Comparing the Cauchy and Gaussian (Normal) density functions

F. Masci, 6/22/2013

1. Relating the location and scale parameters

The Cauchy distribution has no finite moments, i.e., mean, variance etc, but it can be normalized and that's it. When its parameters correspond to a symmetric shape, the "sort-of-mean" is found by symmetry, and since the Cauchy has no (finite) variance, that can't be used to match to a Gaussian either. However, one can compare the Cauchy to a Gaussian such that the modes (peaks) are the same ($1/\pi$ in the example shown Figure 1).

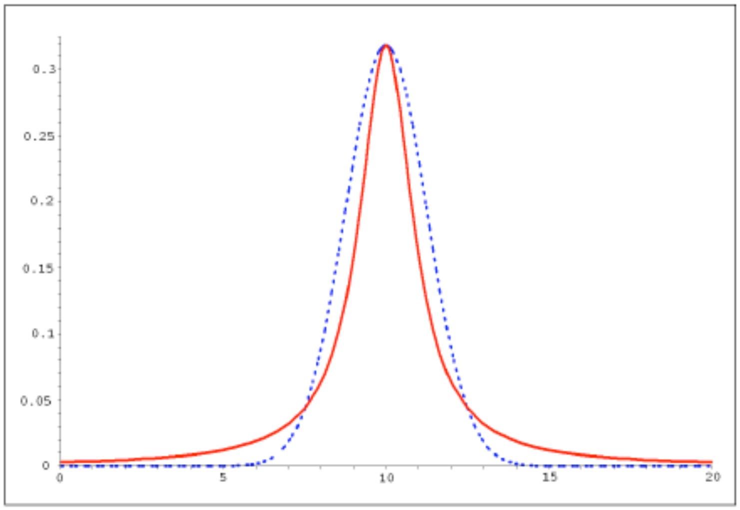

Figure 1: Solid red curve is a Cauchy density function with z_0=10 and b=1. The dashed curve is a Gaussian with the same peak as the Gaussian ($1/\pi$) with mean=10 and variance = $\pi/2$. The Cauchy has heavier tails.

The terminology uses the b and z_0 parameters to define the Cauchy density function:

$$p(z) = \frac{b/\pi}{(z - z_0)^2 + b^2}$$

Given a Cauchy (or Lorentzian) is integrable, you can define probabilities or quantile ranges that correspond to a certain probability. Hence, you can find the relationship between the Cauchy scale parameter "b" and the sigma of a Gaussian such that they contain the same mass (probability) within some quantile (or confidence) interval of interest. For a given probability p, the quantile functions for the Gaussian (z_g) and Cauchy (z_c) are:

$$z_g = \mu + \sigma\sqrt{2}\,\mathrm{erf}^{-1}(2p-1)$$
$$z_c = z_0 + b\tan\left[\pi(p-0.5)\right]$$
$$\text{where } 0 \le p \le 1$$

Therefore, if you want the Cauchy parameters that would give the same p-confidence interval as you would get from a Gaussian, simply equate the above. Not sure why anyone would want to do this in practice.

2. The sampling distribution of the mean for a Cauchy population

There's something we usually take for granted but never think about deeply – basically the distribution of the mean of a set of N independent measurements drawn from a population with finite σ will have standard-deviation "$\sigma/\sqrt{N}$". As N increases, this distribution approaches "normality" (the Central Limit Theorem).

So, regardless of the underlying population, it only needs to have a finite variance for the $\sigma/\sqrt{N}$ rule to hold. Therefore, if one is drawing samples from a Cauchy population and *naively* computes the sample mean and σ, they should never see $1/\sqrt{N}$ behavior! This is because the Cauchy distribution has no finite variance. In fact, there will be diminishing returns as N increases because more of the Cauchy tails will be sampled. The latter will inflate the sample σ more than what can be compensated by any $\sqrt{N}$ diminution.

Figure 2 compares the standard deviation of the sample mean for sample sizes $N = 1,2,3\ldots100$ computed from 10000 simulated samples for each N drawn from a Gaussian (normal), Uniform, and Cauchy population. This is a log-log plot so that "$\sigma/\sqrt{N}$" behavior will be represented by a straight line with slope = -1/2. As expected, there is no $1/\sqrt{N}$ reduction in the mean for any sample drawn from a Cauchy population.

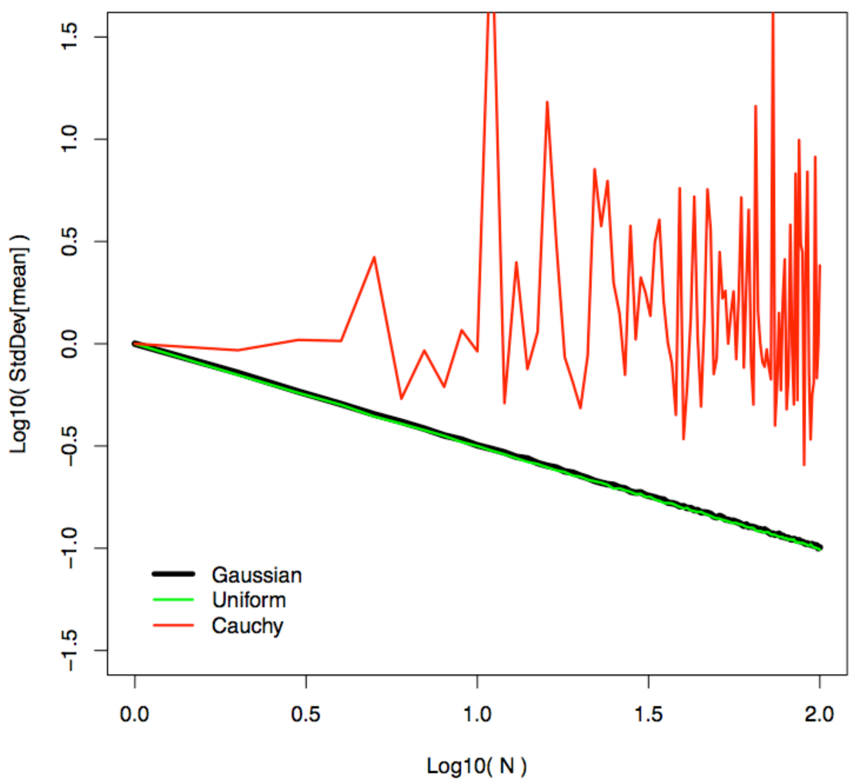

Figure 2: standard-deviation of the sample mean for sample sizes $N = 1,2,3\ldots100$ drawn from three popular distributions. All estimates are scaled to have standard-deviation = 1 at sample size $N = 1$.