

Occurrence and quantity of precipitation can be modelled simultaneously

Peter K Dunn*

April 28, 2004

Abstract

Many statistical models exist for modelling precipitation. One difficulty is that two issues need to be addressed: the probability of precipitation occurring, and then the quantity of precipitation recorded. This paper considers a family of distributions for modelling the quantity of precipitation including those observations in which exactly no precipitation is recorded. Two examples are then discussed showing the distributions model the precipitation patterns well.

Keywords: Precipitation models; Poisson distribution; gamma distribution; Tweedie distribution; generalized linear models.

*Department of Mathematics and Computing, Faculty of Sciences, University of Southern Queensland, Toowoomba, Queensland, 4350 AUSTRALIA. email; dunn@usq.edu.au; phone: +61 7 4631 5527; fax: +61 7 4631 5550.

1 Introduction

Many researchers have studied precipitation modelling to enable a better understanding of erosion, runoff and pollutant transport, for example. Many simulation studies need to have a good model for precipitation which are then used in the design and management of irrigation systems, farm management systems and water supplies. They can also be used in computer simulations such as APSIM (McCown *et al.*, 1995).

As a consequence of the large number of subsequent applications, many characteristics of precipitation can be modelled statistically. Models have been developed, for example, for the amount of precipitation that falls, the length of dry and wet spells, and the number of rain days per month. Precipitation can be modelled on numerous time scales, from daily precipitation to monthly and seasonal. For some applications, it is necessary to understand the number of rain events in a given day and the intensity of the precipitation in each (see, for example, Connolly *et al.*, 1998).

One difficulty with modelling precipitation is that precipitation is continuous with exact zeros (when no precipitation is recorded). Most statistical models have difficulty with this mixture of discrete and continuous distributions.

In this paper, a precipitation model is considered which simultaneously models the discrete and continuous features of precipitation, with a sensible interpretation of the parameters. In the next section, brief notes are made about some models already proposed for precipitation modelling. In Section 3, a family of distributions is presented for precipitation modelling with exact zeros. Some notes on parameter

estimation are given in Section 4. Section 5 gives two examples, and some comments are made in Section 6

2 Current approaches to precipitation modelling

Many different methods have been tried for modelling both the continuous and discrete components of precipitation. Glasbey & Nevison (1997) and Glasbey & Durban (2001), for example, discuss a model using a monotonic transformation to define a latent Gaussian variable; zero rainfall then corresponds to censored values below some threshold. Alternatively, the modelling can be split into two parts: the initial task is to identify when actual precipitation is recorded, and then the amount of this precipitation is modelled. This second approach, reviewed here, forms the starting point for this paper. This review is in no way a complete review of the vast body of research, but is intended to present the broad flavour of models that have been used.

2.1 Models for rain days

The modelling of rain days (days on which precipitation has been recorded) has often been based on Markov chains. The simplest has two-states ('rain' and 'no rain') and is first order (the rain probability depends only on the previous time period). An obvious extension is to extend the order of the Markov chains (see, for example, Coe & Stern (1982)). Other models include Markov chains of 'hybrid order', in which first-order Markov chains are used for wet spells but higher order dependence is used for dry spells. For the first-order Markov chains, the lengths

of dry and wet spells have been modelled using a geometric distribution. Katz & Parlange (1995) use a two-stage process also: the occurrence of precipitation is based on a first-order Markov process.

Instead of Markov chains, others have used a mixture of geometric distributions or the negative binomial distribution. Feuerverger (1979) uses a logistic regression model for the precipitation probability; similarly, Chandler & Wheater (2002) use a binomial-based generalized linear model.

2.2 Modelling precipitation amounts

Numerous distributions have been used to model precipitation amounts on rain days ('daily intensity'). Because precipitation is generally heavily skewed to the right, distributions such as the gamma and exponential have formed the basis for this work.

Chapman (1998) states that the skewed normal distribution has been used by Nicks & Lane (1989) and the Weibull distribution by Zucchini & Adamson (1984). Katz & Parlange (1995) model precipitation quantity using a power transformation to normality.

More commonly used, however, is the gamma distribution. Allan & Haan (1975), for example, used a special case of the gamma distribution, the exponential, to model precipitation.

Das (1955) considered a truncated gamma distribution for modelling precipitation on all days, not just rain days. He chooses a small interval $(0, \delta)$ and ignores the actual values less than δ . However, knowing the number of observations less

than δ enables maximum likelihood estimators of the parameters to be found. A similar approach is adopted by Wilks (1990).

Daily precipitation is sometimes modelled as an independent, identically distributed process. Sometimes autocorrelation is included. Wilks (1999) states that for daily precipitation, the autocorrelation of (positive) precipitation amounts has been found to be statistically significant, but of little practical significance.

Stern & Coe (1984) use a shifted gamma distribution since precipitation amounts less than some threshold, say δ , cannot be recorded.

The identical distribution assumption can obviously be relaxed. Some researchers have used different gamma distributions depending on whether or not the previous day had rain. Others have used three different gamma distributions: one for (i) single wet days; (ii) the first days in spells of wet days; (iii) the subsequent wet days in spells of wet days. Wilks (1999) shows the mean precipitation generally increases from (i) to (iii), and uses three different gamma distributions with a common shape parameter but different scale parameters. This produces a four-parameter model. Similar models are also studied by Chapman (1998).

Wilks (1999) also considers a mixed exponential distribution. This is a mixture of two different exponential distributions with different parameters, resulting in a three-parameter model.

For all these models, it is not uncommon for the parameters to vary throughout the year. This can be modelled using different parameters for each month or each season. This results in a large number of parameters. Other approaches include modelling parameter variation with a polynomial or a Fourier series (Stern & Coe,

1984).

Feuerverger (1979) also uses a gamma distribution, conditional on precipitation being recorded. A logistic regression model is used to model the probability of precipitation, with a gamma for the precipitation on rain days. The precipitation is then modelled using the joint density. Feuerverger uses the cloud seeding data of Bethwaite *et al.* (1966). Chandler & Wheater (2002) use a gamma-based generalized linear model.

3 The Tweedie model

The discussion in Section 2.2 suggests using the gamma distribution (and variations) for modelling precipitation on rain days is very common. On this basis, the following argument is presented.

Assume any precipitation event i results in an amount of precipitation R_i , and that each R_i has a gamma distribution $\text{Gam}(-\alpha, \gamma)$. (In the parameterization used here, the mean is $-\alpha\gamma$ and variance $-\alpha\gamma^2$.) The negative value for α is used to be consistent with that used elsewhere in this paper. Then, assume the number of precipitation events in any one day, say N , has a Poisson distribution. Note this implies there will be days with no precipitation events (when $N = 0$). The total daily precipitation, Y , can be found as the Poisson sum of the gamma random variables so that

$$Y = R_1 + R_2 + \cdots + R_N,$$

where N has a Poisson distribution with mean λ .

An identical argument can be applied to monthly precipitation, when R_i could refer to the precipitation recorded on any one day, and Y the total monthly precipitation. The generalization to longer time scales follows.

The distribution of total precipitation Y can be deduced by working with cumulant generating functions (Smyth, 1996) by noting that Y given N has the gamma distribution $\text{Gam}(-N\alpha, \gamma)$. The resulting distribution has been called a compound Poisson distribution (Bar-Lev & Stramer, 1987, Feller, 1968, Jørgensen & Paes de Souza, 1994, Smyth & Jørgensen, 1999), a compound gamma distribution (Johnson & Kotz, 1970), or a Poisson–gamma distribution (Smyth, 1996). The resulting probability function is complicated, and can be written as

$$\log f_p(y; \mu, \phi) = \begin{cases} -\lambda & \text{for } y = 0 \\ -y/\gamma - \lambda - \log y \\ \quad + \log W(y, \phi, p) & \text{for } y > 0. \end{cases}$$

where $\gamma = \phi(p-1)\mu^{p-1}$ and $\lambda = \mu^{2-p}/[\phi(2-p)]$, and W is an example of Wright's generalized Bessel function (Wright, 1933), but cannot be written in terms of more common Bessel functions. W can be expressed as the infinite summation

$$W(y, \phi, p) = \sum_{j=1}^{\infty} \frac{y^{-j\alpha} (p-1)^{\alpha j}}{\phi^{j(1-\alpha)} (2-p)^j j! \Gamma(-j\alpha)}$$

where $\alpha = (2-p)/(1-p)$. The mean of the distribution is μ and the variance is $\text{var}[Y] = \phi\mu^p$. In these formulae, $1 < p < 2$ is the index which determines which Poisson–gamma distribution is used. Importantly, the probability of recording no

precipitation is given by

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

The Poisson–gamma distributions belong to the class of distributions known as the Tweedie family of distributions, named by Jørgensen (1987, 1997) after Tweedie (1984). These distributions have a variance of the form $\text{var}[Y] = \phi\mu^p$ for $p \notin (0, 1)$. For $1 < p < 2$, the distributions have supported on the non-negative real numbers, and the distribution are the Poisson–gamma distributions just discussed; when $\phi = 1$, as p tends to 1 from above the distribution tends to a series of spikes located on the integers and the probabilities are those of the Poisson distribution. As p tends to 2 from below, the distribution tends to a gamma distribution. For $p > 2$ the distributions are continuous for positive Y . The distributions for which $p < 0$ are not considered here; the distributions are continuous on the entire real axis. A referee suggested these distributions may be of use in modelling temperatures. The normal distribution is a special case of the Tweedie family with $p = 0$. No distributions exists for $0 < p < 1$ (Jørgensen, 1987). The Tweedie distributions for which $p > 2$ can also be useful in precipitation modelling as these distributions have a similar shape to the gamma but are more right skewed. Figure 1 show the density function of some Tweedie distributions for various p .

Mathematically, the distributions are best analysed using the (μ, ϕ, p) parametrization; climatologically, parametrization in terms of $(\lambda, \gamma, \alpha)$ is more useful. In this parameterization, λ refers to the mean number of precipitation events per month; γ to the shape of the precipitation events; and $-\alpha\gamma$ to the mean quantity of precipita-

tion per event. Using this parameterization, it is possible to investigate finer-scale climatological structure when only coarse-scale (eg monthly) data are available.

There are some important properties of the Tweedie distributions that make them particularly appealing for use in precipitation modelling:

- There is some intuitive appeal for the models, considering total precipitation as a sum of precipitation on smaller time-scales (outlined above).
- These distributions belong to the exponential family of distributions (McCullagh & Nelder)), upon which generalized linear models are based. Consequently, there is a framework already in place for fitting models based on the Tweedie distributions, and for diagnostic testing. In addition, covariates can be incorporated into the modelling procedure; this, however, is out of the scope of the current paper.
- They provide a mechanism for understanding the fine-scale structure in coarse-scale data.

Despite these attractive features, using the distributions is not straightforward. Indeed, no closed forms exist for evaluating the density function or cumulative distribution functions except in special cases; numerical methods are required, such as evaluating an infinite series or an evaluating infinite oscillating integrals. However, programs for evaluating the density and cumulative distribution function for these distributions exist that provide fast and accurate algorithms for almost all parameter values (Dunn & Smyth, 2003). In addition, these programs allow the computation of quantile residuals (Dunn & Smyth, 1996) in diagnostic analysis.

These residuals prove superior to the more often used deviance and Pearson residuals in situations where responses contain discrete responses. (Quantile residuals use the minimum amount of randomization on the cumulative density scale for discrete point(s) to produce continuous, Gaussian residuals, and avoid distracting clumps in residual plots corresponding to the discrete points.)

In related work, Seigel (1975) and Seigel (1985) use the non-central χ^2 -distribution, a special case of the Tweedie distributions with $1 < p < 2$, to model snowfall with exact zeros. Jørgensen (1987) uses Tweedie distributions in a precipitation example.

4 Estimation

In this section, methods for estimating the parameters of Tweedie distributions are considered. The Tweedie family is a three parameter family of distributions in μ (the mean), $\phi > 0$ (the dispersion parameter) and p . In the generalized linear model context, μ corresponds to predicted values and can be estimated using standard algorithms (see McCullagh & Nelder (1989) for details). Since we restrict ourselves to using no covariates, μ can be estimated by the sample mean. Importantly however, maximum likelihood estimates for μ can be found based on only the first two moments, even in the full generalized linear model case. This means only the first two moments of the distribution are necessary for maximum likelihood estimation of the linear predictor based on the distributional assumption.

The maximum likelihood estimation of ϕ is more difficult; complicated algorithms are available for maximum likelihood estimation of ϕ (see Dunn & Smyth

(2001) and Dunn & Smyth (2003)). An alternative method of estimation is to use the mean deviance estimator of ϕ ; this is closely related to using the quasi-likelihood in place of the true likelihood, which has been shown by Dunn & Smyth (2003) to have difficulties for continuous data with exact zeros. In this paper, the maximum likelihood estimate of ϕ is used throughout.

Estimating the maximum likelihood value of p is performed using a profile (log-) likelihood plot which requires the computation of the density. This is difficult, though computational programs are available (Dunn & Smyth, 2003). For a given fixed value of p , estimates of μ and ϕ can be computed as above, and the log-likelihood computed. The value of p for which the log-likelihood is maximum is chosen as the maximum likelihood value. Some examples are shown in the next section. Nominal confidence intervals for p can also be found, using that $2[\log L(\hat{p}) - \log L(p_0)]$ has, asymptotically, a χ_1^2 distribution, where p_0 is the true parameter value.

An alternative to using an accurate density evaluation is the saddlepoint approximation to the Tweedie densities (Jørgensen, 1997). This simpler solution may produce estimates that are very inaccurate. An example of the problems that may result is given in Dunn & Smyth (2003).

5 Examples

In this section, two examples are given of how the Tweedie distributions can be used for precipitation modelling. First, monthly precipitation at Charleville is considered with a different Tweedie distribution for each month. Secondly, Melbourne

daily and monthly precipitation is examined.

5.1 Charleville monthly precipitation

Consider the total monthly precipitation recorded at Charleville, Queensland, Australia, from 1882 to 1994. There are 113 observations for each month. For the months of January, November and December, precipitation was recorded every year, while other months have years with no recorded precipitation. The maximum likelihood estimates of p and ϕ for each month are shown in Table 1. A typical profile likelihood plot for selecting the value of p is shown in Figure 2. Note that for November, $\hat{p} = 2.30$, which is not too dissimilar to a gamma distribution ($p = 2$). In this case, the nominal 95% confidence interval almost includes $p = 2$, which suggests that perhaps only slight improvements can be made over using the gamma distribution. However, the other cases suggest significant improvements can be made over the gamma distribution.

Table 1 shows the maximum likelihood estimators of the mathematical parameters; Table 2 shows the parameters expressed in terms of λ (the mean number of precipitation events per month), γ (the shape of the rainfall gamma distribution) and $-\alpha\gamma$ (the amount of rain per event). Note that for January and November, this interpretation is nonsense since $p > 2$. Nonetheless, this interpretation is enlightening when $1 < p < 2$; for example, the model estimates 6.7 precipitation events in June with a mean rainfall of 42mm.

To assess the quality of the fitted distributions, quantile residuals can be used (Dunn & Smyth, 1996). One feature of quantile residuals is they have an ex-

act standard normal distribution (apart from sampling error) provided the correct distribution is used. A typical QQ-plot of these quantile residuals are shown in Figure 3. In all cases, the plots show the distributions model the total monthly precipitation well. (In contrast, using deviance or Pearson residuals makes this decision difficult as the residuals corresponding to exact zeros form distinct and distracting lines in the plots.)

5.2 Melbourne daily precipitation

Data for Melbourne daily precipitation has been taken from the Time Series Data Library (Hyndman (2001), and originally from the Australian Bureau of Meteorology). The data give the daily precipitation in Melbourne, Victoria, Australia from 1981 to 1990. For the purpose of demonstration, the daily precipitation from the months of April have been chosen.

The data used consists of 300 daily observations, of which 197 are zeros (no precipitation). A profile likelihood plot for estimating p for the daily data, (Figure 4) shows the maximum likelihood estimate of p is approximately 1.61.

Figure 5 shows a QQ-plot of the quantile residuals from fitting a model to the daily precipitation for April. The plot indicates that the distribution fits the data well.

6 Comments and Discussion

The Tweedie family of distributions has proven to be useful for modelling precipitation on a daily and monthly time-scale. The main appeal of the distributions

is they model the precipitation amounts when precipitation is recorded as well as when precipitation has not been recorded. That is, separate models are not necessary for modelling the occurrence of precipitation days and the amount of precipitation. The model parameters can also have a useful interpretation.

Further, there is a theoretical justification for using these distributions, since for $1 < p < 2$ they can be seen as a Poisson sum of gamma distributions. The number of precipitation events have been modelled using a Poisson distribution and the precipitation amounts with a gamma distribution.

The possibility of using the distributions in more advanced applications is implied since the distributions belong to the exponential family of distributions which form the basis of generalized linear models.

7 Acknowledgments

The author thanks Queensland Department of Primary Industries for the data. The comments of a reviewer are most gratefully acknowledged; they improved the flow, interpretation and understanding of the paper.

8 References

Allan DM, Haan CT. 1975. Stochastic simulation of daily rainfall. Research Report No. 82, Water Resources Institute, University of Kentucky.

Bar-Lev SK, Stramer O. 1987. Characterizations of natural exponential families with power variance functions by zero regression properties. *Probability and Related*

Fields. **76**: 509–522.

Bethwaite FD, Smith EJ, Warburton JA, Heffernan KJ. 1966. Effects of seeding isolated cumulus clouds with silver iodide. *Journal of Applied Meteorology*. **5**: 513–520.

Chandler RE, Wheater HS. 2002. Analysis of rainfall variability using generalized linear models: A case study from the west of Ireland. *Water Resources Research*. **38**(10), 1192–1202.

Chapman, T. 1998. Stochastic modelling of daily rainfall: the impacts of adjoining wet days on the distribution of rainfall amounts. *Environmental Modelling and Software*. **13**: 317–324.

Coe R, Stern RD. 1982. Fitting models to daily rainfall data. *Journal of Applied Meteorology*. **21**: 1024–1031.

Connolly RD, Schirmer J, Dunn PK. 1998. A daily rainfall disaggregation model. *Agricultural and Forest Meteorology*. **92**: 105–117.

Das SC. 1955. The fitting of truncated Type III curves to daily rainfall data. *Australian Journal of Physics*. **8**: 298–304.

Dunn PK, Smyth GK. 1996. Randomized quantile residuals. *Journal of Computational and Graphical Statistics*, **5**: 236–244.

Dunn PK, Smyth GK. 2001. Tweedie family densities: methods of evaluation. *Proceedings of the 16th International Workshop on Statistical Modelling*, Odense, Denmark, 2–6 July.

Dunn PK, Smyth GK. 2003. Series evaluation of Tweedie exponential dispersion model densities. Submitted.

- Feller W. 1968. *An Introduction to Probability Theory and its Applications*, Volume I, third edition. John Wiley and Sons: New York.
- Feuerverger A. 1979. On some methods of analysis for weather experiments. *Biometrika*, **66**: 655–658.
- Durban M, Glasbey CA. 2001. Weather modelling using a multivariate latent Gaussian model. *Agricultural and Forest Meteorology*. **109**: 187–201.
- Glasbey CA, Nevison, IM. 1997. Rainfall modelling using a latent Gaussian variable. In: Gregoire, TG, Brillinger DR, Duggle PJ, Russek-Cohen E, Warren WG, Wolfinger RD (Eds.). *Modelling Ongitudinal and SPatially COrelated Data: Methods, Applications, and Future Directions*. No. 122 in *Lecture Notes in Statistics*, Springer, New York, 233–242.
- Hyndman R. 2001. *The Time Series Data Library*, located at <http://www-personal.buseco.monash.edu.au/hyndman/TSDL/>. Accessed 15 November 2001.
- Johnson NL, 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, Series B*. **49**: 127–162.
- Jørgensen B. 1997. *The Theory of Dispersion Models*, Chapman and Hall: London.
- Jørgensen B, Paes de Souza MC. 1994. Fitting Tweedie’s compound Poisson model to insurance claims data. *Scandinavian Actuarial Journal*. **1**: 69–93.
- Katz RW, Parlange MB. 1995. Generalizations of chain-dependent processes:

applications to hourly precipitation. *Water Resources Research*. **31**: 1331–1341.

McCown RL, Hammer GL, Hargreaves JNG, Holzworth D, Huth NI. 1995. AP-SIM: an agricultural production system simulation model for operational research. *Mathematics and Computers in Simulation*. **39**: 225–231.

McCullagh P, Nelder JA (1989). *Generalized Linear Models*, second edition, Chapman and Hall: London.

Nicks AD, Lane LJ. 1989. Chapter 2: Weather Generator. In *Water Erosion Prediction Project*, National Erosion Research Laboratory, editors L. J. Lane and M. A. Nearing, West Lafayette, 2.1–2.19.

Seigel AF. 1979. The noncentral chi-squared distribution with zero degrees of freedom and testing for uniformity. *Biometrika*, **66**: 381–386.

Seigel AF. 1985. Modelling data containing exact zeros using zero degrees of freedom. *Journal of the Royal Statistical Society B*, **47**: 267–271.

Smyth GK. 1996. Regression analysis of quantity data with exact zeros. *Proceedings of the Second Australia–Japan Workshop on Stochastic Models in Engineering, Technology and Management*. Technology Management Centre, University of Queensland, 572–580.

Smyth GK, Jørgensen B. 1999. Fitting Tweedie’s compound Poisson model to insurance claims data: dispersion modelling. In *Proceedings of the 52nd Session of the International Statistical Institute*, Helsinki, Finland, August 10–18, Contributed Paper Meeting 68: Statistics and insurance.

Stern RD, Coe R. 1984. A model fitting analysis of daily rainfall data (with discussion). *Journal of the Royal Statistical Society, Series A*, **147**: 1–34.

Tweedie MCK. 1984. An index which distinguishes between some important exponential families. *Statistics: Applications and New Directions. Proceedings of the Indian Statistical Institute Golden Jubilee International Conference*, editors J. K. Ghosh and J. Roy, Calcutta: Indian Statistical Institute, 579–604.

Wilks DS. 1990. Maximum likelihood estimation for the gamma distribution using data containing zeros. *Journal of Climate*, **3**: 1495–1501.

Wilks DS. 1999. Interannual variability and extreme-value characteristics of several stochastic daily precipitation models. *Agricultural and Forest Meteorology*, **93**: 153–169.

Wright EM. 1933. On the coefficients of power series having essential singularities. *Journal of the London Mathematical Society*, **8**: 71–79.

Zucchini W, Adamson PT. 1984. The occurrence and severity of droughts in South Africa, 91/1/94, Water Research Commission, Pretoria.

Table 1: Maximum likelihood estimates of p and ϕ for fitting Tweedie distributions to total monthly precipitation at Charleville.

Month	\hat{p}	$\hat{\phi}$	Month	\hat{p}	$\hat{\phi}$
Jan	2.4	0.0068	Jul	1.6	4.8
Feb	1.7	2.7	Aug	1.6	5.3
Mar	1.7	3.8	Sep	1.7	4.5
Apr	1.6	5.9	Oct	1.7	2.4
May	1.5	5.1	Nov	2.3	0.016
Jun	1.6	3.9	Dec	1.9	0.0084

Table 2: The maximum likelihood estimates from Table 1 for Charleville, reparameterized in terms of λ (the mean number of precipitation events per month), γ (the shape of the rainfall gamma distribution) and $-\alpha\gamma$ (the amount of rain per event). Note that for January and November, this interpretation is nonsense as $p > 2$.

Month	$\hat{\lambda}$	$\hat{\gamma}$	$-\hat{\alpha}\hat{\gamma}$	Month	$\hat{\lambda}$	$\hat{\gamma}$	$-\hat{\alpha}\hat{\gamma}$
Jan	-25.3	91.0	-26	Jul	4.6	102.3	62
Feb	9.4	150.9	73	Aug	4.1	68.7	49
Mar	5.8	244.2	102	Sep	3.7	138.7	58
Apr	4.1	126.9	80	Oct	7.6	114.5	45
May	6.2	58.0	51	Nov	-49.5	40.4	-8
Jun	6.7	56.5	42	Dec	3027.1	3.1	0.2




Figure 1: Some Tweedie densities function. The polygons indicate the discrete probability of $Y = 0$ when $1 < p < 2$. In each case, the mean and variance are fixed at 1.




Figure 2: The profile likelihood functions showing the maximum (log-) likelihood values of p for July precipitation totals in Charleville. The horizontal line indicates approximate 95% confidence intervals for p .

July rainfall




Figure 3: The QQ-plot of the quantile residuals after fitting a Tweedie distribution to July precipitation totals in Charleville. An ideal plot would show the points falling on the solid line, which corresponds to the standard normal distribution.




Figure 4: The profile likelihood functions showing the maximum likelihood values of p for daily Melbourne precipitation in April. The horizontal dashed line indicates approximate 95% confidence intervals for p . The maximum likelihood estimate of p is 1.61.

Apr rainfall




Figure 5: A QQ-plot of the quantile residuals after fitting a Tweedie distribution to daily precipitation totals in Melbourne. An ideal plot would show the points falling on the solid line, which corresponds to the standard normal distribution.