David Orrell

Why risk models don't work

As economists reconsider perceptions of risk management, there are many models to discard

One of the more memorable statements to come out of the recent crunch was during the summer of 2007, when the CFO at Goldman Sachs reported "seeing things that were 25-standard deviation moves, several days in a row." The phrase "standard deviation" is often thrown around when discussing risk, or evaluating a company or country's chance of default. But what does it really mean?

Standard deviation is a measure of the width of a normal distribution, or bell curve. About 68 percent of the values fall within a standard deviation of the mean, and 95 percent are within two standard deviations. IQ tests are designed to give an average score of 100 with a standard deviation of about 15. So all but five percent of the population has an IQ in the range 70 to 130 (the mean plus or minus two standard deviations). Normal distribution is popular because it can be widely applied. The height of adult males in the US, for example, follows a normal distribution with an average of 70 inches, and a standard deviation of three inches. Now, seeing a 25-standard deviation price change, as Goldman claimed, is statistically equivalent to seeing a man who is 12ft tall. Unless there's some undiscovered basketball star hidden away in a third-world village somewhere, I think it's safe to say that's impossible. The problem is not with the observations of the markets – those extreme events did happen – but with the use of the standard deviation. For while the normal distribution is a useful statistical tool for many situations, it unfortunately doesn't work well in finance.

Scaled distribution

Instead of following a normal distribution, with a well-defined mean, price changes in the stock market are better described by a so-called power-law distribution. Such a distribution is scale-free, i.e. there is no typical scale. A good way to illustrate this, is with the surface of the moon.

They don't have financial crashes on the moon, but they do have craters. It turns out that the sizes of craters follow exactly the same kind of scale-free distribution as markets. There are many small craters, and fewer large ones, but it is impossible to define typical size.

In fact, one obvious consequence of a scale-free distribution is that you can't tell by looking at it what the scale is. Markets have the same property – remove the time scale from a plot of price changes, and it is hard to know whether it is over a decade, a year, or a month.

The problem with our current risk models is that they assume that price changes follow a normal distribution, with at best small modifications. Consider for example the Gaussian copula models, used to value the Collateralised Mortgage Obligations which played a key role in the sub-prime debacle. Hint: Gaussian distribution is a different name for normal distribution. The Value at Risk (VaR) models used by banks and firms are also usually based on the normal distribution, or variants thereof. That's why people still describe risk in terms of standard deviations.

But because financial markets are like moon craters, instead of IQ scores or human heights, the standard deviation is a meaningless concept. You can calculate it for a particular sample, but it doesn't tell you much about the real risk. In particular, the normal distribution drastically underestimates the probability of extreme events.

Such events are sometimes known as "black swans" after Nassim Taleb's book. However it isn't the case that markets are normal most of the time (white swans) and every now and then something really

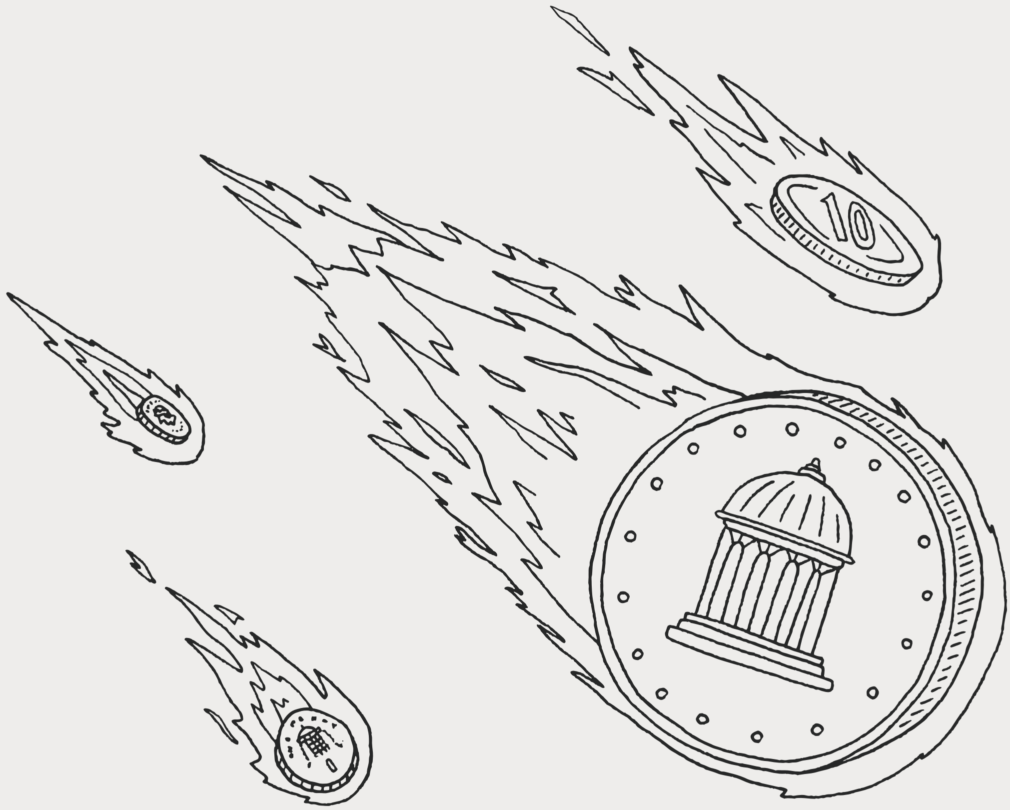

weird happens (a black swan, or a 12-foot basketball player). As with craters on the moon, extreme events aren't freaks – they are part of the landscape. The largest and the smallest events are all part of the same distribution, and are generated by the same processes. It's just that, in our everyday lives, we only notice them when they're big.

Re-thinking risk

So how can we revise our risk models to better reflect reality? One approach is to use power-law distributions to stress-test portfolios and see how they will respond to extreme events. However, a problem with the power law is that it is very hard to determine the parameters which give the correct distribution. By definition, the most important (i.e. extreme) events only happen rarely, and these are precisely the ones you need to fit the model. And using older historical data only makes sense if you really believe that the stockmarket hasn't changed much since the 1930s.

The power law also isn't much use at predicting the future, because it only tells us about the distribution of sizes, and not the timing. Earthquake magnitudes follow a power-law distribution too, but we have no idea when the next one will strike. Nor do we know when an asteroid will next strike the moon.

The real message to come out of this is that, while we can improve risk models by adding in a margin of error, mathematical formulae will never be able to fully capture risk. Just because an equation says it is safe, doesn't mean it is.

Another problem with risk formulae, is that if they are widely adopted (as with VaR) then they can actually end up influencing the market and creating systemic risk. However that is a story for another column.

We can improve risk models by adding in a margin of error. But mathematical formulae will never be able to fully capture risk

By the way, the diameter of that large crater, Korolev, is 437 km, a little larger in area than England. In terms of moon craters, it's a Great Depression – and just as statistically impossible, if you believe the normal distribution. Something to remember, next time you hear the standard deviation being used to describe risk. ◇

David Orrell is a mathematician and author. His next book is Economyths: Ten Ways That Economics Gets It Wrong. His column features regularly on www.worldfinance.com