Tweedie Family Densities: Methods of Evaluation

Peter K. Dunn[1,2] and Gordon K. Smyth[1]

[1] University of Queensland, St Lucia, Q 4072, Australia.
[2] University of Southern Queensland, Toowoomba, Q 4350, Australia.

Abstract: Two numerical evaluation methods are considered for Tweedie family densities which are then used to assess the accuracy of the saddlepoint approximation. This has implications for the distribution of the residual deviance in generalized linear models. Other applications include residual analysis through quantile residuals.

Keywords: Tweedie densities, saddlepoint approximation, quantile residuals

1 Introduction

The Tweedie family of densities are a versatile family of distributions belonging to the class of exponential dispersion models (EDMs) (see Jørgensen, 1997). For a random variable Y which follows an EDM, we can write the density in one of two ways:

$$\begin{aligned} p_Y(y;\mu,\phi) &= a_p(y,\phi)\exp\{[y\theta-\kappa(\theta)]/\phi\} \\ &= b_p(y,\phi)\exp\{-d(y,\mu)/(2\phi)\} \qquad (1) \end{aligned}$$

where $\mu = \mathrm{E}[Y] = \kappa'(\theta)$ is the mean, $\phi > 0$ is the dispersion parameter, $d(y,\mu)$ is the unit deviance, θ is the canonical parameter, and $\kappa(\theta)$ is the cumulant function. The functions $a_p(y,\phi)$ and $b_p(y,\phi)$ cannot generally be written in closed form apart from some special cases listed below.
The variance is given by $\mathrm{var}[Y] = \phi V(\mu)$ where $V(\mu) = \kappa''(\theta)$ is the variance function viewed as a function of the mean μ. Tweedie family densities are characterised by power variance functions of the form $\mathrm{V}[\mu] = \phi\mu^p$, where $p \in (-\infty, 0] \cup [1,\infty)$ is the index determining the distribution (Jørgensen, 1997). The family includes discrete and continuous densities, as well

as mixed densities. Special cases include the normal ($p = 0$), Poisson ($p = 1$), gamma ($p = 2$) and inverse Gaussian ($p = 3$) distributions. For $1 < p < 2$, the distributions are continuous for $y > 0$ and have a discrete mass at $y = 0$ such that

$$\Pr(Y = 0) = \exp(-\lambda) = \exp\left\{-\mu^{2-p}/[\phi(2-p)]\right\}.$$

Tweedie densities are not known in closed form (apart from the special cases given), but are instead known by their relatively simple cumulant generating function (cgf). The cgf is given by

$$K(t) = [\kappa(\theta + \phi t) - \kappa(\theta)]/\phi,$$

where $\kappa(\theta)$ is the cumulant function,

$$\kappa(\theta) = \begin{cases} \left[\{(1-p)\theta + 1\}^{(2-p)/(1-p)} - 1\right]\Big/(2-p) & \text{for } p \neq 2 \\ -\log(1-\theta) & \text{for } p = 2. \end{cases}$$

2 Two methods of evaluation

Tweedie distributions could be used far more frequently in data analysis if methods of evaluation were more readily available. Being EDMs, the Tweedie distributions fit the generalized linear model framework of Nelder and Wedderburn (1972). Examples of Tweedie models used in generalized linear modelling are given in Smyth (1996) and Gilchrist and Drinkwater (1999). We consider two numerical methods here for evaluation of the Tweedie family densities.

2.1 Inversion of the cgf

Given the simple form of the cgf, one method of evaluation is to use Fourier inversion to invert the cgf using

$$p_Y(y; \mu, \phi) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \exp\{K(it) - ity\}\, dt \qquad (2)$$

where $i = \sqrt{-1}$. In the case $1 < p < 2$, we need to consider the continuous conditional density $Y|Y > 0$ for which we obtain

$$p_{Y|Y>0}(y) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \left\{\frac{M_Y(it) - \exp(-\lambda)}{1 - \exp(-\lambda)}\right\} \exp(-ity)\, dt \qquad (3)$$

and this integration is more difficult. In both cases, the infinite oscillating integral is evaluated by converting it into a series by determining the zeros of the integrand and integrating between them. Analytical analysis of the integrand assists in locating the required zeros and ensuring that the algorithms are known to converge. The convergence is made faster and more reliable using an acceleration technique called the W-transformation due to Sidi (1988) and implemented using the W-algorithm of Sidi (1982).

The cumulative distribution function can also be evaluated using similar techniques after directly integrating (2) or (3) with respect to y. In both cases, convergence is more rapid than for the density function.

2.2 Evaluating an infinite series

Jørgensen (1997) (among others) gives two series expressions for evaluating the densities: one for $1 < p < 2$ and one for $p > 2$. For quick and accurate evaluation, the series is summed only over those terms in the series which contribute significantly to the sum. This is done by treating the index of the summation as continuous and then using Stirling's approximation for the gamma functions to approximate the terms in the series. By differentiating with respect to the index of the summation and equating to zero we can find the value of the index for which the terms are a maximum; we then evaluate the approximate terms of the series on either side of the maximum to locate upper and lower bounds on the index over which to sum. The exact terms in the series are then summed over these values of the index.

The series and the inversion methods work best in different parts of the parameter space (see Table 1). The series approach also has difficulty as p gets very close to 2. Plots of the densities are given in Figure 1 for $1 < p < 2$ and Figure 2 for $p > 2$. Note that the distribution approaches the Poisson as $p \to 1$ and the gamma as $p \to 2$.

	'Small' y	'Large' y
$1 < p < 2$	both work well	both are OK
$p > 2$	inversion best	series best

TABLE 1. The performance of the two evaluation methods.

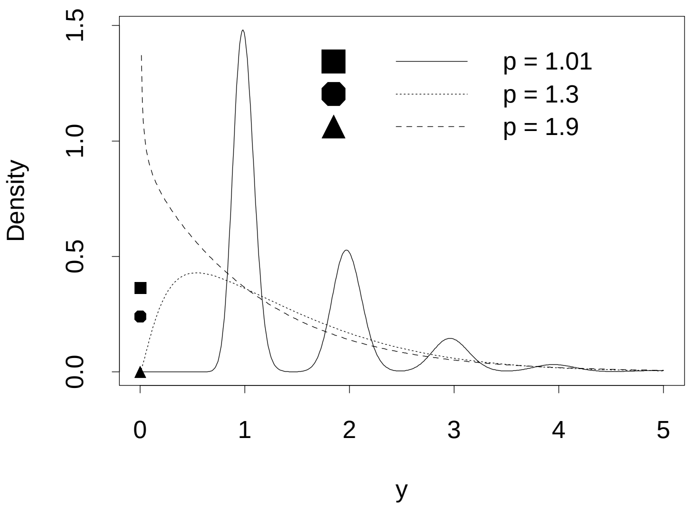

FIGURE 1. Some Tweedie densities with $1 < p < 2$. In all cases, $\mu = 1$ and $\text{var}[Y] = 1$. The solid shapes represent $\Pr(Y = 0)$.

3 The saddlepoint approximation

The saddlepoint approximation can be used to approximate the Tweedie densities. The part of the density that cannot be written in closed form is replaced by a simple analytic expression. Under the saddlepoint approximation we have

$$p(y; \mu, \phi) = [2\pi\phi y^p]^{-1/2} \ \exp\{-d(y, \mu)/(2\phi)\} \{1 + O(\phi)\}$$

as $\phi \to 0$ for the Tweedie densities. The ratio of this approximation to the form of the density (1) is $\rho = b_p(y, \phi)\sqrt{2\pi\phi y^p}$. It can be shown that

$$f_p(y; \mu, \phi) = \frac{1}{y} b_p(1, \xi) \exp\{-d(y, \mu)/(2\phi)\}$$

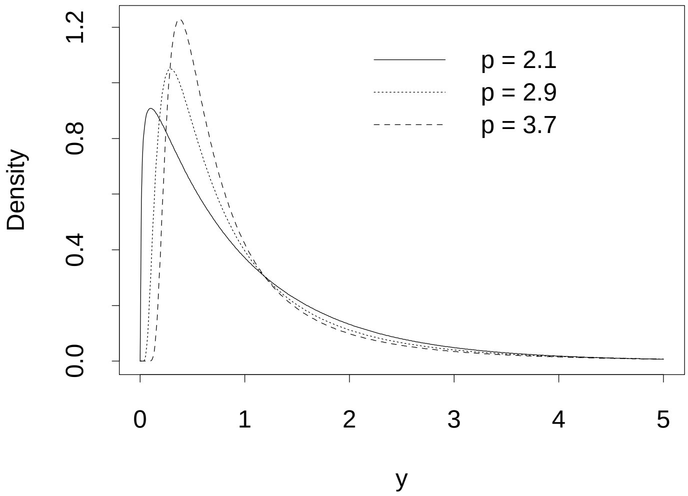

FIGURE 2. Some Tweedie densities with $p > 2$. In all cases, $\mu = 1$ and $\text{var}[Y] = 1$.

where $\xi = \phi y^{p-2}$, and so the ratio of the density to the saddlepoint approximation can be expressed as

$$\rho = b_p(1, \xi)\sqrt{2\pi\xi}. \tag{4}$$

This shows that ρ is a function of p, not of μ, and is a function of y and ϕ only through ξ. The ratio ρ is for each p an increasing monotonic function of $\xi = \phi y^{p-2}$ for $p > 3$, and a decreasing monotonic function of ξ for $1 < p < 3$ provided p is not close to 1; see Figures 3 and 4. Note that ρ is plotted against $\xi/(1+\xi)$ which (unlike ξ) is bounded on the positive side.

We can use the saddlepoint approximation to create a fast evaluation scheme for evaluating the Tweedie densities as follows: We evaluate very accurately the density on a grid of values given by the roots of a Chebyshev polynomial and then form the ratio ρ from (4). For any necessary evaluation, we then use a fast two-dimensional Chebyshev interpolation scheme to interpolate for any values of the parameters, and hence find ρ. From ρ, we can reconstruct the density. This method can be vectorised in S-Plus, and so is fast as well as accurate. Note that the problem has been reduced from one on four dimensions (y, p, μ and ϕ) to one on two dimensions (p and ξ).

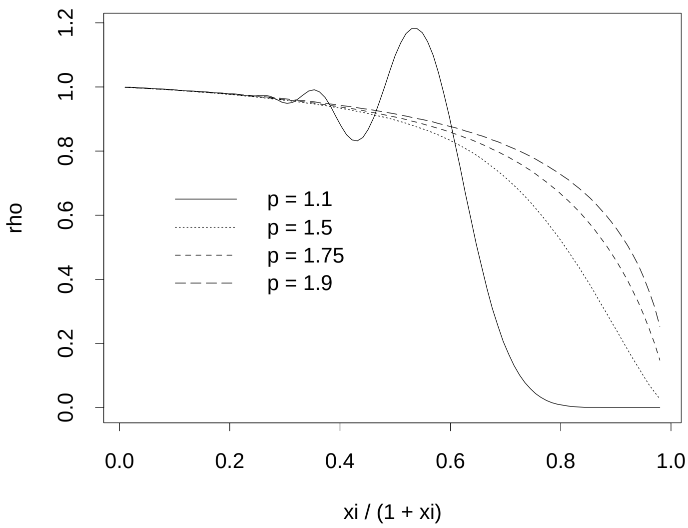

FIGURE 3. The ratio of the exact (numerical) density to the saddlepoint approximation, ρ, for $1 < p < 2$.

4 Applications

As an example, we consider the data used by Box and Cox (1964) from a 3×4 factorial experiment for the survival times of animals. The factors are three poisons and four treatments, and each combination of the two factors is used on four animals. To determine the maximum likelihood estimate of p from the saturated model, the profile (log-) likelihood function is used as shown in Figure 5. From a finer grid, we deduce that $\hat{p} \approx 3.59$ and $\hat{\phi} \approx 0.7937$. Using these estimates, a normal probability plot of quantile residuals (see Dunn and Smyth, 1996) in Figure 6 shows that the fit is adequate. The quantile residuals use the cumulative distribution function, computed using a cgf inversion technique discussed near the end of §2.1. Evaluation of the derivatives with respect to ϕ are computed using a series evaluation technique like that in §2.2.

References

Box, G. E. P., and Cox, D. R. (1964). 'An analysis of transformations (with discussion)' in *J. R. Statistic. Soc.* B, **26**,

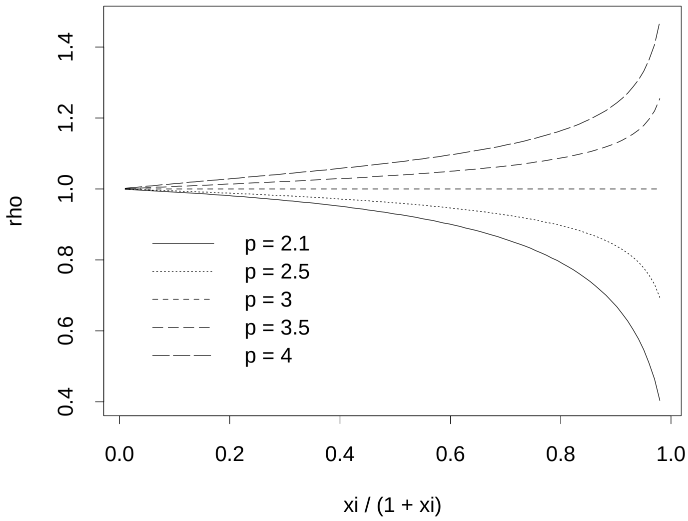

FIGURE 4. The ratio of the exact (numerical) density to the saddlepoint approximation, ρ, for $p > 2$. Note that the inverse Gaussian distribution ($p = 3$) has an exact saddlepoint expression.

211–252.

Dunn, Peter K. and Smyth, Gordon K. (1996). 'Randomized quantile residuals' in *J. Comp. and Graph. Statistics*, **5**(3), 236–244.

Gilchrist, Robert and Drinkwater, Denise (1999). 'Fitting Tweedie Models to Data with Probability of Zero Responses' in *Proc. of the 14th International Workshop on Statistical Modelling*, Graz, Austria, July 19–23.

Jørgensen, Bent (1997). *The Theory of Dispersion Models*, Chapman and Hall, London.

Nelder, J. A. and Wedderburn, R. W. M. (1972). 'Generalized linear models' in *J. R. Statistic. Soc.* A, **135**, 370–384.

Sidi, Avram (1982). 'An algorithm for a special case of a generalization of the Richardson extrapolation process' in *Numerische Mathematik*, **38**, 299–307.

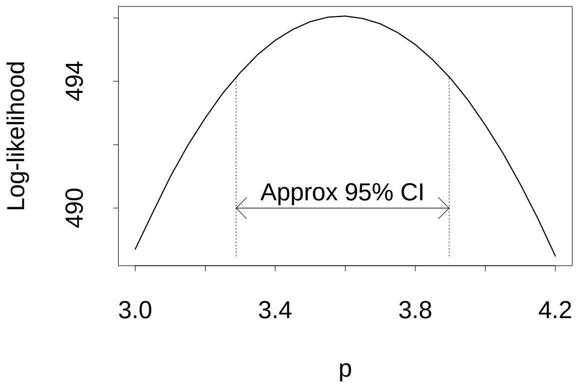

FIGURE 5. The profile log-likelihood function for the poison survival times. p is estimated as 3.59; the 95% confidence interval for p is approximately $(3.29, 3.90)$.

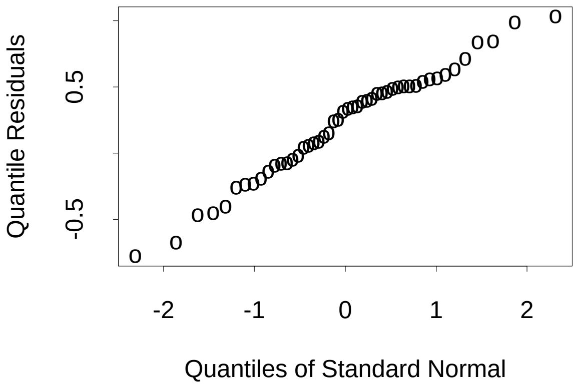

FIGURE 6. The normal probability plot of the quantile residuals applied to the survival times of the poison data, using $\hat{p} = 3.59$.

Sidi, Avram (1988). 'A user-friendly extrapolation method for oscillatory infinite integrals' in *Math. Comp.*, **51**(183), 249–266.

Smyth, Gordon K. (1996). 'Regression analysis of quantity data with exact zeros' in *Proc. Second Australia-Japan Workshop on Stochastic Models in Engineering, Technology and Management*. Technology Management Centre, University of Queensland, 572–580.