

Evaluation of Tweedie exponential dispersion model densities by Fourier inversion

Peter K. Dunn

e-mail: dunn@usq.edu.au

Australian Centre for Sustainable Catchments and
Department of Mathematics and Computing
University of Southern Queensland
Toowoomba Queensland 4350 Australia

Gordon K. Smyth

Bioinformatics Division

Walter and Eliza Hall Institute of Medical Research
Melbourne, Vic 3050, Australia

May 9, 2007

Abstract

The Tweedie family of distributions is a family of exponential dispersion models with power variance functions $V(\mu) = \mu^p$ for $p \notin (0, 1)$. These distributions do not generally have density functions that can be written in closed form. However, they have simple moment generating functions, so the densities can be evaluated numerically by Fourier inversion of the characteristic functions. This paper develops numerical methods to make this inversion fast and accurate. Acceleration techniques are used to handle oscillating integrands. A range of analytic results are used to ensure convergent computations and to reduce the complexity of the parameter space. The Fourier inversion method is compared to a series evaluation method and the two methods are found to be complementary in that they perform well in different regions of the parameter space.

Keywords: compound Poisson distribution; generalized linear models; numerical integration; numerical acceleration; power variance function

1 Introduction

It is well known that the density function of a statistical distribution can be represented as an integral in terms of the characteristic function for that distribution (Abramowitz and Stegun, 1965, 26.1.10). This relationship is a special case of the Fourier inversion

theorem. There are many distributions for which the moment generating function has a analytic expression while the density function does not. In such cases it is natural to try evaluate the Fourier inversion numerically in order to compute the density function or cumulative probability function. Although natural and obvious, this approach has been relatively little used in practice because it is difficult to implement with reliable numeric accuracy, involving as it does an infinite integral of a rapidly oscillating integrand. This article demonstrates that the approach can be brought to a successful conclusion by combining advanced integration methods with some analytic analysis of the integrand. The approach is used to evaluate the densities of an important class of distributions generated by exponential families.

An exponential dispersion model (EDM) is a two-parameter family of distributions consisting of a linear exponential family with an additional dispersion parameter. EDMs are important in statistics because they are the response distributions for generalized linear models (McCullagh and Nelder, 1989). EDMs were established as a field of study in their own right by Jørgensen (1987, 1997), who undertook a detailed study of their properties.

An EDM is characterized by its variance function $V()$, which describes the mean–variance relationship of the distribution when the dispersion is held constant. If Y follows an EDM distribution with mean μ and variance function $V()$ then

$$\text{var}(Y) = \phi V(\mu)$$

where ϕ is the dispersion parameter. Of special interest are EDMs with power mean–variance relationships, $V(\mu) = \mu^p$ for some p . Following Jørgensen (1987, 1997), we call these Tweedie models. The class of Tweedie models includes most of the important distributions commonly associated with generalized linear models including the normal ($p = 0$), Poisson ($p = 1$), gamma ($p = 2$) and the inverse Gaussian ($p = 3$) distributions. Although the other Tweedie model distributions are less well known, Tweedie models exist for all values of p outside the interval $(0, 1)$.

All Tweedie distributions with $p \geq 1$ have non-negative support and strictly positive means, $\mu > 0$. The Tweedie models for $p > 2$ are generated by stable distributions and are continuous with strictly positive support. The Tweedie model distributions for $1 < p < 2$ can be represented as Poisson mixtures of gamma distributions. They are mixed distributions with mass at zero but are otherwise continuous on the positive

reals. The distributions with $p < 0$ are unusual in that they have positive means but support on the whole real line. These distributions seem to have limited practical application so we do not consider them further in this article.

Tweedie models are natural candidates for modelling continuous positive quantity data with an arbitrary measurement scale because they are the only EDMs which are closed under re-scaling. The distributions with $1 < p < 2$ are especially appealing for modelling quantity data when exact zeros are possible. Dunn and Smyth (2005) gave a survey of published applications showing that Tweedie distributions have been used in a diverse range of fields including actuarial studies, assay analysis, survival analysis, time studies, expense studies, consumption studies, ecology and meteorology. Recent applications include fisheries (Candy, 2004) and rainfall prediction (Dunn, 2004). In these and many other applications in which physical quantities are measured, the occurrence of continuous data with exact zeros is common. In such contexts, Tweedie models often are useful candidate distributions and would be more frequently used in practice if high quality numeric computations were readily available. Dunn and Smyth (2005) also give two detailed data analyses, one for $1 < p < 2$ and another for $p > 2$, demonstrating the usefulness of the distributions and methodology on real data.

Apart from the well-known distributions with $p = 0, 1, 2,$ or 3 , none of the Tweedie models have density functions with explicit analytic forms. This complicates the use of these distributions in statistical modelling. In particular, it prevents their use with likelihood based estimation, testing or diagnostic procedures. On the other hand, Tweedie models do have simple, analytic moment generating functions. The purpose of this article is to provide fast, accurate computation of the Tweedie densities by Fourier inversion of the characteristic functions. Our aim is to develop algorithms which will compute the Tweedie density functions to a relative accuracy of 10^{-10} in 64-bit double precision arithmetic for all parameter values.

All previous attempts to evaluate the Tweedie densities have used infinite series expansions for the densities (Jørgensen and Paes de Souza, 1994; Jørgensen, 1997, Section 4.2; Gilchrist, 2000; Dunn and Smyth, 2005). The series expansion actually arises itself from Fourier inversion of the characteristic function, after applying a Taylor series expansion to the main exponential term in the integrand (Feller, 1971, page 582). Dunn and Smyth (2005) presented the first rigorous study of the accuracy of the series

expansion and showed that it is not a practical numeric strategy for all parameter values. For $1 < p < 2$, the number of terms necessary for accurate evaluation becomes arbitrarily large for p near 2, y large or ϕ small. For $p > 2$, subtractive cancellation in floating point arithmetic prevents accurate summation of the series for p close to 2, y small or ϕ small,

Our approach in this article is to treat the inversion of the characteristic function directly as a numerical integration problem. A number of probability identities are derived which make it necessary only to evaluate the density at the mean, $y = \mu$, and in fact only at $y = \mu = 1$. Amongst other advantages, this ensures that the value of the integral is never very small and increases the accuracy of the integration.

The characteristic function produces highly oscillatory infinite integrals; the top panel of Figure 1 shows an example. Specialised extrapolation methods are used to evaluate the oscillating integral using an economical number of terms (Sidi, 1982b). These advanced integration methods turn out to be crucial for numerical accuracy as well as economy. As far as we know, ours is the first application of such methods to compute probability functions. As part of this procedure, strategies are developed for locating the zeros of the integrand.

We find that the numerical inversion strategy evaluates the density functions to the desired 10^{-10} accuracy over a very wide range of parameter values, and is also computationally modest. The inversion approach turns out to be somewhat complementary to the series approach as it performs best for y large when $p < 2$ and for small y for $p > 2$, whereas the series approach is the opposite. In this way the inversion method “plugs the holes” left by the series method. Together, the two strategies enable the Tweedie densities to be accurately evaluated for any region of the parameter space.

In the next section, the properties of Tweedie densities are outlined. The use of Fourier inversion to evaluate the densities is discussed in Section 3, followed by a discussion on numerical integration and acceleration methods in Section 4. Section 5 derives strategies for finding the zeros of the oscillating integrand. Implementation details are given in Section 6. Sections 7 and 8 evaluate the accuracy and computational complexity of the algorithms, comparing comparing two alternative acceleration schemes and comparing the series and Fourier inversion methods. Conclusions and discussion follow in Section 9.

2 The Tweedie densities

2.1 Exponential dispersion models

EDMs have probability density functions or probability mass functions of the form

$$f(y; \mu, \phi) = a(y, \phi) \exp \left[\frac{1}{\phi} \{y\theta - \kappa(\theta)\} \right] \quad (1)$$

for suitable known functions $\kappa(\cdot)$ and $a(\cdot)$ (Jørgensen, 1997). The canonical parameter θ belongs to the open interval satisfying $\kappa(\theta) < \infty$ and the dispersion parameter ϕ is positive. The function $\kappa(\cdot)$ is called the cumulant function of the EDM because, if $\phi = 1$, the derivatives of κ give the successive cumulants of the distribution. In particular, the mean of the distribution is $\mu = \kappa'(\theta)$ and the variance is $\phi\kappa''(\theta)$. The mapping from θ to μ is invertible, so we may write $\kappa''(\theta) = V(\mu)$ for a suitable function $V(\cdot)$, called the variance function of the EDM.

EDMs have a simple form for the moment generating function, a fact which we exploit in this paper. The moment generation function is $M(t) = \int \exp(ty)f(y; \mu, \phi) dy$. Substituting (1) into $M(t)$ and completing the integral shows that the cumulant generating function is

$$K(t) = \log M(t) = [\kappa(\theta + t\phi) - \kappa(\theta)]/\phi. \quad (2)$$

There is another form of the probability function which more convenient than (1) for some purposes. Differentiating $\log f$ shows that $f(y; \mu, \phi)$ is maximized with respect to μ at $\mu = y$. This calculation assumes that the support for y is contained in the domain of μ , which is true for the EDMs considered in this article. Write $t(y, \mu) = y\theta - \kappa(\theta)$. Then the unit deviance, $d(y, \mu) = 2\{t(y, y) - t(y, \mu)\}$, can be viewed as a distance measure, satisfying $d(y, y) = 0$ and $d(y, \mu) > 0$ for $y \neq \mu$. For example, unit deviance of the normal distribution is $d(y, \mu) = (y - \mu)^2$. The probability function can be re-written in terms of the deviance as

$$f(y; \mu, \phi) = b(y, \phi) \exp \left\{ -\frac{1}{2\phi} d(y, \mu) \right\} \quad (3)$$

where $b(y, \phi) = f(y; y, \phi)$. Following Jørgensen (1997), we call this the *dispersion model* form of the probability function.

2.2 Tweedie models

In this article we are interested in EDMs with variance functions of the form $V(\mu) = \mu^p$ for some $p \geq 1$. The cumulant function $\kappa(\theta)$ can be found for Tweedie EDMs by equating $\kappa''(\theta) = d\mu/d\theta = \mu^p$ and solving for κ . Without loss of generality we can choose $\kappa(\theta) = 0$ and $\mu = 1$ at $\theta = 0$. This gives

$$\theta = \begin{cases} \frac{\mu^{1-p} - 1}{1-p} & p \neq 1 \\ \log \mu & p = 1 \end{cases} \quad (4)$$

with inverse

$$\mu = \tau(\theta) = \{\theta(1-p) + 1\}^{1/(1-p)}$$

for $p \neq 1$, and

$$\kappa(\theta) = \begin{cases} \frac{\mu^{2-p} - 1}{2-p} & p \neq 2 \\ \log \mu & p = 2. \end{cases} \quad (5)$$

Note that our definitions for θ and $\kappa(\theta)$ are continuous in p as well as in θ .

The above expression for $\kappa(\theta)$ implies that the cumulant generating function (2) has a simple analytic form, but it remains that neither $a()$ in (1) nor $b()$ in (3) have closed form expressions. Finding numeric approximations for these functions is the aim of this article. The only exceptions where $a()$ and $b()$ can be obtained analytically are the well known distributions at $p = 1$, $p = 2$ and $p = 3$ and at $y = 0$ for $1 < p < 2$. For $1 < p < 2$ there is mass at zero equal to $f(0; \mu, \phi) = \exp[-\mu^{2-p}/\{\phi(2-p)\}]$.

2.3 A re-scaling identity

A fundamental property of Tweedie model densities is that they are closed under re-scaling. Consider the transformation $Z = cY$ for some $c > 0$ where Y follows a Tweedie model distribution with mean μ and variance function $V(\mu) = \mu^p$. Finding the cumulant generating function for Z reveals that it follows a Tweedie distribution with the same p , with mean $c\mu$ and dispersion $c^{2-p}\phi$. Meanwhile, the Jacobian of the transformation is $1/c$ for all $y > 0$. Putting these two facts together gives the extremely useful *rescaling identity*

$$f(y; \mu, \phi) = cf(cy; c\mu, c^{2-p}\phi) \quad (6)$$

for all p , $y > 0$ and $c > 0$. As far as we know, this is the first statement of this identity in the EDM literature. This identity will allow us to select y and parameter values

which are favourable for numeric evaluation, and to obtain the density at other values by rescaling.

3 Fourier inversion

3.1 Overall strategy

The Fourier inversion theorem (Abramowitz and Stegun, 1965, 26.1.10) allows a continuous EDM probability function to be written in terms of its cumulant generating function as

$$f(y; \mu, \phi) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp\{K(it) - ity\} dt \quad (7)$$

where $i = \sqrt{-1}$. Our strategy is to evaluate $f(y; \mu, \phi)$ by numerical evaluation of the integral. However the use of Tweedie density properties allows us to simplify the problem considerably before we resort to numerical integration.

It turns out that the integral (7) doesn't need to be evaluated for every possible combination of y , μ and ϕ . For example, it would be sufficient to evaluate $a(y, \phi) = f(y; 1, \phi)$ by numeric integration, after which $f(y; \mu, \phi)$ could be evaluated for any μ from (1). Hence we only need $\mu = 1$. Alternatively, it would be sufficient to evaluate $b(y, \phi) = f(y; y, \phi)$ by numeric integration, after which $f(y; \mu, \phi)$ could be evaluated from (3). Hence we only need $\mu = y$. The choice of $\mu = 1$, corresponding to $\theta = \kappa(\theta) = 0$, is convenient because it simplifies the integrand.

There are several methods by which $f(y; \mu, \phi)$ can be evaluated for any y , μ , ϕ and p while inverting the cumulant generating function only at $\mu = 1$. Here we list three methods in increasing order of sophistication. The second and third methods use the rescaling identity (6). The third method uses the dispersion model form of the density as well:

Method 1: Evaluate a . Compute $a(y, \phi) = f(y; 1, \phi)$ by Fourier inversion and substitute into (1).

Method 2: Rescale μ to 1. Compute $I = a(y/\mu, \phi/\mu^{2-p}) = f(y/\mu; 1, \phi/\mu^{2-p})$ by Fourier inversion and use $f(y; \mu, \phi) = I/\mu$. This method uses (6) with $c = 1/\mu$.

Method 3: Rescale y to 1 and evaluate b . Compute $I = a(1, \phi/y^{2-p}) = f(1; 1, \phi/y^{2-p})$ by Fourier inversion, use $b(y, \phi) = I/y$, and substitute into (3). This method uses (6) with $c = 1/y$.

Section 7 compares the three methods in terms of computational load and numeric accuracy, showing that, although all three methods are equivalent in exact arithmetic, they are substantially different in terms of floating point errors. Only Method 3 fulfils our aim of ten significant figure accuracy. For the moment we note that Method 3 has theoretical advantages. For one, it is the simplest from a programming point of view, because it requires Fourier inversion only at $y = \mu = 1$ rather than for general y . The crucial motivation however for Method 3 is that the integral I evaluated by Fourier inversion is almost always larger under Method 3 than the other two methods. Under Method 3, numerical integration is always used to evaluate a density at its mean value, i.e., at $y = \mu$, which is close to the mode of the distribution. This should be an advantage because numerical integration is an additive process, involving a summation error which is roughly constant. It should follow that the larger the value of the integral being evaluated, the smaller the relative error.

The remainder of this section and the following two sections of this article develop methodology for evaluating $a(y; \phi)$ by Fourier inversion for any p , y and ϕ . This enables any of Methods 1–3 above to be used to obtain $f(y; \mu, \phi)$.

3.2 The case $p > 2$

When $\mu = 1$ the inversion formula (7) simplifies to

$$a(y, \phi) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp k(t) dt, \quad (8)$$

where $k(t) = \kappa(it\phi)/\phi - ity$. Extracting the real and imaginary components of the integrand and changing the limits of integration gives

$$a(y, \phi) = \frac{1}{\pi} \int_0^{\infty} \exp \Re k(t) \cos \Im k(t) dt. \quad (9)$$

Expressions for $\Re k(t)$ and $\Im k(t)$ are derived in the Appendix. The integrand is, for large t , an exponentially damped cosine oscillating about zero.

3.3 The conditional density for $1 < p < 2$

The Tweedie models with $1 < p < 2$ are mixed distributions with mass at $y = 0$, and this prevents (7), which is for continuous densities, from being used directly. Fourier inversion is instead used to evaluate the density of the conditional distribution of Y

given $Y > 0$. Write $\pi_0 = \Pr(Y = 0 | \mu = 1) = \exp[-1/\{\phi(2-p)\}]$. The moment generating function of the conditional distribution at $\mu = 1$ is related to that of the full distribution by

$$M^*(t) = \frac{M(t) - \pi_0}{1 - \pi_0}.$$

The inversion formula for the conditional distribution at $\mu = 1$ is therefore

$$\begin{aligned} a^*(y, \phi) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} M^*(it) \exp(-ity) dt \\ &= \frac{1}{2\pi(1 - \pi_0)} \int_{-\infty}^{\infty} \{\exp k(t) - \pi_0 \exp(-ity)\} dt. \end{aligned} \quad (10)$$

Extracting the real and imaginary components and changing the limits of integration gives

$$a^*(y, \phi) = \frac{1}{\pi(1 - \pi_0)} \int_0^{\infty} \{\exp \Re k(t) \cos \Im k(t) - \pi_0 \cos(ty)\} dt. \quad (11)$$

Expressions for $\Re k(t)$ and $\Im k(t)$ are derived in the Appendix. The non-conditional density is recovered for $y > 0$ by

$$a(y, \phi) = a^*(y, \phi)(1 - \pi_0).$$

4 Integrating oscillating functions

4.1 The W -transformation

The integrands in (9) and (11) are highly oscillatory, and special purpose methods are needed if numerical integration is to be successful. A popular strategy is to find the zeros of the integrand, integrate between successive zeros, and sum the resulting series. Nolan (1997) and Lambert and Lindsey (1999) adopt this approach in related work on stable distributions. The difficulty with this straightforward approach is that the series may converge very slowly. An arbitrarily large number of terms may be necessary, in which case rounding error can overtake precision and accuracy is lost; see the top panel of Figure 1.

—INSERT FIGURE 1 ABOUT HERE—

An improved strategy, which can decrease computation and overcoming subtractive cancellation, is to use an acceleration method which extrapolates the series based on a limited number of terms (Rabinowitz 1992; Evans 1993; Krommer and Überhuber

1998). Although a number of extrapolation methods are available, little use has been made of them so far for evaluating probability density functions, perhaps because the best known methods are not particularly well suited to summing terms of alternating sign.

Sidi (2003, Chapter 11) discusses a number of advanced extrapolation methods designed for infinite oscillatory integrals. Most of the methods described require substantial asymptotic analysis of the integral as the variable of integration tends to infinity. We adopt the modified W -transformation method of Sidi (1988), which has the advantage of remaining effective while keeping the need for asymptotic analysis to a minimum. We describe it briefly here.

Consider an infinite integral

$$I = \int_0^{\infty} f(t) dt$$

where $f()$ is an oscillating function, and define

$$F(x) = \int_0^x f(t) dt.$$

For any integer $r \geq 0$, let x_0, \dots, x_{r+2} be a set of successive zeros of the integrand.

Write

$$w(x_j) = \int_{x_j}^{x_{j+1}} f(t) dt,$$

for the integrals between the zeros. The modified W -transformation of Sidi (1988) is defined by the system of $r + 2$ linear equations

$$W_r = F(x_j) + w(x_j) \sum_{i=0}^r \frac{v_i}{x_j^i} \quad 0 \leq j \leq r + 1 \quad (12)$$

where W_r and the v_i are unknown constants. The solution for W_r is the approximation to I . Sidi (1982) developed an efficient algorithm, called the W -algorithm, for solving the system of equations using a small number of arithmetic operations and minimal storage space.

4.2 Convergence criterion

Subject to rounding error, W_r becomes an increasingly accurate approximation to I as r increases. Convergence is detected by comparing the most recent estimate of the integral with the previous two estimates. The estimate of the relative error is

$$\text{RelErr}_r = \frac{|W_r - W_{r-1}| + |W_r - W_{r-2}|}{W_r},$$

similar to that used by Piessens et al. (1983). When $\text{RelErr}_r < 10^{-10}$, the algorithm stops and reports W_r as the best estimate of the integral.

This convergence criterion requires least $r = 2$ before an estimate of the relative error can be established. This means that five zeros and five finite integrals are evaluated before convergence is tested. For the continuous densities $p > 2$, the W -transformation is invoked starting from x_0 as the first positive zero of the integral in (9). For the mixed densities $1 < p < 2$, the integrand in (11) often remains irregular in shape for the first few zeros, so the W -transformation scheme is invoked from the fourth zero instead of the first. This means a minimum of eight regions of integration are used before convergence is tested in this case.

5 Analytic root finding

5.1 Exact vs asymptotic zeros

In this section, methods are developed for finding the zeros of the integrals in (9) and (11) for use in Sidi's W -transformation (12). First we consider why exact zeros are necessary.

The zeros of the integrands occur at the zeros of $\cos \Im k(t)$. The Appendix shows that $\Im k(t) \rightarrow ty$ as $t \rightarrow \infty$, so for large t the zeros of the integrals are approximately equal to the zeros of $\cos(ty)$. The zeros of $\cos(ty)$ are simply $z_j = \pi/2 + j\pi/y$ for $j = 0, 1, \dots$. There is no theoretical reason why Sidi's method (12) could not be implemented using the asymptotic zeros z_j instead of the exact zeros x_j . This would avoid the need to find the exact zeros.

In Section 7, we show that the exact zeros give better accuracy. The accuracy obtainable using the asymptotic zeros declines to unacceptable levels for small y when $p > 0$, making it necessary to determine the exact zeros.

5.2 $p > 2$

Consider now how to locate the zeros of the oscillating integrand in (9). Our strategy is to use Newton's method to find each zero, carefully choosing the starting value in each case so that convergence of the iteration to the desired solution is guaranteed.

The zeros occur at

$$\Im k(t) = \frac{\pi}{2} + m\pi \tag{13}$$

for integer values of m . From (17) in the Appendix, see that $\Im k''(t) < 0$ for all t , so $\Im k(t)$ is a convex function. In addition, $k(0) = 0$ and $\Im k'(0) = 1 - y$ so $\Im k(t)$ starts at the origin and initially tends upwards if $y < 1$ or downwards if $y > 1$. Also, $\Im k'(t) \rightarrow -y$ as $t \rightarrow \infty$ so $\Im k(t)$ is asymptotically linear decreasing.

Since $\Im k(t)$ is convex, it has only one local maximum on $[0, \infty)$. Let k_{\max} be this maximum value and let t_{\max} be the corresponding value of t . If $y \geq 1$, k_{\max} is achieved at $t_{\max} = 0$ and all the zeros occur as $\Im k(t)$ decreases, i.e., at $m = -1, -2, \dots$. If $y < 1$, there are two possible scenarios. Either $k_{\max} < \pi/2$, in which case the zeros occur at $m = -1, -2, \dots$ as for $y \geq 1$, or $k_{\max} \geq \pi/2$, in which case one or more zeros will occur with $m \geq 0$. Let m_{\max} be the largest value of m satisfying (13). If $k_{\max} \geq \pi/2$, then the zeros occur successively at $m = 0, \dots, m_{\max}$ as $\Im k(t)$ increases followed by $m = m_{\max}, m_{\max} - 1, \dots$ as $\Im k(t)$ decreases.

Finding the zeros when $y \geq 1$. In this case $\Im k(t)$ is monotonic decreasing on $t \geq 0$. The convexity of $\Im k(t)$ ensures that Newton's method is globally convergent for any solution of (13) from any non-negative starting value. If the starting value is to the right of the solution, then Newton's method converges monotonically to the solution. If not, the iteration takes one step to the right and then converges monotonically.

The slope $\Im k'(t)$ varies from $1 - y$ to $-y$, so the first zero is bracketed between $t = \pi/(2y)$ and $t = \pi/\{2(y - 1)\}$. We start the Newton iteration from $t = \pi/(2y)$ to find the first zero and from the previous zero for subsequent zeros. In each case the first Newton step will be smaller than the distance between the two previous zeros and thereafter convergence will be monotonic.

Finding the zeros when $y < 1$. The key here is to find k_{\max} and t_{\max} . If $k_{\max} < \pi/2$ then we revert to the method for $y \geq 1$ but with starting value $t = t_{\max} + \pi/(2y)$. So consider the case with $k_{\max} \geq \pi/2$ and $m_{\max} \geq 0$.

A lower bound for the first zero is $t = \pi/\{2(1 - y)\}$ and Newton's method will converge monotonically to it from this starting value. Other zeros on the upswing of $\Im k(t)$, i.e., zeros with $m = 1$ to m_{\max} , can be found by starting Newton's method from the previous zero. The iteration converges monotonically in each case.

A problem is to find a good starting value for the first zero on the down-swing. One method is to treat $\Im k(t)$ as roughly symmetric about its peak and start Newton's method the same distance to the right of the peak as the previous zero was to the left.

Convexity of $\mathfrak{S}k(t)$ ensures eventual convergence of the iteration. Subsequent zeros on the down-swing are found by starting Newton's method from the previous zero as was done for $y \geq 1$.

Finding k_{\max} when $y < 1$. Some care is required in solving $\mathfrak{S}k'(t) = 0$ because $\mathfrak{S}k'(t)$ is not everywhere convex. Note that $\mathfrak{S}k^{(3)}(t) = 0$ only once in $[0, \infty)$ at

$$t_{\text{infl}} = \frac{1}{\phi(1-p)} \tan \left\{ \frac{\pi}{2} \left(\frac{1-p}{2p-1} \right) \right\}. \quad (14)$$

This is the point of inflection where $\mathfrak{S}k'(t)$ changes from convex to concave. This value can be used to start Newton's method to solve $\mathfrak{S}k'(t) = 0$. The iteration either descends the concave part of the function to the solution or ascends the convex part to the solution, with monotonic convergence in either case.

When y is small, t_{\max} tends to be much larger than t_{infl} . For example at $p = 2.3$, $\mu = 1$, $\phi = 1$ and $y = 0.001$, $t_{\max} = 1589$ while $t_{\text{infl}} = 2.5$. So an alternative starting point is useful when t_{\max} may be large. For large t , $\cos \zeta \approx -[(1-p)t\phi]^{-1}$. Substituting this into (16) and solving $\mathfrak{S}k'(t) = 0$ gives the approximation

$$\hat{t}_{\max} = -\frac{1}{\phi(1-p)} \left\{ \frac{1}{y} \cos \left(-\frac{\pi}{2(1-p)} \right) \right\}^{p-1}. \quad (15)$$

For the example parameter values above, $\hat{t}_{\max} = 1588$, an excellent approximation.

It is easy to locate t_{infl} and \hat{t}_{\max} relative to t_{\max} by examining $\mathfrak{S}k'(t)$ at the two values. Our strategy is simply to start the Newton iteration for t_{\max} from whichever of t_{infl} or \hat{t}_{\max} is closer to it, unless they are on opposite sides of t_{\max} in which case t_{infl} is used. This improves the starting value without altering the monotonic convergence.

5.3 $1 < p < 2$

Consider now the problem of locating the zeros of the integrand in (11). Since $\cos \mathfrak{S}k(t)$ approaches $\cos(ty)$ as $t \rightarrow \infty$, the zeros of the integrand are asymptotically those of $\cos(ty)$, that is, the zeros will eventually occur arbitrarily close to $(\pi + 2m\pi)/(2y)$ for integer m .

It proves difficult to set up monotonically convergent Newton iterations when $1 < p < 2$ as we did for $p > 2$. Instead we find an interval of values for t known to bracket each successive zero, then use a modified version of Newton's method with bound checking to locate the zero within that interval. The modified algorithm follows an

idea described by Press et al (1996, §9.4). Newton’s method is used while the iterations remain within the *a priori* bounds. If the iteration attempts to leave the bracketing interval, then the algorithm switches to the bisection method to refine the location of the zero until Newton’s method can be used again.

The bracketing interval is found for each zero by stepping out from the previous zero an interval proportional to the previous inter-zero distance, then checking that the integrand changes sign between the end points. Since the zeros are known to become equi-spaced for t large, this procedure causes no problems in practice.

6 Quadrature

The finite integrals $F(x_j)$ and $w(x_j)$ required as input to the Sidi extrapolation algorithm (12) are computed using 512-point Gaussian quadrature. The Gaussian abscissa and weights were generated using the algorithm given by Davis and Rabinowitz (1975, Appendix 2).

The initial region of integration for the case $1 < p < 2$ can be quite irregular; consequently the region between $t = 0$ and the first zero is actually divided into 20 panels, and the area in each panel evaluated using the 512-point Gaussian integrator.

In the case $p > 2$ especially, the integrand is asymptotically a damped cosine, suggesting the use of Gaussian quadrature with cosine weights rather than normal weights, also called Gauss–cos integration (Evans, 1993). Our experiments showed that the Gauss–cos rule did not always evaluate the finite integrals to full machine precision because the integrand often doesn’t closely resemble a cosine shape for the initial integration regions (bottom panel of Figure 1). There was no reliable method of determining when the Gauss–cos integrator would be satisfactory, and so the normal Gaussian integration rule is used throughout. This is still fast and accurate, as well as more reliable.

7 Accuracy and computational complexity

The accuracy of the Fourier inversion evaluation algorithms can be examined in two cases for which exact densities are available. The special cases are $p = 3$, which is the inverse Gaussian distribution (Johnson and Kotz, 1970, §15.3), and $p = 1.5$ which

intersects with the non-central χ^2 distribution on zero degrees of freedom studied by Seigel (1979, 1985). The χ_0^2 distributions are Tweedie distributions with $p = 1.5$ and $\phi = 4/\sqrt{\mu}$. The χ_0^2 densities can be expressed in terms of the modified Bessel function of the first kind of order 1, implemented in the R function `bessel1`.

First we compare the three rescaling strategies described in Section 3.1 for evaluating densities via $\mu = 1$. In Figure 2, the relative errors of the three methods are shown for computing the density with $p = 3$, $\mu = 1.4$ and $\phi = 0.74$. Method 3 enjoys a substantial advantage in relative accuracy for values of y much larger or smaller than μ . There are only small intervals of y values where the relative accuracy of the other two methods are comparable. This confirms the theoretical argument posited in Section 3.1 in favour of Method 3. From this point we adopt Method 3 as our preferred rescaling method, and all the results which follow are for Method 3.

— INSERT FIGURE 2 ABOUT HERE —

Method 3 uses Fourier inversion to compute $f(1; 1, \xi)$ with $\xi = \phi/y^{2-p}$. For any p the performance of the method depends only on ξ . Larger values of ξ generally correspond to integrands with slower decay and the faster oscillations. This suggests that accuracy may deteriorate for ϕ large, for y large for $p > 2$, or for y small for $1 < p < 2$. This is qualitative pattern that we expect. In practice, the use of Sidi extrapolation improves the accuracy of the inversion enormously and causes the simple qualitative pattern in terms of ξ to be less apparent.

Table 1 compares the accuracy of Sidi’s extrapolation method with exact and asymptotic zeros for evaluating inverse Gaussian densities. Exact zeros give better relative accuracy than asymptotic zeros except for very large y values. Accuracy using asymptotic zeros is poor for small y . Accuracy using exact zeros declines slowly as ξ increases, but remains acceptable even for quite large values. Curiously, accuracy using approximate zeros actually improves as ξ increases. Table 2 shows a similar comparison for the non-central χ_0^2 distribution. Here the difference is smaller but still in favour of exact zeros. In this case there is no noticeable deterioration with ξ .

— INSERT TABLE 1 ABOUT HERE —

— INSERT TABLE 2 ABOUT HERE —

Tables 1 and Table 2 also give the number of regions of integration required to achieve the convergence criterion and nominal accuracy. In both cases using the exact zeros converges more quickly, although the difference is not great. The number of integration regions required using exact zeros is explored more widely in Table 3 for $p > 2$, and in Table 4 for $1 < p < 2$. The number of regions required is usually less than 20, resulting in a fast computation. For some parameter values the upper limit of 101 regions is reached. These are cases in which the accumulation of floating point errors prevents the nominal accuracy from being achieved.

— INSERT TABLE 3 ABOUT HERE —

— INSERT TABLE 4 ABOUT HERE —

For values of p other than those tested here the number of integration regions and the value of ξ give a guide to accuracy. For large values of p the upper limit for the number of integration regions is reached increasing quickly for large y , suggesting that full accuracy is not being reached. The form of ξ leads us to expect that the difference in accuracy between large and small y will become more pronounced as p increases. This seems to be born out in Table 3 in which the counts increase with p for larger y but decrease with p for small y . It appears that good accuracy can be maintained for $y < \mu$ even for large values of p .

8 Comparing the series and inversion methods

In an earlier publication (Dunn and Smyth, 2005), we explored a different method for evaluating Tweedie densities which did not involve numerical integration. That method consists of summing an infinite series arising from a Taylor expansion of the characteristic function. The series expansion approach is simple and explicit, but becomes prohibitive for some parameter values because the number of terms required to evaluate the density to a given accuracy increases without bound. The number of terms required depends on their rate of decay. Dunn and Smyth (2005) showed that, under exact arithmetic, the number of series expansion terms required for any desired relative accuracy is roughly proportional to $j_{\max} = y^{2-p}/(\phi|2-p|)$. For $p < 2$ the terms in the series are positive so the main issue is computational complexity. For

$p > 2$ the terms alternate in sign so the number of terms equates to error as subtractive cancellation errors accumulate. The expression for j_{\max} shows the series method must fail for y large for $p < 2$ and for small y for $p > 2$. The series method also fails for p near 2. This section compares the series method with the inversion method of the current article.

Figure 3 compares the relative accuracy of the inversion and series methods for the inverse Gaussian distribution with $\mu = 1.4$ and $\phi = 0.74$. The series expansion gives close to full double-precision accuracy for y greater than about 0.1 but deteriorates rapidly for smaller y . For y less than about 0.1, either inversion method is preferred over the series method. The asymptotic zero acceleration method is inferior to exact zeros except for until about $y > 20$. The Fourier inversion method with exact zeros is the most consistent in accuracy. It can be seen return about ten significant figures accuracy over the whole range of y values.

— INSERT FIGURE 3 ABOUT HERE

Figure 4 gives a similar comparison for the non-central χ_0^2 with $p = 1.5$, $\mu = 4$ and $\phi = 2$. In this case the series method is always superior in terms of relative accuracy. This is not unexpected as $p = 1.5$ is exactly halfway between the discontinuities $p = 1$ and $p = 2$ and the series is expected to perform best here. The error of the series method may in fact be somewhat under-estimated here, because the Bessel function, taken to be the exact density, is itself computed using a series expansion.

— INSERT FIGURE 4 ABOUT HERE

What Figure 4 does not show is the number of terms necessary for accurate evaluation. For the series method the number of terms required increases without bound as p approaches 2, $y \rightarrow \infty$ or $\phi \rightarrow 0$ and the inversion solution becomes the only practical option. For example at $p = 1.9999$, $\phi = 0.01$, $\mu = 1$ and $y = 100$ the series method requires 17,212 terms while the inversion method uses the pre-set minimum of 8 integration regions. Note $\xi = 0.009995$ here. Table 5 compares the series and inversion methods as p approaches 2 from above for $\mu = \phi = 1$. The series method fails entirely for p less than 2.07 while the inversion method converges to the correct gamma density value at $p = 2$.

— INSERT TABLE 5 ABOUT HERE

It can be seen that the parameter domains in which the two numerical methods perform best are to a large extent complementary. When $1 < p < 2$, the inversion method performs excellently for large y but gradually deteriorates in accuracy as y decreases. Meanwhile, the series method performs excellently for small y but becomes increasingly computationally expensive as y increases. Hence the inversion method is preferred for large y and the series method for small y . For $p > 2$ the situation is reversed. Here the inversion method is excellent for small y but loses accuracy as y increases, while the series method is excellent for large y but deteriorates in both accuracy and computational complexity as y decreases. Hence the inversion method is preferred for small y and the series method for large y . The question is where to draw the line between the two. A simple guideline which has proved effective is to cut on values of ξ . Use the inversion method in preference to the series method if $\xi < 1$ when $p > 2$, since subtractive cancellation causes the series method to eventually fail in this neighbourhood. For $1 < p < 2$, where the series does not suffer subtractive cancellation, use the inversion method in preference to the series method if $\xi < 0.01$.

9 Discussion and conclusions

This article describes the effective use of Fourier inversion to accurately compute the density functions for Tweedie densities with $p > 1$. The method shows excellent relative accuracy for a wide range of parameter values.

Probability identities were derived which allow the densities to be evaluated by way of the dispersion model form of the density (3). This allows the densities to be computed for all parameter values while using Fourier inversion only at $y = \mu = 1$. This strategy not only simplifies the computations but allows good relative accuracy to be maintained even in the extreme tails of the density functions. The use of the re-scaling identity (6) is critical for achieving our aim of ten-figure relative accuracy, because it allows us to evaluate the density near its mode, where Fourier inversion has the best relative accuracy. Other strategies are shown to be inadequate.

This article also demonstrates the use of advanced acceleration methods to evaluate the infinite oscillating integrals that arise in Fourier inversion. The acceleration methods not only economize on computations but also prove crucial to achieving good

accuracy for many parameter values. These methods could effectively be used in other applications of Fourier inversion in statistics. Two extrapolation methods were implemented. The use of exact zeros of the integrand produces better relative accuracy at the expense of more analytical and computational effort whereas the use of asymptotic zeros is simpler and quicker but the relative accuracy of the results suffer. Given the results presented here, the loss of speed using the exact zero acceleration scheme is in general more than offset by the gain in accuracy.

The algorithms presented in this paper are implemented in the `tweedie` software package for R (R Development Core Team, 2005). The Fourier inversion component is implemented in FORTRAN for speed. Double precision arithmetic was used for all calculations. The function `dtweedie.inversion` implements the algorithms discussed in this paper.

This article considers density functions. The Fourier inversion method can also be adapted to compute the Tweedie cumulative distribution functions (Dunn, 2001). The function `ptweedie.inversion` in the `tweedie` package implements the inversion method for the cumulative distribution function.

The Fourier inversion method is found to complement the series expansion method of evaluation of Dunn and Smyth (2005) in that the two methods perform best in different regions of the parameter space. Simple guidelines are provided to choose between the two methods. It is possible to set up an effective interpolation scheme which blends the inversion and series methods to provide comprehensive evaluation of Tweedie densities across the parameter space (Dunn, 2001). This work makes use of saddlepoint approximations to aid in the blending process and will be published separately. In the meantime it is available in the function `dtweedie` in the `tweedie` package.

References

- Abramowitz, M., and Stegun, I. A. (eds.) 1965. A handbook of mathematical functions. Dover Publications, New York.
- Candy, S. G. 2004. Modelling catch and effort data using generalized linear models, the Tweedie distribution, random vessel effects and random stratum-by-year effects.

- CCAMLR Science. 11: 59–80.
- Davis, P. J. and Rabinowitz, P. 1975. *Methods of Numerical Integration*. Academic Press, New York.
- Dunn, P. K. 2001. Likelihood-based inference for Tweedie exponential dispersion models. Unpublished PhD Thesis, University of Queensland.
- Dunn, P. K. 2004. Occurrence and quantity of precipitation can be modelled simultaneously. *International Journal of Climatology*. 24: 1231–1239.
- Dunn, P. K. and Smyth, G. K. 2005. Series evaluation of Tweedie exponential dispersion models densities. *Statistics and Computing*, 15: 267–280
- Evans, G. 1993. *Practical Numerical Integration*. John Wiley and Sons, New York.
- Feller, W. 1971. *An Introduction to Probability Theory and its Applications, Volume II*, second edition. John Wiley and Sons, New York.
- Gilchrist, R. 2000. Regression models for data with a non-zero probability of a zero response. *Communications in Statistics—Theory and Methods*. 29: 1987–2003.
- Hasegawa, T. and Sidi, A. 1996. An automatic integration procedure for infinite range integrals involving oscillatory kernels. *Numerical Algorithms*. 13: 1–19.
- Johnson, N. L. and Kotz, S. 1970. *Continuous Univariate Distributions—I*. Houghton Mifflin Company, Boston.
- Jørgensen, B. 1987. Exponential dispersion models (with discussion). *Journal of the Royal Statistical Society B*. 49: 127–162.
- Jørgensen, B. 1997. *The Theory of Dispersion Models*. Chapman and Hall, London.
- Jørgensen, B. and Paes de Souza, M. C. 1994. Fitting Tweedie’s compound Poisson model to insurance claims data. *Scandinavian Actuarial Journal*. 1: 69–93.
- Krommer, A. R. and Überhuber, C. W. 1998. *Computational Integration*. Society for Industrial and Applied Mathematics, Philadelphia.
- Lambert, P. and Lindsey, J. K. 1999. Analysing financial returns using regression models based on non-symmetric stable distributions. *Journal of the Royal Statistical*

- Society C. 48: 409–424.
- McCullagh, P. and Nelder, J. A. 1989. Generalized Linear Models, second edition. Chapman and Hall, London.
- Nolan, J. P. 1997. An algorithm for evaluating stable densities in Zolotarev's (M) parameterization. *Mathematical and Computer Modelling*. 29: 229–233.
- Piessens, R., de Doncker-Kapenga, E., Überhuber, C. W. and Kahaner, D. K. 1983. *Quadpack: A Subroutine Package for Automatic Integration*. Springer-Verlag, Berlin.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. 1996. *Numerical Recipes in FORTRAN 77: The Art of Scientific Computing*, second edition. Cambridge University Press, Cambridge.
- R Development Core Team. 2005. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org>.
- Rabinowitz, P. 1992. Extrapolation methods in numerical integration. *Numerical Algorithms*. 3: 17–28.
- Seigel, A. F. 1979. The noncentral chi-squared distribution with zero degrees of freedom and testing for uniformity. *Biometrika*. 66: 381–386.
- Seigel, A. F. 1985. Modelling data containing exact zeros using zero degrees of freedom. *Journal of the Royal Statistical Society B*. 47: 267–271.
- Sidi, A. 1982. An algorithm for a special case of a generalization of the Richardson extrapolation process. *Numerische Mathematik*. 38: 299–307.
- Sidi, A. 1988. A user-friendly extrapolation method for oscillatory infinite integrals. *Mathematics of Computation*. 51: 249–266.
- Sidi, A. 2003. *Practical extrapolation methods: theory and applications*. Cambridge University Press, Cambridge.

Appendix: derivatives of $k(t)$

To successfully evaluate the integrals in Sections 3.2 and 3.3, we need to understand the behaviour of $k(t)$, in particular its real and imaginary components.

First, recall that $\tau(\theta) = \mu$ and note that $V(\mu) = \mu^p$ implies $\tau'(\theta) = \tau^p(\theta)$. The following results are found for $p \neq 2$:

$$\begin{aligned} k(t) &= \frac{\tau^{2-p}(it\phi) - 1}{\phi(2-p)} - ity; \\ k'(t) &= i\tau(it\phi) - iy; \\ k''(t) &= -\phi\tau^p(it\phi); \\ k^{(3)}(t) &= -ip\phi^2\tau^{2p-1}(it\phi). \end{aligned}$$

Some further notation is useful as this point. Write $\alpha = (2-p)/(1-p)$ and $\zeta = \tan^{-1}\{(1-p)t\phi\}$. Note $-\pi/2 < \zeta < 0$ for $p > 1$. Also $1 + (1-p)it\phi = \exp(i\zeta)/\cos(\zeta)$ and $\cos \zeta = \{1 + (1-p)^2t^2\phi^2\}^{-1/2} > 0$. We can now re-write

$$\tau(it\phi) = \frac{\exp\{i\zeta/(1-p)\}}{(\cos \zeta)^{1/(1-p)}}$$

and hence

$$\begin{aligned} \Re k(t) &= \frac{1}{\phi(2-p)} \frac{\cos(\zeta\alpha)}{(\cos \zeta)^\alpha} - \frac{1}{\phi(2-p)}; \\ \Im k(t) &= \frac{1}{\phi(2-p)} \frac{\sin(\zeta\alpha)}{(\cos \zeta)^\alpha} - yt; \\ \Im k'(t) &= \frac{\cos\{\zeta/(1-p)\}}{(\cos \zeta)^{1/(1-p)}} - y; \end{aligned} \tag{16}$$

$$\Im k''(t) = -\phi \frac{\sin\{\zeta p/(1-p)\}}{(\cos \zeta)^{p/(1-p)}}. \tag{17}$$

We can now confirm that the infinite integrals in Sections 3.2 and 3.3 do have finite values. Some algebra shows $\Re k(t) \rightarrow -\infty$ as $t \rightarrow \infty$ when $p > 2$, so the integral in (9) converges. When $1 < p < 2$, we find that $\Re k(t) \rightarrow \pi_0$ as $t \rightarrow \infty$ and $\Im k(t) \rightarrow ty$ so the two terms in the integrand of (11) asymptotically cancel each other out, showing that (11) converges as well. Notice that neither of the individual integrand terms in (11) give convergent integrals if treated alone.

Table 1: Accuracy of Fourier inversion for evaluating inverse Gaussian densities with $p = 3$, $\mu = 1.4$ and $\phi = 0.74$. Results are given for Sidi extrapolation with *exact* and with *asymptotic* zeros. The table gives \log_{10} relative error, the number of integration regions required, and the value of $\xi = \phi/y^{2-p}$ which indexes the degree of difficulty.

y	Exact Density	Exact zeros		Asymptotic zeros		ξ
		Log- Error	No. Ints	Log- Error	No. Ints	
0.001	1.4×10^{-289}	-9.1	5	-4.5	21	0.00074
0.01	5.5×10^{-27}	-9.5	6	-3.4	17	0.0074
0.05	0.00014	-11	9	-5.1	17	0.037
0.1	0.043	-10	10	-6	16	0.074
0.5	0.75	-10	13	-3.4	17	0.37
1	0.44	-10	14	-4.4	18	0.74
2	0.15	-10	15	-6.1	18	1.48
3	0.067	-9.9	16	-5.4	18	2.22
4	0.032	-9.7	16	-3.4	18	2.96
5	0.017	-9.5	16	-2.4	18	3.7
6	0.0094	-9.4	17	-2.2	18	4.44
7	0.0053	-9.4	17	-2.9	18	5.18
8	0.0031	-9.3	17	-5	18	5.92
9	0.0019	-9.3	17	-6.1	18	6.66
10	0.0011	-9.2	17	-6.8	18	7.4
15	0.00011	-9.0	17	-8.4	18	11.1
20	1.3×10^{-5}	-8.9	17	-9	18	14.8
50	1.1×10^{-10}	-8.6	18	-9.9	18	37
100	1.3×10^{-18}	-8.5	18	-9.4	16	74
250	1.1×10^{-41}	-8.2	18	-9.2	17	185
500	1.5×10^{-79}	-8.1	18	-9.1	17	370
750	3.1×10^{-117}	-8.0	18	-9.1	17	555
1000	7.4×10^{-155}	-7.9	101	-9.1	17	740

Table 2: Accuracy of Fourier inversion for evaluating non-central χ_0^2 densities with $p = 1.5$, $\mu = 4$ and $\phi = 2$. Results are given for Sidi extrapolation with *exact* and with *asymptotic* zeros. The table gives the density value from the Bessel function, the \log_{10} relative error, the number of integration regions required, and the value of $\xi = \phi/y^{2-p}$ which indexes the degree of difficulty.

y	Bessel Form	Exact zeros		Asymptotic zeros		ξ
		Log- Error	No. Ints	Log- Error	No. Ints	
0.001	0.135335	-11	13	-8.4	12	63.2
0.01	0.135335	-9.9	13	-9.0	12	20
0.05	0.135321	-9.8	13	-9.6	12	8.94
0.1	0.13528	-9.8	13	-10	12	6.32
0.5	0.134039	-9.9	13	-11	12	2.83
1	0.130567	-9.9	13	-11	12	2.00
2	0.119232	-10	13	-10	12	1.41
3	0.104744	-9.9	13	-10	12	1.15
4	0.0893754	-9.9	13	-10	12	1.00
5	0.0745335	-9.9	12	-10	12	0.894
6	0.0610078	-9.9	12	-10	12	0.816
7	0.0491654	-9.9	13	-10	12	0.756
8	0.0391003	-9.8	13	-9.9	12	0.707
9	0.0307413	-9.8	13	-9.9	12	0.667
10	0.0239277	-9.8	13	-9.9	12	0.632
15	0.00608121	-9.8	13	-9.9	12	0.516
20	0.00134336	-9.7	13	-9.9	12	0.447
50	3.8×10^{-8}	-9.7	14	-9.8	14	0.283
100	1.1×10^{-16}	-9.7	14	-9.0	17	0.200
1000	4.5×10^{-194}	-9.6	10	-9.6	15	0.0632

Table 3: The number of integration regions required to evaluate various Tweedie densities with $p > 2$. The mean $\mu = 1$ in all cases and the exact zero acceleration algorithm is used. Five regions of integrations is the minimum possible and 101 is the maximum.

p	ϕ	y						
		0.001	0.01	1	5	10	100	1000
2.001	10	13	13	13	13	13	13	13
2.001	1	11	11	11	11	11	11	11
2.001	0.1	10	10	10	10	10	10	10
2.001	0.01	6	6	6	6	6	6	6
2.01	10	14	14	14	14	14	14	14
2.01	1	13	13	13	13	13	13	13
2.01	0.1	9	9	10	10	10	10	10
2.01	0.01	6	6	6	6	6	6	6
2.5	10	12	15	17	17	17	18	18
2.5	1	8	10	15	16	16	17	17
2.5	0.1	5	6	10	12	12	15	16
2.5	0.01	101	5	6	8	8	10	12
3	10	7	10	17	18	18	101	101
3	1	5	7	15	17	17	18	101
3	0.1	101	5	10	13	15	17	18
3	0.01	5	5	7	9	10	15	17
5	10	5	5	17	101	101	101	101
5	1	5	5	14	101	101	101	101
5	0.1	5	5	11	17	101	101	101
5	0.01	5	5	8	15	17	101	101

Table 4: The number of integration regions required to evaluate various Tweedie densities with $1 < p < 2$. The mean is $\mu = 1$ in all cases and the exact zero acceleration algorithm is used. Eight regions of integrations is the minimum possible and 101 is the maximum.

p	ϕ	y						
		0.001	0.01	1	5	10	100	1000
1.01	10	11	11	101	15	17	17	17
1.01	1	11	11	17	18	17	17	21
1.01	0.1	11	101	17	16	17	21	22
1.01	0.01	101	17	17	20	21	22	23
1.5	10	12	12	12	12	12	12	13
1.5	1	12	12	12	12	13	16	15
1.5	0.1	12	12	16	15	15	17	20
1.5	0.01	13	16	17	19	20	21	22
1.7	10	12	12	13	13	13	14	14
1.7	1	13	13	15	18	18	16	16
1.7	0.1	16	18	16	16	16	16	17
1.7	0.01	15	16	17	18	19	20	21
1.9	10	14	14	14	14	14	14	15
1.9	1	15	15	15	18	18	16	14
1.9	0.1	14	17	16	16	16	16	16
1.9	0.01	16	17	17	17	17	18	19
1.99	10	13	13	13	13	13	13	14
1.99	1	13	13	14	18	18	16	15
1.99	0.1	16	16	16	16	16	16	16
1.99	0.01	17	17	17	17	17	17	17

Table 5: Computed densities from the inversion and series evaluation algorithms as $p \rightarrow 2$ from above, for $\mu = \phi = 1$. The series method failes entirely for $p < 2.07$ while the inversion method converges to the correct gamma density value at $p = 2$ of 0.3678794.

p	Computed density		Terms required	
	Inversion	Series	Inversion	Series
2.00001	0.367880	0.000000	11	5444
2.00005	0.367881	0.000000	11	2438
2.00010	0.367882	> 100	11	1724
2.00020	0.367885	0.000000	11	1220
2.00030	0.367888	0.000000	11	998
2.00040	0.367891	0.000000	11	864
2.00050	0.367894	> 100	11	774
2.00060	0.367897	> 100	11	706
2.00070	0.367900	> 100	11	654
2.00080	0.367903	0.000000	11	612
2.00090	0.367905	> 100	11	578
2.00100	0.367908	> 100	11	548
2.00500	0.368024	> 100	12	246
2.01000	0.368169	> 100	13	176
2.05000	0.369342	> 100	14	73
2.07000	0.369935	0.370289	14	61
2.10000	0.370832	0.370832	14	52
2.15000	0.372344	0.372344	14	43
2.20000	0.373874	0.373874	14	39




Figure 1: Two example integrands. Top: Showing the slow decay possible for the integrand. The integrand is computed for $p = 3$, $\phi = 2$ and $y = 10$, showing the slowly decaying nature of the integrand. The thick solid line is the integrand; the dashed lines are the envelope of the integrand. The shaded regions have areas of approximately 0.064407 and -0.061380 respectively. Bottom: Showing the irregular initial integrands sometimes possible when $1 < p < 2$. The integrand is computed for $p = 1.02$, $\phi = 0.5$ and $y = 0.01$.




Figure 2: Relative errors of the three Fourier inversion re-scaling strategies mentioned in Section 3.1 to compute densities with $p = 3$, $\mu = 1.4$ and $\phi = 0.74$. Method 3, which transforms all densities to $y = \mu = 1$, gives smaller relative errors. All computations are performed using Sidi extrapolation with exact zeros.




Figure 3: Relative errors of inversion and series methods for evaluating inverse Gaussian densities with $\mu = 1.4$ and $\phi = 0.74$. The thick solid line is the inversion method using Sidi extrapolation with *exact* zeros, the dashed line is inversion with *asymptotic* zeros, and the dotted line is the series method. The thin solid horizontal line represents the target accuracy of 10^{-10} .

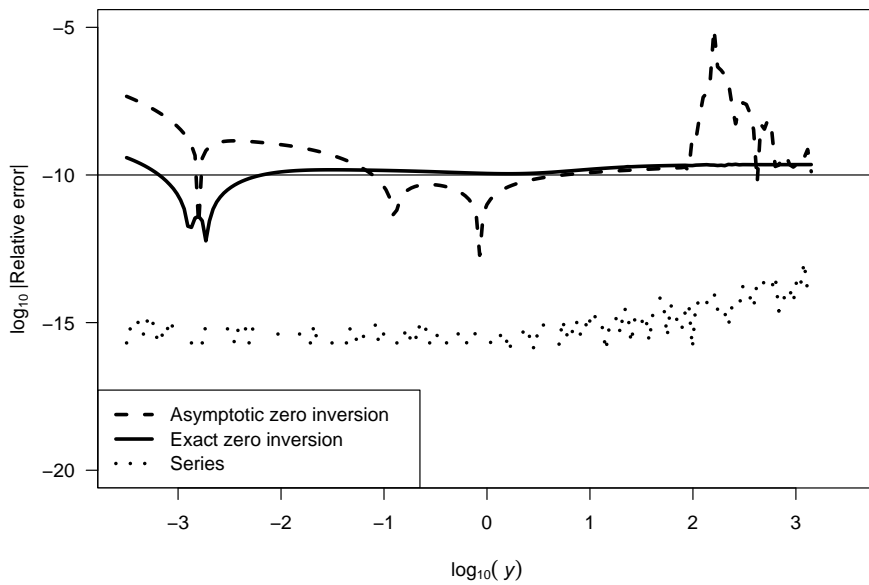


Figure 4: Relative errors of series and inversion methods for evaluating non-central χ_0^2 densities with $p = 1.5$, $\mu = 4$ and $\phi = 2$. The thick solid line is the inversion method using *exact* zeros, the dashed line is the modified method with *asymptotic* zeros, and the dotted line is the series method. The thin solid horizontal line represents the target accuracy of 10^{-10} .