section 5. Section 6 describes the classical conditional and marginal likelihood approach.

2. NOTATION AND ORTHOGONAL PARAMETERS

We assume our measurement(s) y can be modelled as coming from a probability distribution with density or mass function $f(y; \theta)$, where $\theta = (\psi, \lambda)$ takes values in R^d . We assume ψ is a one-dimensional parameter of interest, and λ is a vector of nuisance parameters. If there is interest in more than one component of θ , the methods described here can be applied to each component of interest in turn. The likelihood function is

$$L(\theta) = L(\theta; y) = c(y)f(y; \theta); \quad (1)$$

it is defined only up to arbitrary multiples which may depend on y but not on θ . This ensures in particular that the likelihood function is invariant to one-to-one transformations of the measurement(s) y . In the context of independent, identically distributed sampling, where $y = (y_1, \dots, y_n)$ and each y_i follows the model $f(y; \theta)$ the likelihood function is proportional to $\prod f(y_i; \theta)$ and the log-likelihood function becomes a sum of independent and identically distributed components:

$$\ell(\theta) = \ell(\theta; y) = \sum \log f(y_i; \theta) + a(y). \quad (2)$$

The maximum likelihood estimate $\hat{\theta}$ is the value of θ at which the likelihood takes its maximum, and in regular models is defined by the score equation

$$\ell'(\hat{\theta}; y) = 0. \quad (3)$$

The observed Fisher information function $j(\theta)$ is the curvature of the log-likelihood:

$$j(\theta) = -\ell''(\theta) \quad (4)$$

and the expected Fisher information is the model quantity

$$i(\theta) = E\{-\ell''(\theta)\} = \int -\ell''(\theta; y)f(y; \theta)dy. \quad (5)$$

If y is a sample of size n then $i(\theta) = O(n)$.

In accord with the partitioning of θ we partition the observed and expected information matrices and use the notation

$$i(\theta) = \begin{pmatrix} i_{\psi\psi} & i_{\psi\lambda} \\ i_{\lambda\psi} & i_{\lambda\lambda} \end{pmatrix} \quad (6)$$

and

$$i^{-1}(\theta) = \begin{pmatrix} i^{\psi\psi} & i^{\psi\lambda} \\ i^{\lambda\psi} & i^{\lambda\lambda} \end{pmatrix}. \quad (7)$$

We say ψ is *orthogonal* to λ (with respect to expected Fisher information) if $i_{\psi\lambda}(\theta) = 0$. When ψ is scalar a transformation from (ψ, λ) to $(\psi, \eta(\psi, \lambda))$ such that ψ is orthogonal to η can always be found (Cox and Reid, [1]). The most directly interpreted consequence

of parameter orthogonality is that the maximum likelihood estimates of orthogonal components are asymptotically independent.

Example 1: ratio of Poisson means Suppose y_1 and y_2 are independent counts modelled as Poisson with mean λ and $\psi\lambda$, respectively. Then the likelihood function is

$$L(\psi, \lambda; y_1, y_2) = e^{-\lambda(1+\psi)} \psi^{y_2} \lambda^{y_1+y_2}$$

and ψ is orthogonal to $\eta(\psi, \lambda) = \lambda(\psi + 1)$. In fact in this example the likelihood function factors as $L_1(\psi)L_2(\eta)$, which is a stronger property than parameter orthogonality. The first factor is the likelihood for a binomial distribution with index $y_1 + y_2$ and probability of success $\psi/(1 + \psi)$, and the second is that for a Poisson distribution with mean η .

Example 2: exponential regression Suppose $y_i, i = 1, \dots, n$ are independent observations, each from an exponential distribution with mean $\lambda \exp(-\psi x_i)$, where x_i is known. The log-likelihood function is

$$\ell(\psi, \lambda; y) = -n \log \lambda + \psi \sum x_i - \lambda^{-1} \sum y_i \exp(\psi x_i) \quad (8)$$

and $i_{\psi\lambda}(\theta) = 0$ if and only if $\sum x_i = 0$. The stronger property of factorization of the likelihood does not hold.

3. LIKELIHOOD INFERENCE WITH NO NUISANCE PARAMETERS

We assume now that θ is one-dimensional. A plot of the log-likelihood function as a function of θ can quickly reveal irregularities in the model, such as a non-unique maximum, or a maximum on the boundary, and can also provide a visual guide to deviance from normality, as the log-likelihood function for a normal distribution is a parabola and hence symmetric about the maximum. In order to calibrate the log-likelihood function we can use the approximation

$$r(\theta) = \text{sign}(\hat{\theta} - \theta) [2\{\ell(\hat{\theta}) - \ell(\theta)\}]^{1/2} \sim N(0, 1), \quad (9)$$

which is equivalent to the result that twice the log likelihood ratio is approximately χ_1^2 . This will typically provide a better approximation than the asymptotically equivalent result that

$$\hat{\theta} - \theta \sim N(0, i^{-1}(\theta)) \quad (10)$$

as it partially accommodates the potential asymmetry in the log-likelihood function. These two approximations are sometimes called first order approximations because in the context where the log-likelihood is $O(n)$, we have (under regularity conditions) results such as

$$\Pr\{r(\theta; y) \leq r(\theta; y^0)\} = \Pr\{Z \leq r(\theta; y^0)\} \quad (11)$$

$$\{1 + O(n^{-1/2})\}$$

Table I The p -values for testing $\mu = 0$, i.e. that the number of observed events is consistent with the background.

upper p -value	0.0005993
lower p -value	0.0002170
mid p -value	0.0004081
$\Phi(r^*)$	0.0003779
$\Phi(r)$	0.0004416
$\Phi\{(\hat{\theta} - \theta)j^{1/2}\}$	0.0062427

where Z follows a standard normal distribution. It is relatively simple to improve the approximation to third order, i.e. with relative error $O(n^{-3/2})$, using the so-called r^* approximation

$$r^*(\theta) = r(\theta) + \{1/r(\theta)\} \log\{q(\theta)/r(\theta)\} \sim N(0, 1) \quad (12)$$

where $q(\theta)$ is a likelihood-based statistic and a generalization of the Wald statistic $(\hat{\theta} - \theta)j^{1/2}(\hat{\theta})$; see Fraser [2].

Example 3: truncated Poisson

Suppose that y follows a Poisson distribution with mean $\theta = b + \mu$, where b is a background rate that is assumed known. In this model the p -value function can be computed exactly simply by summing the Poisson probabilities. Because the Poisson distribution is discrete, the p -value could reasonably be defined as either

$$\Pr(y \leq y^0; \theta) \quad (13)$$

or

$$\Pr(y < y^0; \theta), \quad (14)$$

sometimes called the upper and lower p -values, respectively.

For the values $y^0 = 17$, $b = 6.7$, Figure 1 shows the likelihood function as a function of μ and the p -value function $p(\mu)$ computed using both the upper and lower p -values. In Figure 2 we plot the *mid* p -value, which is

$$\Pr(y < y^0) + (1/2)\Pr(y = y^0). \quad (15)$$

The approximation based on r^* is nearly identical to the mid- p -value; the difference cannot be seen on Figure 2. Table 1 compares the p -values at $\mu = 0$. This example is taken from Fraser, Reid and Wong [3].

4. PROFILE AND ADJUSTED PROFILE LIKELIHOOD FUNCTIONS

We now assume $\theta = (\psi, \lambda)$ and denote by $\hat{\lambda}_\psi$ the restricted maximum likelihood estimate obtained by

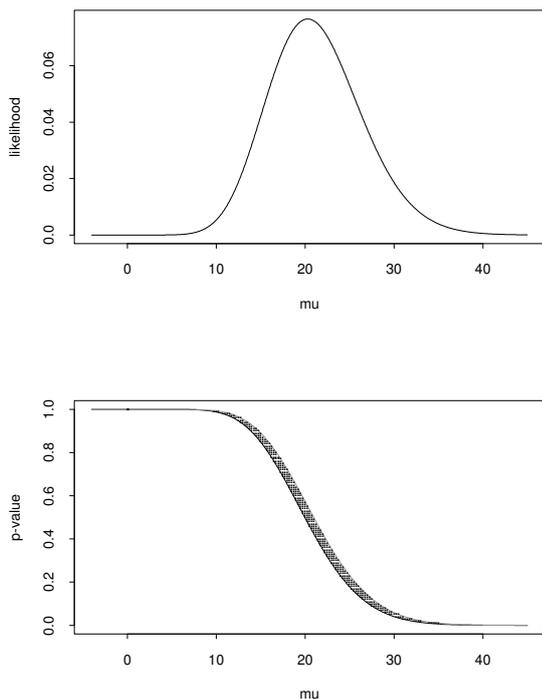


Figure 1: The likelihood function (top) and p -value function (bottom) for the Poisson model, with $b = 6.7$ and $y^0 = 17$. For $\mu = 0$ the p -value interval is $(0.99940, 0.99978)$.

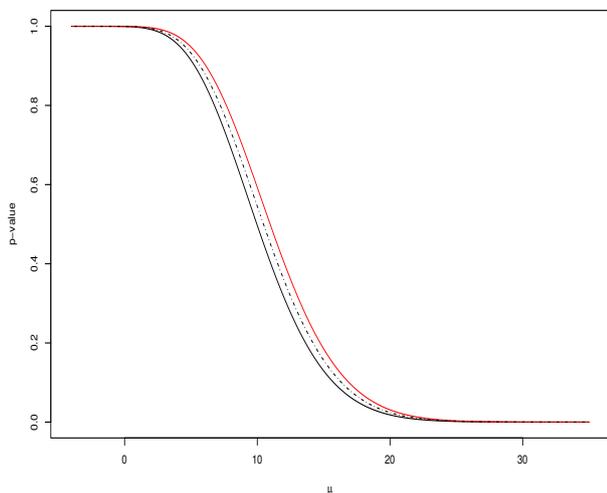


Figure 2: The upper and lower p -value functions and the mid- p -value function for the Poisson model, with $b = 6.7$ and $y^0 = 17$. The approximation based on $\Phi(r^*)$ is identical to the mid- p -value function to the drawing accuracy.

THAT001

maximizing the likelihood function over the nuisance parameter λ with ψ fixed. The profile likelihood function is

$$L_p(\psi) = L(\psi, \hat{\lambda}_\psi); \tag{16}$$

also sometimes called the concentrated likelihood or the peak likelihood. The approximations of the previous section generalize to

$$r(\psi) = \text{sign}(\hat{\psi} - \psi)[2\{\ell_p(\hat{\psi}) - \ell_p(\psi)\}]^{1/2} \sim N(0, 1), \tag{17}$$

and

$$\hat{\psi} - \psi \sim N(0, \{i^{\psi\psi}(\theta)\}^{-1}). \tag{18}$$

These approximations, like the ones in Section 3, are derived from asymptotic results which assume that $n \rightarrow \infty$, that we have a vector y of independent, identically distributed observations, and that the dimension of the nuisance parameter does not increase with n . Further regularity conditions are required on the model, such as are outlined in textbook treatments of the asymptotic theory of maximum likelihood. In finite samples these approximations can be misleading: profile likelihood is too concentrated, and can be maximized at the ‘wrong’ value.

Example 4: normal theory regression Suppose $y_i = x'_i\beta + \epsilon_i$, where $x_i = (x_{i1}, \dots, x_{ip})$ is a vector of known covariate values, β is an unknown parameter of length p , and ϵ_i is assumed to follow a $N(0, \psi)$ distribution. The maximum likelihood estimate of ψ is

$$\hat{\psi} = \frac{1}{n} \sum (y_i - x'_i\hat{\beta})^2 \tag{19}$$

which tends to be too small, as it does not allow for the fact that p unknown parameters (the components of β) have been estimated. In this example there is a simple improvement, based on the result that the likelihood function for (β, ψ) factors into

$$L_1(\beta, \psi; \bar{y})L_2\{\psi; \sum (y_i - x'_i\hat{\beta})^2\} \tag{20}$$

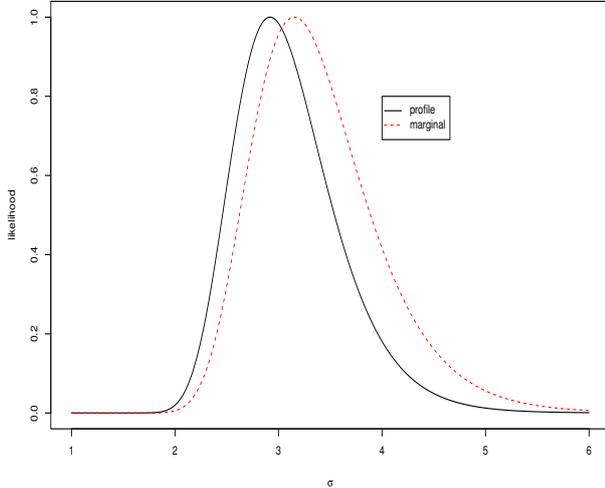
where $L_2(\psi)$ is proportional to the marginal distribution of $\sum (y_i - x'_i\hat{\beta})^2$. Figure 3 shows the profile likelihood and the marginal likelihood; it is easy to verify that the latter is maximized at

$$\hat{\psi}_m = \frac{1}{n-p} \sum (y_i - x'_i\hat{\beta})^2 \tag{21}$$

which in fact is an unbiased estimate of ψ .

Example 5: product of exponential means Suppose we have independent pairs of observations y_{1i}, y_{2i} , where $y_{1i} \sim \text{Exp}(\psi\lambda_i)$ $y_{2i} \sim \text{Exp}(\psi/\lambda_i)$, $i = 1, \dots, n$. The limiting normal theory for profile likelihood does not apply in this context, as the dimension

Figure 3: Profile likelihood and marginal likelihood for the variance parameter in a normal theory regression with 21 observations and three covariates (the "Stack Loss" data included in the Splus distribution). The profile likelihood is maximized at a smaller value of ψ , and is narrower; in this case both the estimate and its estimated standard error are too small.



of the parameter is not fixed but increasing with the sample size, and it can be shown that

$$\hat{\psi} \rightarrow \frac{\pi}{4}\psi \quad (22)$$

as $n \rightarrow \infty$ (Cox and Reid [4]).

The theory of higher order approximations can be used to derive a general improvement to the profile likelihood or log-likelihood function, which takes the form

$$\ell_a(\psi) = \ell_p(\psi) + \frac{1}{2} \log |j_{\lambda\lambda}(\psi, \hat{\lambda}_\psi)| + B(\psi) \quad (23)$$

where $j_{\lambda\lambda}$ is defined by the partitioning of the observed information function, and $B(\psi)$ is a further adjustment function that is $O_p(1)$. Several versions of $B(\psi)$ have been suggested in the statistical literature: we use the one defined in Fraser [5] given by

$$B(\psi) = -\frac{1}{2} \log |\varphi'_\lambda(\psi, \hat{\lambda}_\psi) j_{\varphi\varphi}(\hat{\psi}, \hat{\lambda}) \varphi'_\lambda(\psi, \hat{\lambda}_\psi)|. \quad (24)$$

This depends on a so-called canonical parametrization $\varphi = \varphi(\theta) = \ell_{;V}(\theta; y^0)$ which is discussed in Fraser, Reid and Wu [6] and Reid [7].

In the special case that ψ is orthogonal to the nuisance parameter λ a simplification of $\ell_a(\psi)$ is available as

$$\ell_{CR}(\psi) = \ell_p(\psi) - \frac{1}{2} \log |j_{\lambda\lambda}(\psi, \hat{\lambda}_\psi)| \quad (25)$$

which was first introduced in Cox and Reid (1987). The change of sign on $\log |j|$ comes from the orthogonality equations. In i.i.d. sampling, $\ell_p(\psi)$ is $O_p(n)$, i.e. is the sum of n bounded random variables, whereas $\log |j|$ is $O_p(1)$. A drawback of ℓ_{CR} is that it is not invariant to one-to-one reparametrizations of λ , all of which are orthogonal to ψ . In contrast $\ell_a(\psi)$ is invariant to transformations $\theta = (\psi, \lambda)$ to $\theta' = (\psi, \eta(\psi, \lambda))$, sometimes called interest-respecting transformations.

Example 5 continued In this example ψ is orthogonal to $\lambda = (\lambda_1, \dots, \lambda_n)$, and

$$\ell_{CR}(\psi) = -(3n/2) \log \psi - (2/\psi) \Sigma \sqrt{(y_{1i} y_{2i})}. \quad (26)$$

The value that maximizes ℓ_{CR} is 'more nearly consistent' than the maximum likelihood estimate as $\hat{\psi}_{CR} \rightarrow (\pi/3)\psi$.

5. P-VALUES FROM PROFILE LIKELIHOOD

The limiting theory for profile likelihood gives first order approximations to p -values, such as

$$p(\psi) \doteq \Phi(r_p) \quad (27)$$

and

$$p(\psi) \doteq \Phi\{(\hat{\psi} - \psi) j_p^{1/2}(\hat{\psi})\} \quad (28)$$

although the discussion in the previous section suggests these may not provide very accurate approximations. As in the scalar parameter case, though, a much better approximation is available using $\Phi(r^*)$ where

$$r^*(\psi) = r_p(\psi) + 1/\{r_p(\psi)\} \log\{Q(\psi)/r_p(\psi)\} \quad (29)$$

where Q can also be derived from the likelihood function and a function $\varphi(\theta, y^0)$ as

$$Q = (\hat{\nu} - \hat{\nu}_\psi) \hat{\sigma}_\nu^{-1/2}$$

where

$$\begin{aligned} \nu(\theta) &= e_{\hat{\psi}}^T \varphi(\theta), \\ e_{\hat{\psi}} &= \psi_{\varphi'}(\hat{\theta}_\psi) / |\psi_{\varphi'}(\hat{\theta}_\psi)|, \\ \hat{\sigma}_\nu^2 &= |j_{(\lambda\lambda)}(\hat{\theta}_\psi)| / |j_{(\theta\theta)}(\hat{\theta})|, \\ |j_{(\theta\theta)}(\hat{\theta})| &= |j_{\theta\theta}(\hat{\theta})| |\varphi_{\theta'}(\hat{\theta})|^{-2}, \\ |j_{(\lambda\lambda)}(\hat{\theta}_\psi)| &= |j_{\lambda\lambda}(\hat{\theta}_\psi)| |\varphi_{\lambda'}(\hat{\theta}_\psi)|^{-2}. \end{aligned}$$

The derivation is described in Fraser, Reid and Wu [6] and Reid [7]. The key ingredients are the log-likelihood function $\ell(\theta)$ and a reparametrization $\varphi(\theta) = \varphi(\theta; y^0)$, which is defined by using an approximating model at the observed data point y^0 ; this approximation in turn is based on a conditioning argument. A closely related approach is due to Barndorff-Nielsen; see Barndorff-Nielsen and Cox [8, Ch. 7], and the two approaches are compared in [7].

Table II Employment of men and women at the Space Telescope Science Institute, 1998–2002 (from *Science* magazine, Volume 299, page 993, 14 February 2003).

	Left	Stayed	Total
Men	1	18	19
Women	5	2	7
Total	6	20	26

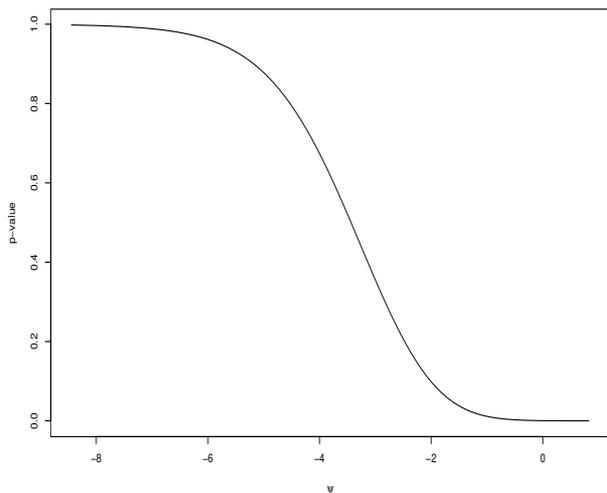


Figure 4: The p -value function for the log-odds ratio, ψ , for the data of Table II. The value $\psi = 0$ corresponds to the hypothesis that the probabilities of leaving are equal for men and women.

Example 6: comparing two binomials Table 2 shows the employment history of men and women at the Space Telescope Science Institute, as reported in *Science* Feb 14 2003. We denote by y_1 the number of males who left and model this as a Binomial with sample size 19 and probability p_1 ; similarly the number of females who left, y_2 , is modelled as Binomial with sample size 7 and probability p_2 . We write the parameter of interest

$$\psi = \log \frac{p_1(1-p_2)}{p_2(1-p_1)}. \quad (30)$$

The hypothesis of interest is $p_1 = p_2$, or $\psi = 0$. The p -value function for ψ is plotted in Figure 4. The p -value at $\psi = 0$ is 0.00028 using the normal approximation to r_p , and is 0.00048 using the normal approximation to r^* . Using Fisher's exact test gives a mid p -value of 0.00090, so the approximations are anticonservative in this case.

THAT001

Example 7: Poisson with estimated background Suppose in the context of Example 3 that we allow for imprecision in the background, replacing b by an unknown parameter β with estimated value $\hat{\beta}$. We assume that the background estimate is obtained from a Poisson count x , which has mean $k\beta$, and the signal measurement is an independent Poisson count, y , with mean $\beta + \mu$. We have $\hat{\beta} = x/k$ and $\text{var}\hat{\beta} = \beta/k$, so the estimated precision of the background gives us a value for k . For example, if the background is estimated to be 6.7 ± 2.1 this implies a value for k of $6.7/(2.1)^2 \doteq 1.5$. Uncertainty in the standard error of the background is ignored here. We now outline the steps in the computation of the r^* approximation (29).

The log-likelihood function based on the two independent observations x and y is

$$\ell(\beta, \mu) = x \log(k\beta) - k\beta + y \log(\beta + \mu) - \beta - \mu \quad (31)$$

with canonical parameter $\varphi = (\log \beta, \log(\beta + \mu))'$.

Then

$$\varphi_{\theta'}(\theta) = \frac{\partial \varphi(\theta)}{\partial \theta'} = \begin{pmatrix} 0 & 1/\beta \\ 1/(\beta + \mu) & 1/(\beta + \mu) \end{pmatrix}, \quad (32)$$

$$\varphi_{\theta'}^{-1} = \begin{pmatrix} -\beta & \beta + \mu \\ -\beta & 0 \end{pmatrix} \quad (33)$$

from which

$$\psi_{\varphi'} = (-\beta, \beta + \mu). \quad (34)$$

Then we have

$$\chi(\hat{\theta}) = \frac{-\hat{\beta}_\mu \log(\hat{\beta}) + (\hat{\beta}_\mu + \mu) \log(\hat{\beta} + \hat{\mu})}{\sqrt{\{\hat{\beta}_\mu^2 + (\hat{\beta}_\mu + \mu)^2\}}} \quad (35)$$

$$\chi(\hat{\theta}_\psi) = \frac{-\hat{\beta}_\mu \log(\hat{\beta}_\mu) + (\hat{\beta}_\mu + \mu) \log(\hat{\beta}_\mu + \mu)}{\sqrt{\{\hat{\beta}_\mu^2 + (\hat{\beta}_\mu + \mu)^2\}}} \quad (36)$$

$$|j_{(\theta\theta)}(\hat{\theta})| = y_1 y_2 = k/\hat{\beta}(\hat{\beta} + \hat{\mu}) \quad (37)$$

$$|j_{(\lambda\lambda)}(\hat{\theta}_\psi)| = \frac{y_1(\hat{\beta}_\mu + \mu)^2 + y_2 \hat{\beta}_\mu^2}{(\hat{\beta}_\mu + \mu)^2 + \hat{\beta}_\mu^2} \quad (38)$$

and finally

$$Q = \frac{\left\{ (\hat{\beta}_\mu + \mu) \log \left(\frac{\hat{\beta} + \hat{\mu}}{\hat{\beta}_\mu + \mu} \right) - \hat{\beta}_\mu \log \frac{\hat{\beta}}{\hat{\beta}_\mu} \right\}}{\{k\hat{\beta}(\hat{\beta} + \hat{\mu})\}^{1/2} \{k\hat{\beta}(\hat{\beta}_\mu + \mu)^2 + (\hat{\beta} + \hat{\mu})\hat{\beta}_\mu^2\}^{1/2}}. \quad (39)$$

The likelihood root is

$$\begin{aligned} r &= \text{sign}(Q) \sqrt{2\{\ell(\hat{\beta}, \hat{\mu}) - \ell(\hat{\beta}_\mu, \mu)\}} \quad (40) \\ &= \text{sign}(Q) \sqrt{2\{k\hat{\beta} \log\{\hat{\beta}/\hat{\beta}_\mu\} + (\hat{\beta} + \hat{\mu}) \log\{(\hat{\beta} + \hat{\mu})/(\hat{\beta}_\mu + \mu)\} - k(\hat{\beta} - \hat{\beta}_\mu) - \{\hat{\beta} + \hat{\mu} - (\hat{\beta}_\mu + \mu)\}\}}. \quad (41) \end{aligned}$$

The third order approximation to the p -value function is $1 - \Phi(r^*)$, where

$$r^* = r + (1/r) \log(Q/r). \quad (42)$$

Figure 5 shows the p -value function for μ using the mid- p -value function from the Poisson with no adjustment for the error in the background, and the p -value function from $1 - \Phi(r^*)$. The p -value for testing $\mu = 0$ is 0.00464, allowing for the uncertainty in the background, whereas it is 0.000408 ignoring this uncertainty.

The hypothesis $Ey = \beta$ could also be tested by modelling the mean of y as $\nu\beta$, say, and testing the value $\nu = 1$. In this formulation we can eliminate the nuisance parameter exactly by using the binomial distribution of y conditioned on the total $x + y$, as described in example 1. This gives a mid- p -value of 0.00521. The computation is much easier than that outlined above, and seems quite appropriate for testing the equality of the two means. However if inference about the mean of the signal is needed, in the form of a point estimate or confidence bounds, then the formulation as a ratio seems less natural at least in the context of HEP experiments. A more complete comparison of methods for this problem is given in Linnemann [8].

6. CONDITIONAL AND MARGINAL LIKELIHOOD

In special model classes, it is possible to eliminate nuisance parameters by either *conditioning* or *marginalizing*. The conditional or marginal likelihood then gives essentially exact inference for the parameter of interest, if this likelihood can itself be computed exactly. In Example 1 above, L_1 is the density for y_2 conditional on $y_1 + y_2$, so is a conditional likelihood for ψ . This is an example of the more general class of linear exponential families:

$$f(\underline{y}; \psi, \lambda) = \exp\{\psi s(\underline{y}) + \lambda' t(\underline{y}) - c(\psi, \lambda) - d(\underline{y})\}; \quad (43)$$

in which

$$f_{cond}(s | t; \psi) = \exp\{\psi s - C_t(\psi) - D_t(s)\} \quad (44)$$

defines the conditional likelihood. The comparison of two binomials in Example 6 is in this class, with ψ as defined at (30) and $\lambda = \log\{p_2/(1 - p_2)\}$. The difference of two Poisson means, in Example 7, cannot be formulated this way, however, even though the Poisson distribution is an exponential family, because the parameter of interest ψ is not a component of the canonical parameter.

It can be shown that in models of the form (43) the log-likelihood $\ell_a(\psi) = \ell_p(\psi) + (1/2) \log |j_{\lambda\lambda}|$ approximates the conditional log-likelihood $\ell_{cond}(\psi) =$

$\log f_{cond}(s | t; \psi)$, and that

$$p(\psi) = \Phi(r^*) \quad (45)$$

where

$$\begin{aligned} r^* &= r_a + \frac{1}{r_a} \log\left(\frac{Q}{r_a}\right) \\ r_a &= \pm[2\{\ell_a(\hat{\psi}_a) - \ell_a(\psi)\}]^{1/2} \\ Q &= (\hat{\psi}_a - \psi)\{j_a(\hat{\psi})\}^{1/2} \end{aligned}$$

approximates the p -value function with relative error $O(n^{-3/2})$ in i.i.d. sampling. An asymptotically equivalent approximation based on the profile log-likelihood is

$$p(\psi) = \Phi(r^*) \quad (46)$$

where

$$\begin{aligned} r^* &= r_p + \frac{1}{r_p} \log\left(\frac{Q}{r_p}\right) \\ r_p &= \pm[2\{\ell_p(\hat{\psi}) - \ell_p(\psi)\}]^{1/2} \\ Q &= (\hat{\psi} - \psi)\{j_p(\hat{\psi})\}^{1/2} \frac{|j_{\lambda\lambda}(\psi, \hat{\lambda}_\psi)|^{1/2}}{|j_{\lambda\lambda}(\hat{\psi}, \hat{\lambda})|^{1/2}}. \end{aligned}$$

In the latter approximation an adjustment for nuisance parameters is made to Q , whereas in the former the adjustment is built into the likelihood function. Approximation (46) was used in Figure 3.

A similar discussion applies to the class of transformation models, using marginal approximations. Both classes are reviewed in Reid [9].

Acknowledgments

The authors wish to thank Anthony Davison and Augustine Wong for helpful discussion. This research was partially supported by the Natural Sciences and Engineering Research Council.

References

- [1] D.R. Cox and N. Reid, "Parameter Orthogonality and Approximate Conditional Inference", *J. R. Statist. Soc. B*, **47**, 1, 1987.
- [2] D.A.S. Fraser, "Statistical Inference: Likelihood to Significance", *J. Am. Statist. Assoc.* **86** 258, 1991.
- [3] D.A.S. Fraser, N. Reid and A. Wong, "On Inference for Bounded Parameters", [arXiv: physics/0303111](https://arxiv.org/abs/physics/0303111), v1, 27 Mar 2003. to appear in *Phys. Rev. D*.
- [4] D.R. Cox and N. Reid, "A Note on the Difference between Profile and Modified Profile Likelihood", *Biometrika* **79**, 408, 1992.

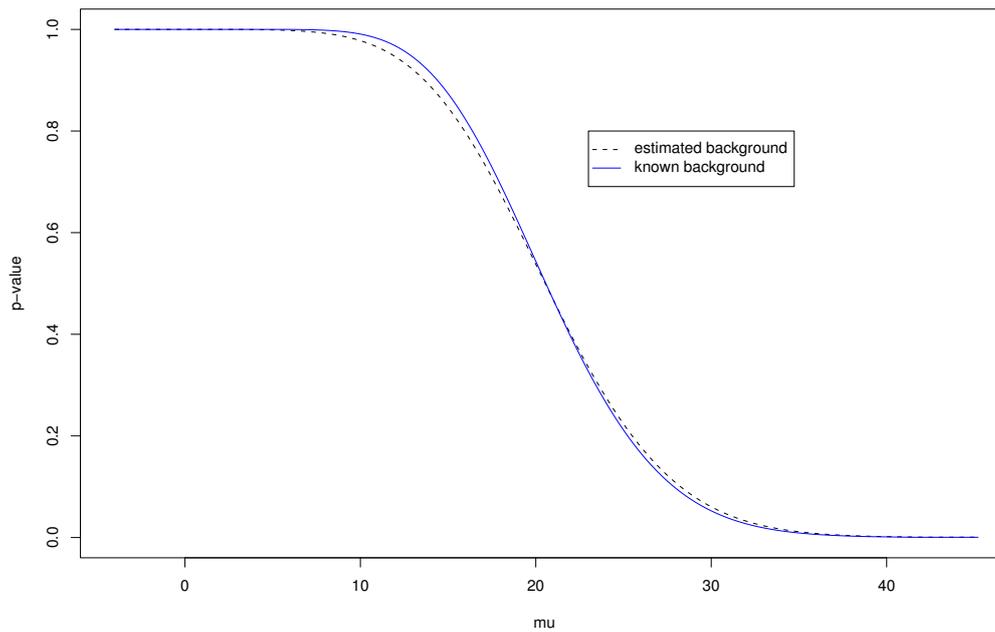


Figure 5: Comparison of the p -value functions computed assuming the background is known and using the mid- p -value with the third order approximation allowing a background error of ± 1.75 .

- [5] D.A.S. Fraser, “Likelihood for Component Parameters”, *Biometrika* **90**, 327, (2003).
- [6] D.A.S. Fraser, N. Reid and J. Wu, “A Simple General Formula for Tail Probabilities for Frequentist and Bayesian Inference”, *Biometrika* **86**, 246, 1999.
- [7] N. Reid, “Asymptotics and the Theory of Inference”, *Ann. Statist.*, to appear, 2004.
- [9] J. T. Linnemann, “Measures of significance in HEP and astrophysics”, *Phystat 2003*, to appear, 2004.
- [9] N. Reid, “Likelihood and Higher-Order Approximations to Tail Areas: a Review and Annotated Bibliography”, *Canad. J. Statist.* **24**, 141, 1996.