5.13: The Folded Normal Distribution

The General Folded Normal Distribution

Introduction

The *folded normal distribution* is the distribution of the absolute value of a random variable with a normal distribution. As has been emphasized before, the normal distribution is perhaps the most important in probability and is used to model an incredible variety of random phenomena. Since one may only be interested in the magnitude of a normally distributed variable, the folded normal arises in a very natural way. The name stems from the fact that the probability measure of the normal distribution on $(-\infty, 0]$ is “folded over” to $[0, \infty)$. Here is the formal definition:

Suppose that Y has a normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$. Then $X = |Y|$ has the *folded normal distribution* with parameters μ and σ.

So in particular, the folded normal distribution is a continuous distribution on $[0, \infty)$.

Distribution Functions

Suppose that Z has the standard normal distribution. Recall that Z has probability density function ϕ and distribution function Φ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}, \quad z \in \mathbb{R} \tag{5.13.1}$$

$$\Phi(z) = \int_{-\infty}^z \phi(x)\,dx = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} e^{-x^2/2}\,dx, \quad z \in \mathbb{R} \tag{5.13.2}$$

The standard normal distribution is so important that Φ is considered a special function and can be computed using most mathematical and statistical software. If $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$, then $Y = \mu + \sigma Z$ has the normal distribution with mean μ and standard deviation σ, and therefore $X = |Y| = |\mu + \sigma Z|$ has the folded normal distribution with parameters μ and σ. For the remainder of this discussion we assume that X has this folded normal distribution.

X has distribution function F given by

$$F(x) = \Phi\left(\frac{x-\mu}{\sigma}\right) - \Phi\left(\frac{-x-\mu}{\sigma}\right) = \Phi\left(\frac{x-\mu}{\sigma}\right) + \Phi\left(\frac{x+\mu}{\sigma}\right) - 1 \tag{5.13.3}$$

$$= \int_0^x \frac{1}{\sigma\sqrt{2\pi}} \left\{\exp\left[-\frac{1}{2}\left(\frac{y+\mu}{\sigma}\right)^2\right] + \exp\left[-\frac{1}{2}\left(\frac{y-\mu}{\sigma}\right)^2\right]\right\} dy, \quad x \in [0, \infty) \tag{5.13.4}$$

Proof

For $x \in [0, \infty)$,

$$F(x) = \mathbb{P}(X \le x) = \mathbb{P}(|Y| \le x) = \mathbb{P}(|\mu + \sigma Z| \le x) = \mathbb{P}(-x \le \mu + \sigma Z \le x) \tag{5.13.5}$$

$$= \mathbb{P}\left(\frac{-x-\mu}{\sigma} \le Z \le \frac{x-\mu}{\sigma}\right) = \Phi\left(\frac{x-\mu}{\sigma}\right) - \Phi\left(\frac{-x-\mu}{\sigma}\right) \tag{5.13.6}$$

which gives the first expression. The second expression follows since $\Phi(-z) = 1 - \Phi(z)$ for $z \in \mathbb{R}$. Finally, the integral formula follows from the form of Φ given above and simple substitution.

We cannot compute the quantile function F^{-1} in closed form, but values of this function can be approximated.

Open the special distribution calculator and select the folded normal distribution, and set the view to CDF. Vary the parameters and note the shape of the distribution function. For selected values of the parameters, compute the median and the first and third quartiles.

X has probability density function f given by

$$f(x) = \frac{1}{\sigma}\left[\phi\left(\frac{x-\mu}{\sigma}\right) + \phi\left(\frac{x+\mu}{\sigma}\right)\right] \tag{5.13.7}$$
$$= \frac{1}{\sigma\sqrt{2\pi}}\left\{\exp\left[-\frac{1}{2}\left(\frac{x+\mu}{\sigma}\right)^2\right] + \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right]\right\}, \quad x \in [0,\infty) \tag{5.13.8}$$

Proof

This follows from differentiating the CDF with respect to x, since $F'(x) = f(x)$ and $\Phi'(z) = \phi(z)$.

Open the special distribution simulator and select the folded normal distribution. Vary the parameters μ and σ and note the shape of the probability density function. For selected values of the parameters, run the simulation 1000 times and compae the empirical density function to the true probability density function.

Note that the folded normal distribution is unimodal for some values of the parameters and decreasing for other values. Note also that μ is not a location parameter nor is σ a scale parameter; both influence the shape of the probability density function.

Moments

We cannot compute the the mean of the folded normal distribution in closed form, but the mean can at least be given in terms of Φ. Once again, we assume that X has the folded normal distribution with parmaeters $\mu \in \mathbb{R}$ and $\sigma \in (0,\infty)$.

The first two moments of x are

1. $\mathbb{E}(X) = \mu[1 - 2\Phi(-\mu/\sigma)] + \sigma\sqrt{2/\pi}\exp(-\mu^2/2\sigma^2)$
2. $\mathbb{E}(X^2) = \mu^2 + \sigma^2$

Proof

From the definition, we can assume $X = |\mu + \sigma Z|$ where Z has the standard normal distribution. Then

$$E(X) = \mathbb{E}(|\mu + \sigma Z|) = \mathbb{E}(\mu + \sigma Z; Z \ge -\mu/\sigma) - \mathbb{E}(\mu + \sigma Z; Z \le -\mu/\sigma) \tag{5.13.9}$$
$$= \mathbb{E}(\mu + \sigma Z) - 2\mathbb{E}(\mu + \sigma Z : Z \le -\mu/\sigma) = \mu - 2\mu\Phi(-\mu/\sigma) - 2\sigma\mathbb{E}(Z; Z \le -\mu/\sigma) \tag{5.13.10}$$

So we just need to compute the last expected value. Using the change of variables $u = z^2/2$ we get

$$\mathbb{E}(Z; Z \le -\mu/\sigma) = \int_{-\infty}^{-\mu/\sigma} z\frac{1}{\sqrt{2\pi}}e^{-z^2/2}\,dz = -\int_{\mu^2/2\sigma^2}^{\infty}\frac{1}{\sqrt{2\pi}}e^{-u}\,du = -\frac{1}{\sqrt{2\pi}}e^{-\mu^2/2\sigma^2} \tag{5.13.11}$$

Substituting gives the result in (a). For (b), let Y have the normal distribution with mean μ and standard deviation σ so that we can take $X = |Y|$. Then $\mathbb{E}(X^2) = \mathbb{E}(Y^2) = \text{var}(Y) + [\mathbb{E}(Y)]^2 = \sigma^2 + \mu^2$.

In particular, the variance of X is

$$\text{var}(X) = \mu^2 + \sigma^2 - \left\{\mu\left[1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right] + \sigma\sqrt{\frac{2}{\pi}}\exp\left(-\frac{\mu^2}{2\sigma^2}\right)\right\}^2 \tag{5.13.12}$$

Open the special distribution simulator and select the folded normal distribution. Vary the parameters and note the size and location of the mean±standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

Related Distributions

The most important relation is the one between the folded normal distribution and the normal distribution in the definition: If Y has a normal distribution then $X = |Y|$ has a folded normal distribution. The folded normal distribution is also related to itself through a symmetry property that is perhaps not completely obvious from the initial definition:

For $\mu \in \mathbb{R}$ and $\sigma \in (0,\infty)$, the folded normal distribution with parameters $-\mu$ and σ is the same as the folded normal distribution with parameters μ and σ.

Proof 1

The PDF is unchanged if μ is replaced with $-\mu$.

Proof 2

Suppose that Y has the normal distribution with mean μ and standard deviation σ so that $|Y|$ has the folded normal distribution with parameters μ and σ. Then $-Y$ has the normal distribution with mean $-\mu$ and standard deviation σ so that $|-Y|$ has the folded normal distribution with parameters $-\mu$ and σ. But $|Y| = |-Y|$.

The folded normal distribution is also closed under scale transformations.

Suppose that X has the folded normal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ and that $b \in (0, \infty)$. Then bX has the folded normal distribution with parameters $b\mu$ and $b\sigma$.

Proof

Once again from the definition, we can assume $X = |Y|$ where Y has the normal distribution with mean μ and standard deviation σ. But then $bX = b|Y| = |bY|$, and bY has the normal distribution with mean $b\mu$ and standard deviation $b\sigma$.

The Half-Normal Distribution

When $\mu = 0$, results for the folded normal distribution are much simpler, and fortunately this special case is the most important one. We are more likely to be interested in the magnitude of a normally distributed variable when the mean is 0, and moreover, this distribution arises in the study of Brownian motion.

Suppose that Z has the standard normal distribution and that $\sigma \in (0, \infty)$. Then $X = \sigma|Z|$ has the *half-normal distribution* with *scale parameter* σ. If $\sigma = 1$ so that $X = |Z|$, then X has the *standard half-normal distribution*.

Distribution Functions

For our next discussion, suppose that X has the half-normal distribution with parameter $\sigma \in (0, \infty)$. Once again, Φ and Φ^{-1} denote the distribution function and quantile function, respectively, of the standard normal distribution.

The distribution function F and quantile function F^{-1} of X are

$$F(x) = 2\Phi\left(\frac{x}{\sigma}\right) - 1 = \int_0^x \frac{1}{\sigma}\sqrt{\frac{2}{\pi}} \exp\left(-\frac{y^2}{2\sigma^2}\right) dy, \quad x \in [0, \infty) \tag{5.13.13}$$

$$F^{-1}(p) = \sigma \Phi^{-1}\left(\frac{1+p}{2}\right), \quad p \in [0, 1) \tag{5.13.14}$$

Proof

The result for the CDF follows from the CDF of the folded normal distribution with $\mu = 0$. Recall that $\Phi(-z) = 1 - \Phi(z)$ for $z \in \mathbb{R}$. The result for the quantile function follows from the result for the CDF and simple algebra.

Open the special distribution calculator and select the folded normal distribution. Select CDF view and keep $\mu = 0$. Vary σ and note the shape of the CDF. For various values of σ, compute the median and the first and third quartiles.

The probability density function f of X is given by

$$f(x) = \frac{2}{\sigma}\phi\left(\frac{x}{\sigma}\right) = \frac{1}{\sigma}\sqrt{\frac{2}{\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right), \quad x \in [0, \infty) \tag{5.13.15}$$

1. f is decreasing with mode at $x = 0$.
2. f is concave downward and then upward, with inflection point at $x = \sigma$.

Proof

The formula for f follows from differentiating the CDF above. Properties (a) and (b) follow from standard calculus.

Open the special distribution simulator and select the folded normal distribution. Keep $\mu = 0$ and vary σ, and note the shape of the probability density function. For selected values of σ, run the simulation 1000 times and compare the emprical density function to the true probability density function.

Moments

The moments of the half-normal distribution can be computed explicitly. Once again we assume that X has the half-normal distribution with parameter $\sigma \in (0, \infty)$.

For $n \in \mathbb{N}$

$$\mathbb{E}(X^{2n}) = \sigma^{2n} \frac{(2n)!}{n! 2^n} \tag{5.13.16}$$

$$\mathbb{E}(X^{2n+1}) = \sigma^{2n+1} 2^n \sqrt{\frac{2}{\pi}} n! \tag{5.13.17}$$

Proof

As in the definition, we can take $X = \sigma |Z|$ where Z has the standard normal distribution. The even order moments of X are the same as the even order moments of σZ. These were computed in the section on the normal distribution. For the odd order moments we again use the simple substitution $z = x^2/2$ to get

$$\mathbb{E}(X^{2n+1}) = \sigma^{2n+1} \int_0^\infty x^{2n+1} \sqrt{\frac{2}{\pi}} e^{-x^2/2} \, dx = \sigma^{2n+1} 2^n \sqrt{\frac{2}{\pi}} \int_0^\infty z^n e^{-z} \, dz = \sigma^{2n+1} 2^n \sqrt{\frac{2}{\pi}} n! \tag{5.13.18}$$

In particular, we have $\mathbb{E}(X) = \sigma \sqrt{2/\pi}$ and $\text{var}(X) = \sigma^2 (1 - 2/\pi)$

Open the special distribution simulator and select the folded normal distribution. Keep $\mu = 0$ and vary σ, and note the size and location of the mean$\pm$standard deviation bar. For selected values of σ, run the simulation 1000 times and compare the mean and standard deviation to the true mean and standard deviation.

Next are the skewness and kurtosis of the half-normal distribution.

Skewness and kurtosis

1. The skewness of X is

$$\text{skew}(X) = \frac{\sqrt{2/\pi}(4/\pi - 1)}{(1 - 2/\pi)^{3/2}} \approx 0.99527 \tag{5.13.19}$$

2. The kurtosis of X is

$$\text{kurt}(X) = \frac{3 - 4/\pi - 12/\pi^2}{(1 - 2/\pi)^2} \approx 3.8692 \tag{5.13.20}$$

Proof

Skewness and kurtosis are functions of the standard score and so do not depend on the scale parameter σ. The results then follow by letting $\sigma = 1$ and using the standard computational formulas for skewness and kurtosis in terms of the moments of the half-normal distribution.

Related Distributions

Once again, the most important relation is the one in the definition: If Y has a normal distribution with mean 0 then $X = |Y|$ has a half-normal distribution. Since the half normal distribution is a scale family, it is trivially closed under scale transformations.

Suppose that X has the half-normal distribution with parameter σ and that $b \in (0, \infty)$. Then bX has the half-normal distribution with parameter $b\sigma$.

Proof

As in the definition, let $X = \sigma |Z|$ where Z is standard normal. Then $bX = b\sigma |Z|$.

The standard half-normal distribution is also a special case of the chi distribution.

The standard half-normal distribution is the chi distribution with 1 degree of freedom.

Proof

If Z is a standard normal variable, then Z^2 has the chi-square distribution with 1 degree of freedom, and hence $|Z| = \sqrt{Z^2}$ has the chi distribution with 1 degree of freedom.