Appendix 9.1

What Are the Problems with Value at Risk?

Introduction

Value at Risk (VaR) is a statistical measure of market risk,[1] which is used by institutions (mainly financial institutions) that want to determine the vulnerability of their actively traded portfolios to changes in market prices and returns. Calculating VaR could be a difficult chore, but for many companies, it is as easy and as quick as sending a short e-mail message because services, like Bloomberg, have built-in programs that can value a company's portfolio and calculate VaR statistics in just a few seconds.[2]

VaR is a compelling statistic because it is easy to understand, transmits a very strong message about a company's vulnerability to changes in the market prices of financial assets, and is well understood throughout the risk management community. Of course, there are problems with using the past to predict the future.

How Is VaR Calculated?

There are three main ways to calculate VaR: the correlation (variance/covariance) method, historical method, and simulation method.[3] The *correlation method* is the quickest and easiest to use because all that is needed to make these estimates are the mean, standard deviation, and cross-correlations among the asset returns in a portfolio. Most of these figures are available electronically through RiskMetrics.[4] Implicit in the use of the correlation method is the assumption that the distribution of all possible fluctuations in prices and yields is normally distributed (see Exhibit A9.1.1).

[1]Notice that VaR is a one-sided test. It is only worried about potential losses and not potential gains.
[2]Bloomberg offers this service to its customers, and its estimates are based on the RiskMetrics variance and covariance estimates.
[3]The three methods mentioned here are the major ways to calculate VaR. There are other, less well used, methods, such as Closed Form VaR and Delta-Gamma VaR.
[4]VaR was started by J.P. Morgan under the name RiskMetrics. RiskMetrics is credited with popularizing VaR, and it has made the job of calculating this statistic easier by providing free access to the financial data needed for the analyses. See http://www.riskmetrics.com. Accessed 20 January 2008.

Exhibit A9.1.1 Standard Normal Distribution

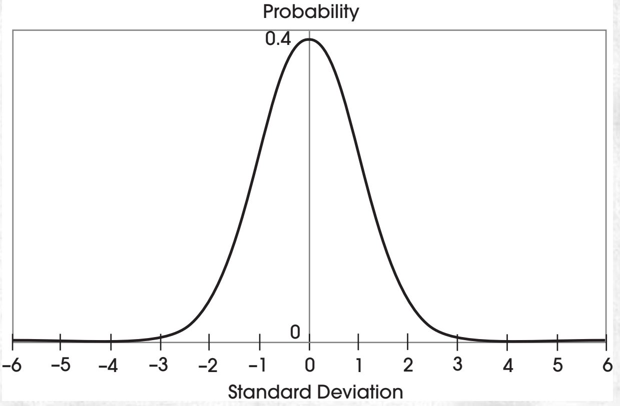

One problem with using a standard normal distribution as the basis for this analysis is that the real world is not so easily described. Like the drunk looking for his keys under the lamppost, this distribution is used because that is (statistically) where the light is. Analysts hope the true state of the world will be close enough to the standard normal distribution to make predictions based on this assumption credible.

Ask yourself, is the world depicted by Exhibit A9.1.1 the one in which *we* live? Over the past 100 years, events have occurred multiple times (e.g., stock market crashes in 1929, 1987, and 1997, two World Wars, the Value Jet and Swissair plane crashes, Three-Mile Island and Chernobyl nuclear accidents, and the Challenger and Columbia space shuttle disasters) that a normal distribution would predict should happen, at most, once in a millennium.[5] If we lived in a normally distributed world, 99.9% of everything that happened in life would occur about three standard deviations from the "average" humdrum day.[6] The problem is the real world has fat tails, which means that many more good things and bad things happen in real life than a standard normal distribution would predict. Therefore, the existence of fat tails causes analyses, which are based on the normal distribution, to underestimate the true level of risk.

The *historical method* is popular because it makes no assumptions about the distribution of future price movements. Rather, this method uses actual historical prices to calculate how the current portfolio's value would have changed over a given period (e.g., a day or a week) in the past. Once the results are

[5]For an enlightening discussion of the relatively high incidence of disasters happening in highly complex organizations despite very low probabilities of their occurrences, see Charles, Perrow, *Normal Accidents: Living With High-Risk Technologies*, Princeton University Press, NJ, 1999.

[6]More specifically for one-tailed tests, 99.9% would be 2.56 standard deviations, 99% would be 2.33 standard deviations, 95% would be 1.65 standard deviations, and 83.5% would be 1.00 standard deviation to the left of the mean.

calculated, they are rank-ordered and cut-off points (e.g., 99% level or 95% level) are identified. The major disadvantage of the historical method is that it is time consuming and begs the question whether the sample period chosen for the analysis is appropriate (e.g., not too long or too short).

The *simulation method* builds a model of the relationships among the market values in a portfolio and then runs repeated and random iterations (scenarios) of the model to determine the most likely outcomes. Clearly the most labor intensive of all three methods, the simulation method is also the most flexible, especially for multifaceted portfolios. Nevertheless, in spite of its flexibility and sophistication, it has not been shown to be the most accurate.

In general, if money and time were not issues, then it would be wise to estimate VaR using all three methods, but money and time are almost always issues; so, most companies use just one method.

Problems With VaR

Value at risk can help to control risks, but improper use of these estimates can lead to poor risk-management decisions, and focusing too narrowly on this one measure alone can lead to serious mistakes. A full understanding of VaR implies knowing what the measure shows and what it hides. Among the major problems are the following:

1. *VaR provides no information about the maximum losses that could occur.*

VaR provides information on the *minimum* that a company might expect to lose at threshold levels of confidence (e.g., 99% or 95% level of confidence). It says nothing about the *maximum* that could be lost. One way to get an idea of the risks in the uncharted territory of the distribution tails is to stress test VaR valuations by bombarding them with the worst of all price and return possibilities to see what could happen. Another technique is to *back test* the VaR results to validate that what has actually happened is in sync with what VaR predicted. In this way, results that were clear outliers could be identified and precautions could be taken. Another alternative is to derive VaR statistics using distributions that are not normally distributed (e.g., distributions that have fat tails), but perhaps the best advice is to always interpret VaR conclusions with a heavy dose of common sense.

2. *Volatility may not be constant over time.*

VaR estimates that use the Correlation Method assume the standard deviation (i.e., volatility) is constant over time; so, distant past variations are weighted with the same importance as nearby variations. The problem with this assumption is that statistical evidence indicates that the standard deviations of returns on assets are not constant over time, and many can vary considerably, which means that more accurate measures of volatility might be obtained by finding the statistically most significant past changes and weighing them more heavily than others.[7]

3. *Discrete movements of prices and structural shifts in the economy reduce the relevance of VaR statistics.*

VaR assumes the world moves in very small steps from one equilibrium to another. When the changes are significant, VaR results deteriorate, and the larger the steps, the greater the deterioration. The Crashes of 1929 and 1987, *Tequila* (Peso) Crisis, terrorist attacks on the Twin Towers, Argentine and Brazilian economic crises, Venezuelan leadership debacle, Asian Tiger crisis, and Russian debt default are just a few examples of when world markets took discontinuous, quantum stumbles. As well, structural shifts in the economy weaken VaR results because future changes in prices and returns are not as closely correlated to historical patterns and relationships.

4. *VaR assumes the assets in a portfolio can be easily liquidated.*

VaR assumes that the assets in a portfolio can be liquidated at will and without substantially moving market prices in disadvantageous directions. Illiquidity increases risks, and these risks are not reflected in the VaR estimates.

5. *VaR does not account for credit risk.*

None of the major methods used to calculate VaR includes an imputed value for credit risk. To the extent counterparties might not repay or might delay paying, or to the extent a large failure might set off a capital market meltdown, this omission could be serious because potential losses are underestimated.

[7]A popular method of addressing this problem is to calculate time-varying volatilities using a method called GARCH, which is an abbreviation for Generalized Auto-Regressive Conditional Heteroscedacity. For example, a GARCH estimate of today's volatility would combine yesterday's prediction of volatility and yesterday's squared return. Therefore, if yesterday's return and volatility were large, the estimate for today's volatility would be large. The technique for estimating today's volatility and then calculating VaR using GARCH is complicated, but, in general, it uses the same method as constant standard deviation. The difference is that calculations of constant standard deviation assume there is no systematic relationship between past and present volatilities. GARCH uncovers these relationships, if they exist, and allows the model to predict future volatilities based on past volatilities. In spite of its sophistication, GARCH is still only an extrapolation, which means it does little to predict future inflection points or discrete changes in economic activity. As well, if the analysis is for a long period of time, GARCH results tend to converge to the constant volatility result; so, you are back to square one. See Tim Bollerslev, "Generalizing Autoregressive Conditional Heteroskedacity," *Journal of Econometrics*, 31 (1986), 307–327, and see Robert F. Engle, "Autoregressive Conditional Heteroskedacity with Estimates of the Variance of U.K. Inflation," *Econometrica*, 50 (1982), 987–1008.