

Opinion

10 March 2009

To VaR or not to VaR?



Value-at-Risk (VaR) is appropriate and effective for its proper purpose - but it addresses only one of the two key challenges of financial risk management, argues David Rowe.

The current economic crisis has produced considerable soul searching on many fronts. This has been especially evident in the field of financial risk management. Particular attention has been paid to the subject of Value-at-Risk (VaR) as a tool for risk estimation and reporting. This discussion has special relevance for the insurance industry since VaR is an element of the proposed Solvency II framework.

Some have condemned VaR as a fundamentally flawed concept that has been a contributor to risk rather than an effective means for managing it. To see why I take issue with this view requires a bit of history. Before the early 1990s, traders were constrained by a complex array of micro-limits on:

- Total gross open positions
- Gross open positions at a variety of fixed maturities
- Total absolute open positions summed across all maturity buckets
- Limits across multiple option types on
 - Delta
 - (Negative) gamma
 - Vega (possibly asymmetric limits on long and short vega)

Market risk policy committees faced a steady stream of recurring requests for increases in these limits. Unfortunately there was no meaningful way for such committees to weigh the amount of risk implicit in the existing structure of these limits let alone what further risk was implied by the requested increases. An experienced risk manager named Aaron Brown captured it well when he recently told me that VaR was the first effective means for communicating risk implications between traders and general managers. That was the voice of someone who had experience in the pre-VaR world and recognizes what an advance VaR represented.

Not the "worst case loss"

Like so many advances, however, VaR became over-hyped and over-used. Financial risk managers must bear some of the blame for the ensuing criticism that VaR created a false sense of security among senior managers and watchdogs. For *much* too long many were prepared to use the sloppy shorthand of calling VaR the "worst case loss." The belief seemed to be that everyone would realize this was not to be taken literally, that what was *really* meant was "a loss that will only be exceeded two or three times a year." Unfortunately, that belief was often unfounded. The problem with calling VaR a "worst case loss" is that many people will take you at your word. General business executives and the general public will frequently believe exactly what is said, that VaR is the most that can be lost in any given day. A far better alternate shorthand description is to call VaR "the minimum twice-a-year loss." This terminology conveys two things. It indicates the approximate rarity of the stated loss threshold being breeched but it also begs the right question, namely "How big could the loss be on those two days a year?" The answer is, VaR says *nothing* about what lurks beyond the 1% threshold.

Another problem is that VaR fell victim to a corrosive feedback process. This point has been well made by my SunGard colleague Till Guldemann, sometimes called the "father of VaR¹". When VaR became the standard metric for measuring and monitoring the limits on market risk taken by traders, both individually and collectively, they had no choice but to comply. Traders who repeatedly and willfully exceed their institutionally established limits will ultimately be fired.

"VaR says nothing about what lurks beyond the 1% threshold."

Nevertheless, traders still want to make their returns. It is hardly a big leap to realize that one way of doing this is to pile on risk in the tail of the loss distribution. This opens a firm up to what are known as Low Probability High Impact events. Because the probability of an occurrence falls below the usual 1% VaR threshold, such positions have little or no impact on a VaR-based risk measure. The corrosive feedback effect is that the widespread use of VaR as a control metric has encouraged exactly the type of risk-taking that VaR fails to measure, namely exposure to extreme events. Hence VaR doesn't just fail to address the most extreme losses, it actually encourages behavior that increases their magnitude.

In my view, the problem is that risk managers should feel responsible for two quite different tasks. I characterize these as:

1. Making sure one's institution does not die a death of a thousand cuts in normal market conditions.
2. Protecting against the potentially lethal impact of a truly catastrophic set of external events.

Despite all the criticism that has been levelled at VaR, it works remarkably well as a tool to fulfill the first of these two tasks. As we see currently, however, the second task addresses a very real concern that cannot be ignored. It can only be addressed by a series of actions that loosely travel under the title of stress testing. The appropriate share of risk management

effort allocated to these two tasks is open to debate and is unlikely to be the same for all institutions. Nevertheless, a significant amount of resources, including staff time, computing resources, and senior management attention is required if the dangers of catastrophic events are to be avoided.

More stress-testing needed

It is clear from my experience that far too little attention has been paid to stress testing and scenario analysis in many financial institutions despite the millions of dollars devoted to data gathering and preparation of VaR estimates. In some cases, much of the work that has been done could be leveraged for the purpose of catastrophe avoidance. In others, however, shortcuts have been taken to prepare VaR estimates with reasonable accuracy while avoiding the cost necessary to gather the level of detailed information needed for effective stress analysis.

In essence, financial risk managers have two tasks but often they have attended only to the first of these. In part this may be the result of restrictive institutional mandates. Beyond that, however, I suspect it has to do with the much more amorphous and ill-defined character of addressing catastrophic risk. Too often this is viewed as requiring no more than some minor tweaking to existing VaR analysis. Some argue for use of VaR excess (the average of all the losses beyond the 1% threshold) instead of VaR itself. Others argue for inclusion of distributions with much fatter tails or with jump diffusion² characteristics. In my view these approaches offer little that is helpful. They may yield bigger risk numbers, but do little or nothing to enrich an organization's understanding of its potentially lethal vulnerabilities.

"The corrosive feedback effect is that widespread use of VaR as a control metric has encouraged exactly the type of risk-taking that VaR fails to measure, namely exposure to extreme events."

A successful stress testing and scenario analysis program must be fundamentally different in kind from traditional VaR estimation. One approach I have advanced is called the "stress testing trident." It involves three broadly different lines of attack. These are:

1. simulating the market's greatest disasters ("the market's greatest hits")
2. defining and then stressing the most serious current vulnerabilities ("the Achilles heel approach")
3. using imagination based on social, geopolitical and economic inputs to formulate plausible crisis scenarios for investigation ("the imagination approach")

The first approach involves defining stress scenarios that replicate the relative changes in all applicable market variables for selected historical events. Typical events to be included might be, to name a few:

- The October 1987 stock market crash in the US
- Britain's forced withdrawal from the European Exchange Rate Mechanism (ERM) in September 1992
- Selected dates during the Asian currency crisis of 1997/98

- The Russian debt crisis and devaluation of August 1998

The big advantage of this exercise is that no one can defend the position that "this scenario could never happen" because it has. The big drawback is that the market movements being simulated usually have nothing to do with the vulnerabilities of the current trading positions. While such simulations may alleviate the anxiety of some who lived through the trauma of these events, they represent a scatter shot approach that is not guaranteed to highlight current "worst case" losses.

"Pessimisation"

The second approach involves what I have termed "pessimisation" and others call constrained loss maximization. The idea is to examine the existing portfolio in a systematic way to define its particular vulnerabilities and then construct stress scenarios that exploit these vulnerabilities to the full³. Among other things, this type of exercise can reveal cases where traders are systematically writing large volumes of out-of-the-money options. Often this will not become obvious in standard VaR results without analyzing the market scenarios that generate losses beyond the 1% cut-off point.

The third approach is to use subjective assessment of current socioeconomic and geopolitical conditions to define dangerous scenarios. This requires thinking through both the initial and potential secondary effects of a hypothetical disaster. Like the "market's greatest hits", this approach fails to tailor the scenarios to current vulnerabilities. In contrast, however, it is forward looking and driven by current external conditions.

An exercise based on imagination also is useful in forcing an assessment of secondary implications that may not be immediately obvious. Too much of what passes for stress testing boils down to comparative static analysis. Shift the yield curve by 300 or 400 basis points and take all credit ratings down two or three notches and see what happens. Real crises, such as the one we are living through at the moment, don't unfold that way. Rather they play out dynamically over time, with one shock triggering something that leads to something else in an unfolding series of cause and effect linkages.

Thinking through how such a series of causal relationships might unfold requires the combined input of many different disciplines involving a wide range of staff from a variety of functions across the institution. This involvement of people with a wide range of differing perspectives is an important strength of the approach. Done well it both enriches the range of linkages to be considered and stimulates thought about how to respond. Such advance thinking about consequences and potential responses can facilitate faster reaction in the midst of an actual crisis when speed is of the essence⁴.

Effective stress testing is a bit like trying to cure the common cold. Like infectious viruses, crises come in too many varieties to allow a single silver bullet solution. Nevertheless, applying all three approaches described here can do a great deal to limit exposure to a crisis and to respond more effectively when a crisis does occur.

The lessons

So what is the answer to the question "To VaR or not to VaR?" I conclude that VaR is an essential and effective tool for meeting the first of risk management's two tasks, namely avoiding a death of a thousand cuts. The danger comes in forgetting about the second task (or devoting far too little time and too few resources to it.) VaR is effectively silent about what "lurks beyond the 1% threshold." Only a messy, qualitative, judgmental and somewhat unsatisfying process of grappling institutionally with potential crisis scenarios and their impact can set the stage for prompt action when low probability but very high impact events materialize.

David M. Rowe, Ph.D., is EVP for Risk Management, SunGard. He is a monthly columnist for Risk magazine and contributes to many other publications. See:

<http://www.sungard.com/enterpriserisk>

e-mail: david.rowe@sungard.com

References

1. Till was at JP Morgan when they developed their VaR system in the early 1990s and was instrumental in persuading the bank to make its data and methodology freely available on the internet, thereby hastening the adoption of VaR as a standard approach.
2. A jump diffusion process can be thought of as a combination of two distributions, one well behaved and quite "predictably random" and a second latent distribution which is dormant most of the time but occasionally erupts in a change that swamps the behaviour of the well behaved distribution.
3. See Rowe, D., *Building Pessimized Scenarios*, Risk, January 2006, p. 90.
4. For an excellent example of stress testing based on imagination see, *Financial Risk Outlook 2007*, UK Financial Services Authority, February 1, 2007 available at http://www.fsa.gov.uk/pubs/plan/financial_risk_outlook_2007.pdf.