Chapter 3

A Primer on Value at Risk

Michael Minnich
Vice President
Capital Market Risk Advisors, Inc.

INTRODUCTION

The accelerating trend towards measuring and monitoring risk at a firm-wide level has increased the focus on Value at Risk (VAR) and the need for a consistent firm-wide approach. Although VAR is only one of many both quantitative and qualitative factors that should be incorporated into a cohesive risk measurement and risk management approach, it is an extremely important one. Exhibit 1 shows the key components of a risk management framework.

Practitioners, regulators, and academics have embraced VAR, and many view VAR as a vital component of current "best" practices in risk measurement. VAR can be defined as the maximum loss a portfolio is expected to incur over a specified time period, with a specified probability.[1] Exhibit 2 illustrates this point. The shaded area to the left of the graph shows that there is a 5% probability that a sample portfolio of actually traded securities will lose more than $8.2 million ($100 million – $91.8 million) over the next one week period.

One of the most important aspects of VAR is that unlike scenario analysis or stress testing which shows what loss would occur given a certain scenario, VAR actually assigns a probability to a dollar amount of loss occurring. This probability and its corresponding loss amount (5% and $8.2 million in the example above) are not associated with any particular event, but encompass any event that would cause such a loss.[2] It is important to remember that VAR is not the *maximum* loss that will occur, but only a loss level threshold that will be pierced some percentage of the time (5% in the example above). The actual loss that occurs could be much higher than the VAR.

---

[1] In this chapter the term probability is used for purposes of clarity only. "Confidence level" is a more appropriate term that statisticians or mathematicians would prefer.

[2] Any loss figure generated from a VAR calculation can only anticipate losses resulting from events encompassed by the model chosen and that fall within the assumptions made. For example, a VAR that only measures losses due to market risk (such as Riskmetrics) will not capture credit losses. Likewise, most VAR models will not capture losses due to volatility movements.

The author would like to thank his colleagues at Capital Market Risk Advisors, Inc. for their valuable suggestions and help.

Exhibit 1: Key Components of a Risk Management Framework

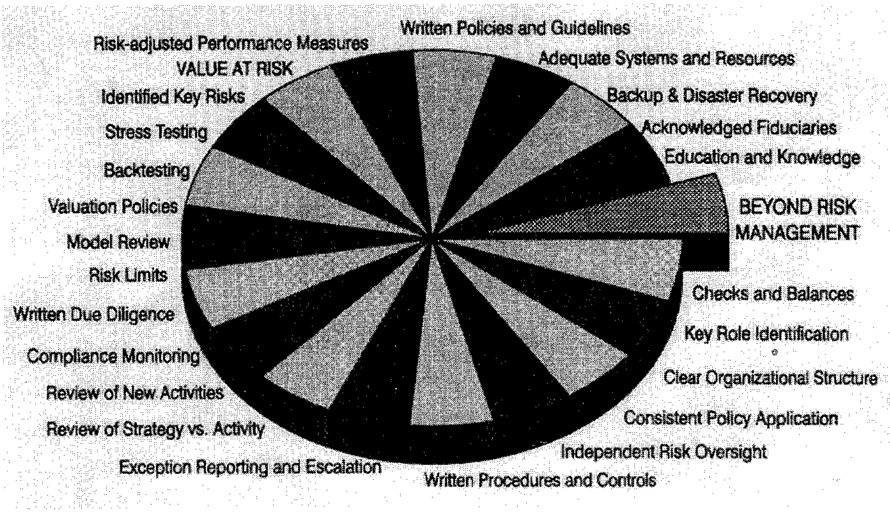

Copyright, Capital Market Risk Advisors, Inc.

Exhibit 2: Illustration of Value at Risk Probability Distribution of Portfolio Value

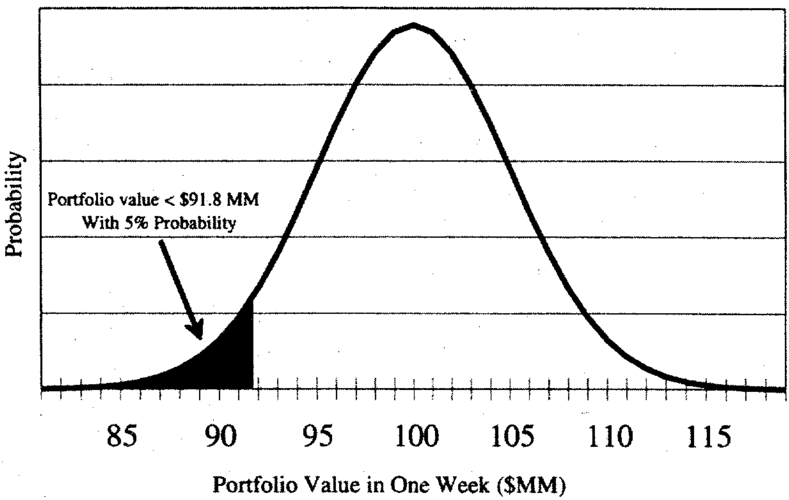

Despite all of the support for VAR, it is not a sufficient parameter for risk measurement. VAR is after all only one number from a rich distribution of return information. VAR should be used in conjunction with other risk measures such as scenario testing, stress testing, and other asset/business specific risk measures.

TYPES OF VAR

One of the most difficult aspects of calculating VAR is selecting from among the many types of VAR methodologies and their associated assumptions. Depending on the organization, some of these decisions can be straightforward and clear, but more often than not, certain trade-offs will be made.

There are three main categories of VAR whose primary distinction is the type of calculations performed. These categories are summarized in Exhibit 3. Below we will illustrate how each of these VARs is calculated along with the strengths and weaknesses of each approach. But first we will consider the necessary decisions and assumptions that need to be made regardless of which of the three methodologies is chosen.

VAR CHOICES

The following are decisions that must be decided prior to calculating VAR: (1) time horizon; (2) confidence interval; (3) data series; (4) mapping/selecting relevant risk factors; and, (5) option valuation. Care should be taken when considering each, as the choices made can change not only the actual number,[3] but also the uses and meaning of the VAR number itself.

Exhibit 3: Main Categories of VAR

Methodology	Calculations involved
Variance-Covariance	Volatility and correlation matrix Matrix algebra to arrive at VAR
Monte Carlo Simulation	Volatility and correlation matrix Monte Carlo simulation to generate portfolio return distribution
Historical Simulation	Historical Data set Simulate portfolio using historical returns as actual return distribution

[3] Tanya Styblo Beder has clearly demonstrated that the different choices can change the VAR by up to 14 times for some portfolios. See Tanya Styblo Beder, "VAR: Seductive but Dangerous," *Financial Analysts Journal* (September-October 1995), pp. 12-24.

Exhibit 4: Considerations in Selecting a Time Horizon

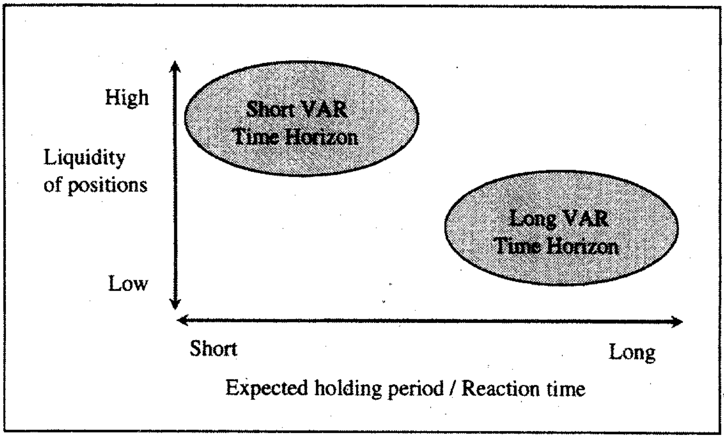

Time Horizon

The choice of time horizon depends upon the objectives of the portfolio and liquidity of its positions. Typically for trading and market making operations, VAR is computed using a one day, one week, or a two week time horizon. However, longer time horizons are often used for proprietary trading firms, institutional investors, and corporations.

Two of the most important considerations for selecting a time horizon are the liquidity of the instruments and the expected holding period. Exhibit 4 illustrates that firms who hold highly liquid securities (such as swaps) for short amounts of time are better suited for a short VAR time horizon, while firms who hold illiquid securities (such as real estate) with a long expected holding period would be wise to choose a longer VAR time horizon.

Confidence Interval

The confidence interval defines the percentage of time that the firm should not lose more than the VAR amount. Commonly used confidence intervals range from 90% to 99%. The actual choice of confidence interval is not as important as understanding the implications of the choice and ensuring that limits are set accordingly. The Bank for International Settlement and the Derivatives Policy Group recommend a confidence level of 99%, while research shows that 95% performs best under back-testing due to "fat tails." The term "fat tails" refers to the fact that large market moves occur more frequently than what would occur if market returns were normally distributed. Despite this fact, many market practitioners and academics assume that returns are normally distributed. This assumption, although usually insignificant, can cause problems in VAR calculations on some asset classes.

Exhibit 5: Normal versus "Fat Tails" Distribution

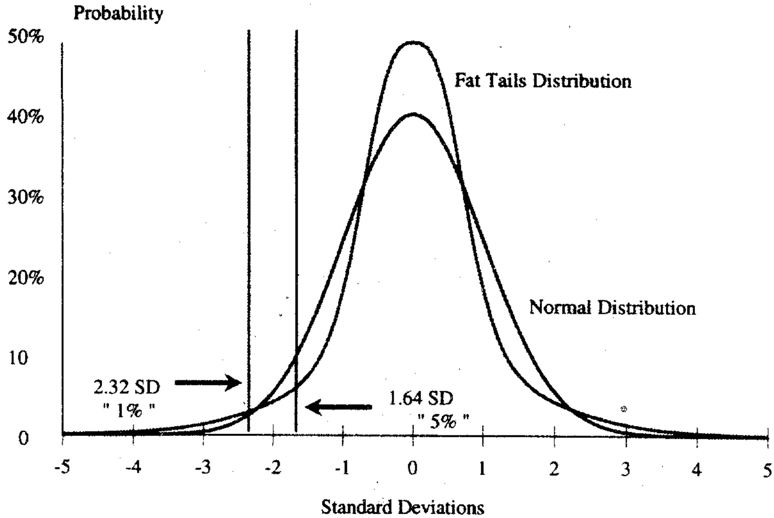

Exhibit 5 illustrates how "fat tails" can cause problems in calculating VAR at higher confidence levels — the VAR should be further to the left for a distribution with "fatter tails" than a normal distribution. As discussed later in this chapter, historical VAR can capture "fat tailed" effects of a return distribution.

Data Series

By most accounts, VAR is fairly data intensive. The choice of historical, implied, or other ways to determine security relationships is important, but typically there is very little choice. Some argue that using implied correlations and volatilities results in a better predictor of risk than historical correlations and volatilities, but very little implied data are available. As a result, historical data sets have become commonplace in VAR calculations, thereby driving the need for additional data decisions to be made.

How much historical data should be used? Longer periods of data have a richer return distribution while shorter periods allow the VAR to react more quickly to changing market events. Three to five years of historical data are typical. In addition, the role of outliers in the data set needs to be considered. Should the Persian Gulf War effects on the oil market be excluded when looking at historical data? Some market participants believe that it should not be excluded because this event reflects real history and adds to the "fat tailed" richness of a data series. Others argue that it should be excluded because the border of its inclusion versus exclusion could imply very different VARs. For instance, consider if a firm were using 10 years of historical data to calculate VAR on an equity portfolio. In October 1997 when the 1987 crash data point falls out of the data set, a risk manager will most likely see a decrease in VAR that has nothing to do with the firm's actual risk, but just the exclusion of one data point due to the passage of time.

One common method that solves both of the above issues is to use exponentially weighted data. Exponential weighting gives more recent data more weight, allowing the VAR to react to changing market conditions quickly. Exhibit 6 illustrates this fact using the Persian Gulf War period as an example. In addition, all of the available historical data are used and hence there is not a discrete trailing point of historical data that would cause the 1987 crash problem mentioned above.[4]

Exhibit 7 shows how equally weighted and exponentially weighted VAR can actually move in different directions due to data issues. Depicted in Exhibit 7 is the VAR of a constantly rebalanced emerging markets portfolio over the course of a 2-year period following the Mexican peso crises. Note that the exponential VAR increases in March 1996, while the equally weighted VAR decreases dramatically during the same period. The increase in the exponentially weighted VAR is caused by the few weeks prior to March 1996 being more volatile than average, while the equally weighted decrease is due solely to the peso crash data point falling out of its one year data set.

Exhibit 6: Exponential versus Unweighted Volatility

S&P 500 Volatility during the Persian Gulf Crisis

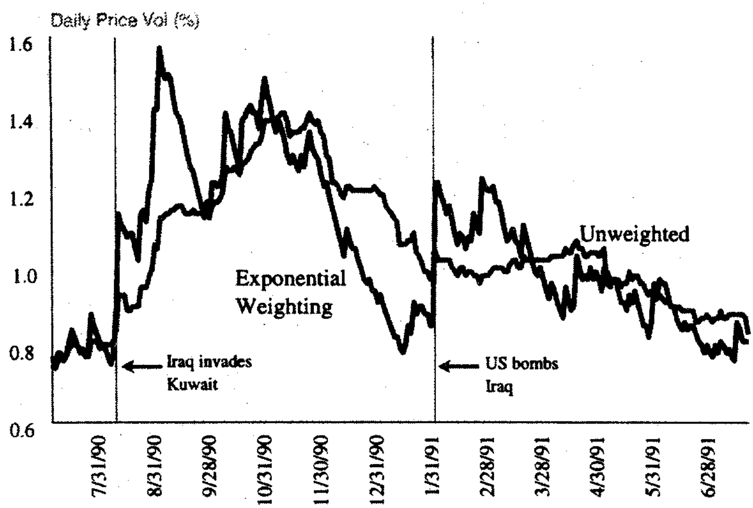

[4] Exponential weighting can be approximated in a historical VAR context simply by giving more recent data points more weight in the historical distribution. Depending on how this is implemented, there could be a discrete drop off point in the data series, but the effect will be diminished due to the exponential weighting.

Exhibit 7: $100 Million in an Emerging Markets Portfolio
VAR Stability

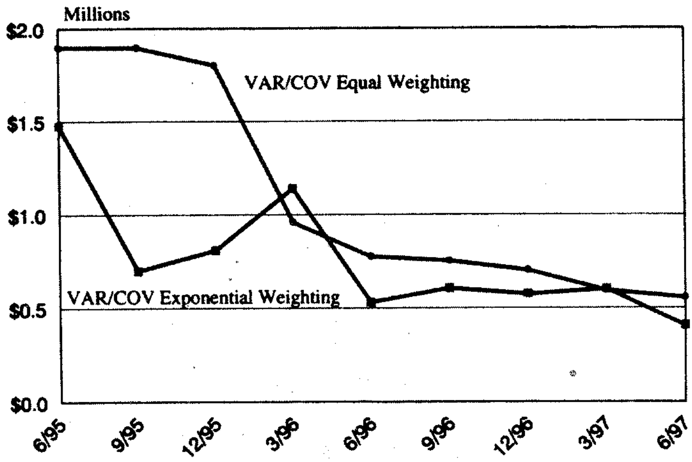

Because slightly different VAR methodologies can yield not only different VARs, but also different relative VAR movements, it is important to either monitor the spread relationships between the VAR calculations on an actual portfolio or at least a sample portfolio. This practice will help avoid placing too much dependence on a single VAR number within an organization.

Mapping/Selecting Relevant Risk Factors

To the extent that historical price series for every position is not available or is too monumental a data problem, assumptions as far as mapping securities into security groups or relating security risk to risk factors need to be made. These assumptions can be considered a balancing act between mapping to such an extent that the ultimate result may not be indicative of the actual portfolio and being so exact that the calculation could take a large team of people several days to compute. The key is to stress test in what areas exactness really matters and then be less exact on the others. This decision is highly dependent on the firm's business and the types of securities it trades.

Option Valuation

The variance-covariance VAR methodologies are not very good at capturing the non-linear behavior of options. So depending on how far in or out of the money an option position is, a VAR number can over or under estimate the real VAR. Depending on the amount of options held, certain second order limits may need to be established or multipliers added to approximate an accurate VAR number. As the next section mentions, Monte Carlo Simulation VAR is considered the best methodology for calculation VAR involving non-linear securities such as options.

Exhibit 8: Portfolio Volatility Example

Position	Treasury Strip	Equity Fund
Units	\$100 million	1 million shares
Value	\$61.78 million	\$60.00 million
Daily Vol	\$0.26 million	\$0.38 million
Correlation	25%	

$$Vol_{Portfolio} = \sqrt{Vol_{Strip}^2 + Vol_{Equity}^2 + 2 \times Corr \times Vol_{Strip} \times Vol_{Equity}}$$

$$Vol_{Portfolio} = \sqrt{0.26^2 + 0.38^2 + 2 \times 0.25 \times 0.26 \times 0.38}$$

$$Vol_{Portfolio} = \sqrt{0.2614}$$

$$Vol_{Portfolio} = \$0.51 \text{ million}$$

$$5\%\ VAR = 1.645 \times Vol_{Portfolio} = \$0.84 \text{ million}$$

COMPUTING VAR

Now that we have explored the various VAR choices that need to be tackled, we can show the actual calculations necessary in each of the three major types of VAR. Exhibit 8 illustrates the components of the sample portfolio used in each of the following VAR calculations.

Variance-Covariance VAR

Variance-covariance VAR in its simplest form involves finding the expected volatility in a portfolio of two securities and then multiplying by a factor that is selected based on the desired confidence level. For two securities, the VAR is:

$$VAR = MV \times \text{Factor} \times \sqrt{Wt_1^2 \times Vol_1^2 + Wt_2^2 \times Vol_2^2 + 2 \times Corr_{12} \times Wt_1 \times Wt_2 \times Vol_1 \times Vol_2}$$

where

$Wt_1 + Wt_2 = 1$
MV = total market value of portfolio
Factor = confidence level specific factor derived from cumulative normal distribution
Wt_1 = weight of security 1
Wt_2 = weight of security 2
Vol_1 = volatility of security 1
Vol_2 = volatility of security 2
$Corr_{12}$ = correlation between security 1 and security 2

Exhibit 8 illustrates this calculation for two securities. Note that the factor of 1.645 corresponds to a confidence level of 95%. A factor of 2.326 corresponds to a 99% confidence level. These factors are a function of the normal distribution of security returns that was illustrated in Exhibit 2.

The mathematical relationship can get arbitrarily complex as more and more securities are added. To avoid added complexity, VAR for more than two securities is represented in terms of matrix algebra. The first step in calculating variance-covariance VAR is computing the volatilities and correlations for the selected risk factors. The correlation matrix and volatilities are then combined into a covariance matrix. Next the dollar risk weights for each risk factor are calculated. These weights are arrived at through the mapping process (in the simplest case they are the market value amounts in each asset). Matrix algebra is then used to calculate the portfolio volatility. The portfolio volatility is then scaled by a factor according to the selected confidence level.

The weaknesses of the variance-covariance approach are:

1. It assumes all risk factors are normally distributed. This assumption causes the "fat tails" problem as illustrated in Exhibit 5. Research has shown that this problem is often not significant at the 95% confidence level, but can cause problems at the 99% confidence level.
2. It does not capture non-linear payout functions (e.g., options, callable bonds).
3. It does not capture time decay or time dependency of delta.
4. It assumes a static portfolio.

Monte Carlo Simulation VAR

Monte Carlo simulation VAR follows similar steps. The first step is to calculate the correlation and volatility matrix for the risk factors. Then these correlations and volatilities are used to drive a random number generator to compute changes in the underlying risk factors. Next the resulting values are used to re-price each portfolio position and determine trial gain or loss. This process is repeated for each random number generation and re-priced for each trial. Exhibit 9 illustrates the first 10 trials for the illustration from Exhibit 8. The results are then ordered such that the loss corresponding to the desired confidence level can be determined.

The greatest benefit to Monte Carlo simulation VAR is the ability to use pricing models to revalue non-linear securities for each trial. In this way, the non-linear effects of options that were missed in the variance-covariance VAR can now be captured.

The computation involved in Monte Carlo simulation VAR can be immense for large portfolios, so careful consideration should be paid to the cost versus benefit of calculating a simulation VAR. For instance if a portfolio is composed of linear (non-option) components, there is little added benefit to using

Monte Carlo simulation VAR compared to variance-covariance VAR. Exhibit 10 illustrates how the results from the Monte Carlo simulation VAR approaches the variance-covariance VAR as more and more trials are added.

Monte Carlo simulation, having its roots as a random number generator, is exposed to sampling error. There is the risk of running too few simulations to adequately capture the distribution and this could result in an inferior answer. Fortunately, calculation methods exist to estimate how far off a simulation is so that we can decide whether or not to run more trials.

Exhibit 9: Monte Carlo Simulation Trials

(Values in $ Millions)

Trial	Strip Value	Equity Value	Portfolio	Gain/Loss
1	61.73	59.80	121.53	–0.25
2	62.09	60.01	122.10	0.32
3	61.79	59.64	121.43	–0.35
4	61.61	60.36	121.96	0.18
5	61.79	60.26	122.05	0.27
6	61.79	60.04	121.83	0.05
7	61.65	60.38	122.03	0.25
8	61.65	59.94	121.59	–0.18
9	61.81	60.44	122.25	0.47
10	61.71	59.38	121.09	–0.69

Exhibit 10: Monte Carlo Simulation Results

Portfolio Gain/Loss

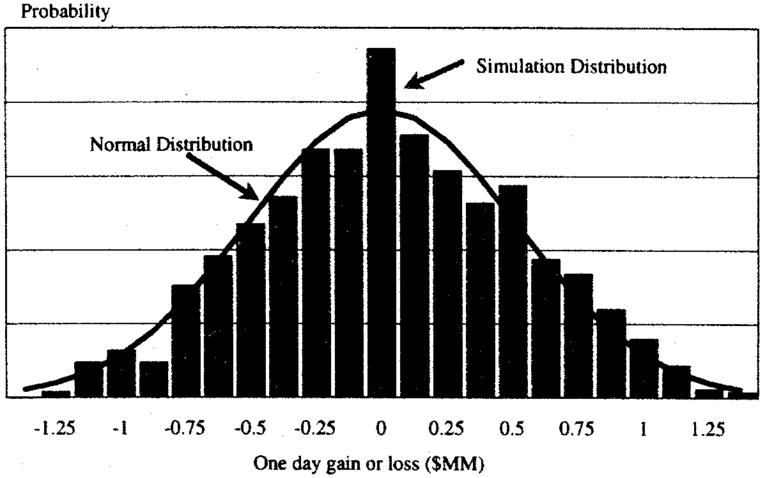

5% VAR = $0.81 million

To analyze the sample error inherent in Monte Carlo simulation, consider the VAR of a portfolio which is between $5 and $6 million with a confidence of 95%. This is simply estimating the range around the estimated VAR in which the real VAR may lie with a certain confidence. This confidence interval can be computed by assuming a binomial distribution with probability equal to the confidence interval (95% in our example) of having a real VAR below the simulated VAR. Using 1,000 trials, we know that the 95% VAR may be viewed as the 50th ranked trial. Using the binomial distribution around the 50th trial, we can compute that with 95% confidence interval around the 50th trial occurs on trials 37 and 64. Therefore in the previous example, the confidence interval for our VAR is $0.75 million – $0.86 million, corresponding to the 64th and 37th ranked trials, respectively.

Historical VAR

Unlike the other two methodologies, historical VAR does not depend on calculated correlations and volatilities. Instead it uses historical data of actual price movements to determine the actual portfolio distribution. In this way the correlations and volatilities are implicitly handled. In fact the most important advantage of historical VAR is that the "fat-tailed" nature of a security's distribution is preserved since there is no abstraction to a correlation and volatility matrix.

Historical VAR is calculated by mapping a portfolio into a historical price distribution. The gains and losses are then added across the portfolio for each day and then ranked in order. The return corresponding to the desired confidence level is then selected as the VAR — meaning that if the desired confidence level were 95% and there were 1,000 data points, we would select the 50th lowest return (50 = (100% – 95%) × 1,000). Exhibit 11 illustrates the historical distribution for the strip and stock example.

One important consideration for computing a historical VAR is the length of the historical period. The historical period should be long enough to form a reliable estimate of the distribution, but short enough to avoid "paradigm shifts."

COMPARISON OF THREE MAJOR VAR TECHNIQUES

Exhibit 12 summarizes the pros and cons of each of the methodologies discussed in this chapter. Recognizing that two different departments with a VAR of $1 million can have very different risk profiles depending on both the nature of their business and the calculation methodology chosen by each is the first step in improving an organization's use of VAR. VAR is a valuable tool, but must be consistently applied to be a meaningful measuring stick.

Exhibit 11: Historical Simulation Results
Portfolio Gain/Loss

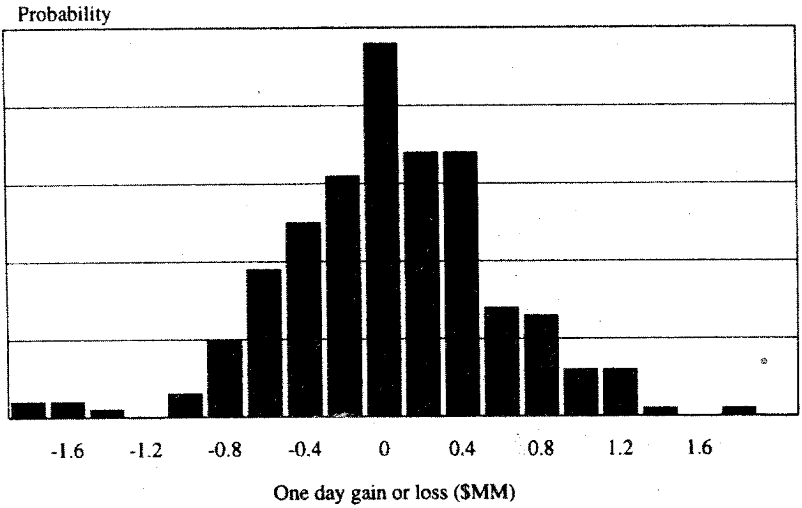

5% VAR = $0.80 million

Exhibit 12: Summary of Pros and Cons of Each VAR Methodology

Method	Pros	Cons
Variance/ Covariance	Easy to understand. Least computationally intensive. Industry standard.	May misstate non-linear risks. "Fat tails" problem. Computationally intensive.
Monte Carlo Simulation	Accommodates any statistical assumptions about risk factors. Can fully capture non-linear risks.	Sampling error.
Historical Simulation	Naturally addresses the "fat tails" problem. Performs well under back-testing. Can fully capture non-linear risks.	Relies on history. Computationally intensive. Data intensive.