Wild Moments in Statistics

Fat and long tails in portfolio management

André Koch
Stachanov Solutions & Services b.v.

http://www.stachanov.com
Business Software Development and Consulting

Wild Moments in Statistics

Nowadays, most risk managers are well aware of the danger of applying thin-tailed distributions to asset performance prone to extreme events. The fact that the widely used normal distribution tends to assign a near zero probability to events four to five standard deviations from the mean is well-known. Portfolio and value-at-risk (VaR) models no longer assume normality for the price movements of the underlying assets. The relentless battering from scholars and publicists like Mandelbrot and Taleb brought home the message that risk management is not about "normal" events but about what is happening in the distribution's tails, where extreme events reveal themselves.

Definition

There is confusion on the definition of what a long, heavy, or fat tail is. We can even wonder how to define a tail. Although this makes an interesting discussion, for the purpose of this paper we can steer away from these problems since most of us will recognise the phenomena when we see them. A parallel can be drawn with the designation of "richness". There are many ways to define this state; however, we usually are well aware of who is rich and who is not. For convenience' sake, we typically settle for a demarcation such as the top ten or twenty per cent of the ranked incomes earned in a country.

Income distribution

Ranking individual income and wealth and assigning the relative probability is a classic example of a long tailed graph: the Pareto distribution. Vilfredo Pareto (1848-1923) was an Italian sociologist and economist who studied the uneven distribution of wealth and income of Italy. He observed that 20% of the Italian citizens possessed 80% of the land. This 80%/20% principle is widely used to describe other unevenly distributed observations.

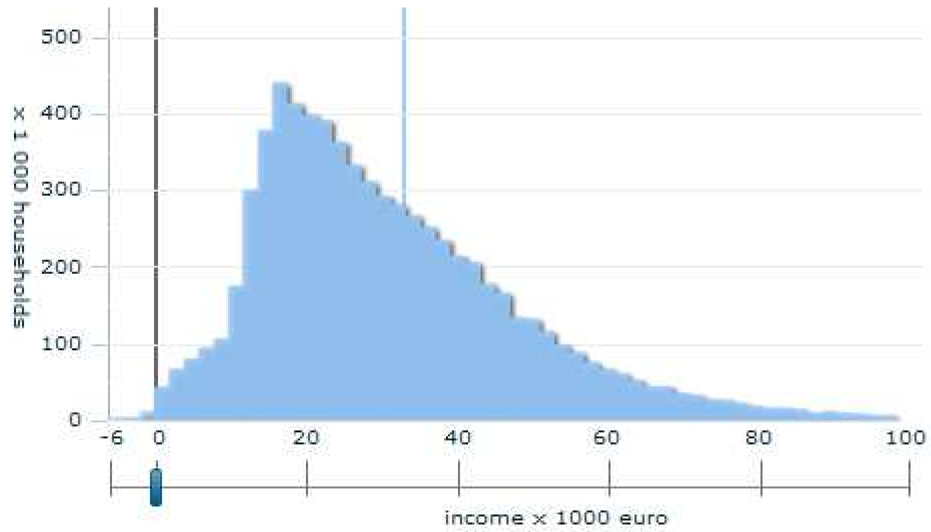

Figure 1: Distribution of spendable household income in The Netherlands in 2010; the blue line showing the average income (source: CBS)

If we look at spendable income per household in The Netherlands we can model this distribution with a Pareto probability density function (PDF).

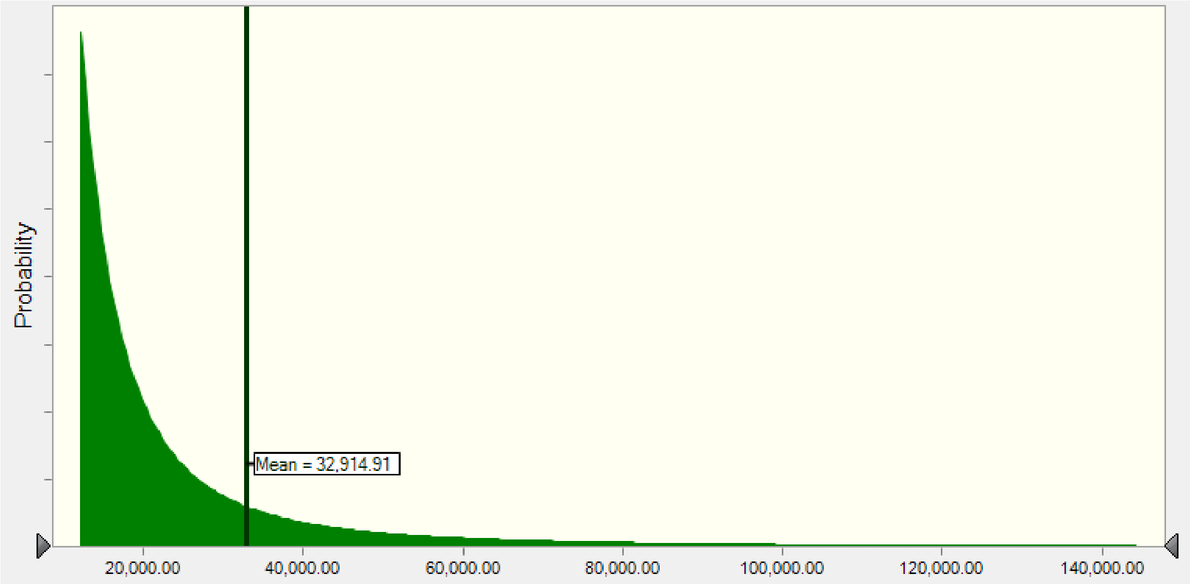

Figure 2: Pareto distribution of spendable household income in The Netherlands in 2010

The Pareto shows the income on the horizontal axis and the probability or likelihood to earn at least such an income on the vertical axis. The area under the curve indicates the percentage of the households. Looking at this income distribution, we can make a couple of observations.

1. The distribution has one tail at the right-hand side. The Pareto function is one-tailed. In our model there is a hard limit on the left presented by zero euros. On the right, the limit is not in sight. Theoretically, there is a ceiling if we know the income of the household with the highest earnings. In risk management, where losses are monitored, the tail will typically appear on the left. By flipping the sign from + to – we can also model left tail events. Some distributions, such as the normal distribution and the student's t, are two-tailed distributions.
2. We can see that the distribution is skewed to the right; with the mass of the graph on the left and the tail pointing to the right. Skew, in statistical terms, measures the degree of symmetry of a distribution. A completely symmetrical distribution, such as a normal distribution, has a skewness of zero. The Pareto distribution in the graph has a strong positive skewness. Intuitively, we see that if the mass of the distribution is pushed to the left, the graph's tail at the right-hand side becomes longer and is lifted up. Skewness is an indicator of long or fat tails.
3. We also notice that the distribution has a long tail indicating that there are households with very high spendable incomes. Although, only a small part of the population consists of high-earners, their numbers cannot be ignored. Contrary to the normal distribution, which in effect does not allow for occurrences beyond four or five standard deviations from the mean, the Pareto displays values that are far from the centre. The more we move to the right, the smaller the likelihood becomes that there are still households making that much money.

Fat and heavy tails

We mentioned that we would recognise a long tail when we see it. Some probability functions include fat tails, others do not. Below in figure 3, the graph zooms in on the tails of a normal, a lognormal, a student's t, and a Pareto distribution. It is clear that the normal distribution does not display a tail, while the latter three distributions do.

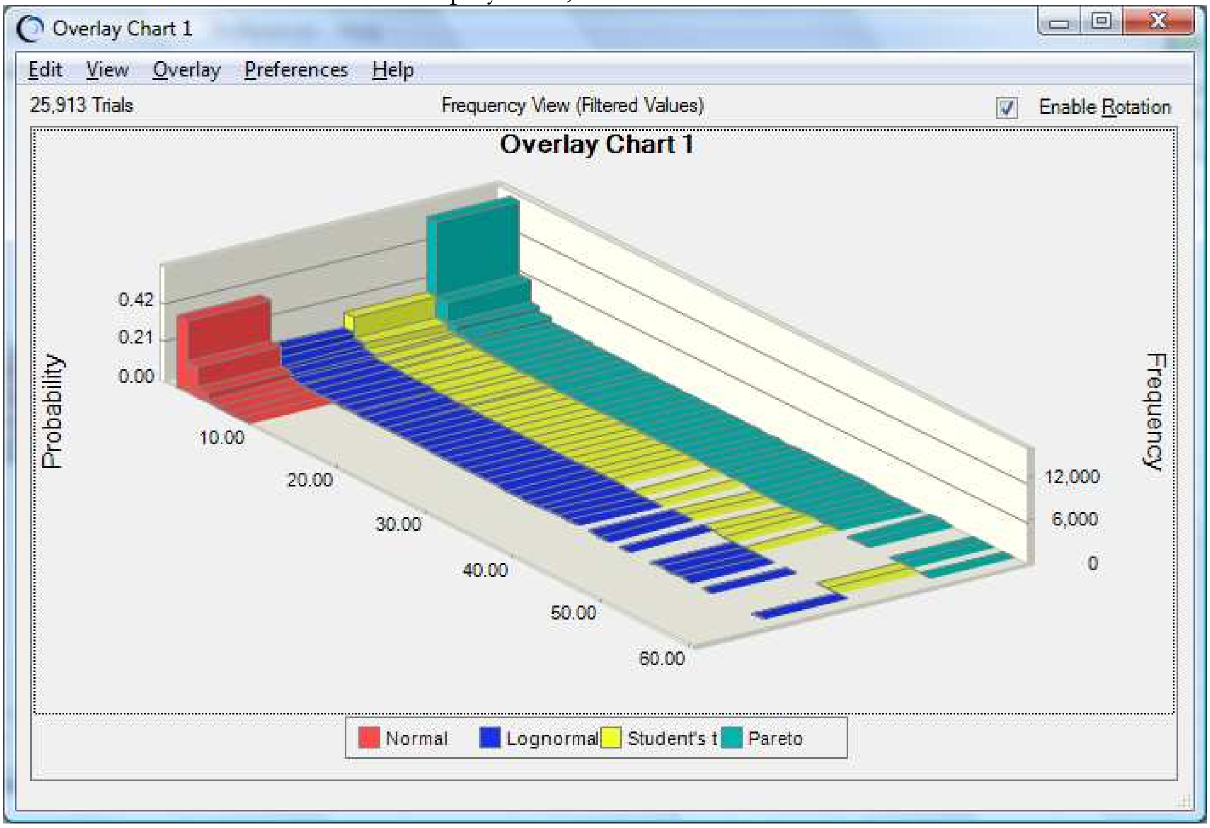

Figure 3: Tails of four distributions compared

All three of them have long tails, and the Pareto seems to be a bit fatter.
The study of fat tailed distributions is, for obvious reasons, a well-studied topic in risk management. Risk managers are not so much interested in what is "normal" but are interested in what happens in extreme cases or in tail events. Naturally, we sense that fat tail distributions reserve a large part of their surfaces for the tails since this is where the trouble lurks. The more of the distribution's surface is in the tail, the higher the probability of tail events becomes. Fat tailed probability density functions (PDF) demonstrate a low speed of moving toward the horizontal X-axis. Therefore, this acceleration factor determines the fatness of the tail. Let's take a look at the following function:

$$k \times \frac{1}{x^{\alpha}}$$ Where k and α are constants > 0

From the formula above, it becomes clear that with larger α's the denominator increases and the value of the fraction will accelerate toward zero. Smaller values for alpha will produce fat tails. The alpha (α) is referred to as the tail parameter or tail index. Working with Excel and using scatter charts we can quickly get a feel for the behaviour of the distribution with various alphas. When we define, for instance, a Pareto distribution, we have to provide the alpha.

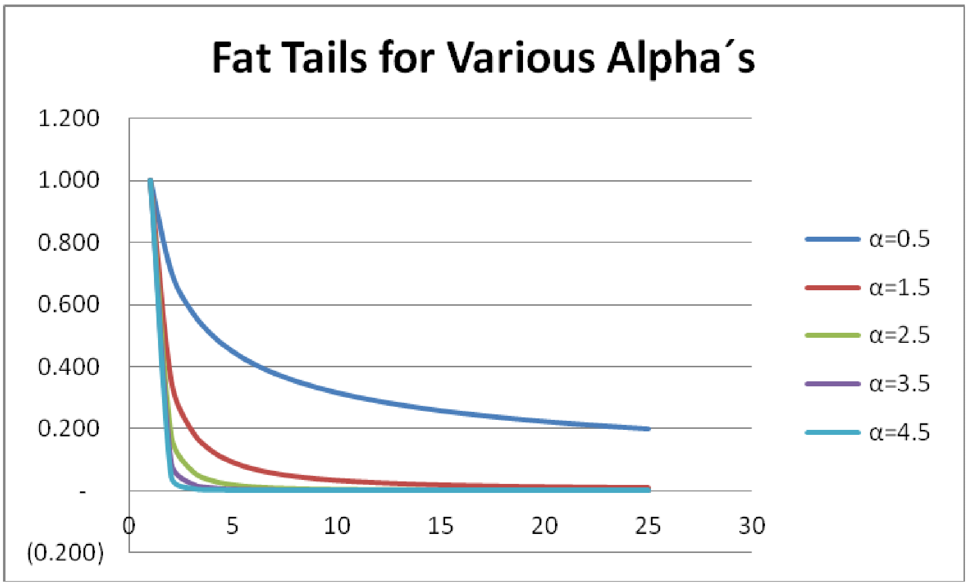

Figure 4: Fat tails for various alphas; smaller alphas produce fatter tails

Fat tailed distributions pose specific problems to portfolio managers and, in extreme cases, might annul the benefits of risk diversification and portfolio management.

Portfolio management

Modern portfolio theory is built on the assumption that combining assets in a portfolio will reduce the risk and that this portfolio benefit can be quantified. For years, portfolio theory has reached beyond traditional instruments such as stocks and bonds and is now applied in many areas including loan and insurance portfolios. Apart from the expected return, default, or claim, the standard deviation of the underlying instruments is a key input parameter. In combination with the correlations of the components, the overall riskiness of the portfolio can be calculated. Therefore, it is crucial that the standard deviations of the portfolio inputs are known and can be relied upon. Fat tails might put a spoke in the portfolio wheel of fortune. Let's take a look.

In statistics, "moments" denote the statistical tools used for characterising and analysing data sets. The first moment is the average of the data set, the second moment the variance (which equals the standard deviation squared), and the third moment is the degree of symmetry or skewness. The fourth moment is kurtosis which measures the pointiness of a graph. A standard normal distribution has a kurtosis of three. Therefore, a kurtosis in the range from zero up to three indicates a graph flatter than a Gaussian curve. A kurtosis larger than three produces a pointed chart. I will not discuss the calculation of the moments, but it is worth mentioning that all four formulas contain a factor that measures the deviation of the expected, average value. This value is raised to the power one, two, three, or four corresponding to the four statistical 'moments'. So, the kurtosis, for example, contains an exponentiation with a power four. It is easy to see that these exponents may accelerate the calculation values and result in large numbers. The kurtosis is the most sensitive value in this respect, followed by the skewness, variance, and the mean. Hence, the moments are in reversed order, from fourth to first, in terms of sensitivity to extreme values. Kurtosis and skewness are the canary birds in the coal mines warning for danger yet to come.

If we return to our Pareto distribution of the spendable income in The Netherlands and use Oracle Crystal Ball to do the statistical analysis we notice that both skewness and the kurtosis are very high: ≈ 47 and ≈ 4227 respectively. These values are far from the zero and three of a standard normal distribution.

Statistic	Assumption values
Trials	100,000
Base Case	0.00
Mean	32,029.25
Median	18,564.24
Mode	---
Standard Deviation	83,822.60
Variance	7,026,228,501.75
Skewness	46.68
Kurtosis	4,226.84
Coeff. of Variability	2.62
Minimum	11,950.03
Maximum	10,909,753.12
Mean Std. Error	265.07

Figure 5: Statics of the Pareto distribution of spendable household income in The Netherlands 2010

High values for the third and fourth moments might, as we have discussed earlier, be the harbingers of trouble ahead. Trouble, for the risk manager, translates to fat tailed distributions with relatively high probabilities for extreme events. The Latin saying goes *in cauda venum*, the poison is in the tail. Therefore, let's take a closer look at the tail behaviour of the distributions discussed.

Tail behaviour

Instead of doing the statistical works for the complete distribution, we could also choose to limit our analysis to just the tail. To keep things simple, we define the tail as the area of the graph starting at the 90th percentile.

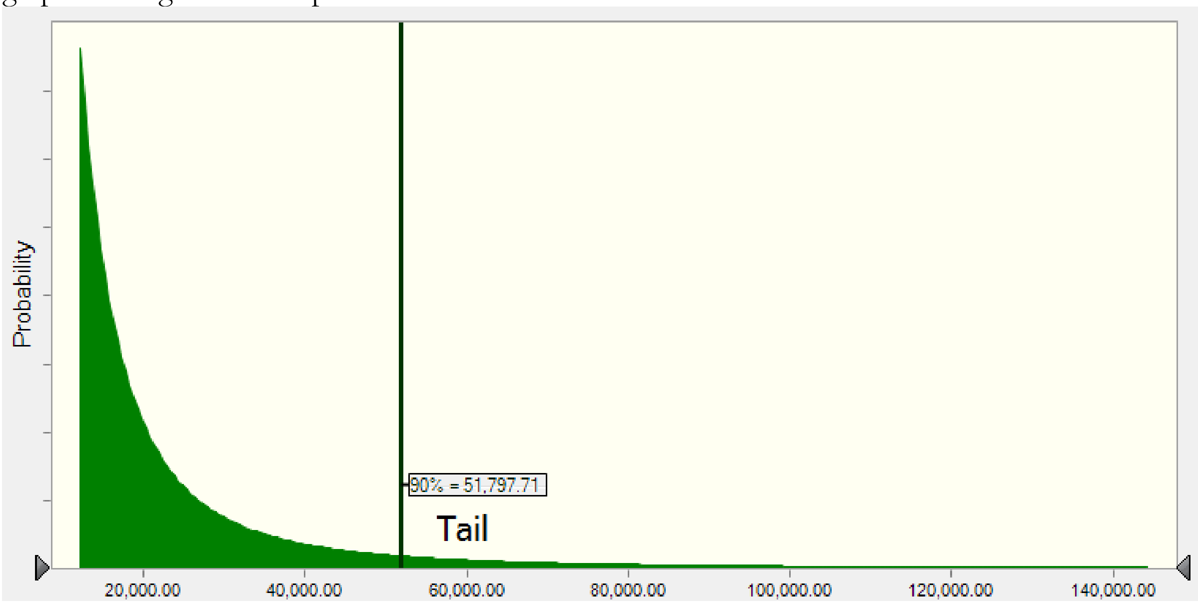

Figure 6: Tail of a Pareto PDF defined as the area beyond the 90th percentile

For the just tail, we can repeat the calculation of the four moments: average, variance, skewness, and kurtosis. Analysing the tail with our statistical tool kit brings about some singularities for fat-tailed distributions. As it turns out, for heavy or fat tailed distributions one or more moments are undefined. Undefined, in our case, means that the kurtosis cannot be determined. Of course, the formula still works and we get a result for the kurtosis calculation, but as we include more observations, does not stabilise, on the contrary: it goes berserk.

Large numbers

This phenomenon conflicts with the law of large numbers, which teaches that the more observations are included in the calculation, the more stable and reliable the result will become. Say, we would like to assess the average height of male adult football supporters in a stadium. It is understood that there is no need measuring all fifty thousand supporters and that a sufficiently large sample will do. As the sample size observed becomes larger, the mean value will get closer to the real mean of all fifty thousand. The larger the sample, the more stable the number. Once we have measured, for instance, a thousand adults, adding another thousand measures will not change the average and other moments much. Remarkably, this is not the case for our tail analysis of heavy and fat tailed distributions. In the case of heavy and fat tails, their statistical moments go wild and refuse to solidify.

Extreme values

A tool commonly used in extreme-value theory is the mean excess plot. Earlier, we have defined, maybe somewhat arbitrarily, the tail as the part of the graph beyond the 90th percentile. Continuing our football supporters' example, we call someone tall when he is in this quantile. The mean excess is the average of the tail which provides the mean extra height of what we have classified as a tall supporter. The extra height is the excess height that comes on top of the length at the starting point of the tail. If we start moving the threshold from the 90th percentile to the 91st, 92nd, and so on, we can expect this excess height to decline. So, there is a diminishing average excess height. This is not the case, however, for fat tails. Here, the mean excess is accelerating. The same goes for the excess plot of the standard deviation. As we move further down the tail, the situation becomes wilder and wilder.

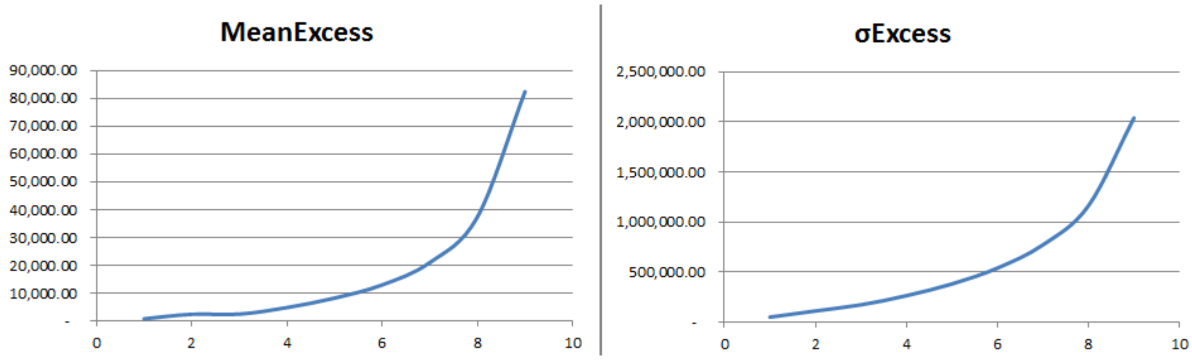

Figure 7: Excess plots for mean and standard deviation for super fat Pareto tail

This is bad news for risk managers. If the variance and standard deviation cannot be determined, the traditional portfolio theory cannot be applied. After all, the standard deviations of the components of a portfolio are crucial input for the portfolio model. The wild statistical moments wipe out the idea of diversification.

Seeking alpha?

Above we discussed the α-parameter, defining the fatness of a tail. The uncontrolled behaviour of the various statistical moments is linked to the α-value. Distributions with an alpha above four the moments are stable. Below four, however, the kurtosis starts floating. Below three, the skewness follows suit, and under two, the variance and standard deviation lose solid grounds. Finally, for alphas smaller than one even the mean becomes unanchored. This way, the alpha presents a classification method for the fatness of distribution tails.

In the investment world, *seeking alpha* has become a winged phrase and it is also the name of one of the most popular investment sites. Alpha, in this case, represents the value with which a portfolio outperforms the exchange. The alpha discussed in this paper is of a different nature: an indicator of the fatness of the tails and of risk. Given the serious impact fat tails can have on the portfolio management approach chosen, it seems recommendable to watch the tail alpha as well.

It is time for a new investment website called *tailing alpha*. The domain name is still available!

Amsterdam, Jan 23rd, 2012

Stachanov Solutions & Services b.v.

André Koch

Stachanov Group, Amsterdam
T +31 20 509 1010
E info@stachanov.com

About Stachanov Group

Comparatively small, but very fine and well-respected engineering company is serving AAA-customers since 1997.
Focus on large database projects for the public sector (e.g. UWV, Ontwikkelingsbedrijf Rotterdam), and (risk-) modelling projects for both public (e.g. Havenbedrijf Rotterdam) and private sectors (e.g. all Dutch banks).
Most work is on project basis, but we also have hosting customers (e.g. Euronext).
ISO-certified since 2004, we embrace: Innovation, Professionalism & Operational Excellence.