



SOCIETY OF ACTUARIES

Article from:

The Financial Reporter

May 1999 – Issue 39

A Framework for Managing Surplus

continued from page 7

4. Differences in asset class specific expenses
5. Differences in RBC requirements
6. Cash flow differences between the two securities
7. Transaction costs
8. Capital gains tax implications
9. Impacts on interest crediting rates, if applicable. There is also the issue of who (i.e., policyholders or shareholders) should benefit from the transaction, and to what degree.
10. Impact on interest maintenance reserve (IMR) and any applicable statutory accounting considerations
11. Impact on GAAP accounting results
Note that realized gains in liability portfolios go to GAAP surplus and do not remain within the liability portfolio. For example, realizing gains on assets supporting a fixed liability effectively advances the timing of GAAP operating income but changes its character into net income instead of operating income. The future GAAP operating income will be lower and the margins in the GAAP reserves will be lower. If the realized gains are too large, then the liability portfolio has negative GAAP margins that would result in loss recognition. This is the worst situation, because not only has the future operating income been converted into net income, but the realization of gains beyond the point of a zero margin results in negative operating income via the loss recognition.
12. Rating agency issues, if any.

David N. Becker, FSA, PhD, CFA, is Vice President and Chief Actuarial Officer, Lincoln National Life Insurance Company, Fort Wayne, Ind. This article appeared in The North American Actuarial Journal, and is being repeated here because of its pertinence.

Value-at-Risk—*an Overview*

(Part Two of Two)

by Glyn Holton

(Editor's Note: Issue 38 of The Financial Reporter, January 1999 contains Part One of this article, dealing with a definition of VAR, a simple model of it, key factors in VAR and discussion of linearity and non-linearity aspects.)

Simulating VAR

Faced with non-linear portfolios, we must discard the linearity and normality assumptions of delta-normal VAR and consider alternative approaches to estimating VAR. The basic problem of estimating VAR, however, remains the same. We consider a set of key factors whose behavior we can describe statistically. We have a portfolio price function that relates those key factors to the portfolio's price. Somehow, we must translate these two pieces of information into an estimate of the portfolio's VAR. In this section, we consider the problem as one of solving an integral equation.

Suppose we wish to estimate 95% VAR for a portfolio. The portfolio's VAR is the bound on a 95% confidence interval for ΔP . As suggested by Exhibit 1, this can be expressed as an integral:

$$95\% = \int_{-VAR}^{\infty} p \, d\Delta P \quad [21]$$

where p is the probability density function for ΔP .

In [21] we are not actually solving for the value of the integral. Instead, we are solving for the value VAR that makes it 95%. If no closed form solution exists for [21], we consider numerical methods of integration. In doing so, we face a problem called the "curse of dimensionality." This arises because, although [21] is presented as a one-dimensional integral, it is in fact an m -dimensional integral—both p and ΔP are functions of the m key factors.

$$\{V_k\}_{k=1}^m$$

Most techniques of numerical integration entail dividing the area of integration into subparts, performing some simple calculations on each subpart, and summing the results.

A problem in multi-dimensions is that, as the number of dimensions grows, so does the number of (multi-dimensional) rectangles used. For example, in the one-dimensional case, the area of integration $[a,b]$ might be divided into 100 subparts. In the two-dimensional case, the area of integration has the form $[a,b] \times [c,d]$. If both the intervals $[a,b]$ and $[c,d]$ are divided into 100 subparts, there are going to be $100^2 = 10,000$ rectangles to evaluate.

In the 50-dimensional case, that number grows to 100^{50} . Reducing the number of subparts into which each interval is divided does not help. In the 50-dimensional case, if each interval were divided into just two subparts, this would translate into $2^{50} = 1,125,899,906,842,620$ rectangles.

This is the "curse of dimensionality." It is a problem that causes most techniques of numerical integration to fail when applied to high-dimensional problems. It is an issue with VAR because many portfolios are exposed to tens or hundreds of key factors—each one adding a dimension to the problem.

Monte Carlo simulation is a form of numerical integration that avoids the curse of dimensionality. Using the numerical approach outlined above, the integral is approximated as:

$$\int f(x) dx \approx \sum_{i=1}^{z^n} A_i \quad [23]$$

where z^n is the total number of rectangles, and A_i is the area (volume) of the i^{th} rectangle. Because of the sheer number of rectangles involved, we do not directly calculate this sum. Instead, we note that

[23] can be rewritten:

$$\int f(x)dx \approx z^n \left(\frac{\sum_{i=1}^{z^n} A_i}{z^n} \right) = z^n \text{Avg}(A_i) \tag{24}$$

where $\text{Avg}(A_i)$ is the average area of all the rectangles.

That average area can be estimated by selecting a sub-sample of the rectangles and computing their average area. The solution of Monte Carlo simulation is to make the selection randomly.

A question that remains is: How many scenarios will be necessary to reasonably estimate an integral using Monte Carlo simulation? In fact, Monte Carlo simulation breaks the curse of dimensionality. It can be shown that, *independent of the number of dimensions to a problem*, the convergence error of Monte Carlo simulation is inversely proportional to the square root of the number of random scenarios used. For example, suppose that a simulation uses 1,000 scenarios and has precision of 8%.¹ The very same analysis would be precise to within 4% if 4,000 scenarios were used. This would be true no matter how many dimensions the problem had.

A Probability Transformation

When we estimate VAR for a portfolio, there are two broad inputs:

1. Historical data for key factors
2. The portfolio price function

These two inputs characterize the two components of market risk that VAR must incorporate—the two pieces of the VAR “puzzle” that we identified in Section 3. Historical data captures the uncertainty of financial markets. The portfolio price function describes the portfolio’s exposures to that uncertainty.

Neither input appears in the VAR integral [21]. The integral depends upon the one-dimensional probability distribution p for ΔP . Somehow we must infer p from the historical data and the portfolio price function.

Historical data tells us nothing

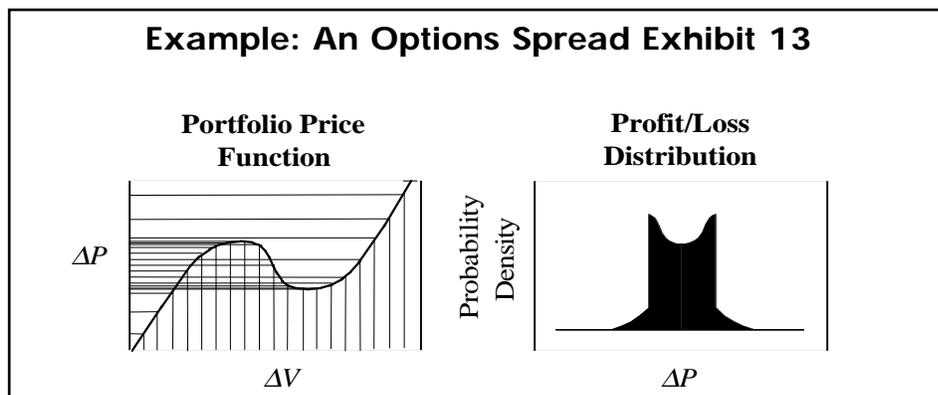
directly about p . Rather, it provides information about the m -dimensional probability distribution for the key factors, which we denote q .

The portfolio price function also tells us nothing directly about p . However, as a transformation from the m -dimensional space of the key factors to the one-dimensional space of the portfolio’s value, it relates p to q . If we somehow apply the transformation to the entire m -dimensional probability distribution q we will obtain the one-dimensional probability distribution p .

In attempting this transformation, we face two challenges:

how the portfolio price function transforms the (normal) distribution q of the underlier into the more complex distribution p shown on the right. By observing how values of ΔP cluster in the left graph, we can infer the appearance of the probability distribution in the right graph.

Exhibit 13 illustrates in one dimension how complex the task of inferring the probability distribution p may be. After all, the portfolio price function may have multiple local maxima and minima as well as multiple inflection points. If the portfolio contains exotic derivatives, it may even have jump discontinuities. The task of inferring p for a



1. Applying the portfolio price function as a transformation to the probability distribution q is a complex mathematical problem.
2. We don’t even know the probability distribution q . We have to decide what inferences to make about that distribution based upon available historical data.

As we shall see, Monte Carlo simulation provides a solution to the first problem. The second problem can be addressed in different ways. Monte Carlo VAR and historical VAR are two forms of Monte Carlo simulation that differ only in how they address this second problem.

Starting with the first problem, let’s consider an example. Exhibit 13 describes a portfolio consisting of a long-short options position in a normally distributed underlier V . The portfolio price function is illustrated on the left, and the probability distribution p is illustrated on the right.

In the left graph, evenly spaced values for the underlier ΔV have been mapped into values of ΔP . This illustrates

portfolio with thousands of positions exposed to hundreds of key factors is potentially staggering.

Exhibit 13, however, also suggests a solution. By mapping a range of values for DV into corresponding values for ΔP , we were able to infer the nature of the distribution p . We can systematize such an approach by selecting a broad sample of scenarios for the and valuing the portfolio under each scenario using the portfolio price function. The histogram of results for ΔP will be a discrete approximation to the probability distribution p from which we can estimate VAR.

If we actually use this approach to estimate VAR for a portfolio, one challenge is deciding how to select the set of scenarios. It is not sufficient to merely select a large number of scenarios. We must make the selection in a manner that will not bias the results—we need a representative selection. One possible solution is to select scenarios that are evenly spaced as was done in Exhibit 13. In higher dimensions, however, this

Value At Risk—An Overview (Part Two of Two)

continued from page 9

approach succumbs to the curse of dimensionality. An alternative is to select the scenarios randomly. Obviously, this is the solution of Monte Carlo simulation which we developed in Section 8.

Accordingly, in this section and the previous section, we have addressed two fundamental challenges in estimating VAR for non-linear portfolios. For both challenges, a solution has been Monte Carlo simulation. In summary, the two distinct problems that Monte Carlo simulation has solved have been:

1. The curse of dimensionality which we face in numerically solving the integral [21]
2. The probability transformation of applying the portfolio price function to q to infer p

When we use Monte Carlo simulation for estimating VAR, we can do so in one of two ways:

1. We can draw our scenarios from an m -dimensional uniform distribution and then weight each scenario to reflect the probability distribution of the key factors, or
2. We can draw the scenarios from the probability distribution of the key factors and weight the scenarios uniformly.

Either approach represents a valid implementation of Monte Carlo simulation. In Sections 12 and 13, we will introduce two different implementations of Monte Carlo simulation for estimating VAR. These are the techniques of Monte Carlo VAR and historical VAR. Both are implemented according to the second of the above two approaches.

Statistical Error

Because we don't know the probability distribution q , we must make inferences about it based upon historical data. In the case of delta-normal VAR, these inferences take the form of a set of standard deviations and

correlations. In the case of Monte Carlo simulation, the inferences can take different forms. The end result, however, must be a set of scenarios. Monte Carlo VAR and historical VAR are both forms of Monte Carlo simulation. They differ only in how they utilize historical data in selecting those scenarios to represent q . Both approaches entail two general types of error:

1. **Error arising from how scenarios are selected:** We must select scenarios in a manner that reflects the characteristics of the distribution q .
2. **Error arising from the number of scenarios selected:** We must select sufficiently many scenarios to adequately reflect the distribution q .

The difference between these is the difference between quality and quantity—selecting the right scenarios vs. selecting enough scenarios. The first type of error arises in different ways, some of which

are unique to either historical VAR or Monte Carlo VAR. The second type of error impacts both historical VAR and Monte Carlo VAR in exactly the same way. It is called convergence error.

We can reduce convergence error simply by using more scenarios. Error relating to *how* we select scenarios, on the other hand, cannot be reduced in this way. If the manner in which we select scenarios has some form of bias, simply selecting more scenarios can not eliminate that bias.

Error relating to how we select scenarios is often a subjective notion. This is because there is no "true" distribution q . People can have legitimate differences of opinion about the nature of that distribution—and consequently about any bias that may exist in how we select scenarios. Convergence error is more objective. Using statistical techniques, we can usually quantify the convergence error of a simulation.

Monte Carlo VAR

The approach of Monte Carlo VAR is to

randomly generate scenarios based upon some assumed joint probability distribution for the ΔV_k . Historical market data is used to infer statistical characteristics such as standard deviations and correlations for the assumed distribution.

For example, for a given portfolio we might assume that each key factor is lognormally distributed with a mean equal to today's value for that factor. Recent market data would then be analyzed to infer a standard deviation for each key factor as well as a correlation for each pair of key factors.

Once an assumed joint distribution is specified, standard techniques for generating correlated random numbers are used to select a set of scenarios. In this way, the selected scenarios are literally drawn from the assumed distribution. They reflect the statistical characteristics—standard deviations and correlations—inferred from the historical data.

This approach to selecting scenarios entails four sources of error:

1. **Assumed distribution:** The standard distribution we assume for key factors may imperfectly reflect the "true" distribution q .
2. **Sampling error:** Because we estimate standard deviations and correlations from a limited set of historical data, those "sample" standard deviations and correlations will only approximately reflect the "true" standard deviations and correlations of q .
3. **Non-stationary:** Because market conditions are non-stationary, the historical data upon which we base standard deviation and correlation estimates may imperfectly reflect today's market conditions.
4. **Imperfect random number generation:** Imperfections in the random number generator we use for selecting scenarios may introduce a bias.

Because Monte Carlo VAR depends upon the inference of standard deviations and correlations from historical data, it is similar to delta-normal VAR. Its sampling error and error from market

non-stationary are identical to those of delta-normal VAR. One can be addressed by using as much historical data as possible. The other can be addressed by using only the most recent data. As with delta-normal VAR, some compromise must be achieved to balance the two.

In addition to error relating to how scenarios are selected, Monte Carlo VAR also entails convergence error. However, there is no theoretical limit to the number of scenarios that can be used with Monte Carlo VAR. Accordingly, this error can be made as small as available computing technology will permit.

We can calibrate a portfolio to determine the number of scenarios required to achieve a desired degree of convergence. For example, suppose an organization wants to simulate the VAR of its portfolio with only 4% convergence error. To find the required number of scenarios, the organization calculates Monte Carlo VAR on the portfolio 50 times, using 1,000 random scenarios in each simulation. The resulting 50 VAR estimates are then gathered and their standard deviation is calculated.

Suppose the standard deviation is 8%. This means that simulation can measure the portfolio's VAR with 8% convergence error using 1,000 scenarios. Because the convergence error of Monte Carlo simulation is inversely proportional to the square root of the number of scenarios used, the same portfolio will require 4,000 scenarios to achieve a convergence error of 4%.

Historical VAR

Like Monte Carlo VAR, historical VAR must somehow select a set of scenarios to reflect the unknown distribution q . The approach of historical VAR is to draw scenarios directly from historical data. For each date represented in the historical data, the one-day return for each of the key factors is calculated. A scenario is constructed by applying those returns to today's values for the key factors. This approach to selecting scenarios entails two sources of error. Both arise from market non-stationary:²

1. **Non-stationary:** Because market conditions are non-stationary, the

historical data upon which we base standard deviation and correlation estimates may imperfectly reflect today's market conditions.

2. **Distortions from assuming market stationarity:** Distortions occur because historical data is treated as arising from a stationary (fixed) probability distribution as opposed to one that has varied over time.

The first source of error also arose with delta-normal VAR and Monte Carlo VAR. The second is new. Its most obvious effect is that heteroscedasticity (non-constant volatility) is mistaken for leptokurtosis (fat tails to a distribution). For this reason, historical VAR tends to overstate the effects of leptokurtosis. Monte Carlo VAR, by comparison, uses standard distributions such as the normal distribution or lognormal distributions to model q . Accordingly, it tends to understate the effects of leptokurtosis.

While Monte Carlo VAR and historical VAR introduce different errors in how they select scenarios, their convergence errors behave identically. This is because the two methodologies differ only in how they specify random scenarios—not in how they use those scenarios. For a given portfolio, the number of scenarios needed to achieve a given degree of convergence will be the same irrespective of whether Monte Carlo VAR or historical VAR is used to generate those scenarios.³ When you calibrate a portfolio for Monte Carlo VAR, the same result applies to historical VAR.

For example, suppose that Monte Carlo VAR is used to calibrate a portfolio to determine that a 8% convergence error can be achieved with 1,000 scenarios. If 2% convergence error were desired, Monte Carlo VAR could achieve that result using 16,000 scenarios. Historical VAR could not match that convergence. If a year (252 trading days) of historical scenarios were used, the convergence error of historical VAR would be 16%. Achieving 2% convergence error with historical VAR would require 63 years of data.

Historical VAR is fairly easy to implement. However, the significant convergence error associated with historical VAR can limit the technique's appeal in many situations.

Conclusion

Value at risk is a powerful measure of market risk. In theory, it is applicable to all portfolios and all sources of market risk.

The challenge of estimating a portfolio's VAR lies in integrating the market information contained in the standard deviations and correlations of key factors with the portfolio information contained in the portfolio price function. For simple portfolios that exhibit linear price behavior, this can be accomplished using the method of delta-normal VAR.

If a portfolio contains options or other positions that exhibit non-linear price behavior, VAR may be estimated using Monte Carlo simulation. Two particular implementations of Monte Carlo simulation for VAR are the techniques of Monte Carlo VAR and historical VAR.

All three VAR techniques presented here entail error relating to statistical inference. In addition, the simulation techniques Monte Carlo VAR and historical VAR entail convergence error.

Glyn Holton is an independent consultant specializing in financial risk management. His practice is called Contingency Analysis and is based in Boston, Mass. His web site is www.contingencyanalysis.com.

References

- 1) Duffie, D. and J. Pan. "An Overview of Value at Risk." *Journal of Derivatives*, Spring 1997, pp. 7-49
- 2) Niederreiter, Harald. "Random Number Generation and Quasi-Monte Carlo Methods." *Society for Industrial and Applied Mathematics*, Philadelphia, 1992
- 3) RiskMetrics Technical Document, 4th ed. JP Morgan, 1997