Article from:

Risks and Rewards Newsletter

March 1998 – Issue No. 30

Measuring Financial Risk

The Layperson's Introduction to Value at Risk

by Barry Schachter

Editor's Note: *The following article originally appeared in the August 1997 issue of* Financial Engineering News *and is reprinted with permission. Free subscriptions to this publication are available by visiting the* Financial Engineering News *web site at www.fenews.com.*

Value at risk (VAR) is much on the minds of risk managers and regulators these days, because of the promise it holds for improving risk management. It is common to hear the question asked, could VAR have prevented Barings, or Orange County, or Sumitomo? No answer to questions of that sort will be attempted here. Instead, this article will take a normative approach. My purpose is more modest, namely, to provide the reader with some background by describing VAR and its evolving role in risk management. Because of its technical nature, it is customary to begin any discussion of VAR with a definition. I offer three equivalent definitions:

- A forecast of a given percentile, usually in the lower tail, of the distribution of returns on a portfolio over some period, similar in principle to an estimate of the expected return on a portfolio, which is a forecast of the 50th percentile.
- An estimate of the level of loss on a portfolio which is expected to be equaled or exceeded with a given, small probability.
- A number invented by purveyors of panaceas for pecuniary peril intended to mislead senior management and regulators into false confidence that market risk is adequately understood and controlled.

The Quest for the "Holy Scale"

Folklore (if it is fair to attribute as folklore that which only dates back five years) tells us that VAR was developed to provide a single number which could encapsulate information about the risk in a portfolio, could be calculated rapidly (by 4:15), and could communicate that information to nontechnical senior managers. Tall order, and not one that could be delivered upon with compromises.

Modern portfolio theory (MPT), as taught in business schools, tells us that the risk in a portfolio can be proxied by the portfolio standard deviation, a measure of spread in a distribution. That is, standard deviation is all you need to know in order to (1) encapsulate all the information about risk that is relevant, and (2) construct risk-based rules for optimal risk "management" decisions. (The more technically proficient will please forgive my playing somewhat fast and loose with the theory in the interests of clarity.) Strangely, when applied to the quest for the Holy Scale, standard deviation loses its appeal found in MPT. First, managers think of risk in terms of dollars of loss, whereas standard deviation defines risk in terms of deviations (!), either above or below, expected return and is therefore not intuitive. Second, in trading portfolios deviations of a given amount below expected return do not occur with the same likelihood as deviations above, as a result of positions in options and option-like instruments, whereas the use of standard deviation for risk management assumes symmetry.

An alternative measure of risk was therefore required. Why not measure the spread of returns, then, by estimating the loss associated with a given, small probability of occurrence? Higher spread or risk should mean a higher loss at the given probability. Then senior management can be told that there is a 1 in 100, say, chance of losing *X* dollars over the holding period. Not only is this intuitively appealing, but it's easy to show that when returns are normally distributed (symmetric), the information conveyed is exactly the same as where standard deviation is employed, it's just that the scale is different. This approach can be consistent with MPT. It seems then that perhaps the Holy Scale has been found in VAR.

The Slip 'twixt Cup and Lip

It's perhaps too easy to criticize efforts to implement the VAR concept. It takes some courage to venture into unfamiliar terrain and missteps are inevitable. The VAR paradigm is still evolving (as is that of financial risk management in general), and experimentation should be encouraged. To speak of "best practices" is surely premature.

The general approaches to VAR computation have fallen into three classes—parametric, historical simulation, and Monte Carlo. Parametric VAR is most closely tied to MPT, as the VAR is expressed as a multiple of the standard deviation of the portfolio's return. Historical simulation expresses the distribution of portfolio returns as a bar chart or histogram of hypothetical returns. Each hypothetical return is calculated as that which would be earned on today's portfolio if a day in the history of market rates and prices were to repeat itself. The VAR then is read from this histogram. Monte Carlo also expresses returns as a histogram of hypothetical returns. In this case, the hypothetical returns are obtained by choosing at random from a given distribution of price and rate changes estimated with historical data. Each of these approaches has strengths and weaknesses.

"The general approaches to VAR computation have fallen into three classes—parametric, historical simulation, and Monte Carlo."

The parametric approach has as its principal virtue speed in computation. The quality of the VAR estimate degrades with portfolios of nonlinear instruments. Departures from normality

continued on page 18, column 1

Measuring Financial Risk
continued from page 17

in the portfolio return distribution also represent a problem for the parametric approach. Historical simulation (my personal favorite) is free from distributional assumptions, but requires the portfolio be revalued once for every day in the historical sample period. Because the histogram from which the VAR is estimated is calculated using actual historical market price changes, the range of portfolio value changes possible is limited. Monte Carlo VAR is not limited by price changes observed in the sample period, because revaluations are based on sampling from an estimated distribution of price changes. Monte Carlo usually involves many more repricings of the portfolio than historical simulation and is therefore the most expensive and time-consuming approach.

Rule or Tool?

It seems that VAR is being used for just about every need: risk reporting, risk limits, regulatory capital, internal capital allocation, and performance measurement. Yet, VAR is not the answer for all risk management challenges. No theory exists to show that VAR is the appropriate measure upon which to build optimal decision rules. VAR does not measure "event" (for example, market crash) risk. That is why portfolio stress tests are recommended to supplement VAR. VAR does not readily capture liquidity differences among instruments. That is why limits on both tenors and option greeks are still useful. VAR doesn't readily capture model risks, which is why model reserves are also necessary.

Because VAR does not capture all relevant information about market risk, its best use is as a tool in the hands of a good risk manager. Nevertheless, VAR is a very promising tool—one that will continue to evolve rapidly because of the intense interest in it by practitioners, regulators, and academics.

Barry Schachter is Vice President, Market Risk Portfolio Manager at Chase Manhattan Bank. He maintains a web page of VAR resources at http://pw2.netcom.com/~bschacht/varbiblio.html.

Seven Quantitative Insights into Active Management—Part 5

Data Mining Is Easy

by Ronald N. Kahn

Editor's Note: *The following article originally appeared in the Winter 1998 issue of the Horizon, a publication of BARRA, Inc., and is reprinted with permission.*

Why is it that so many strategies look great in backtests and disappoint upon implementation? Backtesters always have 95% confidence in their results, so why are investors disappointed far more than 5% of the time? It turns out to be surprisingly easy to search through historical data and find patterns that don't really exist.

To understand why data mining is easy, we must first understand the statistics of coincidence. Let's begin with some noninvestment examples. Then we will move on to investment research.

The Statistics of Coincidence

Several years ago Evelyn Adams won the New Jersey state lottery twice in four months. Newspapers put the odds of that happening at 17 trillion to 1, an incredibly improbable event. A few months later, two Harvard statisticians, Percy Diaconis and Frederick Mosteller, showed that a double win in the lottery is not a particularly improbable event. They estimated the odds at 30 to 1. What explains the enormous discrepancy in these two probabilities?

It turns out that the odds of Evelyn Adams winning the lottery twice are in fact 17 trillion to 1. But that result is presumably of interest only to her immediate family. The odds of someone, somewhere, winning two lotteries— given the millions of people entering lotteries every day—are only 30 to 1. If it wasn't Evelyn Adams, it could have been someone else.

Coincidences appear improbable only when viewed from a narrow perspective. When viewed from the correct (broad) perspective, coincidences are no longer so improbable. Let's consider another noninvestment example: Norman Bloom, arguably the world's greatest data miner.

Norman died a few years ago in the midst of his quest to prove the existence of God through baseball statistics and the Dow Jones average. He argued that "BOTH INSTRUMENTS are in effect GREAT LABORATORY EXPERIMENTS wherein GREAT AMOUNTS OF RECORDED DATA ARE COLLECTED AND PUBLISHED" (capitalization Bloom's). As but one example of thousands of his analyses of baseball, he argues that the fact that George Brett, the Kansas City third baseman, hit his third home run in the third game of the playoffs, to tie the game at 3–3, could not be a coincidence—it must prove the existence of God. In the investment arena, he argued that the Dow's 13 crossings of the 1,000 line in 1976 mirrored the 13 colonies which united in 1776—which also could not be a coincidence. (He pointed out, too, that the 12th crossing occurred on his birthday, deftly combining message and messenger.) He never took into account the enormous volume of data—in fact, an entire New York Public Library's worth—he searched through to find these coincidences. His focus was narrow, not broad.

With Norman's passing, the title of world's greatest living data miner has been left open. Recently, however, Michael Drosnin, author of *The Bible Code,* seems to have filled it.

The importance of perspective to understanding the statistics of coincidence was perhaps best summarized by, of all people, Marcel Proust—who often showed keen mathematical intuition:

> The number of pawns on the human chessboard being less

continued on page 19, column 1