

# 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

**C**onventional 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

**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**

and computationally simpler than the Gaussian, which makes us wonder why it was not implemented before.

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 kilograms, 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 varia-

bles 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:

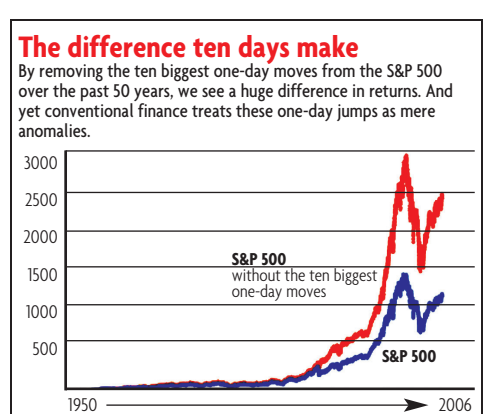
## Probability of exceeding:

- 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 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



**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](http://www.ft.com/masteringuncertainty)  
Note: As sponsors of this publication, Ernst & Young had no input in the selection of its contributors or editorial content



**The Gaussian and fractal models: observations and consequences**

- 1 By itself, no single number can characterise uncertainty and risk but, as we have seen, we can still have a handle on it so long as we can have a table, a chart and an open mind.
- 2 In the Gaussian world, standard tables show that 67 per cent of the observations fall between -1 and +1 sigma. Outside of this, sigma loses its significance. With a scalable distribution, you may have 80 per cent, 90 per cent, even 99.99 per cent of observations falling between -1 and +1 sigmas. In fractals, the standard deviation is never a “typical” value and may even be infinite!
- 3 When assessing the effectiveness of a given financial, economic or social strategy, the observation window needs to be large enough to include substantial deviations, so one must base strategies on a long time frame.
- 4 You are far less diversified than you assume. Because the market returns in the very long run will be dominated by a small number of investments, you need to mitigate the risk of missing these by investing as broadly as possible. Very broad passive indexing is far more effective than active selection.
- 5 Projections of deficits, performance and interest rates are marred with extraordinarily large errors. In many budget calculations, US interest rates were projected to be 5 per cent for 2001 (not 1 per cent); oil prices were projected to be close to \$22 a barrel for 2006 (not \$62). Like prices, forecast errors follow a fractal distribution.
- 6 Option pricing models, such as Black-Scholes-Merton, are strongly grounded in the bell curve in their representation of risk. The Black-Scholes-Merton equation bases itself on the possibility of eliminating an option’s risk through continuous dynamic hedging, a procedure incompatible with fractal discontinuities.
- 7 Some classes of investments with explosive upside, such as venture capital, need to be favoured over those that do not have such potential. Technology investments get bad press; priced appropriately (in the initial stages) they can deliver huge potential profits, thanks to the small, but significant, possibility of a massive windfall.
- 8 Large moves beget large moves; markets keep in memory the volatility of past deviations. A subtle concept, fractal memory provides an intrinsic way of modelling both the clustering of large events and the phenomenon of regime switching, which refers to phases when markets move from low to high volatility.

Nick Lowndes

of variables such as financial markets. For example, a level described as a 22 sigma has been exceeded with the stock market crashes of 1987 and the interest rate moves of 1992.

The key here is to note how the frequencies in the preceding list drop very rapidly, in an accelerating way. The ratio is not invariant with respect to scale.

Let us now look more closely at a fractal, or scalable, distribution using the example of wealth. We find that the odds of encountering a millionaire in Europe are as follows:

- Richer than 1 million: 1 in 62.5
- Richer than 2 million: 1 in 250
- Richer than 4 million: 1 in 1,000
- Richer than 8 million: 1 in 4,000
- Richer than 16 million: 1 in 16,000
- Richer than 32 million: 1 in 64,000
- Richer than 320 million: 1 in 6,400,000

This is simply a fractal law with a “tail exponent”, or “alpha”, of two, which means that when the number is doubled, the incidence goes down by the square of that number – in this case four. If you look at the ratio of the moves, you will notice that this ratio is invariant with respect to scale.

If the “alpha” were one, the incidence would decline by half when the number is

doubled. This would produce a “flatter” distribution (fatter tails), whereby a greater contribution to the total comes from the low probability events.

- Richer than 1 million: 1 in 62.5
- Richer than 2 million: 1 in 125
- Richer than 4 million: 1 in 250
- Richer than 8 million: 1 in 500
- Richer than 16 million: 1 in 1,000

We have used the example of wealth here, but the same “fractal” scale can be used for stock market returns and many other variables. Indeed, this fractal approach can prove to be an extremely robust method to identify a portfolio’s vulnerability to severe risks. Traditional “stress testing” is usually done by selecting an arbitrary number of “worst-case scenarios” from past data. It assumes that whenever one has seen in the past a large move of, say, 10 per cent, one can conclude that a fluctuation of this magnitude would be the worst one can expect for the future. This method forgets that crashes happen without antecedents. Before the crash of 1987, stress testing would not have allowed for a 22 per cent move.

Using a fractal method, it is easy to extrapolate multiple projected scenarios. If your worst-case scenario from the past data was,

say, a move of -5 per cent and, if you assume that it happens once every two years, then, with an “alpha” of two, you can consider that a -10 per cent move happens every eight years and add such a possibility to your simulation. Using this model, a -15 per cent move would happen every 16 years, and so forth. This will give you a much clearer idea of your risks by expressing them as a series of possibilities.

You can also change the alpha to generate additional scenarios – lowering it means increasing the probabilities of large deviations and increasing it means reducing the probabilities. What would such a method reveal? It would certainly do what “sigma” and its siblings cannot do, which is to show how some portfolios are more robust than others to an entire spectrum of extreme risks. It can also show how some portfolios can benefit inordinately from wild uncertainty.

Despite the shortcomings of the bell curve, reliance on it is accelerating, and widening the gap between reality and standard tools of measurement. The consensus seems to be that any number is better than no number – even if it is wrong. Finance academia is too entrenched in the paradigm to stop calling it “an acceptable approximation”.

Any attempts to refine the tools of modern portfolio theory by relaxing the bell curve

assumptions, or by “fudging” and adding the occasional “jumps” will not be sufficient. We live in a world primarily driven by random jumps and tools designed for random walks address the wrong problem. It would be like tinkering with models of gases in an attempt to characterise them as solids and call them “a good approximation”.

While scalable laws do not yet yield precise recipes, they have become an alternative way to view the world, and a methodology where large deviation and stressful events dominate the analysis instead of the other way around. We do not know of a more robust manner for decision-making in an uncertain world.

**Benoit Mandelbrot** is Sterling professor emeritus of mathematical sciences at Yale University. He is the author of “Fractals and Scaling in Finance” (Springer-Verlag, 1999) and, with Richard L Hudson, of “The (Mis)behaviour of Markets” (Profile, 2005). [benoit.mandelbrot@yale.edu](mailto:benoit.mandelbrot@yale.edu)  
**Nassim Nicholas Taleb** is a veteran derivatives trader and Dean’s professor in the sciences of uncertainty at the University of Massachusetts, Amherst. He is also the author of “Fooled by Randomness” (Random House, 2005) and “The Black Swan” (forthcoming). [www.fooledbyrandomness.com](http://www.fooledbyrandomness.com)