
The Tangent Exponential Model

A. C. Davison

Ecole Polytechnique Fédérale de Lausanne

N. Reid

University of Toronto

10.1 Introduction

In a series of papers starting in the late 1980s, D. A. S. Fraser, N. Reid and coworkers developed the tangent exponential model for higher-order likelihood inference. This chapter aims to explain the motivation and justification for this model and to describe how it is used to compute accurate approximations. The literature on this is not entirely transparent, as the argument evolved over numerous articles (Fraser, 1988, 1990, 1991, 2004; Fraser and Reid, 1988, 1993, 1995, 2001; Cakmak et al., 1994; Fraser et al., 1999). We give a heuristic account of this construction, for the most part skating over the technical details, and provide an annotated bibliography as a road map through the literature. The high accuracy of the resulting approximations has been verified empirically both in numerous articles and in books such as Brazzale et al. (2007), Chapter 8 of which overlaps with the account here.

Approximations based on the tangent exponential model build on the theory of conditional and marginal inference, and on Laplace and saddlepoint approximations. Chapter 12 of Davison (2003) defines some basic notions and derives some of the results presented here. Section 11.3.1 of that book contains an account of the Laplace method for integrals and related approximations for cumulative distribution functions. Fuller accounts may be found in Barndorff-Nielsen and Cox (1989, 1994), McCullagh (1987) and Severini (2000). Jensen (1995) and Butler (2007) provided comprehensive accounts of saddlepoint approximations and their many applications in analysis, probability, and statistics, and a helpful derivation is given in Kolassa (2006).

To fix notation and provide some building blocks that will be useful later, we start by summarising first-order inferential approximations related to the normal distribution and the central limit theorem. We then outline how the principles of sufficiency and ancillarity lead to consideration of significance functions, which are central in this and other frameworks for statistical inference, before discussing how the likelihood function can be viewed as a pivot. This leads to two density approximations, the p^* approximation of Barndorff-Nielsen (1983) and the tangent exponential model density approximation of Fraser (1988). Higher-order approximations to significance functions are then developed; first for linear exponential family models, where the link to saddlepoint approximations is most explicit, and then for general models, where the tangent exponential model plays a crucial role. Some extensions and generalisations are sketched in the concluding section.

10.2 Likelihood and Significance Functions

10.2.1 Background

We consider a vector $Y = (Y_1, \dots, Y_n)^T$ of continuous responses and a statistical model for Y with joint density function $f(y; \theta)$ that depends on a parameter $\theta \in \Theta \subset \mathbb{R}^p$. Vectors throughout are column vectors, and the superscripts T and $^\circ$ denote transpose and a quantity evaluated at the observed data, so y denotes a generic response vector and y° its observed value, and $\hat{\theta}$ a maximum likelihood estimator and $\hat{\theta}^\circ$ the maximum likelihood estimate computed from y° . Continuity of the response distribution plays a crucial role in the general development; some comments on extending the tangent exponential model for discrete responses are given in Section 10.5.2.

The likelihood function is proportional to the density of the observed data regarded as a function of the unknown parameter θ , i.e.,

$$L(\theta) = L(\theta; y^\circ) \propto f(y^\circ; \theta),$$

and measures the relative plausibility of different values of θ as explanations of y° . Weighted by prior information, the likelihood is a main ingredient in Bayesian approaches to inference, and, unweighted, it is key to the “pure likelihood” approach (Edwards, 1972; Royall, 1997). A central issue in using the likelihood function for inference is the distribution of $L(\theta; y)$ in repeated sampling under $f(y; \theta)$. This is needed in order to *calibrate* inferences based on the likelihood function, i.e., to ensure that their stated properties are correct under repeated sampling from the model. Calibration is essential to give inferences objective validity, ideally while respecting basic principles of inference, such as conditionality and sufficiency. One important approach to calibration is through the notion of a significance function, to be developed in Section 10.2.2.

We denote the log likelihood by $\ell(\theta) = \log L(\theta)$, and derivatives by subscripts, such as $\ell_\theta(\theta) = \partial \ell(\theta) / \partial \theta$ and $\ell_{\theta\theta}(\theta) = \partial^2 \ell(\theta) / \partial \theta \partial \theta^T$, which are respectively a column vector and a matrix. The observed information function is $j(\theta) = -\ell_{\theta\theta}(\theta)$. The classical asymptotic theory for likelihood-based inference is derived under the following smoothness conditions on the model (Davison, 2003, Section 4.4.2):

- i. the true value of θ is interior to the parameter space Θ ;
- ii. the densities $\{f(y; \theta) : \theta \in \Theta\}$ are distinct and have common support;
- iii. there is a neighbourhood \mathcal{N} of the true value of θ within which the first three derivatives of $\ell(\theta)$ with respect to θ exist, and for $j, k, l = 1, \dots, p$, $n^{-1} \mathbb{E}\{|\ell_{\theta_j \theta_k \theta_l}(\theta)|\}$ is uniformly bounded for $\theta \in \mathcal{N}$;
- iv. the expected Fisher information matrix $I(\theta) = \mathbb{E}\{j(\theta)\}$ is finite and positive definite, and $I(\theta) = \mathbb{E}\{\ell_\theta(\theta) \ell_\theta^T(\theta)\}$.

Chapter 16 of van der Vaart (1998) gives weaker conditions for the limiting distributional results now described.

When θ is a scalar, this classical theory provides three basic distributional approximations for inference:

$$s = s(\theta; y) = \ell_\theta(\theta) j^{-1/2}(\hat{\theta}) \sim N(0, 1), \quad (10.1)$$

$$t = t(\theta; y) = (\hat{\theta} - \theta) j^{1/2}(\hat{\theta}) \sim N(0, 1), \quad (10.2)$$

$$r = r(\theta; y) = \text{sign}(\hat{\theta} - \theta) [2\{\ell(\hat{\theta}) - \ell(\theta)\}]^{1/2} \sim N(0, 1), \quad (10.3)$$

where the maximum likelihood estimator, $\hat{\theta}$, is assumed to satisfy $\ell_{\theta}(\hat{\theta}) = 0$. The approximations (10.1)–(10.3) are derived from the central limit theorem for $\ell_{\theta}(\theta)$, which is $O_p(n^{1/2})$ in independent sampling, as $E_{\theta}\{\ell_{\theta}(\theta)\} = 0$ and $\text{var}_{\theta}\{\ell_{\theta}(\theta)\} = O(n)$ under conditions (i)–(iv). Their derivations also rely on the consistency of $\hat{\theta}$ for θ , and the convergence of the observed Fisher information to its expectation.

Each of s , t and r depends on both θ and the data y . The distributional approximations in (10.1)–(10.3) are all with reference to the distribution of y under the model $f(y; \theta)$, and $s(\theta; y)$, $t(\theta; y)$ and $r(\theta; y)$ are approximate pivots, i.e., functions of the data and a parameter whose distribution is known, at least approximately. Thus if y is fixed at y° , then $s(\theta; y^{\circ})$, $t(\theta; y^{\circ})$ and $r(\theta; y^{\circ})$ can be used to obtain the significance functions $\Phi(s)$, $\Phi(q)$, or $\Phi(r)$, where $\Phi(\cdot)$ is the distribution function for a standard normal random variable.

Versions of (10.1), (10.2) and (10.3) are also available for vector-valued θ . We write $\theta = (\psi, \lambda)$, where ψ is a scalar parameter of interest and λ is a nuisance parameter. The profile log likelihood $\ell_p(\psi) = \ell(\psi, \hat{\lambda}_{\psi})$ can be used to define analogous pivotal quantities,

$$\begin{aligned} s &= \ell'_p(\psi) J_p^{-1/2}(\hat{\psi}) \sim N(0, 1), \\ t &= (\hat{\psi} - \psi) J_p^{1/2}(\hat{\psi}) \sim N(0, 1), \\ r &= \text{sign}(\hat{\psi} - \psi) [2\{\ell_p(\hat{\psi}) - \ell_p(\psi)\}]^{1/2} \sim N(0, 1), \end{aligned} \tag{10.4}$$

where the prime denotes differentiation with respect to ψ , $\hat{\lambda}_{\psi}$ is the constrained maximum likelihood estimator for λ for ψ fixed, and

$$J_p(\psi) = J_{\psi\psi}(\hat{\theta}_{\psi}) - J_{\psi\lambda}(\hat{\theta}_{\psi}) J_{\lambda\lambda}^{-1}(\hat{\theta}_{\psi}) J_{\lambda\psi}(\hat{\theta}_{\psi}).$$

We use the shorthand $\hat{\theta}_{\psi} = (\psi, \hat{\lambda}_{\psi}^{\text{T}})^{\text{T}}$, and partition the observed information matrix as

$$J(\theta) = \begin{pmatrix} J_{\psi\psi}(\theta) & J_{\psi\lambda}(\theta) \\ J_{\lambda\psi}(\theta) & J_{\lambda\lambda}(\theta) \end{pmatrix},$$

with $J_{\psi\lambda}(\theta) = -\partial^2 \ell(\theta) / \partial \psi \partial \lambda^{\text{T}}$, and so forth.

The normal approximations to the distributions of s , t and r are equivalent to first order, but that for r respects any asymmetry of the log likelihood about its maximum, whereas that for t does not, which suggests that for complex problems these approximations may be inadequate. Their generally poor performance in models with many nuisance parameters has been borne out in empirical work.

The approximations (10.1)–(10.4) are usually derived from Taylor series expansion of the score equation $\ell_{\theta}(\hat{\theta}) = 0$ that defines the maximum likelihood estimate. A more elegant, but more difficult, approach shows that the log likelihood, as a function of θ , converges to the log likelihood for a normal distribution. See, for example Fraser and McDunnough (1984) for a Taylor-series type approach to this and LeCam (1960) or van der Vaart (1998, Ch. 6) for a related approach that LeCam called “local asymptotic normality”. We will return to this in describing the p^* approximation in Section 10.2.3.

In some settings we may be able to construct exact significance functions, using principles of sufficiency and ancillarity, which we now describe.

10.2.2 Significance Functions

There are two ideal settings for exactly calibrating inference on a scalar parameter θ . In the first, Y can be reduced to a scalar minimal sufficient statistic S , giving the *significance function*, sometimes also known as a *p-value function*,

$$p^{\circ}(\theta) = \text{pr}(S \leq s^{\circ}; \theta), \tag{10.5}$$

which without loss of generality we suppose to be decreasing in θ . The quantity $P^\circ(\theta)$ obtained by treating s° as random has a uniform distribution under sampling from the model $f(y; \theta)$, and this allows the calibration of inferences described in Section 10.2.1. For example, the limits of a $(1 - 2\alpha)$ confidence interval $(\theta_\alpha, \theta_{1-\alpha})$ for θ are the solutions of equations $p^\circ(\theta_\alpha) = 1 - \alpha$ and $p^\circ(\theta_{1-\alpha}) = \alpha$. Likewise, the evidence against the hypothesis that $\theta = \theta_0$ with alternative $\theta \neq \theta_0$ may be summarised by the p -value $2 \min\{p^\circ(\theta_0), 1 - p^\circ(\theta_0)\}$. The exact uniform distribution of $P^\circ(\theta)$ under repeated sampling from $f(y; \theta)$ implies that both of these inferential summaries are perfectly calibrated, in the sense that they have exactly their stated properties: the confidence interval contains the true θ with probability $1 - 2\alpha$ and the p -value is uniformly distributed.

In the second ideal setting there is a known transformation from Y to a minimal sufficient statistic (S, A) , where S is scalar and A is *ancillary*, i.e., its distribution does not depend on θ . The role of A can be understood by noting that the density $f(y; \theta)$ factorises as $f(s | a; \theta)f(a)$. We can therefore envisage the data as arising first by generating A and then generating S conditional on A . But as the first step does not depend on θ , the relevant subset of the sample space for inference about θ fixes the observed value a° of A (Cox, 1958). Thus we should base inference for θ on the significance function

$$p^\circ(\theta; a^\circ) = \text{pr}(S \leq s^\circ | A = a^\circ; \theta), \tag{10.6}$$

which provides perfectly calibrated inferences, conditionally on the observed value of the ancillary statistic A . Now the limits of an exact $(1 - 2\alpha)$ confidence interval for θ , the solutions of equations $p^\circ(\theta; a^\circ) = \alpha$ and $p^\circ(\theta; a^\circ) = 1 - \alpha$, depend on a° , thus emphasising how conditioning on the ancillary statistic A affects the precision of inferences.

Expressions (10.5) and (10.6) are functions of θ , evaluated only at the observed data.

Example 10.1

Suppose that Y_1/θ and $Y_2\theta$ are independent gamma variables with unit scale and shape parameter n ; their joint density function is

$$f(y_1, y_2; \theta) = \frac{(y_1 y_2)^{n-1}}{\{\Gamma(n)\}^2} \exp(-y_1/\theta - y_2\theta), \quad y_1, y_2 > 0, \theta > 0,$$

and the minimal sufficient statistic is (Y_1, Y_2) . We set $S = (Y_1/Y_2)^{1/2}$ and $A = (Y_1 Y_2)^{1/2}$, which is ancillary, and note that since $Y_1 = AS$ and $Y_2 = A/S$, we have $y^\text{T} = (y_1, y_2) = (as, a/s)$, so $|\partial(y_1, y_2)/\partial(a, s)| = 2a/s$ and


$$f(s | a; \theta) = \frac{1}{sI(a)} \exp\{-a(s/\theta + \theta/s)\}, \quad a, s > 0, \theta > 0,$$

where $I(a) = \int_{-\infty}^{\infty} \exp(-2a \cosh u) du$ is a normalising constant. The significance function (10.6) is readily obtained by numerical integration of $f(s | a^\circ; \theta)$ over the interval $(0, s^\circ)$.

The log likelihood for this model can be written as

$$\ell(\theta; s, a) \equiv -a(s/\theta + \theta/s), \quad \theta > 0,$$

so the maximum likelihood estimator is $\hat{\theta} = s$ and the observed information is $j(\hat{\theta}) = a/s^2$. Larger values of the ancillary a yield more precise inferences, because the standard error for $\hat{\theta}$, $j^{-1/2}(\hat{\theta}) = s/a^{1/2}$, diminishes as a increases.

**FIGURE 10.1**

Inference for gamma example. Significance functions $\text{pr}(s \leq s^\circ \mid a; \theta)$ (left) and log likelihoods (right) for $a = 3$ (solid) and $a = 6$ (dots). The horizontal lines in the left-hand panel are at 0.025, 0.5 and 0.975; the intersections of the highest and lowest of these with the significance functions give the limits of exact 95% confidence intervals. The horizontal lines in the right-hand panel are at -1.92 and -3.32 , and their intersections with the log likelihoods show the limits of 95% and 99% confidence intervals based on the approximation (10.3).

Figure 10.1 shows the significance and log likelihood functions when $s^\circ = 1.6$ and $a^\circ = 3, 6$. They are more concentrated when $a^\circ = 6$, resulting in shorter confidence intervals. Both log likelihood functions show clear asymmetry, suggesting that normal approximation based on t is badly calibrated and would provide poor inferences; indeed, the symmetric approximation (10.2) could produce negative confidence limits.

Though rarely met in practice, the settings above provide blueprints for more complex situations, in which several complications may arise:

- θ consists of a parameter ψ of interest, often scalar, and a vector λ of nuisance parameters. Even if there is a direct analogue of S , the significance functions (10.5) and (10.6) will typically depend on the unknown λ ;
- the reduction to a minimal sufficient statistic of dimension p applies principally in linear exponential family models. No mapping $y \mapsto (s, a)$ can be found in general, so exact inferences are mostly unavailable; and
- the interest parameter ψ may be a vector. We shall not consider this situation here, but Section 10.7.2 has pointers to related approaches.

Significance functions were emphasised as a primary tool for inference in Fraser (1991). Fraser (2019) described many summaries that are then directly available: confidence limits and p -values, as described above, as well as aspects of fixed-level testing. Power, for example, is reflected in the “steepness” of the function. There is a close connection between significance functions and confidence distributions (Cox, 1958; Efron, 1993; Xie and Singh, 2013).

10.2.3 Likelihood as Pivot

A fruitful approach to improved approximations is to use the log likelihood more directly—Hinkley (1980) referred to the likelihood function itself as a pivot. One reason this provides more precise calibration of our inferences is that it is closely related to a saddlepoint approximation (Daniels, 1954), which can be uncannily accurate. The saddlepoint approximation can be applied to the density of the maximum likelihood estimator, leading to the p^* approximation, also called “Barndorff-Nielsen’s formula” (Barndorff-Nielsen, 1980):

$$p^*(\widehat{\theta} \mid a; \theta) = c|j(\widehat{\theta})|^{1/2} \exp\{\ell(\theta) - \ell(\widehat{\theta})\}, \tag{10.7}$$

where $\ell(\theta) = \ell(\theta; \widehat{\theta}, a)$, with $(\widehat{\theta}, a)$ a transformation of the sample $y = (y_1, \dots, y_n)$ such that a is ancillary. Pivoting this approximate density (10.7), which is supported on \mathbb{R}^p , provides confidence intervals or regions for θ , and since a is ancillary no information about the parameter is lost by conditioning. However to compute the approximation requires calculation of the transformation from y to $(\widehat{\theta}, a)$, which can be very difficult. Moreover, computation of a significance function requires only a good approximation at the observed data point, whereas using (10.7) would require high accuracy for all $(\widehat{\theta}, a)$.

An approach intermediate between using the limiting normal form for the log likelihood and using (10.7) was introduced in Fraser (1988). His *tangent exponential model* approximation to the density $f(y; \theta)$, defined for $y \in \mathbb{R}^n$ and $\theta \in \mathbb{R}^p$, is a model on \mathbb{R}^p that implements conditioning on an approximately ancillary statistic a . Its expression is

$$f_{\text{TEM}}(s \mid a; \theta) = \exp\{s^\top \varphi(\theta) + \ell\{\theta(\varphi); y^o\}\}h(s). \tag{10.8}$$

This has the structure of a linear exponential family model for a constructed sufficient statistic $s \in \mathbb{R}^p$ and constructed canonical parameter $\varphi(\theta) \in \mathbb{R}^p$, with $-\ell(\theta; y^o)$ playing the role of the cumulant generator. If the underlying density is in the exponential family, then $\varphi(\theta)$ is simply the canonical parameter. In more general models the canonical parameter $\varphi(\theta)$ may depend on the data y^o , and is constructed using principles of approximate ancillarity, as described in Section 10.5.1.

In the next section, we show how this tangent exponential model is built on the exact significance functions of Section 10.2.2, but incorporates aspects of direct approximation of the log likelihood.

10.3 Approximate Conditional Inference

10.3.1 Ancillary and Sufficient Directions

Below we describe a general approach to approximate but accurate inference when the dimension p of the parameter vector θ is less than the dimension n of the data y . We assume that the components of y are independent and that the regularity conditions outlined in Section 10.2.1 hold. The tangent exponential model is often said to have “asymptotic properties”, which essentially is shorthand for the assumption that the expansions in Section 10.6 have the behaviour in n indicated there, i.e., the log-density is differentiable in both θ and y up to quartic terms, and the standardised derivatives decrease in powers of $n^{-1/2}$.

For a heuristic development we suppose initially that there is a smooth bijection between y and (s, a) , where s and a have respective dimensions p and $n - p$, and a is ancillary. The conditionality principle implies that inference should be based on the conditional density $f(s | a^\circ; \theta)$, so the reference set for frequentist inference is $\mathcal{A}^\circ = \{y \in \mathbb{R}^n : a(y) = a^\circ\}$, i.e., the p -dimensional manifold of the sample space on which the ancillary statistic equals its observed value a° . The bijection between y and (s, a) implies that \mathcal{A}° can be parametrised in terms of s , at which point its tangent plane \mathcal{T}_s is determined by the columns of the $n \times p$ matrix

$$\frac{\partial y(s, a^\circ)}{\partial s^\top}. \quad (10.9)$$

In particular, the tangent plane \mathcal{T}° to \mathcal{A}° at s° is determined by

$$V = \left. \frac{\partial y(s, a^\circ)}{\partial s^\top} \right|_{s=s^\circ}. \quad (10.10)$$

The columns of V have been called ancillary directions, because they are derived from the ancillary manifold \mathcal{A}° at $s = s^\circ$, but one can argue that this is a misnomer: the p columns of V determine how y changes in the direction of s locally at s° , so they might better be called *sufficient directions*, and we shall use this term below. The ancillary statistic itself varies locally at a° in the $n - p$ directions orthogonal to the columns of V .

Suppose for simplicity that θ and S are scalar. Then

$$\text{pr}(S \leq s^\circ | A = a^\circ; \theta) \propto \int_{-\infty}^{s^\circ} f(s, a^\circ; \theta) ds. \quad (10.11)$$

The log likelihood $\ell(\theta; s, a) = \log f(s, a; \theta)$ is a sum of n contributions and therefore has order n . A change of variables can be used to ensure that s is $O_p(1)$ and then, for $s = s^\circ + O_p(n^{-1/2})$, Taylor series expansion gives

$$\ell(\theta; s, a^\circ) = \ell(\theta; s^\circ, a^\circ) + (s - s^\circ) \left. \frac{\partial \ell(\theta; s, a^\circ)}{\partial s} \right|_{s=s^\circ} + \dots, \quad (10.12)$$

where the first term on the right-hand side is of order n , the second is the product of a term of order $n^{-1/2}$ with one of order n and is therefore of order $n^{1/2}$, and the remainder is $O(1)$. Standard first-order results, such as that leading to a significance function from applying the asymptotic approximation (10.3) to the signed likelihood root

$$r^\circ(\theta) = \text{sign}(\hat{\theta}^\circ - \theta) \left[2 \left\{ \ell(\hat{\theta}^\circ; s^\circ, a^\circ) - \ell(\theta; s^\circ, a^\circ) \right\} \right]^{1/2},$$

use only the first term on the right of (10.12), and have error of order $n^{-1/2}$ for one-sided confidence intervals. We hope to reduce this error to order n^{-1} , so-called second-order inference, by including the second term. Although $r(\theta)$ involves only the value of the log likelihood, an approximation based on both terms in (10.12) also requires the derivative of ℓ with respect to s , a so-called sample space derivative. Approximation (10.12) yields a version of the tangent exponential model (10.8), here specialised to the case where we can identify (s, a) directly.

In Section 10.4, we describe the building blocks for exponential family models, but we first return to the example in Section 10.2.2.

Example 10.1 (ctd)

We previously saw that $y^T(s, a) = (y_1, y_2) = (as, a/s)$, yielding

$$\frac{\partial y(s, a)}{\partial s} = \begin{pmatrix} a \\ -a/s^2 \end{pmatrix}, \quad V = \begin{pmatrix} a^\circ \\ -a^\circ/s^{\circ 2} \end{pmatrix}, \quad \frac{\partial \ell(\theta; s, a)}{\partial s} = -a(\theta^{-1} - \theta s^{-2}).$$

The left-hand panel of Figure 10.2 shows $f(y_1, y_2; \theta)$, the conditional sample space \mathcal{A}° and its tangent space \mathcal{T}° for $a^\circ = 3$ and $s^\circ = 1.6$. The right-hand panel shows the logarithm of the conditional density $f(s | a^\circ; \theta)$ and its tangents at s° for θ equal to 1, to $\hat{\theta}^\circ = 1.6$ and to 2.2. The filled circles show $f(s^\circ | a^\circ; \theta)$ for these values of θ , i.e., the corresponding likelihood values. Notice that $\theta = \varphi b + \{\varphi^2 b^2 + (s^\circ)^2\}^{1/2}$, where $b = (s^\circ)^2 / (2a^\circ)$, can be expressed in terms of $\varphi(\theta) = \partial \ell(\theta; s^\circ, a^\circ) / \partial s$, thus parametrising the model in terms of the slope of the tangent to the log density at y° ; in this model, this is a data-dependent parametrisation, to which we return below. \square

10.3.2 Computation of V

At first sight the definition of V at (10.10) suggests that the mapping $y \mapsto (s, a)$ must be known. In fact this is not the case, as we see if we write

$$V = \left. \frac{\partial y}{\partial s^T} \right|_{y=y^\circ} = \left. \frac{\partial y}{\partial \theta^T} \right|_{y=y^\circ} \times \left. \left(\frac{\partial s}{\partial \theta^T} \right)^{-1} \right|_{y=y^\circ}. \tag{10.13}$$

The second matrix on the right has dimension $p \times p$ and is invertible, so the column space of V is also the column space of the first matrix on the right; both define the same space of




FIGURE 10.2

Conditional inference for gamma example. Left: joint density $f(y; \theta)$ with $\theta = 1$, $y^\circ = (4.8, 1.875)$, giving $a^\circ = 3$ and $s^\circ = 1.6$. The white filled circle shows y° and the black filled circle shows $f(y^\circ; \theta)$. The conditional reference set \mathcal{A}° is shown by the dashed white line and the density $f\{y(s, a^\circ); \theta\}$ on \mathcal{A}° is shown by the solid line. The tangent plane \mathcal{T}° is shown by the dotted white line. Right: Log conditional density $f(s | a^\circ; \theta)$ (curved) and its tangents at $\log f(s^\circ | a^\circ; \theta)$ (straight) for $\theta = \hat{\theta}^\circ$ (solid), $\theta = 1$ (dashes) and $\theta = 2.2$ (dots). The filled circles are at $f(s^\circ | a^\circ; \theta)$.

sufficient directions, but $\partial y/\partial\theta^T$ does not require y to be expressed in terms of (s, a) . Hence V could be right-multiplied by any invertible $p \times p$ matrix of constants without changing the span of the sufficient directions. Below we shall generally take V to be $\partial y/\partial\theta^T$, evaluated at $y = y^\circ$ and $\theta = \hat{\theta}^\circ$.

Although it is not customary to express y as a function of θ , it is natural to do so when considering how a dataset would be simulated. In a regression model, for example, we can write $\theta = (\beta, \sigma)$ and

$$y(\theta) = X\beta + \sigma\varepsilon, \quad (10.14)$$

with the design matrix X and error vector ε fixed, giving

$$V = \left(\frac{\partial y}{\partial\beta^T}, \frac{\partial y}{\partial\sigma} \right) \Big|_{y=y^\circ} = (X, \varepsilon) \Big|_{y=y^\circ} = (X, (y^\circ - X\hat{\beta}^\circ)/\hat{\sigma}^\circ), \quad (10.15)$$

where $\hat{\theta}^\circ = (\hat{\beta}^\circ, \hat{\sigma}^\circ)$ are maximum likelihood estimates computed from the data y° . One way to think about (10.14) is as a quantile function, or structural equation, whereby changes in θ are reflected in changes to y for fixed ε .

An alternative approach to deriving V is to note that if the y_j are independent and have distribution functions $F_j(\cdot; \theta)$, then the pivotal equation $F(y_j; \theta) = u_j$ for some fixed u_j implicitly defines how y_j depends on θ . Total differentiation of this equation with respect to θ^T yields

$$0 = \frac{du_j}{d\theta^T} = \frac{\partial y_j}{\partial\theta^T} \frac{\partial F(y_j; \theta)}{\partial y_j} + \frac{\partial F(y_j; \theta)}{\partial\theta^T},$$

which implies that

$$\frac{\partial y_j}{\partial\theta^T} = - \left\{ \frac{\partial F_j(y_j; \theta)}{\partial y_j} \right\}^{-1} \frac{\partial F_j(y_j; \theta)}{\partial\theta^T}. \quad (10.16)$$

In the case of a regression model, we have

$$F_j(y_j; \beta, \sigma) = F\{(y_j - x_j^T\beta)/\sigma\},$$

where x_j^T is the j th row of X and $\varepsilon_1, \dots, \varepsilon_n \stackrel{\text{iid}}{\sim} F$, which gives (10.15) when evaluated at y° .

In group transformation models, there is no need to invoke the distribution function F_j , because one can write $y_j = g_j(\theta, \varepsilon_j)$ for a known function g_j and an error ε_j whose distribution does not depend on θ . Then V can be computed as $\partial y_j/\partial\theta$ for fixed ε_j , or equivalently

$$\frac{\partial y_j}{\partial\theta^T} = - \left\{ \frac{\partial \varepsilon_j(y_j; \theta)}{\partial y_j} \right\}^{-1} \frac{\partial \varepsilon_j(y_j; \theta)}{\partial\theta^T}.$$

This construction is used in the bivariate normal example of Section 10.5.2. The derivation of the sufficient directions in Fraser and Reid (1995) and Fraser and Reid (2001) builds on the local location model of Fraser (1964); see Section 10.7.3.

The construction in terms of a structural equation such as (10.14) does not apply to discrete models, which require a slightly different treatment. We discuss the construction of the sufficient directions V further in Section 10.5.2.

In the next section, we outline the accurate approximation of significance functions in continuous exponential families.

10.4 Exponential Family Inferences

10.4.1 One-Parameter Case

Consider a continuous one-parameter exponential family with a scalar parameter θ , expressed as a tilted version of a baseline density $f_0(y)$,

$$f(y; \theta) = f_0(y) \exp [n \{s(y)\theta - \kappa(\theta)\}], \tag{10.17}$$

with canonical statistic $ns(y)$, canonical parameter θ and cumulant generator $n\kappa(\theta)$. The marginal density of the sufficient statistic s is of the same form,

$$f(s; \theta) = h_0(s) \exp [n \{s\theta - \kappa(\theta)\}], \tag{10.18}$$

where $h_0(s) = \int 1\{y \in \mathbb{R}^n : s(y) = s\}f_0(y)dy$. Here s is assumed to be an average of n independent observations, so its variation around its mean $\kappa_\theta(\theta)$ is $O_p(n^{-1/2})$. Writing densities explicitly in terms of the sample size n , as in (10.17) or (10.18), is a technical device that helps in keeping track of powers of n in theoretical development, but it does not change the quality of subsequent approximations and is unnecessary when applying them.

Although the integral defining h_0 is not usually available in closed form, the saddlepoint approximation can be used to approximate $f(s; \theta)$ very accurately. And as $\hat{\theta}$ satisfies $s = \kappa_\theta(\hat{\theta})$, this gives an approximation to the density of $\hat{\theta}$ (e.g., Davison, 2003, equation (12.32)):

$$f(\hat{\theta}; \theta) = c \hat{j}^{1/2} \exp \left\{ \ell(\theta; \hat{\theta}) - \ell(\hat{\theta}; \hat{\theta}) \right\} \{1 + O(n^{-1})\}, \tag{10.19}$$

where $\hat{j} = n\kappa_{\theta\theta}(\hat{\theta})$ is the observed information evaluated at $\theta = \hat{\theta}$, and the normalising constant c ensures that (10.19) has unit integral. Our goal below is to use (10.19) to obtain a convenient and general expression for the significance function $\text{pr}(\hat{\theta} \leq \hat{\theta}^\circ; \theta)$, where $\hat{\theta}^\circ$ is the observed value of $\hat{\theta}$.

We first make a monotone change of variable $\hat{\theta} \mapsto r(\theta)$, where

$$r(\theta) = \text{sign}(\hat{\theta} - \theta) \left[2 \left\{ \ell(\hat{\theta}; \hat{\theta}) - \ell(\theta; \hat{\theta}) \right\} \right]^{1/2}$$

The Jacobian of this transformation may be obtained from the derivative

$$\begin{aligned} r(\theta) \frac{\partial r(\theta)}{\partial \hat{\theta}} &= \left. \frac{\partial \ell(\theta; \hat{\theta})}{\partial \theta} \right|_{\theta=\hat{\theta}} + \left. \frac{\partial \ell(\theta; \hat{\theta})}{\partial \hat{\theta}} \right|_{\theta=\hat{\theta}} - \frac{\partial \ell(\theta; \hat{\theta})}{\partial \hat{\theta}}, \\ &= \ell_{;\hat{\theta}}(\hat{\theta}; \hat{\theta}) - \ell_{;\hat{\theta}}(\theta; \hat{\theta}), \end{aligned}$$

where here and below the appearance of a variable after a subscripted semi-colon indicates a sample-space derivative with respect to that variable. Differentiation with respect to $\hat{\theta}$ is necessitated by the change of variable $\hat{\theta} \mapsto r(\theta)$. For the exponential model, we have

$$r(\theta) \frac{\partial r(\theta)}{\partial \hat{\theta}} = n\kappa_{\theta\theta}(\hat{\theta})(\hat{\theta} - \theta) = \hat{j}(\hat{\theta} - \theta),$$

and (10.19) becomes

$$f\{r(\theta); \theta\} = c \frac{r(\theta)}{q(\theta)} \exp \left\{ -r(\theta)^2/2 \right\}, \tag{10.20}$$

where the Wald statistic for θ ,

$$q(\theta) = \hat{j}^{-1/2}(\hat{\theta} - \theta) = \hat{j}^{-1/2} \left\{ \ell_{;\hat{\theta}}(\hat{\theta}; \hat{\theta}) - \ell_{;\hat{\theta}}(\theta; \hat{\theta}) \right\},$$

has an asymptotic standard normal distribution; see (10.2), where q is called t .

As $\theta \rightarrow \hat{\theta}$ we have $r(\theta), q(\theta) \rightarrow 0$, but it is possible to show that, using an abbreviated notation, $q = r + a_1 r^2/n^{1/2} + a_2 r^3/n + \dots$, so there is no singularity in (10.20). Taylor series expansion of the logarithm yields

$$\frac{1}{r} \log \left(\frac{q}{r} \right) = a_1 n^{-1/2} + (a_2 - a_1^2/2) r n^{-1} + O(n^{-3/2}), \quad (10.21)$$

so the square of (10.21) is of order n^{-1} . This is useful in the next step.

It follows from (10.21) that a second change of variable

$$r(\theta) \mapsto r^*(\theta) = r(\theta) + \frac{1}{r(\theta)} \log \left\{ \frac{q(\theta)}{r(\theta)} \right\} \quad (10.22)$$

has Jacobian $1 + (a_2 - a_1^2/2)/n + O(n^{-3/2}) = 1 + O(n^{-1})$. That the coefficient of the $1/n$ term is constant in r is important for renormalisation, discussed in Section 10.4.4. Although the transformation (10.22) may not be strictly monotonic over its entire range, it is monotone to the order considered here. After this transformation and a little more algebra, (10.20) becomes

$$f\{r^*(\theta); \theta\} = c \exp \left\{ -r^*(\theta)^2/2 \right\} \{1 + O(n^{-1})\}. \quad (10.23)$$

We deduce that $c = (2\pi)^{-1/2} \{1 + O(n^{-1})\}$, so the significance function may be written

$$\begin{aligned} \text{pr}(\hat{\theta} \leq \hat{\theta}^\circ; \theta) &= \text{pr} \{r^*(\theta) \leq r^{*\circ}(\theta); \theta\} \\ &= \Phi \{r^{*\circ}(\theta)\} \{1 + O(n^{-1})\}, \end{aligned} \quad (10.24)$$

where $r^{*\circ}(\theta)$ is the value of $r^*(\theta)$ actually observed, i.e., with y and $\hat{\theta}$ replaced by y° and $\hat{\theta}^\circ$. The discussion in Section 10.5.1 then allows inference on θ by treating $r^{*\circ}(\theta)$ as a realisation of a standard normal variable. An alternative to (10.24) with the same order of asymptotic error is the Lugannani–Rice (1980) formula

$$\text{pr} \left(\hat{\theta} \leq \hat{\theta}^\circ; \theta \right) \doteq \Phi \{r^\circ(\theta)\} + \left\{ \frac{1}{r^\circ(\theta)} - \frac{1}{q^\circ(\theta)} \right\} \phi \{r^\circ(\theta)\}, \quad (10.25)$$

where ϕ denotes the standard normal density function and $r^\circ(\theta)$ and $q^\circ(\theta)$ are the ingredients to $r^{*\circ}(\psi)$. Both (10.24) and (10.25) are typically highly accurate, with neither systematically better than the other in applications, but they can become numerically unstable for θ near $\hat{\theta}$ and dealing with this may require some careful programming. In Section 10.4.4 we show that when the density is renormalised to integrate to unity the relative error in (10.24) and (10.25) becomes $O(n^{-3/2})$, as is that in (10.23).

Expression (10.21) shows that the second term of $r^*(\theta)$ is an $O(n^{-1/2})$ correction to the $O(1)$ quantity $r(\theta)$; recall the second term on the right-hand side of (10.12).

10.4.2 Linear Exponential Family

The previous argument generalises to a linear exponential family in which $\theta^\top = (\psi, \lambda^\top)$ consists of a scalar parameter ψ of interest and a nuisance parameter λ of dimension $p - 1$. In this case,

$$f(s; \theta) = h(s) \exp [n \{s^\top \theta - \kappa(\theta)\}],$$

where $s(y)^T = (s_1, s_2)$ is partitioned conformably with θ . In this model the conditional density of s_1 given s_2 does not depend on λ , and the ratio of the saddlepoint approximations to the densities of (s_1, s_2) and of s_2 yields the approximation

$$f(s_1 | s_2; \psi) \doteq c \left\{ \frac{|J_{\lambda\lambda}(\widehat{\theta}_\psi)|}{|J(\widehat{\theta})|} \right\}^{1/2} \exp \left\{ \ell(\widehat{\theta}_\psi) - \ell(\widehat{\theta}) \right\}, \quad (10.26)$$

where $\widehat{\theta}_\psi = (\psi, \widehat{\lambda}_\psi^T)^T$ denotes the maximum likelihood estimator for fixed ψ and $J_{\lambda\lambda}(\theta)$ is the sub-matrix of $J(\theta)$ corresponding to λ . A calculation similar to that leading to (10.23) establishes that apart from a relative error of order n^{-1} , the resulting significance function for inference on ψ , $\text{pr}(S_1 \leq s_1^\circ | S_2 = s_2; \psi)$ is again of form (10.24), now with

$$r(\psi) = \text{sign}(\widehat{\psi} - \psi) \left[2 \left\{ \ell(\widehat{\theta}) - \ell(\widehat{\theta}_\psi) \right\} \right]^{1/2}, \quad (10.27)$$

$$q(\psi) = (\widehat{\psi} - \psi) \left\{ \frac{|J(\widehat{\theta})|}{|J_{\lambda\lambda}(\widehat{\theta}_\psi)|} \right\}^{1/2} \quad (10.28)$$

evaluated at the observed data y° and maximum likelihood estimate $\widehat{\theta}^\circ$. Note that

$$q(\psi) = t(\psi)\rho(\psi, \widehat{\psi}), \quad \rho(\psi, \widehat{\psi}) = \left\{ \frac{|J_{\lambda\lambda}(\widehat{\theta})|}{|J_{\lambda\lambda}(\widehat{\theta}_\psi)|} \right\}^{1/2},$$

with $t(\psi)$ the Wald statistic based on the profile log likelihood defined in (10.4). For derivations of these results see Fraser and Reid (1993) or Davison (2003, Section 12.3.3), for example.

Although expression (10.24) was derived by approximating a conditional distribution, it is also an approximation to the marginal distribution of $r^*(\psi)$, because the normal distribution of $r^*(\psi)$ in (10.24) does not depend on the conditioning variable S_2 .

If the parameter of interest is a linear function of the natural parameter θ of the exponential model, say $\psi = C_1^T \theta$ for some known $p \times 1$ vector C_1 , then we can set $\lambda = C_2^T \theta$, so that $\varphi^T = (\psi, \lambda^T) = \theta^T (C_1, C_2) = \theta^T C$, say, where the $p \times p$ matrix C is invertible, express the exponential family in terms of

$$s_*(y) = C^{-1} s(y), \quad \varphi(\theta) = C^T \theta, \quad \kappa_*(\varphi) = \kappa(C^{-T} \varphi) = \kappa(\theta),$$

and finally apply approximation (10.24) in this reparametrised model.

10.4.3 General Exponential Family

We now extend the normal approximation to the distribution of r^* , developed in the previous section, to the general exponential family

$$f(s; \theta) = h(s) \exp [n \{ s^T \varphi(\theta) - \kappa(\theta) \}], \quad (10.29)$$

where $\varphi(\theta)$ may be a nonlinear function of (ψ, λ) . For this we shall need the analogues of $r(\psi)$ and $q(\psi)$ of (10.27) and (10.28). The likelihood is invariant to reparametrisation, so $r(\psi)$ is unchanged, but as ψ is no longer a component of the canonical parameter φ , we need a new form for $q(\psi)$ using a surrogate for ψ . Taylor series expansion for small $\widehat{\theta} - \widehat{\theta}_\psi$ gives

$$\varphi(\widehat{\theta}) - \varphi(\widehat{\theta}_\psi) = \frac{\partial \varphi(\widehat{\theta}_\psi)}{\partial \theta^T} (\widehat{\theta} - \widehat{\theta}_\psi) + \dots,$$

so if the matrix on the right-hand side is invertible at $\widehat{\theta}_\psi$, which is a condition for θ to be identifiable, then

$$\widehat{\theta} - \widehat{\theta}_\psi = \left\{ \frac{\partial \varphi(\widehat{\theta}_\psi)}{\partial \theta^\top} \right\}^{-1} \left\{ \varphi(\widehat{\theta}) - \varphi(\widehat{\theta}_\psi) \right\} + \dots, \tag{10.30}$$

and as the partial derivatives satisfy

$$I_p = \frac{\partial \theta}{\partial \varphi^\top} \frac{\partial \varphi}{\partial \theta^\top} = \begin{pmatrix} \partial \psi / \partial \varphi^\top & \partial \varphi / \partial \varphi^\top \end{pmatrix} \begin{pmatrix} \partial \varphi / \partial \psi & \partial \varphi / \partial \lambda^\top \end{pmatrix} = \begin{pmatrix} \frac{\partial \psi}{\partial \varphi^\top} \frac{\partial \varphi}{\partial \psi} & \frac{\partial \psi}{\partial \varphi^\top} \frac{\partial \varphi}{\partial \lambda^\top} \\ \frac{\partial \lambda}{\partial \varphi^\top} \frac{\partial \varphi}{\partial \psi} & \frac{\partial \lambda}{\partial \varphi^\top} \frac{\partial \varphi}{\partial \lambda^\top} \end{pmatrix}, \tag{10.31}$$

the first row of the inverse on the right-hand side of (10.30) is $\partial \psi / \partial \varphi^\top$, yielding

$$\widehat{\psi} - \psi = \frac{\partial \psi(\widehat{\theta}_\psi)}{\partial \varphi^\top} \left\{ \varphi(\widehat{\theta}) - \varphi(\widehat{\theta}_\psi) \right\} + \dots.$$

Hence $q(\psi)$ should be based on a local departure of $\widehat{\theta}$ from $\widehat{\theta}_\psi$, or equivalently of $\widehat{\varphi}$ from $\widehat{\varphi}_\psi$, which can be measured through the constructed parameter $\chi(\theta)$ given by

$$\chi(\theta) = u^\top \varphi(\theta), \quad u = \frac{\partial \psi(\widehat{\theta}_\psi) / \partial \varphi}{\left\| \partial \psi(\widehat{\theta}_\psi) / \partial \varphi \right\|},$$

i.e., the orthogonal projection of $\varphi(\theta)$ onto a unit vector u parallel to $\partial \psi(\widehat{\theta}_\psi) / \partial \varphi$; note that $\chi(\theta)$ depends on the data through $\widehat{\theta}_\psi$. The Wald-type measure q from (10.28) for $\chi(\theta)$ is

$$q(\psi) = \text{sign}(\widehat{\psi} - \psi) \left| \chi(\widehat{\theta}) - \chi(\widehat{\theta}_\psi) \right| \left\{ \frac{|J(\widehat{\varphi})|}{|J_{(\lambda\lambda)}(\widehat{\varphi}_\psi)|} \right\}^{1/2}, \tag{10.32}$$

where the determinants are computed in the φ parametrisation,

$$|J(\widehat{\varphi})| = \left| J(\widehat{\theta}) \right| \left| \frac{\partial \varphi(\widehat{\theta})}{\partial \theta^\top} \right|^{-2}, \quad |J_{(\lambda\lambda)}(\widehat{\varphi}_\psi)| = \left| J_{\lambda\lambda}(\widehat{\theta}_\psi) \right| \left| \frac{\partial \varphi^\top(\widehat{\theta}_\psi)}{\partial \lambda} \frac{\partial \varphi(\widehat{\theta}_\psi)}{\partial \lambda^\top} \right|^{-1}, \tag{10.33}$$

and the second factor on the right-hand side of the second expression here stems from the ‘‘area formula’’ (see, for example, Krantz and Parks, 2008, Lemma 5.1.4). An equivalent expression for $q(\psi)$ is obtained by substituting these expressions and simplifying, yielding

$$q(\psi) = \frac{|\varphi(\widehat{\theta}) - \varphi(\widehat{\theta}_\psi)| \left| \partial \varphi(\widehat{\theta}_\psi) / \partial \lambda^\top \right|}{\left| \partial \varphi(\widehat{\theta}) / \partial \theta^\top \right|} \times \left\{ \frac{|J(\widehat{\theta})|}{|J_{\lambda\lambda}(\widehat{\theta}_\psi)|} \right\}^{1/2}. \tag{10.34}$$

If $\varphi^\top = (\psi, \lambda^\top)$, then (10.34) reduces to (10.28).

The equivalence of (10.32) and (10.34) follows by noting that the determinant of the $p \times p$ matrix

$$\begin{pmatrix} \varphi(\widehat{\theta}) - \varphi(\widehat{\theta}_\psi) & \partial \varphi(\widehat{\theta}_\psi) / \partial \lambda^\top \end{pmatrix}$$

is the signed volume $l_1 v_2$ of the parallelepiped generated by its p columns, with l_1 the length of the component of its first column in the direction orthogonal to its other columns and v_2 the volume of the $(p - 1)$ -dimensional parallelepiped these last columns generate (Fraser et al., 1999; Skovgaard, 1996), which is

$$v_2 = \left| \frac{\partial \varphi^\top(\widehat{\theta}_\psi)}{\partial \lambda} \frac{\partial \varphi(\widehat{\theta}_\psi)}{\partial \lambda^\top} \right|^{1/2}.$$

As the vector $\partial \psi(\widehat{\theta}_\psi) / \partial \varphi$ is orthogonal to $\partial \varphi(\widehat{\theta}_\psi) / \partial \lambda^\top$, we have $l_1 = \left| \chi(\widehat{\theta}) - \chi(\widehat{\theta}_\psi) \right|$; the sign of $\widehat{\psi} - \psi$ supplies the sign of $l_1 v_2$.

Example 10.2

To illustrate the development above, consider independent exponential variables y_1 and y_2 with rate parameters $\lambda\psi$ and λ ; here $\theta = (\psi, \lambda)^\top$. The corresponding log likelihood is

$$\ell(\psi, \lambda) = 2 \log \lambda + \log \psi - \lambda(\psi y_1 + y_2), \quad \psi, \lambda > 0,$$

leading to $\widehat{\lambda}_\psi = 2/(\psi y_1 + y_2)$, $\widehat{\lambda} = 1/y_2$, $\widehat{\psi} = y_2/y_1$, $|j(\widehat{\theta})| = 1/(\widehat{\lambda}\widehat{\psi})^2$ and $|j_{\lambda\lambda}(\widehat{\theta}_\psi)| = 2/\widehat{\lambda}_\psi^2$.

The elements of (10.29) are

$$\begin{aligned} \varphi &= \begin{pmatrix} \varphi_1 \\ \varphi_2 \end{pmatrix} = - \begin{pmatrix} \lambda\psi \\ \lambda \end{pmatrix}, \\ s(y) &= \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}, \\ \kappa(\varphi) &= -\log(-\varphi_1) - \log(-\varphi_2), \quad \varphi_1, \varphi_2 < 0, \end{aligned}$$

so $\psi = \varphi_1/\varphi_2$, and fixing ψ is equivalent to forcing φ to lie on a line through the origin of gradient $1/\psi$. Now

$$\frac{\partial \varphi}{\partial \lambda} = - \begin{pmatrix} \psi \\ 1 \end{pmatrix}, \quad \frac{\partial \psi}{\partial \varphi} = \begin{pmatrix} 1/\varphi_2 \\ -\varphi_1/\varphi_2^2 \end{pmatrix} = \frac{1}{\lambda} \begin{pmatrix} -1 \\ \psi \end{pmatrix}.$$

Hence the unit vector in the direction $\partial\psi/\partial\varphi$ is $(1 + \psi^2)^{-1/2}(-1, \psi)^\top$, and u is obtained from this by evaluating it at $\varphi(\widehat{\theta}_\psi)$; this is orthogonal to $\partial\varphi/\partial\lambda$ by construction.

To be concrete, let $y_1 = 1$ and $y_2 = 2$, and consider computing the significance function at $\psi = 1$. In this case $\widehat{\theta}_\psi = (1, 2/3)^\top$, $\widehat{\theta} = (2, 1/2)^\top$, $\varphi(\widehat{\theta}_\psi) = -(2/3, 2/3)^\top$, $\varphi(\widehat{\theta}) = -(1, 1/2)^\top$ and $u = (-1/\sqrt{2}, 1/\sqrt{2})^\top$.

Figure 10.3 shows the log likelihoods, the maximum likelihood estimates and the partial maximum likelihood estimates in the θ and φ parametrisations. The construction of $\chi(\theta) = u^\top \varphi(\theta)$ in terms of u and $\varphi(\theta)$, shown in the right-hand panel, yields $\chi(\widehat{\theta}) = 1/\sqrt{8}$ and $\chi(\widehat{\theta}_\psi) = 0$. The difference $\chi(\widehat{\theta}) - \chi(\widehat{\theta}_\psi)$ in the constructed parameter is a signed distance along the dashed line. As ψ varies, the grey line $\varphi_1 = \psi\varphi_2$ and the vector u , which is orthogonal to that line, also vary. Increasing ψ starting from $\psi = 1$ would move $\varphi(\widehat{\theta}_\psi)$ along the dotted line closer to $\varphi(\widehat{\theta})$, inclining the vector u and the dashed line closer to vertical and reducing significance by decreasing $\chi(\widehat{\theta}) - \chi(\widehat{\theta}_\psi)$. When the grey line passes through the black square representing $\varphi(\widehat{\theta})$, it coincides with the black circle and the open square and then $\chi(\widehat{\theta}) - \chi(\widehat{\theta}_\psi) = 0$. Decreasing ψ would move $\varphi(\widehat{\theta}_\psi)$ along the black dotted line away from $\varphi(\widehat{\theta})$, inclining the vector u and the dashed line further from the vertical, and increasing $\chi(\widehat{\theta}) - \chi(\widehat{\theta}_\psi)$ and thus the significance. \square

10.4.4 Renormalisation

Although the error in (10.24) is ostensibly $O(n^{-1})$, it is actually $O(n^{-3/2})$ in wide generality after renormalisation of the approximation to ensure that it integrates to 1. While in practice this would be done numerically, we note that if the error term in (10.23) takes the form B/n where B does not depend on the variable of integration, then

$$1 = \int f\{r^*(\theta); \theta\} dr^* = \int c\sqrt{(2\pi)}\phi\{r^*(\theta)\}\{1 + B/n + O(n^{-3/2})\} dr^*$$

implies

$$c\sqrt{(2\pi)} = 1 - B/n + O(n^{-3/2}),$$

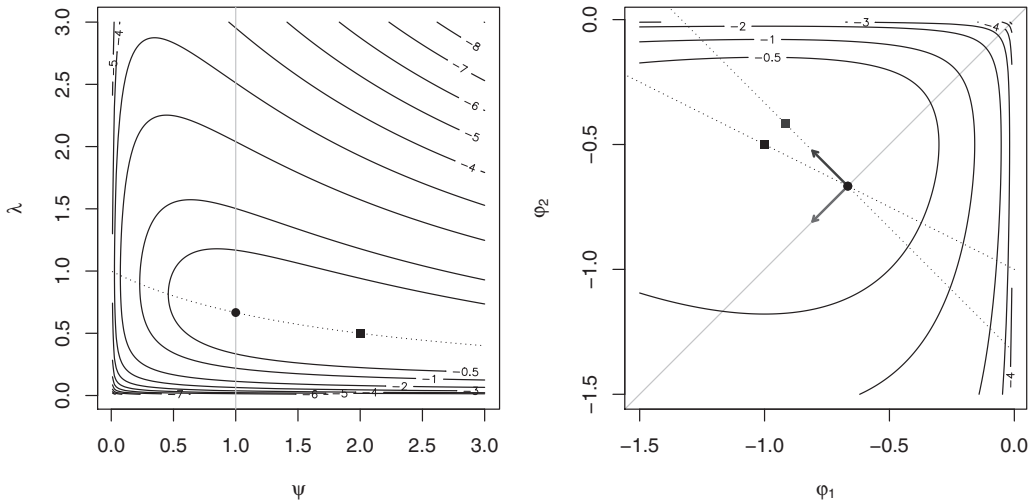


FIGURE 10.3

Log likelihoods for exponential example in θ parametrisation (left) and φ parametrisation (right). The solid grey lines show $\psi = 1$ or equivalently $\varphi_1 = \varphi_2$. The overall maximum likelihood estimates $\hat{\theta}$ and $\varphi(\hat{\theta})$ are shown by black squares, and the partial maximum likelihood estimates $\hat{\theta}_\psi$ and $\varphi(\hat{\theta}_\psi)$, with $\psi = 1$, by black circles. The dotted lines show $\hat{\theta}_\psi$ and $\varphi(\hat{\theta}_\psi)$ as functions of ψ . The arrows in the right-hand panel show the directions of the vectors $\partial\varphi(\hat{\theta}_\psi)/\partial\lambda$ (light) and $\partial\psi(\hat{\theta}_\psi)/\partial\varphi$ (heavy). The difference $\chi(\hat{\theta}) - \chi(\hat{\theta}_\psi)$ appearing in $q(\psi)$ is a distance along the dotted line, between $\chi(\hat{\theta}_\psi)$ (black filled circle) and $\chi(\hat{\theta})$ (open square).

so (10.23) becomes

$$f(r^*; \theta) = (1 - B/n)\phi\{r^*(\theta)\}\{1 + B/n + O(n^{-3/2})\} = \phi\{r^*(\theta)\}\{1 + O(n^{-3/2})\}.$$

As noted below (10.22), in the linear exponential family with $p = 1$ we indeed have B free of r^* . In work as yet unpublished Y. Tang has verified that B is also constant in r^* for the multi-parameter linear exponential family treated in Section 10.4.2, and for regression-scale models (10.14).

If the $O(n^{-1})$ error term depends on $r^*(\theta)$, then under mild conditions

$$B\{r^*(\theta)\} = B(0) + r^*(\theta)B'(s), \quad 0 < |s| < |r^*(\theta)|,$$

and the term $B(0)/n$ cancels with the normalizing constant as above, so we only need consider

$$c\sqrt{(2\pi)}\phi\{r^*(\theta)\}r^*(\theta)B'(s).$$

If $B'(s)$ is constant, then this term integrates to 0 and the norming constant is as before. If $B'(s)$ depends on r^* , it would have to rise very rapidly for the error term to be unbounded, as the normal density function drops rapidly as $|r^*(\theta)|$ increases. Slower, but not constant, dependence of B' on r^* would lead to relative error just $O(1/n)$, not improved by renormalization.

Mild regularity conditions (Daniels, 1956) ensure the validity of this argument, which also applies to Laplace and similar approximations, including those in Sections 10.4.2 and 10.4.3. A similar analysis verifies that the relative error in the Lugannani–Rice (1980) approximation (10.25) is also $O(n^{-3/2})$.

10.5 General Likelihood

10.5.1 Basic Approximation

Section 10.4 described approximations useful for inference in exponential families. We now outline how these may be extended to general models, using the tangent exponential model (10.8). As noted there, the tangent exponential model has the structure of an exponential family model, with canonical parameter φ and cumulant generator $-\ell(\varphi; y^\circ)$. To compare (10.8) to the general exponential family density (10.29), we take the baseline density $h(s)$ to be $h(s; \hat{\theta}^\circ)$, which effectively centers the score variable s so that $s^\circ = 0$. Thus we have

$$f_{\text{TEM}}(s \mid a; \theta) = \exp\{s^\top \varphi(\theta; y^\circ) + \ell(\theta; y^\circ)\} h(s; \hat{\theta}^\circ), \tag{10.35}$$

with canonical variable s and canonical parameter $\varphi(\theta) \equiv \varphi(\theta; y^\circ)$ defined locally at y° , as we now describe.

In an ordinary exponential family, the canonical parameter can be obtained (up to linear transformations) from $\partial \ell(\theta; s) / \partial s$. In a general model, $\varphi(\theta)$ is also constructed by differentiating in the sample space, now using the $n \times p$ matrix V of sufficient directions of Section 10.3. As shown there, projecting onto the space spanned by these directions implements conditioning on an approximate ancillary statistic without requiring an explicit form for that ancillary. In more detail,

$$\varphi(\theta; y^\circ) = \ell_{;V}(\theta; y^\circ), \tag{10.36}$$

where

$$\ell_{;V}(\theta; y^\circ) = \left. \frac{d}{dt} \ell(\theta; y^\circ + Vt) \right|_{t=0} = V^\top \left. \frac{\partial \ell(\theta; y)}{\partial y} \right|_{y=y^\circ} = \sum_{j=1}^n V_j^\top \frac{\partial \ell(\theta; y_j^\circ)}{\partial y_j},$$

with $t = (t_1, \dots, t_p)^\top$; the last equality applies when the y_j make independent contributions $\ell(\theta; y_j)$ to the log likelihood. Note that both $\ell(\theta; y^\circ) = \log f(y^\circ; \theta)$ and $\varphi(\theta; y^\circ)$ are computed from the original model $f(y; \theta)$.

The validity of this local approximation can be established by Taylor series approximation, as described in Section 10.6. Taking (10.35) as our starting point, we simply apply the formulae derived for the general exponential family model in Section 10.4.3, computing $r(\psi)$ at (10.27) and $q(\psi)$ at (10.34). The approximate significance function for an interest parameter ψ is $\Phi\{r^*(\psi)\}$, using (10.22).

The tangent model approximates the original density in a neighbourhood of the observed data point y° ; more precisely the original model and the tangent exponential model have the same observed log-likelihood function, and the same sample-space derivative of the log-likelihood function. This turns out to be enough to ensure that the significance function is accurate to $O(n^{-3/2})$.

The saddlepoint approximation to (10.35) is

$$f_{\text{TEM}}(s \mid a; \theta) \doteq c |j(\hat{\varphi})|^{-1/2} \exp[s^\top \{\varphi(\theta) - \varphi(\hat{\theta}^\circ)\} + \ell(\theta; y^\circ) - \ell(\hat{\theta}^\circ; y^\circ)], \tag{10.37}$$

where $j(\varphi)$ is the observed Fisher information computed in the φ parametrisation as at (10.33), and $\ell(\theta; y^\circ) = \ell\{\theta(\varphi); y^\circ\}$. The validity of this saddlepoint approximation is verified by Taylor-series expansion of the log likelihood $\ell(\theta; s)$ about $\hat{\theta}^\circ$ and s° , described in Section 10.6.

10.5.2 Generalising V

The matrix V of sufficient directions is central to approximations based on the tangent exponential model. Our previous discussion has taken V to equal $\partial y / \partial \theta^T$, evaluated at the observed values y^o of the data y and the maximum likelihood estimate $\hat{\theta}^o$, but this expression presupposes that the individual observations y_j are scalar and can be differentiated with respect to θ ; the latter is not the case for discrete responses. In this section we outline how V can be generalised.

As a preliminary remark, notice that the span of V is invariant to smooth invertible reparametrisation $\theta \mapsto \theta'$, which has the effect of right-multiplying V by an invertible $p \times p$ matrix of constants; see (10.13). It is readily checked that such a multiplication leaves (10.34) unchanged, so V can be computed in whatever parametrisation is simplest.

When the y_j are continuous and independent vectors of possibly different dimensions d_j we can write

$$\varphi(\theta) = \sum_{j=1}^n \frac{\partial y_j^T}{\partial \theta} \Big|_{y_j=y_j^o, \theta=\hat{\theta}^o} \times \frac{\partial \ell(\theta; y_j)}{\partial y_j} \Big|_{y_j=y_j^o} = \sum_{j=1}^n V_j^T \frac{\partial \ell(\theta; y_j)}{\partial y_j} \Big|_{y_j=y_j^o},$$

say, where V_j has dimension $d_j \times p$. This effectively replaces the matrix V by a tensor.

Example 10.3

If $\{(y_{1j}, y_{2j}), j = 1, \dots, n\}$ are independent pairs from a bivariate normal distribution with zero means, unit variances and covariance θ , V can be constructed using the pivotal quantities $z_{1j} = (y_{1j} + y_{2j})^2 / \{2(1 + \theta)\}$ and $z_{2j} = (y_{1j} - y_{2j})^2 / \{2(1 - \theta)\}$, leading to

$$V_j = \left(\frac{y_{2j}^o - \hat{\theta}^o y_{1j}^o}{1 - \hat{\theta}^{o2}}, \frac{y_{1j}^o - \hat{\theta}^o y_{2j}^o}{1 - \hat{\theta}^{o2}} \right), \quad j = 1, \dots, n,$$

and thus to

$$\varphi(\theta) = \ell_{;V}(\theta) = \frac{\theta(t^o - \hat{\theta}^o s^o) - (s^o - \hat{\theta}^o t^o)}{(1 - \theta^2)(1 - \hat{\theta}^{o2})},$$

where $t = \sum (y_{1j}^2 + y_{2j}^2) / 2$ and $s = \sum y_{1j} y_{2j}$. The sufficient statistics (s, t) emerge naturally in the construction of $\varphi(\theta)$. If a preliminary reduction to sufficiency is made, the resulting V is a 2×1 vector instead of a $2n \times 1$ vector as above, though $\varphi(\theta)$ is unchanged. The sample space contours determined by V are illustrated in Reid (2003), and the accuracy of the normal approximation to the distribution of r^* is illustrated in Reid (2005). \square

Computing V for discrete responses is more awkward. In Section 10.3.2 we saw that in the continuous case, total differentiation of the pivot $F_j(y_j; \theta)$ led to the expression (10.16) for V_j , but $\partial F_j(y; \theta) / \partial y = 0$ almost everywhere in the discrete case. To deal with this, note that in a continuous exponential family model with canonical observation y_j , we can write $\partial \ell(\theta; y_j) / \partial y_j = \alpha_j(\theta)$, say, and observe that as

$$\begin{aligned} n^{-1} \varphi(\theta) &= n^{-1} \sum_{j=1}^n \frac{\partial y_j^T}{\partial \theta} \alpha_j(\theta) \\ &= n^{-1} \sum_{j=1}^n \mathbb{E} \left(\frac{\partial y_j^T}{\partial \theta} \right) \alpha_j(\theta) + n^{-1} \sum_{j=1}^n \left\{ \frac{\partial y_j^T}{\partial \theta} - \mathbb{E} \left(\frac{\partial y_j^T}{\partial \theta} \right) \right\} \alpha_j(\theta), \end{aligned}$$

and each term in the final average has mean zero, that term is $O_p(n^{-1/2})$. Thus if the order of integration and differentiation can be interchanged, $n^{-1}\varphi(\theta)$ can be replaced with

$$n^{-1}\tilde{\varphi}(\theta) = n^{-1} \sum_{j=1}^n \left. \frac{\partial \mathbb{E}(y_j^T)}{\partial \theta} \right|_{\theta=\hat{\theta}^\circ} \alpha_j(\theta),$$

at the expense of introducing an $O_p(n^{-1/2})$ error. Hence using $\tilde{\varphi}(\theta)$ in (10.34) does not change the $O(n^{-1})$ error in (10.24). The use of $\partial \mathbb{E}(y_j^T)/\partial \theta$ replaces the sufficient directions with their expectation, which is only tangential to \mathcal{A}° on average, so renormalisation no longer reduces the order of error to $n^{-3/2}$: inference accurate to third order is unavailable.

As the expectation of a discrete response is typically continuous in the parameters, this approach can be applied to discrete exponential family models such as the Poisson, binomial and multinomial. Extension to more general discrete response distributions entails replacing $\alpha_j(\theta)$. Davison et al. (2006) show that one can use a locally-defined score variable w_j , and

$$w_j = \left. \frac{\partial \ell(\theta; y_j)}{\partial \theta} \right|_{\theta=\hat{\theta}^\circ}, \quad V_j = \left. \frac{\partial \mathbb{E}(w_j; \theta)}{\partial \theta^T} \right|_{\theta=\hat{\theta}^\circ}, \quad \varphi(\theta) = \sum_{j=1}^n V_j^T \frac{\partial \ell(\theta; y_j)}{\partial w_j}.$$

Here w_j has dimension $p \times 1$, so V_j is a $p \times p$ matrix that is easily seen to be the contribution from y_j to the expected information matrix, evaluated at $\hat{\theta}^\circ$. The derivative $\partial \ell(\theta; y_j)/\partial w_j$ is most easily computed as $\partial \ell(\theta; y_j)/\partial y_j \times (\partial w_j/\partial y_j)^{-1}$.

Skovgaard (1996) derived a version of r^* that replaces q in (10.22) with a quantity that is computed entirely from cumulants of the log likelihood. The resulting approximation has a relative error that is $O(n^{-1})$ in a large deviation region about the maximum likelihood estimator. This provides highly accurate results far out into the tails of its distribution, which Skovgaard argues may be of more practical value than higher, $O(n^{-3/2})$, accuracy near its mean. Reid and Fraser (2010) show that Skovgaard’s approximation can be related to a tangent exponential model with canonical parameter determined from the derivative of $I(\theta; \hat{\theta}^\circ) = \int \ell(\theta; y) f(y; \hat{\theta}^\circ) dy$. It can also be used for discrete models, and gives the same approximation as Davison et al. (2006) in curved exponential families, but not more generally.

10.6 Derivation of the Tangent Exponential Model

10.6.1 Preliminary Remarks

The expression for the tangent exponential model, (10.8), is concise and emphasises the connection to exponential family models and the role of φ as a canonical parameter, but does not lend itself to ready understanding. It can be derived using Taylor series approximations that can be given explicitly when y and θ are scalar and provide some theoretical illumination. The development for higher dimensions is similar but much more laborious and does not yield additional insights.

We have seen above that the tangent exponential model requires computation of a first derivative in the sample space. We can think of this as assessing how the log likelihood $\ell(\theta; y^\circ)$ changes not only as a function of θ , as is standard in both likelihood and Bayesian inference, but also as a function of y , in a small neighbourhood of the observed data point y° .

Employing this first derivative probes $\ell(\theta; y)$ more deeply than simply using the observed log likelihood function $\ell(\theta; y^\circ)$, but does not involve computing the log likelihood function on its entire domain, i.e., for $(\theta, y) \in \mathbb{R}^p \times \mathbb{R}^n$.

10.6.2 No Nuisance Parameters

Suppose that we have a model $f(s; \theta)$, $s \in \mathbb{R}$ and $\theta \in \mathbb{R}$, and that there is an implicit dependence on n , in the sense that $\ell(\theta; s) = \log f(s; \theta)$ is $O_p(n)$. We will address how to get this reduction in general at the end of this section, but this would be the case for example if s was the sufficient statistic based on a random sample from a linear exponential family model, and it is also the case for models like the Cauchy, where s is the maximum likelihood estimator or any other location-equivariant estimator of θ , based on a sample of size n , and the distribution is conditional on the $(n - 1)$ -dimensional ancillary statistic a .

We first expand the log likelihood $\log f(s; \theta)$ in a Taylor series in both s and θ , about the fixed points s° and $\hat{\theta}^\circ$, giving

$$\begin{aligned} \ell(\theta; s) &= \log f(s; \theta) \\ &= \ell(\hat{\theta}^\circ; s^\circ) + (s - s^\circ)\ell_s(\hat{\theta}^\circ; s^\circ) + (\theta - \hat{\theta}^\circ)\ell_\theta(\hat{\theta}^\circ; s^\circ) \\ &\quad + \frac{1}{2}(s - s^\circ)^2\ell_{ss}(\hat{\theta}^\circ; s^\circ) + (s - s^\circ)(\theta - \hat{\theta}^\circ)\ell_{s\theta}(\hat{\theta}^\circ; s^\circ) \\ &\quad + \frac{1}{2}(\theta - \hat{\theta}^\circ)^2\ell_{\theta\theta}(\hat{\theta}^\circ; s^\circ) + \dots \\ &= \sum_{i,j=0}^{\infty} \frac{1}{i!j!}(\theta - \hat{\theta}^\circ)^i(s - s^\circ)^j b_{ij}, \end{aligned} \tag{10.38}$$

say, where $b_{ij} = \partial^i \partial^j \ell(\hat{\theta}^\circ; s^\circ) / \partial \theta^i \partial s^j$ for $i, j = 0, 1, \dots$; note that $b_{10} = 0$.

Now consider how expansion (10.38) would differ if $f(s; \theta)$ were an exponential family model with canonical parameter θ and sufficient statistic s . In this case

$$\log f(s; \theta) = \ell(\theta; s) = \theta s - \kappa(\theta) - d(s),$$

so the coefficients b_{i0} would be those for a Taylor series expansion of $-d(s)$, the b_{0j} would be those for a Taylor series expansion of $-\kappa(\theta)$, and the only other non-zero term would be $b_{11} = 1$.

Andrews et al. (2005) showed that for any continuously differentiable model $f(s; \theta)$, there exists a transformation $x = x(s)$ and $\varphi = \varphi(\theta)$ such that the expansion of $\log f(x; \varphi)$ has the coefficient array starting with $i = j = 0$ at the top left and terms b_{ij} in row i and column j , given by

$$\begin{pmatrix} b + \frac{3\alpha_4 - 5\alpha_3^2 - 12\gamma}{24n} & \frac{-\alpha_3}{2n^{1/2}} & -\left(1 + \frac{\alpha_4 - 2\alpha_3^2 - 5\gamma}{2n}\right) & \frac{\alpha_3}{n^{1/2}} & \frac{\alpha_4 - 3\alpha_3^2 - 6\gamma}{n} \\ 0 & 1 & 0 & 0 & \\ -1 & 0 & \frac{\gamma}{n} & & \\ -\frac{\alpha_3}{n^{1/2}} & 0 & & & \\ -\frac{\alpha_4}{n} & & & & \end{pmatrix}. \tag{10.39}$$

Only terms up to $O(n^{-1})$ are shown: the terms in the blank spaces are $O(n^{-3/2})$ or smaller, as are the terms implicitly omitted. The constants α_3 , α_4 and γ are the derivatives of $\log f(x; \varphi)$ at $x = 0$ and $\varphi = 0$:

$$\begin{aligned} \alpha_3 &= - \left. \frac{\partial^3 \log f(x; \varphi)}{\partial \varphi^3} \right|_{x=0, \varphi=0}, \\ \alpha_4 &= - \left. \frac{\partial^4 \log f(x; \varphi)}{\partial \varphi^4} \right|_{x=0, \varphi=0}, \\ \gamma &= \left. \frac{\partial^4 \log f(x; \varphi)}{\partial x^2 \partial \varphi^2} \right|_{x=0, \varphi=0}, \end{aligned}$$

and γ is related to the exponential curvature of the model (Efron, 1975). Ignoring terms of $O(n^{-3/2})$, the expansion (10.39) in terms of x and φ is almost that of an exponential family model; the only additional coefficient is the (2, 2) entry γ/n , which adds a term $\gamma\varphi^2x^2/(4n)$ to the log likelihood expansion.

The variables s and θ are both scaled and centered as part of the transformation to x and φ : note that the observed information $-\partial^2 \log f(0; 0)/\partial \varphi^2 = -b_{20} = 1$. The point $(x, \varphi) = (0, 0)$ corresponds to the original point of expansion $(s^\circ, \hat{\theta}^\circ)$, where s° is the observed value and $\hat{\theta}^\circ = \hat{\theta}(s^\circ)$ is the corresponding value of the maximum likelihood estimator.

Another way to write the model given by (10.39) is

$$\log f(x; \varphi) = b_{00} + P_{1n}(x) + P_{2n}(\varphi) + x\varphi + \gamma x^2 \varphi^2 / (4n) + O(n^{-3/2}),$$

where $P_{1n}(x)$ is given by the first row of the array and $P_{2n}(\varphi)$ by its first column, each omitting b_{00} . On examining the elements of (10.39) we see that we can write

$$f(x; \varphi) \propto \phi(x - \varphi) \left\{ 1 + \frac{a_1(x, \varphi)}{n^{1/2}} + \frac{a_2(x, \varphi)}{n} + \gamma \frac{a_3(x, \varphi)}{n} + O(n^{-3/2}) \right\}, \quad (10.40)$$

in terms of the standard normal density function ϕ and suitable polynomials a_1 , a_2 and $a_3(x, \varphi) = (x^2\varphi^2 - x^4 + 5x^2 - 2)/4$. Equation (10.40) can be integrated term by term with respect to x ; an explicit array for the resulting approximate distribution function $F(x; \varphi)$ is given in Andrews et al. (2005). Remarkably, although $F(x; \varphi)$ depends on γ in general, $F(0; \varphi)$ does not depend on γ , because $\phi(x - \varphi)a_3(x; \varphi)$ has integral zero over the negative half-line. As $x = 0$ corresponds to $s = s^\circ$, the significance function $F(s^\circ; \theta)$ does not depend on γ and can be computed using the exponential family version of (10.39), in which $\gamma = 0$.

The tangent exponential model (10.8) is just an invariant version of the simplified expansion (10.40): s is now the transformed variable called x here, and is the corresponding score variable, and φ is by definition $\partial \log f_{\text{TEM}}(0; 0)/\partial x$, the canonical parameter of the exponential model approximation.

The steps in going from the original model to (10.39) are outlined in Andrews et al. (2005) and Cakmak et al. (1998).

The reduction to a single variable s in a scalar parameter model is straightforward if, for example, $y = (y_1, \dots, y_n)$ is a sample from an exponential family model, with density function (10.17), as the log likelihood $\ell(\theta; s)$ then equals $\exp\{\varphi(\theta)s - n c(\theta)\}$, where $s = \sum_{i=1}^n s(y_i)$, and has the dependence on n summarised in (10.39), with $\gamma = 0$.

Similarly, in the case of a sample $y = (y_1, \dots, y_n)$ from a location model, the exact distribution of any location-invariant estimator, say s , of the location parameter θ given the location ancillary statistic $a = (y_1 - s, \dots, y_n - s)$ is

$$f(s | a; \theta) = \exp\{\ell(\theta; s, a)\} \bigg/ \int_t \exp\{\ell(t; s, a)\} dt,$$

and (10.39) is equivalent to the density approximation arising when Laplace's method is applied to the denominator integral.

For more general models, the discussion in Section 10.3 establishes the existence of a conditional distribution on \mathbb{R} that can be determined by finding the $n \times 1$ vector V of sufficient directions. The arguments above show that this conditional distribution, which now has a scalar variable and scalar parameter, is effectively an exponential family for the purpose of approximating the significance function.

10.6.3 Nuisance Parameters

The expansion in (10.39) can be generalised to vector parameters, as in Cakmak et al. (1994) and Fraser and Reid (1993), but the notation is cumbersome, and the various multi-dimensional analogues to $\alpha_3, \alpha_4, \gamma$ are not explicitly available. However, the expansion verifies that the coordinate-free version of the tangent exponential model has the form given at (10.8) and (10.35), with saddlepoint approximation (10.37).

This gives a tangent exponential model on \mathbb{R}^p for inference about θ , which is implicitly conditioned on an approximate ancillary statistic through the use of the $n \times p$ matrix V , and this is now the full model used to obtain an approximate significance function for a scalar parameter of interest ψ .

In this full model, consider fixing ψ , and constructing a new tangent exponential model on \mathbb{R}^{p-1} with parameter λ . We can write, suppressing the conditioning on a ,

$$f_{\text{TEM}}(s; \theta) = f_{1, \text{TEM}}(s_\psi \mid \tilde{a}_\psi; \lambda) f_2(\tilde{a}_\psi), \quad (10.41)$$

where \tilde{a}_ψ is a new approximate ancillary statistic for the model with ψ held fixed. This gives us a one-dimensional distribution for inference about ψ ,

$$f_2(\tilde{a}_\psi) = f_{\text{TEM}}(s; \theta) / f_{1, \text{TEM}}(s_\psi \mid \tilde{a}_\psi; \lambda),$$

and as we know the left-hand side is free of both s_ψ and λ , we can choose $s_\psi = 0$ and $\lambda = \hat{\lambda}_\psi$. Using the saddlepoint form (10.37) of the tangent exponential model in the numerator and denominator yields a model on \mathbb{R} of the form

$$h(s; \psi) = c \exp\{\ell(\hat{\varphi}_\psi; s) - \ell(\hat{\varphi}; s)\} |J_{\varphi\varphi}(\hat{\varphi})|^{-1/2} |J_{(\lambda\lambda)}(\hat{\varphi}_\psi)|^{1/2}, \quad s \in \mathcal{L}_\psi, \quad (10.42)$$

where $\ell(\varphi; s) = s^T \varphi + \ell(\varphi; y^o)$, and \mathcal{L}_ψ is a line in the sample space corresponding to fixing $\hat{\lambda}_\psi$ (and a). Expressing the result with a constraint on s avoids explicit identification of \tilde{a}_ψ : it is enough to know that it exists. As in Section 10.4.3 the information determinants are computed in the φ parameterization; see (10.33).

The right-hand-side of (10.42) has the form of our original tangent exponential model (10.37), with an adjustment factor in the ratio of determinants; note also the similarity to the approximate conditional density (10.26) for linear exponential families. As a result, the approximate significance function is the same as that for general exponential families outlined in Section 10.4.3, with the significance function as in (10.24) or (10.25), with r defined in (10.27), and q defined in (10.34). Once the tangent exponential approximation to the original model has been established, the exponential model formulas apply directly.

It would be natural to partition s into a component related to ψ and one related to λ , and this is how the result is presented in Fraser and Reid (1995, Sec. 6). In later work (Reid and Fraser, 2010; Fraser et al., 2016b) the simpler notation of (10.42) is preferred, with a constraint on s to emphasise that the density is for a variable of the same dimension as the parameter of interest ψ .

10.7 Concluding Remarks

10.7.1 Summary

The tangent exponential model and associated significance function implement inference conditional on an approximate ancillary statistic, followed by marginalization to a pivotal quantity, $r^*(\psi; y)$, for a scalar parameter of interest. This pivot is readily computed using only $\ell(\theta; y^o)$ and $\varphi(\theta; y^o) = \ell_{;V}(\theta; y^o)$, and the full and constrained maximum likelihood estimators $\hat{\theta}$ and $\hat{\theta}_\psi$. Fraser et al. (1999) and Reid (2003) present this “inference algorithm” as two dimension-reduction steps: from a model $f(y; \theta)$ on \mathbb{R}^n to a model $f(s | a; \theta)$ on \mathbb{R}^p , by conditioning, and from this model to another on \mathbb{R} , by marginalizing. The model on \mathbb{R} can be approximated by a simple standard normal distribution for the pivotal quantity $r^*(\psi; y)$, and in continuous models the approximation to the significance function based on $r^*(\psi; y^o)$ has relative error $O(n^{-3/2})$.

The final approximation step is somewhat separate from the development of the model on \mathbb{R} , and follows closely the derivation of the r^* approximation in Barndorff-Nielsen (1986). It can also be applied in other contexts, and in particular to approximation of a Bayesian posterior survivor function, starting from the the Laplace approximation to the posterior marginal density (Tierney and Kadane, 1986).

As our focus here is on the steps leading to the tangent exponential model and their implications for inference, we have not included numerical work indicating the accuracy of the approximations. There are many examples and exercises in Brazzale et al. (2007), in the literature referred to there and in Brazzale and Davison (2008).

There is a close relation between the r^* approximation and the parametric bootstrap; the higher-order properties of the latter are investigated in DiCiccio and Young (2008), Lee and Young (2005), and DiCiccio et al. (2015). To achieve the same order of accuracy it is necessary to bootstrap under the constrained maximum likelihood estimate $(\psi, \hat{\lambda}_\psi)$, which increases the computational burden. Fraser and Rousseau (2008) also consider the relationship between significance functions based on the parametric bootstrap, on r^* , and on Bayesian versions of predictive p -values.

10.7.2 Extensions

If the parameter of interest is a vector, a significance function is not easily obtained unless one can construct a scalar measure of departure such as the log likelihood ratio statistic $w(\psi) = 2\{\ell(\hat{\theta}) - \ell(\hat{\theta}_\psi)\}$, Wald statistic $(\hat{\psi} - \psi)^\top j^{\psi\psi}(\hat{\theta})(\hat{\psi} - \psi)$, or score statistic, each to first order approximately distributed as χ_d^2 . Davison et al. (2014) and Fraser et al. (2016b) use the tangent exponential model as the building block for a directional approach to inference for a d -dimensional parameter ψ which creates a univariate summary, by considering the magnitude of ψ conditional on its direction from a null value ψ_0 . The saddlepoint approximation to $f_{\text{TEM}}(s | a; \theta)$ on this line in the sample space forms the basis for inference. A new scalar-parameter exponential family is constructed from the multi-parameter exponential family model or the approximating tangent exponential model.

Our discussion has presumed that the underlying data are independent, but the geometric motivation in Section 10.3 suggests that the approach should provide improved accuracy more generally. Belzile and Davison (2022) adapt the approach for discrete responses to the inhomogeneous Poisson process, but this is a special case owing to its independence properties. The main difficulty in broader settings is to compute V , and from this the constructed parameter $\varphi(\theta)$. In a time series setting, a series of pivotal quantities may be

generated from the predictive distributions $F(y_j | y_{j-1}, \dots, y_1; \theta)$ for $j = 2, \dots, n$, using martingale differences or a lower triangular square root of the covariance matrix for the response (Fraser et al., 2005; Lozada-Can and Davison, 2010). It is not yet clear whether other decompositions of the covariance matrix would lead to asymptotically equivalent results.

There is an r^* approximation for Bayesian inference, readily obtained from the Laplace approximation, as mentioned above; see also Fraser et al. (1999). This provides a route to examining the discrepancy between posterior survivor functions and significance functions. Equating the two versions of r^* leads to a data-dependent prior that ensures agreement of the significance and survivor functions up to terms of $O(n^{-1})$. The former was emphasised in Fraser (2011) and Fraser et al. (2016a); the latter formed the basis for a discussion of default priors in Fraser et al. (2010b).

10.7.3 A Brief Historical Note

Fraser viewed the dimension-reduction steps in Section 10.7.1 as essentially unique, and consequently the pivotal quantity r_ψ^* not as an arbitrary choice among several potential pivotal quantities, but as the only route to higher-order approximation for a scalar parameter in the presence of nuisance parameters:

This ancillary density is uniquely determined by steps that retain continuity of the model in the derivation of the marginal distribution. It thus provides the unique null density for assessing a value $\psi = \psi_0$, and anyone suggesting a different null distribution would need to justify inserting discontinuity where none was present (Fraser, 2017, Section 4).

The continuity referred to there is the presumption that changes in θ are smoothly related to changes in y and vice-versa, as in a pure location model $f(y - \theta)$. The vectors V determining the tangent plane to the ancillary surface are based on the local location model defined in Fraser (1964). Suppose y_i has density $f(y_i; \theta)$ and cumulative distribution function $F(y_i; \theta)$, $\theta \in \mathbb{R}$. Define a transformation $y_i \mapsto x_i$ by setting

$$x_i = \int^{y_i} -\frac{F_\theta(y; \theta_0)}{F_y(y; \theta_0)} dy,$$

where θ_0 is some fixed value, and write the density of x_i as $g(x_i; \theta)$. Then $g\{x_i - (\theta - \theta_0); \theta_0\}$ coincides with the true density at θ_0 , as does its derivative $g_x(x_i; \theta_0)$. For a sample y_1, \dots, y_n , this local location model has an ancillary statistic $(x_1 - \bar{x}, \dots, x_n - \bar{x})$, and sufficient direction $(1, \dots, 1)$, which transforms back to the sufficient direction

$$V^\top = -\left(\frac{F_\theta(y_1; \theta_0)}{F_y(y_1; \theta_0)}, \dots, \frac{F_\theta(y_n; \theta_0)}{F_y(y_n; \theta_0)}\right)$$

in terms of y_1, \dots, y_n ; see (10.16). As noted in Fraser and Reid (2001), this construction does not give a local location model for the full sample y_1, \dots, y_n because the vector field $V(y)$ is not guaranteed to be integrable. But the expansions in that paper verify that the approximations derived from the tangent exponential model are still valid, as the sufficient directions V describe the same tangent plane as a second-order ancillary statistic that exists under mild regularity conditions. Fraser and Reid (2001, Section 10.3) promised that “the integrability of the $V(y)$ to the required order will be examined elsewhere”, and this was fulfilled in Fraser et al. (2010a).

Fraser viewed as intrinsically linked the construction of the tangent exponential model, the application of the saddlepoint approximation and the construction of significance functions

as key inferential summaries. A first, lengthy, paper written shortly after the simpler developments in Fraser (1988, 1990, 1991) included all these pieces, and was met with some puzzlement by editors and reviewers: one reviewer advised “it should probably be several papers” — a reaction that might be rather unusual nowadays. This led to the asymptotic expansions being the focus of Fraser and Reid (1993), although much of the original draft was published in Fraser and Reid (1995). That latter paper derived the tangent exponential approximation to general models, derived the directional vectors V from a local location model, showed the existence of a second-order ancillary statistic with the same directional vectors, verified that the dimension of this ancillary is fixed as $n \rightarrow \infty$, and derived the r^* approximation in its general form. The construction of the directional vectors was discussed in more detail in Fraser and Reid (2001), which is confusingly referred to in some of his papers as Fraser and Reid (1999).

The annotations in the bibliography below attempt to provide a road map through the most relevant of these papers. Copies of the less readily accessible ones are posted at

<https://utstat.toronto.edu/reid/fraser-papers.html>

Acknowledgements

The work was supported by the Swiss National Science Foundation and the Natural Sciences and Engineering Council of Canada. We thank Léo Belzile, Yanbo Tang and two anonymous referees for helpful comments.

Bibliography

- Andrews, D. A., Fraser, D. A. S. and Wong, A., 2005. Computation of distribution functions from likelihood information near observed data. *Journal of Statistical Planning and Inference*, 134, pp. 180–193. Detailed derivation of Taylor series expansions for density and distribution functions. Corrections to some formulae available at <https://utstat.toronto.edu/dfraser/documents/210Typo-corr.pdf>.
- Barndorff-Nielsen, O. E., 1980. Conditionality resolutions. *Biometrika*, 67, pp. 293–310.
- Barndorff-Nielsen, O. E., 1983. On a formula for the distribution of the maximum likelihood estimator. *Biometrika*, 70, pp. 343–365.
- Barndorff-Nielsen, O. E., 1986. Inference on full or partial parameters based on the standardized signed log likelihood ratio. *Biometrika*, 73, pp. 307–322.
- Barndorff-Nielsen, O. E. and Cox, D. R., 1989. *Asymptotic Techniques for Use in Statistics*. London: Chapman & Hall.
- Barndorff-Nielsen, O. E. and Cox, D. R., 1994. *Inference and Asymptotics*. London: Chapman & Hall.
- Belzile, L. R. and Davison, A. C., 2022. Improved inference on risk measures for univariate extremes. *Annals of Applied Statistics*, 16, pp. 1524–1549.
- Brazzale, A. R. and Davison, A. C., 2008. Accurate parametric inference for small samples. *Statistical Science*, 23, pp. 465–484.

- Brazzale, A. R., Davison, A. C. and Reid, N., 2007. *Applied Asymptotics: Case Studies in Small Sample Statistics*. Cambridge: Cambridge University Press.
- Butler, R. W., 2007. *Saddlepoint Approximations with Applications*. Cambridge: Cambridge University Press.
- Cakmak, S., Fraser, D. A. S., McDunnough, P., Reid, N. and Yuan, X., 1998. Likelihood centered asymptotic model: Exponential and location model versions. *Journal of Statistical Planning and Inference*, 66, pp. 211–222. One of a series of papers giving expansions of likelihood quantities in various versions; here for a scalar parameter and scalar response.
- Cakmak, S., Fraser, D. A. S. and Reid, N., 1994. Multivariate asymptotic model: Location and exponential approximations. *Utilitas Mathematica*, 46, pp. 21–31. Extends the expansion given in 10.6.1 above to p -dimensional parameter and p -dimensional response. A location model version is also developed, which was used in the development of so-called default priors.
- Cox, D. R., 1958. Some problems connected with statistical inference. *Annals of Mathematical Statistics*, 29, pp. 357–372.
- Daniels, H. E., 1954. Saddlepoint approximations in statistics. *Annals of Mathematical Statistics*, 25, pp. 631–650.
- Daniels, H. E., 1956. The approximate distribution of serial correlation coefficients. *Biometrika*, 43, pp. 169–185.
- Davison, A. C., 2003. *Statistical Models*. Cambridge: Cambridge University Press.
- Davison, A. C., Fraser, D. A. S. and Reid, N., 2006. Improved likelihood inference for discrete data. *Journal of the Royal Statistical Society: Series B*, 68, pp. 495–508. Derives an approximation to V based on the derivative of the expected value of the response, as described in Section 10.5.2.
- Davison, A. C., Fraser, D. A. S., Reid, N. and Sartori, N., 2014. Accurate directional inference for vector parameters in linear exponential families. *Journal of the American Statistical Association*, 109, pp. 302–314. Uses the tangent exponential model and a conditional argument to construct approximations to P -values for assessing a vector parameter of interest.
- DiCiccio, T. J., Kuffner, T. A., Young, G. A. and Zaretzki, R., 2015. Stability and uniqueness of p -values for likelihood-based inference. *Statistica Sinica*, 25, pp. 1355–1376.
- DiCiccio, T. J. and Young, G. A., 2008. Conditional properties of unconditional parametric bootstrap procedures for inference in exponential families. *Biometrika*, 95, pp. 747–758.
- Edwards, A. W. F., 1972. *Likelihood*. Cambridge: Cambridge University Press.
- Efron, B., 1975. Defining the curvature of a statistical problem (with applications to second order efficiency). *Annals of Statistics*, 3, pp. 1189–1242.
- Efron, B., 1993. Bayes and likelihood calculations from confidence intervals. *Biometrika*, 80, pp. 3–26.

- Fraser, A. M., Fraser, D. A. S. and Staicu, A.-M., 2010a. Second order ancillary: A differential view from continuity. *Bernoulli*, 16, pp. 1208–1223. Resolves concerns about non-uniqueness of the approximate ancillary determined using the tangent exponential model.
- Fraser, D. A. S., 1964. Local conditional sufficiency. *Journal of the Royal Statistical Society: Series B*, 26, pp. 52–62.
- Fraser, D. A. S., 1988. Normed likelihood as saddlepoint approximation. *Journal of Multivariate Analysis*, 27, pp. 181–193. Studies the relationship of the p^* and saddlepoint approximations for the distribution of the score function. Introduction of the tangent exponential model for scalar y and scalar θ .
- Fraser, D. A. S., 1990. Tail probabilities from observed likelihoods. *Biometrika*, 77, pp. 65–76. Shows the relationship of the p^* and saddlepoint approximations in exponential families, and discusses how this generalizes. Gives an expression for the tangent exponential model when y and θ have the same dimension.
- Fraser, D. A. S., 1991. Statistical inference: Likelihood to significance. *Journal of the American Statistical Association*, 86, pp. 258–265. Based on Fisher Lecture at 1990 Joint Statistical Meetings; proposes significance functions as an encompassing inferential method.
- Fraser, D. A. S., 2004. Ancillaries and conditional inference (with discussion). *Statistical Science*, 19, pp. 333–369. Detailed discussion of conditioning as the main focus for inference, with many examples of exact ancillary statistics and a discussion of the approximate ancillarity underlying the tangent exponential model.
- Fraser, D. A. S., 2011. Is Bayes posterior just quick and dirty confidence? *Statistical Science*, 26, pp. 299–316. Argues that the use of Bayesian inference with convenience priors could be adequate to first order but potentially misleading in finite samples.
- Fraser, D. A. S., 2017. p -values: The insight to modern statistical inference. *Annual Review of Statistics and Its Application*, 4, pp. 1–14. Fraser continued to simplify and refine his approach, here giving a concise but clear account of the material covered in this chapter.
- Fraser, D. A. S., 2019. The p -value function and statistical inference. *The American Statistician*, 73, pp. 135–147. This is in a special issue on the role of “ $p < 0.05$ ” in contributing to a lack of replicability of scientific work. This paper emphasizes the significance, or p -value, function as preferable to a single p -value.
- Fraser, D. A. S., Bédard, M., Wong, A., Lin, W. and Fraser, A. M., 2016a. Bayes, reproducibility and the quest for truth. *Statistical Science*, 31, pp. 578–590.
- Fraser, D. A. S. and McDunnough, P., 1984. Further remarks on asymptotic normality of likelihood and conditional analyses. *Canadian Journal of Statistics*, 12, pp. 183–190.
- Fraser, D. A. S. and Reid, N., 1988. On conditional inference for a real parameter: A differential approach on the sample space. *Biometrika*, 75, pp. 251–264.
- Fraser, D. A. S. and Reid, N., 1993. Third order asymptotic models: Likelihood functions leading to accurate approximations to distribution functions. *Statistica Sinica*, 3, pp. 67–82. Taylor expansion of §10.6.1 first published; also a multivariate version. In the scalar parameter case the r^* approximation based on the tangent exponential model established. Builds on Fraser (1990) where the p^* and tangent exponential model approximations are

connected, for scalar y and scalar θ . Nuisance parameters tackled in linear exponential and location models; suggests using r^* based on either approximate conditional log likelihood or approximate marginal log-likelihood.

- Fraser, D. A. S. and Reid, N., 1995. Ancillaries and third order significance. *Utilitas Mathematica*, 47, pp. 33–53. Most of the main theoretical results first appear here, re-worked in later papers. Uses the Taylor expansion as in Fraser and Reid (1993), §10.6.1 and Andrews et al. (2005) to derive the saddlepoint version of the tangent exponential model. Shows that only second-order ancillarity is needed for third-order inference, following Skovgaard (1986); that V gives a tangent plane for a fixed ancillary to $O(n^{-1/2})$, but can be 'bent' to get ancillary to $O(n^{-1})$ without changing V , i.e., there exists an $O(n^{-1})$ ancillary W , but V is adequate. Shows the relation to local location model in (Fraser, 1964). An elegant argument in Section 3 shows that the dimension of the ancillary statistic does not grow with sample size.
- Fraser, D. A. S. and Reid, N., 2001. Ancillary information for statistical inference. In *Empirical Bayes and Likelihood Inference*, eds S. E. Ahmed and N. Reid, pp. 185–207. New York: Springer. Clarifies the development of approximate ancillarity in Fraser and Reid (1995), later put on a more rigorous footing in Fraser et al. (2010a).
- Fraser, D. A. S., Reid, N., Marras, E. and Yi, G. Y., 2010b. Default priors for Bayesian and frequentist inference. *Journal of the Royal Statistical Society: Series B*, 72, pp. 631–654.
- Fraser, D. A. S., Reid, N. and Sartori, N., 2016b. Accurate directional inference for vector parameters. *Biometrika*, 103, pp. 625–639.
- Fraser, D. A. S., Reid, N. and Wu, J., 1999. A simple general formula for tail probabilities for frequentist and Bayesian inference. *Biometrika*, 86, pp. 249–264. Spells out the dimension reduction from n to p to 1; compares the r^* approximation using the tangent exponential model to the approach of Barndorff-Nielsen (1986). Derives the Bayesian version of r^* . Also considers models where the parameter of interest is defined implicitly via a constraint on the full parameter θ . Provides several numerical examples.
- Fraser, D. A. S., Rekkas, M. and Wong, A., 2005. Highly accurate likelihood analysis for the seemingly unrelated regression problem. *Journal of Econometrics*, 127, pp. 17–33.
- Fraser, D. A. S. and Rousseau, J., 2008. Studentization and deriving accurate p-values. *Biometrika*, 95, pp. 1–16.
- Hinkley, D. V., 1980. Likelihood as approximate pivotal distribution. *Biometrika*, 67, pp. 287–292.
- Jensen, J. L., 1995. *Saddlepoint Approximations*. Oxford: Oxford University Press.
- Kolassa, J., 2006. *Series Approximation Methods in Statistics* (3rd ed.). New York: Springer.
- Krantz, S. G. and Parks, H. R., 2008. *Geometric Integration Theory*. Boston: Birkhäuser.
- LeCam, L., 1960. Locally asymptotically normal families of distributions. *University of California Public Statistics*, 3, pp. 27–98.
- Lee, S. M. S. and Young, G. A., 2005. Parametric bootstrapping with nuisance parameters. *Statistics & Probability Letters*, 71, pp. 143–153.
- Lozada-Can, C. and Davison, A. C., 2010. Three examples of accurate likelihood inference. *American Statistician*, 64, pp. 131–139.

- Lugannani, R. and Rice, S., 1980. Saddlepoint approximation for the distribution of the sum of independent random variables. *Advances in Applied Probability*, 12, pp. 475–490.
- McCullagh, P., 1987. *Tensor Methods in Statistics*. London: Chapman & Hall.
- Reid, N., 2003. Asymptotics and the theory of inference. *Annals of Statistics*, 31, pp. 1695–1731.
- Reid, N., 2005. Asymptotics and the theory of statistics. In *Celebrating Statistics: Papers in Honour of D.R. Cox*, eds A. C. Davison, Y. Dodge and N. Wermuth, pp. 73–88. Oxford: Oxford University Press.
- Reid, N. and Fraser, D. A. S., 2010. Mean loglikelihood and higher-order approximations. *Biometrika*, 97, pp. 159–170. *Compares inference based on r^* using $\varphi(\theta) = \ell_{;V}(\theta; y^o)$ to the method proposed by (Skovgaard, 1996), which can be formulated as the same approximation using a parametrization $\tilde{\varphi}(\theta)$ that is constructed from likelihood cumulants. The Appendix attempts to clarify the derivation of the tangent exponential model.*
- Royall, R. M., 1997. *Statistical Evidence: A Likelihood Paradigm*. London: Chapman & Hall/CRC.
- Severini, T. A., 2000. *Likelihood Methods in Statistics*. Oxford: Clarendon Press.
- Skovgaard, I. M., 1986. Successive improvement of the order of ancillarity. *Biometrika*, 73, pp. 516–519.
- Skovgaard, I. M., 1996. An explicit large-deviation approximation to one-parameter tests. *Bernoulli*, 2, pp. 145–165.
- Tierney, L. and Kadane, J. B., 1986. Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81, pp. 82–86.
- van der Vaart, A. W., 1998. *Asymptotic Statistics*. Cambridge: Cambridge University Press.
- Xie, M. and Singh, K., 2013. Confidence distributions: The frequentist distribution estimator of a parameter. *International Statistical Review*, 81, pp. 3–39.