

# The Tweedie Power Variance Function Family

Stephen J. Mildenhall

2020-10-20

## 1 The Tweedie-Power Variance Function Family

### 1.1 Introduction

Readers who have worked with GLMs will recall seeing an exhibit in the documentation for their software similar to Table 1. It shows the univariate exponential family distributions with a power variance function (PVF)

$$V(\mu) = \mu^p,$$

for different values of  $p$ . This section will answer several questions raised by the table: Why do these distributions appear? Why are they in this order? What about  $0 < p < 1$  or  $p > 2$  or even  $p < 0$ ? And, what ties all of these distributions together?

Table 1: Exponential families with PVFs.

Distribution Name	Variance Function $V(\mu)$
Normal	1
Poisson	$\mu$
Tweedie	$\mu^p, 1 < p < 2$
Gamma	$\mu^2$
Inverse Gaussian	$\mu^3$

This section begins by identifying all PVF distributions from their cumulant generating functions. In the process, it shows that the Tweedie family is a compound Poisson with gamma severity and identifies the Lévy measure that underlies the PVF family. In conjunction with the Tweedie limit theorem at the end of [Part III](#), this measure can be regarded as a kind of **universal severity distribution**. The section goes on to show how the Tweedie interpolates between diversifying and non-diversifying growth models and concludes by illustrating the range of shapes its density can take.

### 1.2 Identifying All Members of The Power Variance Family

For any NEF we know  $\kappa'(\theta) = \tau(\theta) = \mu$  and  $\kappa''(\theta) = \tau'(\tau^{-1}(\mu)) = 1/(\tau^{-1})'(\mu) = V(\mu)$ , by the results in [Part II](#). For the power variance family, we can integrate  $(\tau^{-1})'(\mu) =$

$1/V(\mu) = \mu^{-p}$  to get

$$\theta = \tau^{-1}(\mu) = \frac{\mu^{1-p}}{1-p},$$

provided  $p \neq 1$ . Rearranging determines

$$\kappa'(\theta) = \mu = \{(1-p)\theta\}^{\frac{1}{1-p}}.$$

We can integrate this expression for  $\kappa'(\theta)$  and, for  $p \neq 2$ , obtain

$$\begin{aligned} \kappa(\theta) &= (1-p)^{\frac{1}{1-p}} \theta^{\frac{2-p}{1-p}} \Big/ \frac{2-p}{1-p} \\ &= (1-p)^{\frac{2-p}{1-p}} \frac{\theta^{\frac{2-p}{1-p}}}{2-p} \\ &= \frac{1}{2-p} \{(1-p)\theta\}^{\frac{2-p}{1-p}} \end{aligned}$$

since  $\frac{1}{1-p} + 1 = \frac{2-p}{1-p}$ . Substituting  $\alpha = \frac{2-p}{1-p}$ , and noting  $p = \frac{\alpha-2}{\alpha-1}$ ,  $1-p = \frac{1}{\alpha-1}$ , and  $\frac{1}{2-p} = \frac{\alpha-1}{\alpha}$ , yields

$$\kappa(\theta) = \frac{\alpha-1}{\alpha} \left( \frac{\theta}{\alpha-1} \right)^\alpha.$$

For reference, the most basic stable distribution with tail parameter  $\alpha$  has cumulant generating function  $-\gamma|\theta|^\alpha$  for a scale constant  $\gamma$ , which has the same form.

The two excluded cases,  $p = 1, 2$ , work out as follows. When  $p = 1$ ,  $\alpha = \infty$ . The first integral yields  $\theta = \log(\mu)$  and the second  $\kappa(\theta) = e^\theta$ , giving a Poisson. And when  $p = 2$ ,  $\alpha = 0$ . The first integral yields  $\theta = -\mu^{-1}$  and the second  $\kappa(\theta) = -\log(-\mu)$ , giving a gamma.

One further case will occur, which in a sense corresponds to  $p = \infty$ . It turns out the correct interpretation is  $V(\mu) = e^\mu$ . Then, the first integral yields  $\theta = -e^{-\mu}$  and the second  $\kappa(\theta) = \theta - \theta \log(-\theta)$ . This will give an extreme Cauchy distribution.

To summarize: the cumulant generator for the power variance family is given by

$$\kappa_p(\theta) = \begin{cases} \frac{\alpha-1}{\alpha} \left( \frac{\theta}{\alpha-1} \right)^\alpha & p \neq 1, 2; \alpha = \frac{2-p}{1-p} \\ e^\theta & p = 1; \alpha = \infty \\ -\log(-\theta) & p = 2; \alpha = 0. \\ \theta - \theta \log(-\theta) & p = \infty; \alpha = 1 \end{cases}$$

where the subscript indicates dependence on  $p$ . The normal,  $\alpha = 2$ ,  $p = 0$  is a special case of the first row.

The canonical parameter domain  $\Theta_p$  is the largest interval real numbers for which  $\kappa_p(\theta)$  is finite. Therefore

$$\Theta_p = \begin{cases} (-\infty, \infty) & p = 0, \alpha = 2; p = 1, \alpha = \infty \\ [0, \infty) & p < 0, 1 < \alpha < 2; 0 < p < 1, \alpha > 2 \\ (-\infty, 0) & 1 < p \leq 2, \alpha \leq 0 \\ (-\infty, 0] & 2 < p < \infty; 0 < \alpha < 1 \\ (-\infty, 0] & p = \infty; \alpha = 1. \end{cases}$$

In the last row we allow  $\theta = 0$  because  $\lim_{\theta \uparrow 0} -\theta \log(-\theta) = 0$ .

In general, a member of the NEF generated by  $V$  has cumulant generating function  $K_\theta(t) = \kappa(\theta + t) - \kappa(\theta)$ . Assuming  $p \neq 1, 2$  and substituting for  $\kappa$  produces

$$\begin{aligned} K_\theta(t) &= \frac{\alpha - 1}{\alpha} \left\{ \left( \frac{\theta + t}{\alpha - 1} \right)^\alpha - \left( \frac{\theta}{\alpha - 1} \right)^\alpha \right\} \\ &= \frac{\alpha - 1}{\alpha} \left( \frac{\theta}{\alpha - 1} \right)^\alpha \left\{ \left( 1 + \frac{t}{\theta} \right)^\alpha - 1 \right\}. \end{aligned}$$

Which distributions correspond to these cumulant generating functions? We distinguish the following situations:

1. when  $0 \notin \Theta_p$  we identify  $K_\theta$ , and
2. when  $0 \in \Theta_p$  we identify  $K_0(t) = \kappa(t)$ , with sub-cases:
  - i. the normal and Poisson, where 0 is in the interior of  $\Theta_p$  and
  - ii. the second, fourth and fifth rows where 0 is a boundary point of  $\Theta_p$ . These NEFs are not regular and  $\theta = 0$  corresponds to a distribution with infinite mean that is not part of the family<sup>1</sup>.

Figure 1 provides a handy reference of the relationship between  $p$  and  $-\alpha$  and the corresponding distributions. The analysis below is ordered by  $\alpha$  (along the  $x$ -axis), starting with  $\alpha = 2$ , the normal.

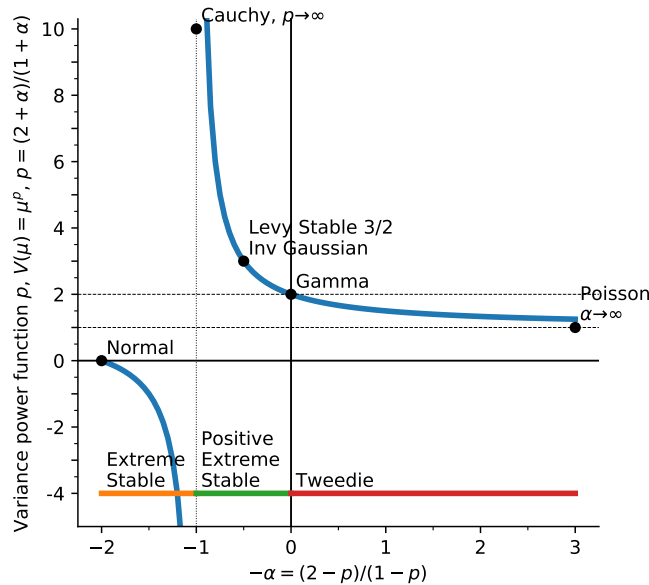


Figure 1: The power variance families:  $p$  plotted against  $-\alpha$ . Horizontal color bars indicate range of  $-\alpha$  values for extreme stable ( $-2 < -\alpha < -1$ ), positive extreme stable ( $-1 < -\alpha < 0$ ) and Tweedie ( $-\alpha > 0$ ) families. No distributions correspond to  $0 < p < 1$ ,  $-\alpha < -2$  on the left.

<sup>1</sup>The Lévy stable  $\alpha = 1/2$  distribution vs. inverse Gaussian distributions when  $p = 3$  is an example of the latter behavior.

### 1.3 Members of the PVF.

- When  $\alpha = 2$  and  $p = 0$ , then  $\theta = 0 \in \Theta_p$  and

$$K_0(t) = \kappa(t) = \frac{1}{2}t^2,$$

which is the cumulant function for the normal distribution. As expected,  $V(\mu) = \mu^0$  is constant. This NEF is regular because 0 is an interior point of  $\Theta_p$ .

- When  $1 < \alpha < 2$  and  $p < 0$ , then  $\theta = 0 \in \Theta_p$  and

$$K_0(t) = \kappa(t) = \frac{\alpha - 1}{\alpha} \left( \frac{t}{\alpha - 1} \right)^\alpha,$$

which is an  $\alpha$ -stable distribution with Lévy density  $|x|^{-\alpha-1}$  on  $x < 0$ . It falls into the compensated IACP, case 3 group, discussed in [Part III](#), and takes any real value, positive or negative, despite only having negative jumps. It has a thick left tail and thin right tail. Its mean is zero, but the variance does not exist. A tilt with  $e^\theta$ ,  $\theta > 0$  makes the left tail thinner, the right tail thicker, and increases the mean. The effect on the right tail is manageable because it is thinner than a normal, [\[1\]](#), [\[2\]](#). As  $\theta \downarrow 0$  the mean decreases to zero and the variance increases, finally becoming infinite when  $\theta = 0$ .

- When  $\alpha = 1$  and  $p = \infty$ , then  $\theta = 0 \in \Theta_p$ . Here we interpret  $V$  to be the exponential variance function and have seen

$$K_0(t) = \kappa(\theta) = \theta(1 - \log(-\theta)),$$

corresponding to an extreme stable distribution with  $\alpha = 1$ , [\[3\]](#) and [\[4\]](#). The NEF exhibits a kind of super-contagion between risks, resulting in an exponential increase in variance with the mean. The distribution is still case 3 and is supported on the whole real line.

- When  $0 < \alpha < 1$  and  $p > 2$ , then  $0 \in \Omega_p$  and

$$K_0(t) = \kappa(t) = -\frac{1 - \alpha}{\alpha} \left( \frac{-t}{1 - \alpha} \right)^\alpha,$$

which is an extreme stable with Lévy distribution  $x^{-\alpha}$ ,  $x > 0$ . The distribution is very thick tailed and does not have a mean ( $\kappa(\theta)$  is not differentiable at 0 because it has a cusp). In the case  $\alpha = 1/2$  we get the Lévy stable distribution and the inverse Gaussian NEF.

- When  $\alpha = 0$  and  $p = 2$ , then  $0 \notin \Omega_p$  and we analyze

$$\begin{aligned} K_\theta(t) &= \kappa(t + \theta) - \kappa(\theta) \\ &= -\log(-(t + \theta)) + \log(-\theta) \\ &= \log\left(1 + \frac{t}{\theta}\right) \\ &= \log\left(1 - \frac{t}{-\theta}\right) \end{aligned}$$

which is a gamma distribution with rate  $-\theta$  and shape 1.

- When  $\alpha < 0$  and  $1 < p < 2$ , then  $0 \notin \Omega_p$  and we analyze

$$K_\theta(t) = \frac{\alpha - 1}{\alpha} \left( \frac{\theta}{\alpha - 1} \right)^\alpha \left\{ \left( 1 + \frac{t}{\theta} \right)^\alpha - 1 \right\}$$

for  $\theta < 0$ . The value  $\theta/(\alpha - 1)$  is positive and  $(1 - t/(-\theta))^\alpha$  is the MGF of a gamma with shape  $-\alpha$  and rate  $-\theta$ . Thus  $K_\theta(t)$  is the cumulant generating function for a CP with expected frequency  $\frac{\alpha-1}{\alpha} \left( \frac{\theta}{\alpha-1} \right)^\alpha > 0$  and gamma severity. Distributions with  $1 < p < 2$  are called **Tweedie** distributions. Finally, we have demonstrated why Tweedie distributions are a compound Poisson with gamma severity!

- When  $\alpha = \infty$  and  $p = 1$ , then 0 is an interior point of  $\Omega_p$  and the NEF is again regular, like the normal. Now

$$K_0(t) = \kappa(t) - \kappa(0) = e^t - 1$$

is the cumulant function for the Poisson distribution.

- Finally, when  $\alpha > 2$  and  $0 < p < 1$ , there is no corresponding NEF. Arguing by contradiction, suppose there is a NEF with  $V(\mu) = \mu^p$ ,  $0 < p < 1$ . Since  $\theta = 0 \in \Theta_p$  the NEF has a generating distribution with variance

$$\kappa''(0) = \left( \frac{\theta}{\alpha - 1} \right)^{\alpha-2} \Big|_{\theta=0} = 0.$$

A distribution with zero variance is degenerate and therefore has a cumulant generating function  $e^{ct}$  for some constant  $c$ . But this contradicts the known form of  $\kappa$ . Therefore, no such NEF can exist. This result was first proved in [5]. Considering a possible NEF with variance function in this range reveals it would be very odd: its variance explodes, relative to the mean, as the mean increases from zero because  $V$  is not differentiable at zero. It is hard to conceive how such a thing would occur, and Jorgensen's result confirms that it cannot.

We have now covered all the possible values for  $p$  and  $\alpha$  and identified all the distributions in the power variance family. The findings are summarized in Table 2. The Support column shows the possible values taken by family members and Means shows the range of possible means.

Table 2: The Tweedie-PVF distributions. For  $\alpha < 2$  they are generated by the jump size density proportional to  $x^{-\alpha-1}e^{\theta x}/x$ . A NEF is regular if the canonical parameter domain  $\Theta$  is open and is steep if  $\Omega = \text{int conv}(S)$ . The extreme Cauchy has  $V(\mu) = e^\mu$ . The Levy stable is a limit case of inverse Gaussian as  $\theta \rightarrow 0$ .

Name	$\alpha$	$p$	Support $S$	Mean $\Omega$	$\Theta_p$	Regu- lar	Steep
Normal	-2	0	$(-\infty, \infty)$	$(-\infty, \infty)$	$(-\infty, \infty)$	Yes	Yes
Extreme stable	$(-2, -1)$	$(-\infty, 0)$	$(-\infty, \infty)$	$(0, \infty)$	$[0, \infty)$	No	No
Extreme Cauchy	-1	$\infty$	$(-\infty, \infty)$	$(-\infty, \infty)$	$(-\infty, 0)$	Yes	Yes
Positive ext. stb.	$(-1, -1/2)$	$(3, \infty)$	$(0, \infty)$	$(0, \infty)$	$(-\infty, 0]$	No	Yes
Inverse Gaussian	-1/2	3	$(0, \infty)$	$(0, \infty)$	$(-\infty, 0]$	No	Yes
Positive ext. stb.	$(-1/2, 0)$	$(2, 3)$	$(0, \infty)$	$(0, \infty)$	$(-\infty, 0]$	No	Yes
Gamma	0	2	$(0, \infty)$	$(0, \infty)$	$(-\infty, 0)$	Yes	Yes

Name	$\alpha$	$p$	Support $S$	Mean $\Omega$	$\Theta_p$	Regu- lar	Steep
Tweedie	$(-\infty, 0)$	$(1, 2)$	$[0, \infty)$	$(0, \infty)$	$(-\infty, 0)$	Yes	Yes
Poisson	$\infty$	1	$\{0, 1, 2, \dots\}$	$(0, \infty)$	$(-\infty, \infty)$	Yes	Yes

We can identify the Lévy measure underlying power variance NEFs as a corollary of this analysis, [6], [7]. In **Part II** we showed that under exponential tilting a Lévy density  $j$  transforms into  $e^{\theta x} j(x)$ . From the classification given in the first section, we see that all PVF exponential distributions except the normal have a jump density proportional to

$$j(x) = x^{-\alpha-1} e^{\theta x}$$

for  $\alpha < 2$  and varying  $\theta$ . In a sense, this distribution can be regarded as a universal severity distribution. When  $\alpha \leq 0$ , the Tweedie range, then  $\theta < 0$  to ensure  $j$  is finite. Then  $\alpha > 0$ ,  $\theta = 0$  defines a legitimate, stable distribution. The common jump density binds all PVF distributions together even though their individual densities are quite different.

Figure 2 shows the completed NEF Circle for the PVF for  $p \notin \{0, 1, 2\}$ . The deviance, log likelihood, and cumulant generator are also correct for  $p = 0$ ,  $\alpha = 2$ , but the density obviously becomes the normal distribution density. The Lévy measure is added at the top to unify the diagram. Series expansions are available for  $c$  and the density, see [3] Section 4.2 or [8] Section 14.

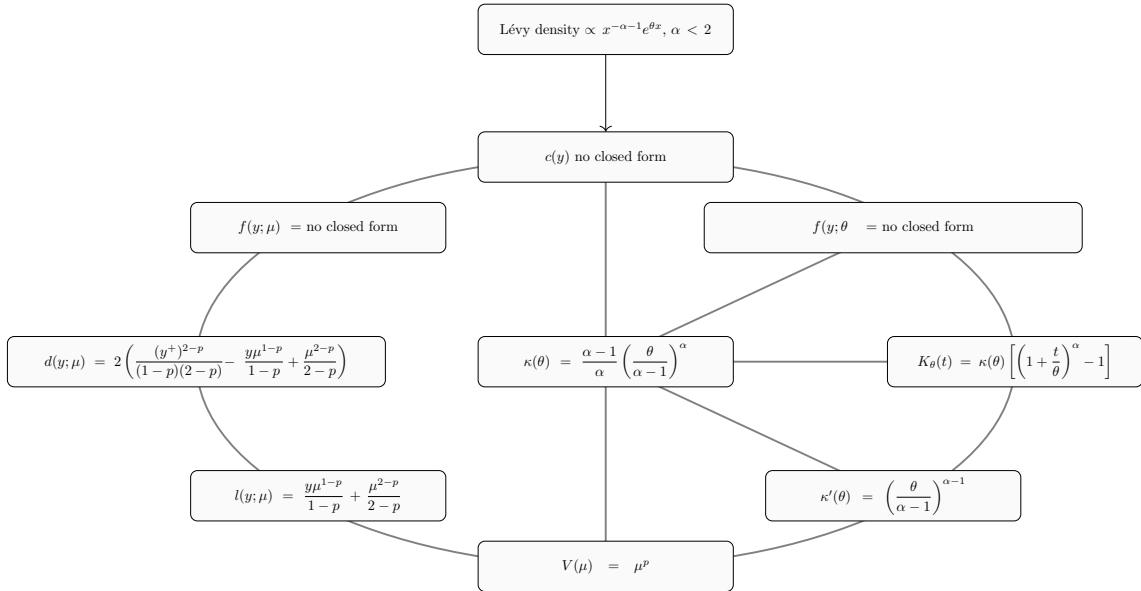


Figure 2: The NEF circle for the PVF families,  $p \notin \{1, 2\}$ ,  $\alpha = (p - 2)/(p - 1)$ ,  $(\alpha - 1)(1 - p) = 1$ . Taking  $\alpha = 2$ ,  $p = 0$  produces the normal distribution.

## 1.4 The Tweedie Compound Poisson Family

The previous section identified the Tweedie family within the PVF as compound Poisson distributions with gamma severity. Tweedie distributions are ideally suited to modeling

losses or pure premiums, where the outcome of no losses has a positive probability, because they are mixed with a mass at 0. Their mean to variance relationship,  $V(\mu) \propto \mu^p$  for  $1 < p < 2$ , spans from a diversifying growth model at  $p = 1$  to the non-diversifying model at  $p = 2$ , as we saw in [Part III](#). Thus the Tweedie family interpolates across the range of a priori reasonable outcomes by varying a family of gamma severity CPs from a Poisson with degenerate severity and  $V(\mu) \propto \mu$ , to a gamma with jump distribution  $\alpha e^{-\beta x}/x$  and  $V(\mu) \propto \mu^2$ . As the parameter  $p$  increases from 1 to 2, the Tweedie moves continuously from modeling diversifiable growth to non-diversifiable. To understand how it achieves this interpolation we introduce a specific parameters, based on parameters for the underlying frequency and severity. We then try to guess how the Tweedie is parameterized based solely on these facts. Obviously, we are cheating because we know the answer, but it instructive to see how close we can get.

The Poisson frequency distribution is easy to understand. It has a single parameter  $\lambda$  equal to the mean, and the variance is also  $\lambda$ . The CV is  $1/\sqrt{\lambda}$ . As  $\lambda$  increases the Poisson becomes relatively more tightly distributed about its mean.

The gamma severity distribution is trickier because it has two common parameterizations. We will use shape and rate parameters,  $\alpha > 0$  and  $\beta > 0$ , with associated gamma density

$$f(x; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \frac{x^\alpha e^{-\beta x}}{x}.$$

The factor  $x$  is written in the denominator to emphasize it is more natural to use  $dx/x$  as a base measure of length on the positive reals than  $dx$ .  $dx/x$  is scale invariant: if  $y = kx$  then  $dy/y = kdx/kx = dx/x$ . The shape parameter is  $\alpha$ . The mean increases with  $\alpha$ —just think about the shape of the density. For  $\alpha > 1$  the gamma distribution has a strictly positive mode at  $(\alpha - 1)/\beta$ .

$\beta$  is a rate parameter. For integer  $\alpha$ , the gamma distribution is the waiting time for  $\alpha$  events when the number of events per unit time is distributed  $\text{Poisson}(\beta)$ , and hence the waiting time between events is exponential with mean  $1/\beta$ . This is reciprocity between the gamma and Poisson, discussed in [Part III](#). When  $\alpha = 1$  the gamma is an exponential, and when  $\alpha = n/2, \beta = 1/2$  it is  $\chi^2$  with  $n$  degrees of freedom. The higher  $\beta$ , the lower the mean: the more frequently trains run (higher departure rate) the less time you wait for one to arrive.

With the shape-rate parameterization, the mean and variance of a  $\text{gamma}(\alpha, \beta)$  are  $\alpha/\beta$  and  $\alpha/\beta^2$  respectively, the CV is  $1/\sqrt{\alpha}$ , and the expected squared value is  $\alpha(\alpha + 1)/\beta^2$ . The skewness is  $2/\sqrt{\alpha}$  showing that the density becomes more symmetric as  $\alpha$  increases.

We can parameterize a Tweedie family distribution as  $Z = \text{CP}(\lambda, \text{gamma}(\alpha, \beta))$ .  $Z$  has mean  $\lambda\alpha/\beta$  and variance  $\lambda\alpha(\alpha + 1)/\beta^2$ . It is more useful if the mean appears as a parameter. To that end, consider the alternative parameters  $p, \mu$  and  $\sigma^2$ , related by  $E[Z] = \mu$  and  $\text{Var}(Z) = \sigma^2\mu^p$ , and  $p$  to be determined.

Let's see if we can solve for the parameters of  $\text{CP}(\lambda, \text{gamma}(\alpha, \beta))$  in terms of  $p, \mu$  and  $\sigma^2$ . We know  $\mu = \lambda\alpha/\beta$  and  $\sigma^2\mu^p = \lambda\alpha(\alpha + 1)/\beta^2$ , but we still need another equation for a unique solution.

Distributions within a NEF are specified by the underlying variance function  $V$ , location  $\mu$ , and dispersion  $\sigma^2$  acting so the variance is  $\sigma^2V(\mu)$ . Dispersion operates analogously to

scale in a normal-distribution world. In  $\text{CP}(\lambda, \text{gamma}(\alpha, \beta))$ ,  $\lambda$  determines location and  $\beta$  (strictly  $1/\beta$ ) determines dispersion, whereas  $\alpha$  will change shape in a more profound way. This suggests that  $p$  should be related to  $\alpha$ , independent of the other parameters. Assuming  $\alpha = \alpha(p)$ , an equation to be specified later, and substituting for the mean and variance gives

$$\lambda = \left( \frac{\alpha + 1}{\alpha \sigma^2 \mu^p} \right) \mu^2$$

and

$$\beta = \left( \frac{\alpha + 1}{\sigma^2 \mu^p} \right) \mu.$$

It is a good exercise to check these formulas give the desired mean  $\mu$  and variance  $\sigma^2 \mu^p$ . The terms involving  $\alpha$  cancel out.

To understand the missing function  $\alpha(p)$  we can look what Table 2 implies about its behavior for  $p$  close to 1 and 2 for a fixed  $\mu$ . Assume  $\sigma^2 = 1$  to simplify the algebra.

- As  $p$  decreases to 1 the Tweedie approaches a Poisson with mean  $\mu$ . At the same time,  $\lambda = \mu^{2-p}(\alpha+1)/\alpha \rightarrow \mu(\alpha+1)/\alpha$  and therefore  $(\alpha+1)/\alpha \rightarrow 1$ , which in turn implies  $\alpha \rightarrow \infty$ . Therefore the mean of the gamma severity  $\alpha/\beta = \alpha \mu^{p-1}/(\alpha+1) \rightarrow 1$  also. The CV of the gamma severity  $1/\sqrt{\alpha} \rightarrow 0$ , and so severity becomes more and more tightly clustered around  $x = 1$ .
- As  $p$  increases to 2 the Tweedie approaches a gamma distribution. Now  $\lambda \rightarrow (\alpha+1)/\alpha$  and  $\beta \rightarrow (\alpha+1)/\mu$ . A gamma distribution is continuous, supported on  $x > 0$ . There is no possibility it takes the value zero. Since  $\Pr(\text{CP}(\lambda) = 0) = e^{-\lambda}$ , corresponding to the Poisson claim count  $N = 0$ , it follows that  $\lambda \rightarrow \infty$ . Therefore  $\alpha \rightarrow 0$  as  $p \rightarrow 2$ .

These two extremes mirror the models of diversifying and non-diversifying growth presented in Part III. When  $p = 1$ ,  $\lambda$  is proportional to  $\mu$  and severity is fixed  $X = 1$ . When  $p = 2$ ,  $\lambda$  is independent of  $\mu$  and mean severity  $\alpha/\beta = \mu\alpha/(\alpha+1)$  increases with  $\mu$ .

Combining the two conditions shows that the unknown function  $\alpha(p)$  satisfies  $\alpha(1) = \infty$  and  $\alpha(2) = 0$ . The simplest positive function with these properties is

$$\bar{\alpha}(p) = \frac{2-p}{p-1},$$

which turns out to be the correct relationship, not only for  $1 \leq p \leq 2$  but for all  $p$ . (Note that  $\bar{\alpha} = -\alpha$  defined in Section 1.2: when discussing stable distributions  $\alpha > 0$  and the jump distribution is  $x^{-\alpha}$ . For the gamma distribution the power of  $x$  appears with a positive parameter. The bar distinguishes the two.) Since  $\bar{\alpha} + 1 = 1/(p-1)$  and  $(\bar{\alpha} + 1)/\bar{\alpha} = 1/(2-p)$ , we can re-write

$$\lambda = \frac{\mu^2}{(2-p)\sigma^2 \mu^p}$$

and

$$\beta = \frac{\mu}{(p-1)\sigma^2 \mu^p}.$$



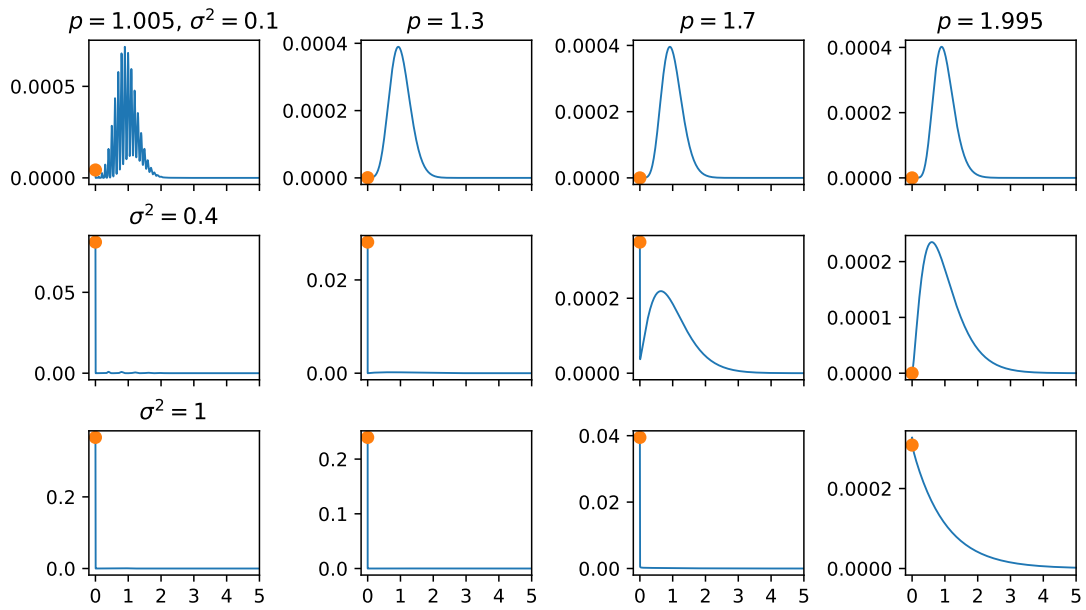


Figure 3: Tweedie family densities with  $\mu = 1$  for various values of  $p$  and  $\sigma^2$ .




Figure 4: Tweedie family log densities with  $\mu = 1$  for various values of  $p$  and  $\sigma^2$ .

## 1.5 Examples of Tweedie Density Functions

The Tweedie distribution is mixed: it has a probability mass at  $x = 0$  and is continuous for  $x > 0$ . Its density can assume a many different shapes, as we now illustrate.

Figure 3 and Figure 4 show sample Tweedie densities for various parameter values on a linear and log scale. The dot shows the probability mass at 0. The distributions transition from near-Poisson,  $p = 1.005$ , on the left to near-gamma,  $p = 1.995$ , on the right. Vertically the dispersion increases from 0.1 to 0.4 to 1.0. The left hand column illustrates how the Tweedie approaches a Poisson. The lower right most plot is close to an exponential distribution.

Table 3: The  $p$ ,  $\mu$  and  $\sigma^2$  Tweedie distribution parameters, with corresponding Poisson  $\lambda$ , gamma  $\alpha, \beta$ , and resulting aggregate variance and CV, probability  $Z = 0$ , expected and CV severity statistics.

$\mu$	$p$	$\sigma^2$	$\lambda$	$\alpha$	$\beta$	Var( $Z$ )	CV( $Z$ )	Pr( $Z = 0$ )	$\alpha/\beta$	CV( $X$ )
1	1.005	0.1	10.0503	199	2000	0.1	0.3162	4.3174e-05	0.0995	0.0708
1	1.005	0.4	2.5125	199	500	0.4	0.6324	0.0810	0.398	0.0708
1	1.005	1	1.0050	199	200	1	1	0.3660	0.995	0.0708
1	1.3	0.1	14.2857	2.3333	33.3333	0.1	0.3162	6.2487e-07	0.07	0.6546
1	1.3	0.4	3.5714	2.3333	8.3333	0.4	0.6324	0.0281	0.28	0.6546
1	1.3	1	1.4285	2.3333	3.3333	1	1	0.2396	0.7	0.6546
1	1.7	0.1	33.3333	0.4285	14.2857	0.1	0.3162	3.3382e-15	0.03	1.5275
1	1.7	0.4	8.3333	0.4285	3.5714	0.4	0.6324	0.0002	0.12	1.5275
1	1.7	1	3.3333	0.4285	1.4285	1	1	0.0356	0.3	1.5275
1	1.995	0.1	2000	0.0050	10.0503	0.1	0.3162	<i>very small</i>	0.0005	14.1067
1	1.995	0.4	500	0.0050	2.5125	0.4	0.6324	7.1245e-218	0.002	14.1067
1	1.995	1	200	0.0050	1.0050	1	1	1.3839e-87	0.005	14.1067

Table 3 shows the underlying parameters, broken into groups by  $p$ . The values of  $\alpha$  and hence CV( $X$ ) only depend on  $p$ . As  $p$  increases to 2,  $\lambda$  increases and the severity becomes thicker tailed and more skewed. The value  $\alpha/\beta$  is the expected severity. Since the overall mean is always 1.0 the expected frequency  $\lambda = \beta/\alpha$ .

## 1.6 Appendix: R and Python Code

There is no closed-form expression for the density or distribution function of a general Tweedie family variable. The R package `tweedie` will compute Tweedie densities. The code below can be used to reproduce Figure 3 and Figure 4} by inputting the appropriate parameters. It also shows how to use the `tweedie.convert` function to reproduce Table 3. R outputs the mean rather than the rate for the second gamma parameter; it is the reciprocal of what we show.

```

# R code for p=1.005 and sigma^2=0.1 graph

library(tweedie)
library(tidyverse)

# Tweedie, phi=sigma^2
mu <- 1
p <- 1.005
phi <- 0.1
twdf = tibble(x = seq(0, 6, length=500),
              fx = dtweedie(y=x, power=p, mu=mu, phi=phi))

ggplot(data=twdf) + geom_line(aes(x=x, y=fx))

# parameters to reproduce Table 7
tbl_tweedies <- crossing(
  mu = 1
  , p = c(1.005, 1.3, 1.7, 1.995)
  , phi = c(0.1, 0.4, 1)
) %>%
  mutate(
    lst_conversions = map2(p, phi, tweedie.convert, mu = 1)
    , lambda = map_dbl(lst_conversions, 'poisson.lambda')
    , alpha = map_dbl(lst_conversions, 'gamma.shape')
    , beta = 1 / map_dbl(lst_conversions, 'gamma.scale')
  ) %>%
  select(-lst_conversions)

tbl_tweedies

```

Figure 3 and Figure 4 were produced using Python, computing the Tweedie density using Fast Fourier Transforms, as shown in the next block of code, see [9]. The second plot,  $p = 1.0005$ , shows just how bizarre the Tweedie density can become. It is not reproduced here.

```

import numpy as np
import pandas as pd

def tweedie(mu, p, sigma2, log2, bs):
    """
    FFT approximation to the Tweedie distribution
    mu = mean of Tweedie with variance function sigma2 mu^p

    log2, bs control FFT discretization
    """

    alpha = (2 - p) / (p - 1)
    lam = mu**(2 - p) / ((2 - p) * sigma2)
    beta = mu**(1 - p) / ((p - 1) * sigma2)

    sev = ss.gamma(alpha, scale=1 / beta)

    ft = np.fft.rfft
    ifft = np.fft.irfft

    # discretize severity
    xs = bs * np.arange(2**log2)

```

```

ps = -np.diff(sev.sf(xs + bs/2), prepend=1)

# convolve with MGF - see Part II
tw = ifft( np.exp( lam * (ft(ps) - 1) ) )

return pd.DataFrame({'tweedie': tw}, index=xs)

# p=1.005, sigma^2=0.1, m
df = tweedie(1., 1.005, 0.1, 16, 1/1024)
df.plot(xlim=[0,5], figsize=(10,8))

# a more spectacular plot (not shown)
df = tweedie(10, 1.0005, 0.1, 16, 1/1024)
df.plot(xlim=[5,15], figsize=(10,8))

```

## References

1. Zolotarev, V.M.: [One-Dimensional Stable Distributions](#), *Translations of Mathematical Monographs*, vol. 65. American Mathematical Society (1986)
2. Carr, P., Wu, L.: The Finite Moment Log Stable Process and Option Pricing. *Journal of Finance*. 58, 753–778 (2003). <https://doi.org/10.1111/1540-6261.00544>
3. Jørgensen, B.: *The theory of dispersion models*. CRC Press (1997)
4. Eaton, M.L., Morris, C., Rubin, H.: On extreme stable laws and some applications. *Journal of Applied Probability*. 8, 794–801 (1971). <https://doi.org/10.2307/3212242>
5. Jørgensen, B.: Exponential Dispersion Models. *Journal of the Royal Statistical Society Series B*. 49, 127–162 (1987)
6. Vinogradov, V.: Properties of certain Lévy and geometric Lévy processes. *Communications on Stochastic Analysis*. 2, (2004). <https://doi.org/10.31390/cosa.2.2.02>
7. Jørgensen, B., Martinez: The Levy-Khinchine Representation of the Tweedie Class. *Brazilian Journal of Probability and Statistics*. 10, 225—233 (1996)
8. Sato, K.-I.: [Levy Processes and Infinitely Divisible Distributions](#). Cambridge University Press, Cambridge (1999)
9. Mildenhall, S.J.: [Correlation and Aggregate Loss Distributions With An Emphasis On The Iman-Conover Method](#). *Casualty Actuarial Society Forum*. Winter, (2005)