SOCIETY OF ACTUARIES

Investment Section

INVESTMENT FALLACIES 2014

A real-world approach to Value at Risk

By Nicholas John Macleod

Introduction

A well-known legal anecdote has it that the barrister Sir Edward Marshall Hall, when asked by a judge, "Is your client familiar with the doctrine of *Res Ipsa Loquitur*?", replied, "My Lord, in the remote hills of County Donegal from where my client hails, they speak of little else."

In many ways that captures the problem faced by investment professionals when they come to consider the calculation of Value at Risk (VaR). Far from being able to provide reasonable estimates based on their day-to-day investment experience, they find themselves confronted with controversy concerning the applicability or otherwise of the Normal distribution, and with proposed remedies that encompass Extreme Value Theory, GARCH models, Copulas, etc., none of which falls within their areas of professional expertise.

The difficulties with such mathematical approaches, however, go beyond their inaccessibility to the practitioner. VaR is concerned with the possibility of unusually large losses, so almost by definition, there is little empirical basis for assessing the nature of potential non-Normal goings-on in the extremes of return distributions. As a result, there are no real grounds for attributing any particular structure to that part of the return distribution, whether by GARCH models, generalized Pareto distributions or anything else; whichever approach we choose, our chosen model will tell us exactly what we told it to tell us.

In this essay I'm going to suggest that estimating VaR is not fundamentally a problem requiring advanced mathematics, and that its solution is not to be found by attempting to divine the nature of hitherto unobserved behavior in the tails of asset return distributions. I'll go further and argue that a common sense approach not only provides more realistic estimates of potential loss, it forces us to identify shortcomings in portfolio structure, and provides practical guidance with respect to what to do about them.

Common-sense Value at Risk

The definition of VaR is straightforward: if I can lose $\$X$ or more with probability p , is my VaR at level p. In terms of return, if a portfolio has a probability p of generating a return less than or equal to R, the VaR at level p is R.

It's clear that VaR is related to the distribution of return. (In fact specifying for every value of is the same thing as specifying the return distribution.) But since it's intended to measure risk, it is generally concerned with low-probability, large losses, events that are typically represented as happenings in the left-hand tail of the return distribution.

For real-world examples, we only have to look back a few years. As the events of 2008 unfolded, equity investors began to see a shift towards more frequent and larger negative returns. Then, in September 2008, the MSCI Europe index lost almost 12%. This was an outlier – a return that did not fall within the range of returns observed during the previous five years. And in October 2008, it happened again, and this time the loss was more than 13%. You wait five years for an extreme event, then two come at once!

Or perhaps not. An investor whose expectations were shaped by the previous five years' performance would certainly have experienced the September and October returns as extreme events. On the other hand, a better-informed investor who realized that the world had changed would have seen them as more or less normal returns within a new and much more hostile environment.

The fact that the outlier of September was followed by an even larger outlier in October supports the latter interpretation. Rather than invoking the statistics of extremes, a more modest view might be that risk arises from our own misjudgments. From that perspective, the returns were not extreme events within a normal environment: they were normal events within a new environment we had failed to reckon with.

How does this help us to estimate Value-at-Risk?

The question of how much a portfolio could lose in hostile conditions boils down to

- What sort of market conditions could harm the portfolio, and what would its returns in those circumstances look like?
- How long would it take to recognize that conditions have become hostile, and how quickly thereafter could action take effect?

Simply contemplating these questions brings VaR back into the real world. How we might make a reasonable estimate of potential loss in practice can be illustrated with an example drawn from real life.

Example

Suppose we expect a bond-substitute portfolio that is invested mainly in low-volatility relative-value hedge funds with limited liquidity to return 4% per year with 4% annual volatility.

The pink line in the chart below represents the anticipated mean return path, and the red line marks the 2-standard-deviation lower boundary under these preliminary expectations. An observed cumulative return path that strays below the boundary would suggest that initial expectations are not being met, so we refer to the red line as the *risk signal boundary.*

The portfolio manager's knowledge of the underlying hedge fund strategies enables him to envisage plausible market conditions in which the portfolio might return -8% annually on 4% volatility. In other words, while it may not be all that likely, there is a set of market conditions that constitutes a realistic "risk" scenario associated with the portfolio. If that scenario materializes, the mean return path would follow the blue line shown in the chart.

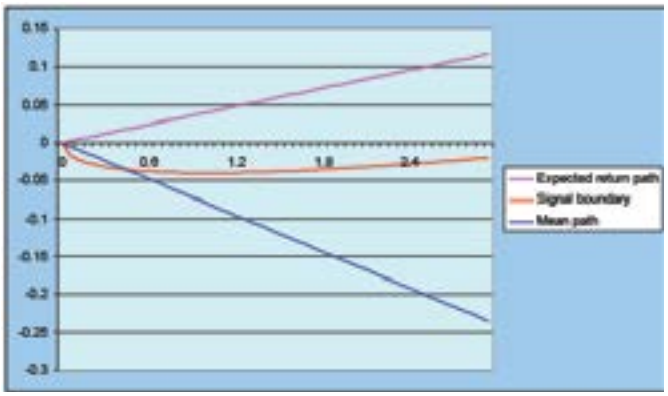

The risk-scenario's mean return path crosses the 2σ signal boundary at about 5 months. At that point we initiate action, but owing to the poor liquidity of the hedge fund investments, it takes effect only after a further year (next chart).

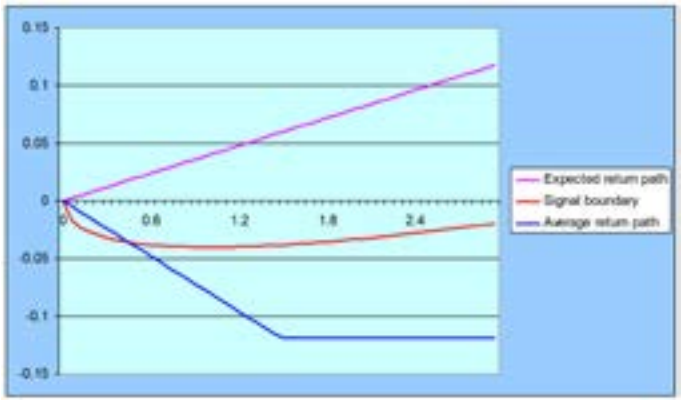

If returns within the -8%; 4% regime follow the average return path, we could expect to lose a total of about 11.5%, made up of 3.5% from inception to signal, plus a further 8% before corrective action takes effect.

But the VaR question has more to do with "How much could I lose *if things go against me?*" than with average loss, so to be safe, the calculation should be based on something worse than the mean path in an unfavorable environment.

The dark blue line on the next chart shows what happens when:

- The actual return path crosses the signal boundary at its lowest point.
- Experience following the signal involves two standard deviations' bad luck[1].

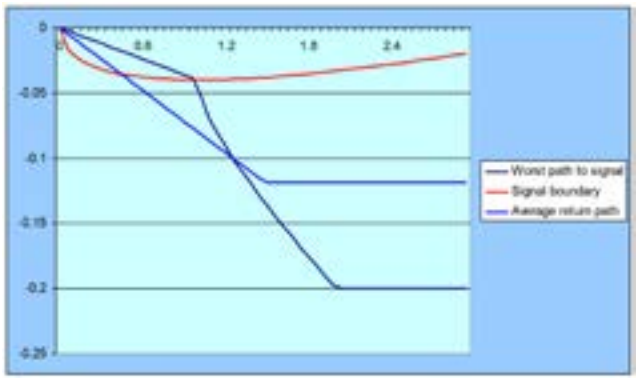

In this case, it takes a year before cumulative return hits the signal boundary — losing 4% in the process — and between Signal and Action we lose a further 16%, for a total loss of 20%, which is clearly unacceptable for a portfolio designed to achieve 4% on 4% volatility[2].

The decomposition of loss into pre-Signal and post-Signal components allows us to identify weaknesses, both in the portfolio, and in the decision-making framework. In terms of the example:

- It doesn't really matter whether the return path crosses the signal boundary at its lower point or somewhere else: the obvious problem is that the time between Signal and effective Action is an entire year, during which we might accrue losses of as much as 16%. The direct solution to that problem is to increase the liquidity of the portfolio.

- We might also find that incurring a loss of around 4% before we even think of taking action is too severe for an ostensibly low-risk portfolio. One solution is to tighten the signal boundary, but that comes at the cost of raising the likelihood of false signals.

- More generally, the idea that a portfolio designed to achieve 4% return on 4% volatility could, under plausible circumstances, annualize at -8% might itself seem inappropriate, in which case the solution is to restructure the portfolio.

Conclusion

The calculation, as described, is obviously a simplification of reality, in that there is typically no single instant at which we decide to take action, and no precise moment at which it takes effect. But the formulation includes the main elements of real risk-management decision procedures, covering questions such as:

- *What market conditions could threaten the portfolio?*
- *How much loss can we tolerate before deciding that observed returns conflict with prior expectations?*
- *Does the liquidity of the portfolio - or our own decision protocol - permit us to act decisively in time to limit losses to an acceptable level?*

These are not easy questions. They require investment expertise and judgment. But they fall squarely within the

---

1 This is not to suggest that return actually follows the path of the dark blue line; it just says that the signal boundary is crossed at its lowest point at time t, and that, at time $t + T$, when action takes effect, things in the meanwhile have gone against us to the tune of two standard deviations.

2 In more conventional terms, a 20% loss over two years on a (4%;4%) portfolio represents a five-standard-deviation event. Even the average loss (with no bad luck) of about 12% is a four-standard deviation event. Each of these numbers is derived here from well-understood, plausible assumptions, but it's hard to imagine that we would have obtained estimates of similar magnitude from a contemplation of the tail risk associated with a (4%;4%) return distribution.

capabilities and experience of investment professionals, and considering them explicitly within the framework of a calculation based on real-world elements helps to restore the responsibility for an investment program to its proper place: in the hands of the program manager.

Nicholas John Macleod, ASA, C.Math., FIMA, is a Principal, Bespoke Financial Modelling, Jersey, Channel Islands. He can be reached at *nicholasjohnmacleod@gmail.com.*

The thoughts and insights shared herein are not necessarily those of the Society of Actuaries, the Investment section of the Society of Actuaries, or corresponding employers of the authors.