A focus on the exceptions that prove the rule

Traditional risk management tools focus on what is normal and consider extreme events as ancillaries. In a world characterised by volatility and uncertainty, **Benoit Mandelbrot** and **Nassim Taleb** argue that this approach is misguided, and propose an alternative methodology where large deviations dominate the analysis

Conventional studies of uncertainty, whether in statistics, economics, finance or social science, have largely stayed close to the so-called "bell curve", a symmetrical graph that represents a probability distribution. Used to great effect to describe errors in astronomical measurement by the 19th-century mathematician Carl Friedrich Gauss, the bell curve, or Gaussian model, has since pervaded our business and scientific culture, and terms like sigma, variance, standard deviation, correlation, R-square and Sharpe ratio are all directly linked to it.

If you read a mutual fund prospectus, or a hedge fund's exposure, the odds are that it will supply you, among other information, with some quantitative summary claiming to measure "risk". That measure will be based on one of the above buzzwords that derive from the bell curve and its kin.

Such measures of future uncertainty satisfy our ingrained desire to "simplify" by squeezing into one single number matters that are too rich to be described by it. In addition, they cater to psychological biases and our tendency to understate uncertainty in order to provide an illusion of understanding the world.

The bell curve has been presented as "normal" for almost two centuries, despite its flaws being obvious to any practitioner with empirical sense. Granted, it has been tinkered with using such methods as complementary "jumps", stress testing, regime switching or the elaborate methods known as GARCH, but while they represent a good effort, they fail to address the bell curve's fundamental flaws.

The problem is that measures of uncertainty using the bell curve simply disregard the possibility of sharp jumps or discontinuities and, therefore, have no meaning or consequence. Using them is like focusing on the grass and missing out on the (gigantic) trees. In fact, while the occasional and unpredictable large deviations are rare, they cannot be dismissed as "outliers" because, cumulatively, their impact in the long term is so dramatic.

The traditional Gaussian way of looking at the world begins by focusing on the ordinary, and then deals with exceptions or so-called outliers as ancillaries. But there is also a second way, which takes the so-called exceptional as a starting point and deals with the ordinary in a subordinate manner – simply because that "ordinary" is less consequential.

These two models correspond to two mutually exclusive types of randomness: mild or Gaussian on the one hand, and wild, fractal or "scalable power laws" on the other. Measurements that exhibit mild randomness are suitable for treatment by the bell curve or Gaussian models, whereas those that are susceptible to wild randomness can only be expressed accurately using a fractal scale. The good news, especially for practitioners, is that the fractal model is both intuitively and computationally simpler than the Gaussian, which makes us wonder why it was not implemented before.

We live in a world of extreme concentration where the winner takes all. Consider how Google grabs much of internet traffic or how Microsoft represents the bulk of software sales

Let us first turn to an illustration of mild randomness. Assume that you round up 1,000 people at random among the general population and bring them into a stadium. Then, add the heaviest person you can think of to that sample. Even assuming he weighs 300kg, more than three times the average, he will rarely represent more than a very small fraction of the entire population (say, 0.5 per cent). Similarly, in the car insurance business, no single accident will put a dent on a company's annual income. These two examples both follow the "Law of Large Numbers", which implies that the average of a random sample is likely to be close to the mean of the whole population.

In a population that follows a mild type of randomness, one single observation, such as a very heavy person, may seem impressive by itself but will not disproportionately impact the aggregate or total. A randomness that disappears under averaging is trivial and harmless. You can diversify it away by having a large sample.

There are specific measurements where the bell curve approach works very well, such as weight, height, calories consumed, death by heart attacks or performance of a gambler at a casino. An individual that is a few million miles tall is not biologically possible, but an exception of equivalent scale cannot be ruled out with a different sort of variable, as we will see next.

Wild randomness

What is wild randomness? Simply put, it is an environment in which a single observation or a particular number can impact the total in a disproportionate way. The bell curve has "thin tails" in the sense that large events are considered possible but far too rare to be consequential. But many fundamental quantities follow distributions that have "fat tails" – namely, a higher probability of extreme values that can have a significant impact on the total.

One can safely disregard the odds of running into someone several miles tall, or someone who weighs several million kilogrammes, but similar excessive observations can never be ruled out in other areas of life.

Having already considered the weight of 1,000 people assembled for the previous experiment, let us instead consider wealth. Add to the crowd of 1,000 the wealthiest person to be found on the planet – Bill Gates, the founder of Microsoft. Assuming that his net worth is close to $80bn, how much would he represent of the total wealth? 99.9 per cent? Indeed, all the others would represent no more than the variation of his personal portfolio over the past few seconds. For someone's weight to represent such a share, he would need to weigh 30m kg.

Try it again with, say, book sales. Line up a collection of 1,000 authors. Then, add the most read person alive, JK Rowling, the author of the Harry Potter series. With sales of several hundred million books, she would dwarf the remaining 1,000 authors who would collectively have only a few hundred thousand readers.

So, while weight, height and calorie consumption are Gaussian, wealth is not. Nor are income, market returns, size of hedge funds, returns in the financial markets, number of deaths in wars or casualties in terrorist attacks. Almost all man-made variables are wild. Furthermore, physical science continues to discover more and more examples of wild uncertainty, such as the intensity of earthquakes, hurricanes or tsunamis.

Economic life displays numerous examples of wild uncertainty. For example, during the 1920s, the German currency moved from three to a dollar to 4bn to the dollar in a few years. And veteran currency traders still remember when, as late as the 1990s, short-term interest rates jumped by several thousand per cent.

We live in a world of extreme concentration where the winner takes all. Consider, for example, how Google grabs much of internet traffic, how Microsoft represents the bulk of PC software sales, how 1 per cent of the US population earns close to 90 times the bottom 20 per cent or how half the capitalisation of the market (at least 10,000 listed companies) is concentrated in less than 100 corporations.

Taken together, these facts should be enough to demonstrate that it is the so-called "outlier" and not the regular that we need to model. For instance, a very small number of days accounts for the bulk of the stock market changes: just ten trading days represent 63 per cent of the returns of the past 50 years.

Let us now return to the Gaussian for a closer look at its tails. The "sigma" is defined as a "standard" deviation away from the average, which could be around 0.7 to 1 per cent in a stock market or 8 to 10 cm for height. The probabilities of exceeding multiples of sigma are obtained by a complex mathematical formula. Using this formula, one finds the following values:

Probability of exceeding:
0 sigmas: 1 in 2 times
1 sigmas: 1 in 6.3 times
2 sigmas: 1 in 44 times
3 sigmas: 1 in 740 times
4 sigmas: 1 in 32,000 times
5 sigmas: 1 in 3,500,000 times
6 sigmas: 1 in 1,000,000,000 times
7 sigmas: 1 in 780,000,000,000 times
8 sigmas: 1 in 1,600,000,000,000,000 times
9 sigmas: 1 in 8,900,000,000,000,000,000 times
10 sigmas: 1 in 130,000,000,000,000,000,000,000 times

and, skipping a bit:

20 sigmas: 1 in 36,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000 times

Soon, after about 22 sigmas, one hits a "googol", which is 1 with 100 zeroes behind it. With measurements such as height and weight, this probability seems reasonable, as it would require a deviation from the average of more than 2m. The same cannot be said

The difference ten days make

By removing the ten biggest one-day moves from the S&P 500 over the past 50 years, we see a huge difference in returns. And yet conventional finance treats these one-day jumps as mere anomalies.

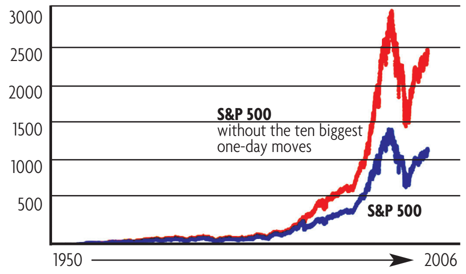

EDITOR Rob Mitchell
DEPUTY EDITOR Ravi Mattu
DESIGN Stephen Stafford
COVER ILLUSTRATION Brett Ryder
BUSINESS DEVELOPMENT MANAGER Kate Harris
FT PARTNERSHIP PUBLICATIONS
Number One Southwark Bridge London SE1 9HL
+44 (0)20 7873 3000
DIRECTOR Richard Foster
www.ft.com/masteringuncertainty
Note: As sponsors of this publication, Ernst & Young had no input in the selection of its contributors or editorial content