Notes from the Vault

Value at Risk: A Valuable Tool That Was Greatly Oversold

David M. Rowe
June 2013

This month I've invited David M. Rowe, one of the speakers at our 2013 Financial Markets Conference, to write about financial risk management. His paper, "Risk Management Beyond VaR," made some important points about value at risk that I think will be appreciated by the readers of Notes from the Vault. *VaR was the main tool to measure risk before the financial crisis, and Rowe explores its limitations in this commentary, to which I've contributed a few minor points.*
—Larry Wall

Value at risk is the most prominent of a set of risk measurement tools to be developed in response to a series of huge, widely publicized losses at large financial firms in the 1980s. Like almost all other risk measures developed over the last 25 to 30 years, value at risk (VaR) relied on classical statistical techniques to measure short-term volatility. Some analysts, such as Nassim Nicholas Taleb, argue that this entire risk measurement enterprise was simply wrongheaded and positively dangerous. I beg to differ. VaR is a good starting point for some risk management discussions. The problems arise when VaR is also the end point for all risk management discussions.

For better or worse, I am old enough to remember the world before VaR. Market risk controls consisted of a complex web of micro position limits. In the fixed income arena, these included:

- Controls on total net duration-adjusted open positions
- Limits on duration-adjusted mismatches at multiple points along the yield curve
- A limit on the sum of the absolute values of such tenor specific mismatches
- Gross position limits
- Issuer concentration limits.

Very importantly, this maze of limits conveyed no instinctive sense of how much risk they allowed. Market risk committees were repeatedly asked for higher limits despite having no real sense of the risks inherent in the limits already in place. In this context, VaR emerged as the first effective communication tool between trading and general management. For the first time it was possible to aggregate risks across very different trading activities to provide some sense of enterprise-wide exposure.

Like all useful innovations, however, VaR had notable weaknesses from the beginning. One weakness was that, inadvertently or deliberately, it was oversold to senior management. Financial risk managers must bear some responsibility for creating a false sense of security among senior managers and watchdogs. For far too long, many were prepared to use the

sloppy shorthand of calling VaR the "worst case loss." A far better alternate shorthand description is to call VaR "the minimum twice-a-year loss." This terminology conveys two things. First, it indicates the approximate rarity of the stated loss threshold being breached. Second, it begs the right question, namely, "How big could the loss be on those two days a year?" To put it bluntly, VaR says nothing about what lurks beyond the 1 percent threshold.

In order to say anything about dangers lying beyond the 1 percent threshold, however, VaR requires some method of estimating the probability of rarely observed events. The standard practice among risk professionals is to use observed returns over some period to estimate a statistical distribution, often the normal distribution. One of the most important lessons Nassim Taleb has driven home is that theoretical statistical methods have an inherent but often unrecognized bias when attention turns to the issue of tail risk.

To be tractable mathematically, statistical distributions, even those with infinite tails, need to have moments that converge. The key characteristic that allows such convergence is rapid attenuation of the probability in the tails. If we measure the standard deviation, skewness or kurtosis of a normal distribution using ever-wider segments of the real number line around zero, these estimates will converge toward limiting values. Absent sufficiently rapid thinning of the tails, moment estimates can diverge indefinitely, being effectively infinite. When we overlay a theoretical distribution on a finite sample, we typically choose a mathematically tractable distribution that "fits" the sample observations we have available based on minimizing some measure such as a squared error penalty function. Thus, by the very act of limiting ourselves to a mathematically tractable distribution, we have implicitly imposed rapidly diminishing probability density in the tails.

Having imposed (or fitted) a theoretical distribution on a finite sample, we then use the tails of that theoretical distribution to make assertions about behavior of the underlying process being examined. As I note in an earlier [article](), Taleb's essential contribution is to hammer home the point that this is both an invalid and a positively dangerous line of reasoning. An important implication of this is that improving our understanding of tail risk will require some difficult cultural changes rather than some minor adjustments to our distributional analysis. Both line and risk managers will need to accept that no single number or even set of numbers can completely summarize a position or a firm's risk exposure.

The key to understanding why no one number can be sufficient is to recognize that the process generating financial returns does not assure that the first, the 1,000th, and the 1 millionth observations are all drawn from the same *underlying stochastic process.* Such a stable process is often a realistic assumption when dealing with physical processes. It is virtually never the case, however, in a social scientific setting. Structural change is the constant bane of econometric forecasters. Such changes are driven by a wide variety of influences, including technological advances, demographic shifts, political upheavals, natural disasters, and, perhaps most importantly, behavioral feedback loops.

Structural change creates a fundamental dilemma for socio-statistical analysis. Classical statistics argues that the more data, the better, since, assuming stochastic stability, this results in smaller estimation errors. For analysis based on time series, however, a larger data set implies incorporation of a greater range of structural changes that undermine the classical assumption of stochastic stability.

This makes it all the more important for risk managers to focus obsessively on what I call "statistical entropy." Like water, information can never rise higher than its source. In the case of information, that source is the set of data on which an analysis is based. In assessing the reliability of any risk estimate, including such things as credit ratings, always start with a review of the volume and quality of the available data. No amount of complex mathematical/statistical analysis can possibly squeeze more information from a data set than it contains initially.

VaR clearly is subject to the constraint implied by "statistical entropy." VaR still can be useful as an easily computed shorthand measure that is employed in day-to-day communication and risk aggregation. VaR is not capable, however, of delivering on managers' need to have a comprehensive understanding of their organization's exposure to loss. An effective process for assessment of loss exposure must be more holistic but also much softer, more amorphous, and less easily defined than the things risk managers do currently. Such a process will require dealing with more unstructured information that is not amenable to precise quantification. Inputs from country risk officers, industry analysts, and macroeconomists must be integrated into regular deliberations about risk. The success of such a process also will require senior managers to abandon the comfortable idea that all forms of risk, including fundamental Knightian uncertainty, can be reduced to a single summary statistic like VaR.

If organizations are to have a reasonable chance to avoid the worst effects of the next crisis, executives and board members must be willing to devote the time and energy to grapple with risk in all its messy multidimensionality.

David M. Rowe is the founder and president of David M. Rowe Risk Advisory. The views expressed here are the author's and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. If you wish to comment on this post, please e-mail atl.nftv.mailbox@atl.frb.org.

http://www.frbatlanta.org/cenfis/pubscf/nftv_1306.cfm