

“Fat Tails”

1 Introduction

The Netherlands learned that it is crucial to have a good understanding of “exceptional” outcomes the hard way, when, in the 1950s, an unexpected weather system caused the dam system protecting a big part of the country to fail. That made “tail events” (that is, supposedly “rare events”) a big item. The latest bank failures in the late 2000s made this very real for a lot of people: the point is not that house prices started suddenly to fall precipitously, but rather that all the “smart” people in the investment banks (the poster child being Lehman Brothers) failed to prepare for such an event. The dynamics of this latter crash are more complex than this, but one contributing factor was that the risk assessment offices in the banks were relying on what is called “VAR”, that is “Value at Risk”, which is an estimate of how big a risk a bank would face with 1% probability – assuming a normal model! Several statistical studies have shown that financial assets, including securities connected to the housing market, are not well described by Gaussian models – the rule seems to be that big events (positive or negative) are much more likely than would be predicted by normal models.

Actually, we were warned: back in the early 1990s a big investment fund (the *Long-Term Capital Management* Fund), managed by no less than two Nobel Prize winners, almost caused a world-wide financial collapse when it went bankrupt after “betting” on assets that went really bad – again, relying on what turned out to be an optimistic evaluation of risk in their investment.

In other words, whenever the potential destruction from a “big” fluctuation is significant (a dam breakdown in the Netherlands, a financial collapse in a financial market, ...), it might be prudent to adopt a more pessimistic model than a “standard” Gaussian one.

2 Comparing with the Normal Distribution

Here is a quick table comparing the (approximate) values of $P[X \geq k]$, for $k = 1, 2, 3$, if X has a standard normal distribution, with the case when X has a “fat tail” distribution, still with mean 0, and variance 1 (For the curious among you, this uses the density function proportional to $\frac{1}{1+x^4}$. The numbers are rounded to the nearest one hundredth)

k	Normal	“Fat Tail”
1	0.16	0.24
2	0.02	0.04
3	0.001	0.01

Table 1.

While the difference might not seem too impressive, it grows fast, in relative terms, as k grows. For example, there is a 0.1% chance of exceeding three standard deviations in the normal case, but that increases tenfold in the example fat tail. Incidentally, our example fat tail distribution is not nearly as “fat” as it could be. The “canonical” fat tail distribution, the Cauchy distribution, goes down, as k increases, so slowly that it does not have a well defined mean or variance at all. For this distribution, the three numbers above are, roughly, respectively, $\frac{1}{\pi} \cdot \frac{\pi}{4} = 0.25, 0.15, 0.1!$

As you can see, what model we adopt impacts our assessment of “outliers” in a big way!

3 Other Consequences

One consequence of having a density that decays slowly, is that not all moments may be defined. If you have some experience in calculus, you will know that for an integral over the whole line (or half a line) to exist, the integrand must approach zero, as $|x|$ grows without bounds, at least as fast as $\frac{1}{|x|^{1+\varepsilon}}$, where $\varepsilon > 0$ is any positive number.

Thus, our density proportional to $\frac{1}{1+x^4}$, will have a mean and a variance, but not a third moment ($\frac{x^3}{1+x^4}$ will go to zero like $\frac{1}{x}$, too slowly). Since the standard proof of the Central Limit Theorem requires four moments to exist, the consequences can be serious.

The Cauchy distribution (the one with density $\frac{1}{\pi} \cdot \frac{1}{1+x^2}$) does not even have a mean, hence, if we are sampling from this distribution, the sample mean will not tend to any finite value, as the sample size increases. We will work a little on this issue, using a *simulated* sample from this distribution.

There are more complex examples, and some have been suggested as models for various important applications. including market prices, Internet traffic, and water flow in complex river systems, presumably better than “standard” models, since they allow for much higher likelihood of extreme events.

4 Exercise

Use simulated data for samples from the Cauchy distribution to experiment:

- what happens to the sample mean, as the sample size increases?
- what happens to the sample variance when the sample size increases?
- If you pretend to ignore that the underlying distribution is very far from normal, what does a “standard” interval estimate, or test, for the mean produce? How does your result change, as you increase the sample size?