LEARNING CURVE EXTREME VALUE VAR*

Derivatives Week, March 6
Aragonés, José.R., Carlos Blanco, and Kevin Dowd

Risk managers are primarily concerned with the risk of low-probability events that could lead to catastrophic losses. Yet traditional VaR methods tend to ignore extreme events and focus on risk measures that accommodate the whole empirical distribution of returns. For example, it is often assumed that returns are normally or lognormally distributed, and little attention is paid to the distribution of the extreme returns we are most concerned about. The danger is then that our models are prone to fail just when they are needed most – in large market moves, when we can suffer very large losses.

One response to this problem is to use stress tests and scenario analyses. These can simulate the changes in the value of our portfolio under hypothesized extreme market conditions. These are certainly very useful. However, they are inevitably limited – we cannot explore all possible scenarios – and by definition give us no indication of the likelihoods of the scenarios considered.

This type of problem is not unique to risk management, but also occurs in other disciplines as well, particularly in hydrology and structural engineering, where the failure to take proper account of extreme values can have devastating consequences. Researchers and practitioners in these areas handle this problem by using Extreme Value Theory (EVT) – a specialist branch of statistics that attempts to make the best possible use of what little information we have about the extremes of the distributions in which we are interested.

EXTREME VALUE THEORY

The key to EVT is the extreme value theorem – a cousin of the better-known central limit theorem – which tells us what the distribution of extreme values should look like in the limit, as our sample size increases. Suppose we have some return observations but do not know the density function from which they are drawn. Subject to certain

* This week's Learning Curve was written by José R. Aragonés from the Universidad Complutense, Madrid; Carlos Blanco of Financial Engineering Associates in Berkeley, CA; and Kevin Dowd from the University of Sheffield. They can be contacted at econ305@sis.ucm.es, carlos@fea.com and k.dowd@shef.ac.uk respectively.

relatively inocuous conditions, this theorem tell us that the distribution of extreme returns converges asymptotically to:

$$H_{\xi,\mu,\sigma}(x) = \begin{cases} \exp(-[1+\xi(x-\mu)/\sigma]^{-1/\xi}) & \text{if} \quad \xi \neq 0 \\ \exp(-e^{-(x-\mu)/\sigma}) & \quad\quad \xi = 0 \end{cases}$$

The parameters μ and σ correspond to the mean and standard deviation, and the third parameter, ξ, gives an indication of the heaviness of the tails: the bigger ξ, the heavier the tail. This parameter is known as the tail index, and the case of most interest in finance is where $\xi > 0$, which corresponds to the fat tails commonly founded in financial return data. In this case, our asymptotic distribution takes the form of a Fréchet distribution.

This theorem tells us that the limiting distribution of extreme returns always has the same form – whatever the distribution of the parent returns from which our extreme returns are drawn. It is important because it allows us to estimate extreme probabilities and extreme quantiles, including VaRs, without having to make strong assumptions about an unknown parent distribution.

EXTREME-VALUE VAR

EVT provides a natural approach to VaR estimation, given that VaR is primarily concerned with the tails of our return distributions. To apply to VaR, we first estimate the parameters of the distribution, and there are a number of standard estimators available. Once we have these, we can plug them into a number of alternative formulas to obtain VaR estimates. To give a simple example, if we want to estimate a VaR that is out of (i.e., more extreme than) our sample range, we can project the tail out from an existing in-sample quantile X_{k+1} – where X_{k+1} is the k+1-th most extreme observation in our sample – and infer the (asymptotic) VaR from the projected tail using the formula:

$$VaR = [CL/k]^{-\hat{\xi}} X_{k+1}$$

where *CL* is the confidence level on which the VaR is predicated. EVT also gives us expressions for the confidence intervals associated with our VaR estimates.

The difference EVT makes to our VaR estimates is illustrated in the Figure. This shows the tail of the West Texas Intermediate (WTI) daily return distribution from 1983 to 1999. The dots indicate the actual extreme return observations, the continuous line by the right vertical axis represents the tail assuming that logarithmic returns follow a normal distribution, and the other continuous line represents the tail of an EV distribution fitted to these data. The message from this Figure is unmistakeable: EV-VaRs are much bigger than normal VaRs, especially at high confidence levels. Assuming normality can thus lead to a major underestimate of our risks.

Figure: EV-VaR vs. Normal VaR

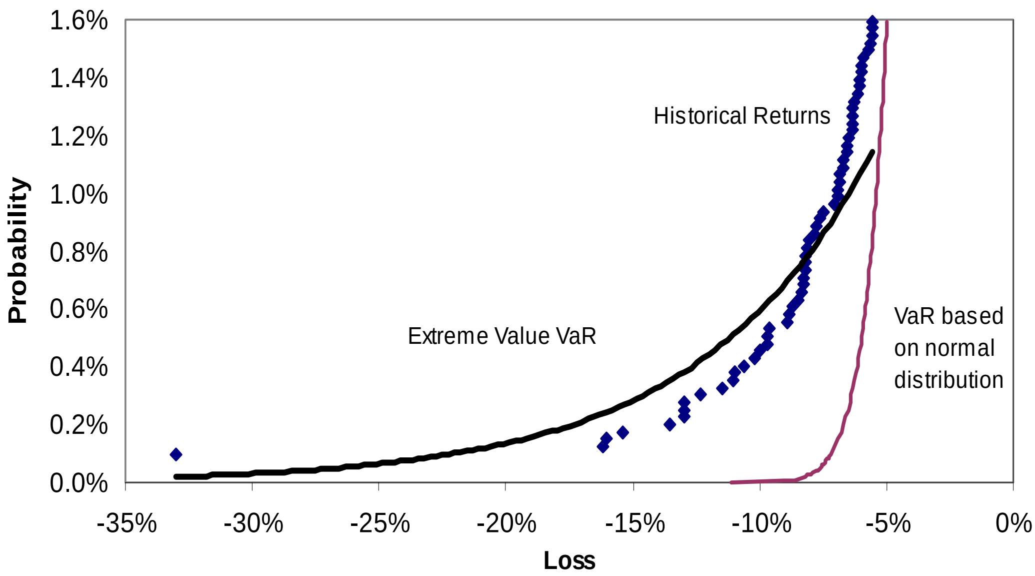

EV-VAR VS. TRADITIONAL VARS

The EV approach to VaR has certain advantages over traditional parametric and non-parametric approaches to VaR.

Parametric approaches estimate VaR by fitting some distribution to a set of observed returns. However, since most observations lie close to the centre of any empirical distribution, traditional parametric approaches tend to fit curves that accommodate the mass of central observations, rather than accommodate the tail observations that are more important for VaR purposes. Traditional parametric approaches also suffer from the drawback that they impose distributions that make no sense for tail estimation and fly in the face of EV theory. By comparison, the EV approach is free of these problems and specifically designed for tail estimation.

Non-parametric or historical simulation approaches estimate VaR by reading off the VaR from an appropriate histogram of returns. However, they lead to less efficient VaR estimates than EV approaches, because they make no use of the EV theory that gives us some indication of what the tails should look like. More importantly, these approaches also have the very serious limitation that they can tell us nothing whatever about VaRs beyond our sample range.

SUMMARY

EVT deals with the frequency and magnitude of very low probability events. Precisely because extreme events are extreme, we have to operate with very small data sets, and this means that our estimates – our quantile estimates, our VaRs, and the estimated probabilities associated with them – are inevitably very imprecise. However, EVT makes the best out of an inherently difficult problem, and therefore marks a significant step forward in VaR estimation.

Of course, as with every other tool in risk management, the successful use of EVT requires an appreciation of strengths and limitations. A lot depends on judgement and experience, and EVT is not an exercise in mechanical number crunching. However, carefully used, EVT can be very useful indeed.