Financial Analysts Journal

A Publication of Association for Investment Management and Research

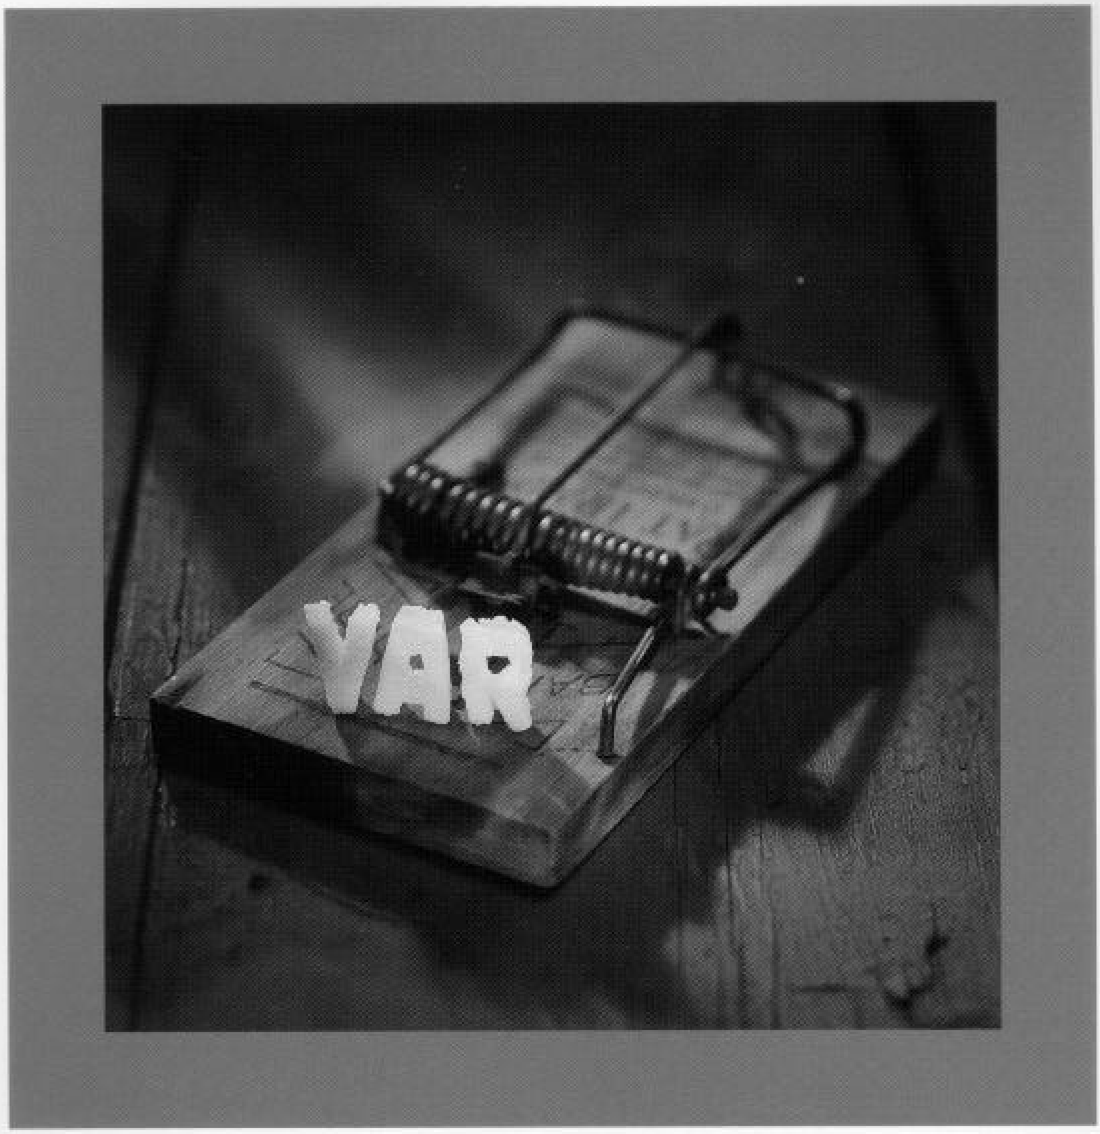

SEPTEMBER/OCTOBER 1995

VAR: Seductive but Dangerous

Tanya Styblo Beder

Value at risk (VAR) has gained rapid acceptance as a valuable approach to risk management. Not all VARs are equal, however. A study of VAR techniques used by dealers and end-users reveals that VAR calculations differ significantly for the same portfolio. VARs are extremely dependent on parameters, data, assumptions, and methodology. Calculation of eight common VARs for three hypothetical portfolios demonstrates the potentially seductive but dangerous nature of any single approach to risk management. In sum, although VAR and other quantitative techniques are necessary aspects of an effective risk-management program, they are not sufficient to control risk.

Value at risk is Wall Street's latest advancement in risk measurement. Simply defined, VAR is an estimate of maximum potential loss to be expected over a given period a certain percentage of the time. Its simplicity is seductive. Used to the extreme, in a single statistic, a firm can measure its exposure to markets worldwide. VAR enables a firm to determine which businesses offer the greatest expected returns at the least expense of risk. When one considers that risk management in the early 1970s consisted almost entirely of the evaluation of credit risk, VAR's power in the context of the galaxy of risks we track, analyze, and manage today is breathtaking to consider.

VAR can be dangerous, however. A review of dozens of dealers' and end-users' VARs revealed radically different approaches to the calculation. In this study, eight common VAR methodologies were applied to three hypothetical portfolios. As illustrated in Figure 1, the magnitude of the discrepancy among these methods is shocking, with VAR results varying by more than 14 times for the same portfolio. These results illustrate the VAR's extreme dependence on parameters, data, assumptions, and methodology.

The implications of these discrepancies for capital adequacy standards are significant, especially given the Basle Committee on Banking Supervision's treatment of VAR in its proposed amendment to The 1988 Basle Capital Accord, "The Supervisory Treatment of Market Risks," published on April 12, 1995. This amendment proposes that dealers use either an internal methodology or a Bank for International Settlements (BIS) standard methodology to compute VAR and that the results be multiplied by a factor of three to determine the amount of capital to be set aside for market risk. Our research indicates that this amount may be too high or too low, depending upon the method used. The need for a uniform VAR methodology or for differing multiplication factors according to the type of VAR is paramount to establish a common ground for comparative purposes.

In our analysis, historical simulations present quite different views of risk relative to Monte Carlo simulations. This difference is attributable to the extreme dependence of historical simulations on the underlying data set and the value of the relative randomness of key variables in Monte Carlo simulations compared with sample-specific values. The results also reveal the exceptional time sensitivity of certain portfolio risks and highlight the potential failure of VAR, even when bolstered by stress testing. In sum, although VAR and stress testing are necessary, they are not sufficient to contain risk.

The differences in common VARs emphasize the fact that no single set of parameters, data, assumptions, and methodology is accepted as the "correct" approach. Even if two firms use the same quantitative technique, they often apply different assumptions in implementing the technique. For example, some firms calculate the global VAR of the firm over a one-day time horizon, using historical data series on markets and a specific set of mathematical models. Others calculate regional VAR of product areas over a monthly or annual time horizon, using random or implied data series

Tanya Styblo Beder is a principal of Capital Market Risk Advisors, Inc., in New York.

Figure 1. Range of VARS: All Simulations

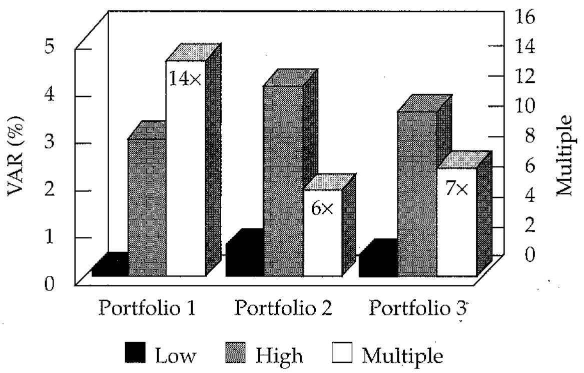

on markets and multiple mathematical models. Depending on the selection of time horizon, data base, and correlation assumptions across instrument/asset classes, the same model may produce widely divergent VAR views for the same portfolio and, therefore, different capital requirements.

THE PORTFOLIOS AND VAR CALCULATIONS

In the remainder of this article, we describe the common VAR calculations and apply them to three hypothetical portfolios. For each methodology presented, VAR is calculated for both one-day (1d) and two-week (2w) time horizons. The first methodology, historical simulation, is performed twice, changing the data base used from the past 100 trading days (Pr100d) to the past 250 trading days (Pr250d). The second methodology, Monte Carlo simulation, also is performed twice, changing the correlation estimates from the JP Morgan RiskMetrics data set to those from the BIS/Basle Committee proposal.[1] Differences in correlation estimates between RiskMetrics and BIS/Basle are significant. RiskMetrics permits correlation across all asset classes, using exponentially weighted daily historical observations. The BIS/Basle proposal permits correlation only within asset classes, not across, effectively forcing the correlation between asset classes to be plus or minus 1, whichever produces the higher estimate of VAR.

The three portfolios were constructed to have increasing complexity in terms of optionality and/or asset class composition and possess properties sought frequently by dealers and end-users. The eight VAR calculations performed for each portfolio are summarized in Table 1.

Portfolio 1

Portfolio 1 consists exclusively of U.S. Treasury strips. It was designed to satisfy three conditions at construction: (1) The duration of the portfolio equals that of the ten-year strip,[2] (2) the portfolio has greater convexity than the ten-year strip, and (3) the portfolio performs at least as well as the ten-year strip under a 100 basis point parallel increase or decrease in the Treasury yield curve or under an inversion of the Treasury yield curve. Table 2 describes the composition of and constraints on Portfolio 1, which consists of a long position in 2-year and 30-year U.S. Treasury strips. The ten-year U.S. Treasury strip, the benchmark, is included in the table for reference. The net investment in Portfolio 1 is $1 million.

The traditional risk measures show that for very small parallel shifts in the Treasury yield curve, Portfolio 1 performs similarly to the ten-year strip. Moreover, for the plus or minus 100 basis point parallel yield curve shifts and inversion, the portfolio's $1 million investment performs slightly better than if it been invested in the ten-year strip.

The VAR analyses of Portfolio 1 reveal quite different risk profiles than older risk measures such as duration, convexity, and scenario analysis. Figure 2 displays the eight common VAR calculations for this portfolio. The VAR results place significantly different degrees of capital at risk both

Table 1. Eight Common VAR Calculations

VAR Approach	Type of Simulation	Data Base/Correlation Assumption	Holding Period
1	Historical	Prior 100 trading days	One day
2	Historical	Prior 250 trading days	One day
3	Monte Carlo	Historical, RiskMetrics correlations	One day
4	Monte Carlo	Historical, BIS/Basle correlations	One day
5	Historical	Prior 100 trading days	Two weeks
6	Historical	Prior 250 trading days	Two weeks
7	Monte Carlo	Historical, RiskMetrics correlations	Two weeks
8	Monte Carlo	Historical, BIS/Basle correlations	Two weeks

Table 2. Portfolio 1: Composition and Constraints

Characteristic	2-Year Strip	30-Year Strip	Total Portfolio	10-Year Strip (Benchmark)
Composition				
Yield[a]	5.91%	6.85%		6.58%
Price[b]	89.12	14.94		52.42
Face amount	$779,778	$2,041,424	$2,281,202	$1,907,670
Purchase amount	$694,964	$305,036	$1,000,000	$1,000,000
Duration and convexity				
Duration	1.712	4.078		5.063
Duration contribution	1.335	8.325	9.660	
Convexity	0.041	1.133		0.514
Convexity contribution	0.032	2.312	2.344	
Scenario analysis				
Yield + 100 bp	6.91%	7.85%		7.58%
Price/yield + 100 bp	87.43	11.38		47.60
Position/yield + 100 bp	$681,775	$232,333	$914,108	$908,051
Yield − 100 bp	4.91%	5.85%		5.58%
Price/yield − 100 bp	90.86	19.64		57.75
Position/yield − 100 bp	$708,743	$401,017	$1,109,490	$1,101,679
Yield curve inversion	7.20%	6.30%		6.58%
Price/inversion	86.95	17.37		52.42
Position/inversion	$678,009	$354,508	$1,032,517	$1,000,000

Note: Prices and yields as of 5/25/95. The maturities of the 2-year, 10-year, and 30-year strips are 5/15/97, 5/15/05, and 8/15/23, respectively.
[a] The market yield for each strip is stated on an actual/365 basis with semiannual compounding.
[b] The price of each strip is stated as a percentage of face amount.

within and between methodologies. The VAR statistics to the right of each bar may be interpreted as follows: Under the assumptions specific to the particular VAR calculation, the probability is 5 percent (1 percent) that the portfolio will suffer a loss greater than or equal to the statistic shown. For the third set of bars, for example, under the assumptions made to perform historical simulation over the prior 250-day period, the probability is 1 percent that a loss equal to or exceeding 1.29 percent of the $1 million portfolio investment will occur over a one-day time horizon.

Figure 3 compares the results of the historical simulation for Portfolio 1 with the results of the

Figure 2. VAR Calculations: Portfolio 1

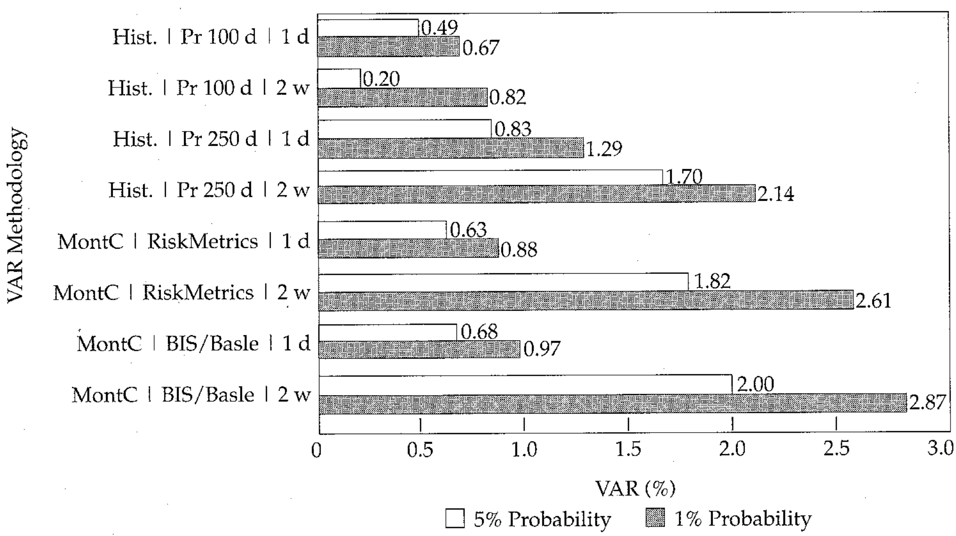

Figure 3. Distributions of One-Day and Two-Week Returns: Portfolio 1

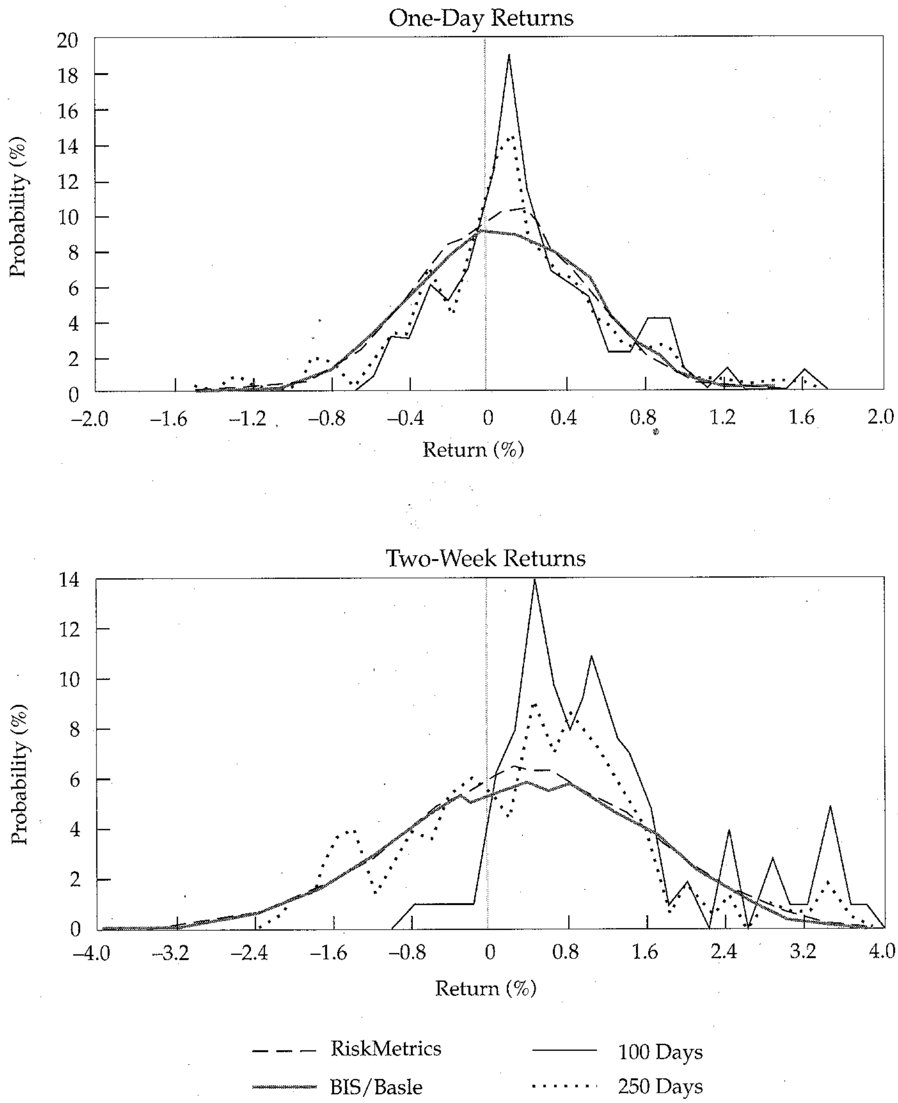

Monte Carlo simulations for one-day and two-week returns.[3] As illustrated by the graphs, the historical simulations present a different view relative to the Monte Carlo simulations. This result is attributable to their extreme dependence on the underlying data set. During the 100-day and 250-day periods included in the historical simulations, the value of Treasury strips largely appreciated. Had a period of rising interest rates been selected, the result would have been the opposite. The danger in basing VAR estimates on relatively short periods of direct historical observations is apparent—history must repeat itself for the results to predict the future. Although historical estimates may be fairly accurate in a trending market, they will be less accurate when the trend changes.

As summarized in the bar charts in Figure 2, the result for the prior 100 trading days, 5 percent probability VAR equals 0.49 percent over a one-day horizon but then drops to 0.20 percent over a two-week period. For all other VAR types, the VAR result increases with the time horizon, as would be expected. This surprising result is explained by the pattern of results during the specific historical periods. Although the average return is positive during the first 100 trading days (the left side of Figure 4), negative returns are more common over the one-day time horizon than over the two-week time

Figure 4. Historical One-Day Returns and Two-Week Returns: Portfolio 1

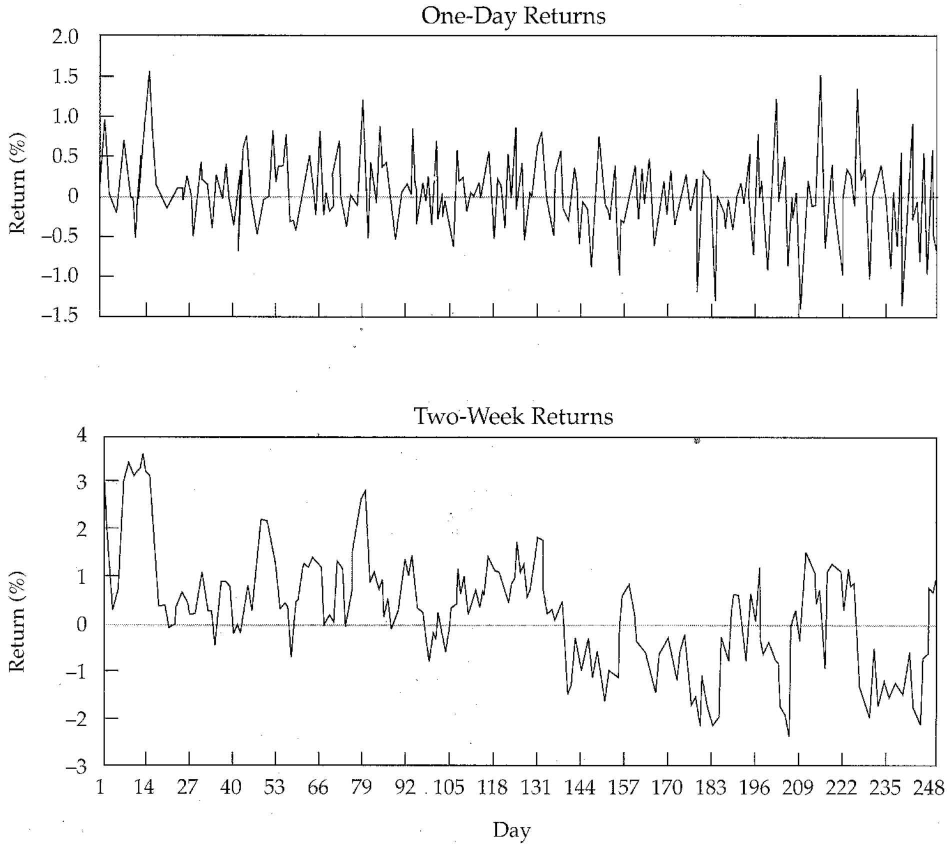

horizon. Thus, VAR is higher for the one-day time horizon than for the two-week time horizon.

Several other conclusions follow from this set of VAR results, as expected given interest rate trends at the time. Monte Carlo simulations indicate higher expected losses than does the 100-day historical simulation but lower expected losses than for the 250-day historical simulation. Historical simulations indicate that increasing the holding period from one day to two weeks decreases the expectation of the highest expected profits and increases the magnitude of expected losses. Note that although Monte Carlo simulations indicate a similar change in the expectation of the highest profits, they predict a larger increase in the magnitude of expected losses (three times versus two times). The difference in VAR driven by the relative randomness of key variables in Monte Carlo versus sample-specific historical simulations is clear.

Time horizon is clearly a crucial parameter in VAR. Firms select quite different time horizons to view their risk. Does the firm wish to analyze its potential capital exposure and expected profit over a short-term or a long-term horizon? In terms of risk–reward appetite, will the firm be satisfied if losses mount for two years but huge profits make up for the losses in the third year? Or, is less erratic performance desired? Often, more-even performance is desired by firms that report public financial information, as well as by funds that must publish daily net asset values. Thus, two firms performing identical VAR calculations, other than selection of time horizon, may have different but not necessarily inconsistent VAR results.

Although a model may produce adequate views of capital at risk on an overnight or weekly basis, it may produce inadequate risk views over time horizons of several months, a year, or longer. For example, the calculation of short-horizon VAR may be misleading for customized or exotic products that cannot be liquidated under the assumed time horizon. To the degree that multiple horizons are required, risk systems rarely are capable of incorporating them in aggregating portfolio risks, creating a limitation in sizing true risk exposures. Note that some firms address this problem by adjusting midmarket valuations.

Yet another challenge is that although longer time horizons may be appropriate for instruments such as illiquid, path-dependent options, some mathematical functions are inaccurate beyond small market moves. For example, many mathematical models are incapable of handling discontinuities such as market gapping, or they require strict assumptions such as linearity to produce accurate information. The April 12, 1995, proposed amendment to the 1988 Basle Capital Accord suggests that firms use a single time horizon of two weeks (ten business days) for VAR calculations.

The selection of data sets is another critical component of VAR. As Portfolio 1 illustrates, alternate data sets may produce vastly different risk views. In our experience, intraday versus end-of-day data often produce contrary views of risk during periods of high volatility. Different risk views will also be created by the use of historical as opposed to market-implied data. Although historical data are most often used to calculate VAR, the length of the historical period selected varies significantly from firm to firm. Several firms and software packages use a 90-day historical time horizon, but many market participants believe that, at a minimum, a one-year data set should be used. The proposed Basle amendment suggests that firms use a one-year-minimum data set for VAR calculations.

Length of time is not the sole criterion to establish regarding the data set. Sampling frequency must be set high enough to ensure that the data set is statistically significant. For example, a one-year data base composed of 12 end-of-month data points may be no more relevant than a data set of 12 points selected through random chance. Often, sampling frequency is also sensitive to the time the data are collected. For example, end-of-day data points are likely to produce a different VAR picture than daily high/low/close data points.

After type, sampling frequency, and length of data base are selected, the VAR user must determine whether to exclude certain data points. For example, should the data set include "outliers" caused by one-time events, market gapping, or other dislocations? Such occurrences are often characterized as extreme but low-probability events. Recent examples are the devaluation of the Mexican peso, the 1987 stock market crashes, and commodity volatility during the Gulf War. Two data bases, distinguished by inclusion of outlier events, are likely to produce different VAR calculations.

Yet another challenge in data set selection is determining whether an outlier event is an indication of structural change in the market. For example, the prepayment patterns for mortgage-based securities in the United States have changed fundamentally during the past few years, driven by mortgage broker activity. Prior to the change, a drop in interest rates had to prevail for several months before home owners refinanced their mortgages. Subsequently, the refinancing lag shortened from months to weeks during the rally that ended with the Federal Reserve's interest rate hike in February 1994. Use of historical prepayment data could thus be misleading in determining the expected life of many mortgage securities.

To reduce dependence on historical data, given that history may not repeat itself, some firms use data sets based on implied market information. Sensitivity of the VAR calculation not only to exclusion of any data points but also to use of implied versus historical data and to the specific time period covered by the data set should be tested to reveal a particular VAR's dependence on such assumptions.

Portfolio 2

Portfolio 2 consists of outright and options positions on the S&P 500 equity index contract. This portfolio was designed to satisfy several conditions at construction: (1) The delta, or price change of the portfolio, equals that of the S&P 500; (2) the gamma, or convexity of the portfolio, is non-negative; and (3) the portfolio significantly outperforms the S&P 500 equity index contract under downward shocks. Table 3 describes the composition of and constraints on Portfolio 2, which consists of a long position in the S&P 500 equity index contract plus long and short options on the same index. As with Portfolio 1, the net investment in Portfolio 2 is $1 million.

The traditional risk measures show that the portfolio outperforms the S&P 500 equity index contract by an approximate factor of five times under a negative 20 percent shock to the S&P 500. This outperformance in a bearish scenario is accomplished at the price of underperformance during a comparable rise in the S&P 500. The portfolio's $1 million investment is preserved in the value of the S&P 500 equity index contract.

Figure 5 summarizes the VAR results for Portfolio 2. Over a one-day horizon, VAR ranges between 0.69 percent to 0.91 percent with a 5 percent

Table 3. Portfolio 2: Composition and Constraints

Instrument	Jun 520	Jun 545	Sep 530	Dec 540	S&P 500	Portfolio
Composition						
Type	Put	Call	Call	Put	Long	
Strike versus market	+20	+45	+30	+40	0	
Price	1.95	0.60	14.90	18.45	528.59	
Number	4,157.40	−28,723.80	19,784.80	11,617.00	945.90	
Purchase amount	$8,107	($17,234)	$294,793	$214,335	$499,999	$1,000,000
Delta and gamma						
Unit delta	−0.239	0.105	0.545	−0.503	1.000	
Delta contribution	−0.001	−0.003	0.011	−0.006	0.001	0.002
Unit gamma	0.023	0.015	0.012	0.009	0.000	
Gamma contribution	0.000	0.000	0.000	0.000	0.000	0.000
Scenario analysis						
S&P + 20%	634.31	634.31	634.31	634.31	634.31	
Price/S&P + 20	0.00	90.02	107.91	0.25	634.31	
Position/S&P + 20	$0	($2,585,725)	$2,134,983	$2,897	$600,001	$152,156
S&P − 20%	422.87	422.87	422.87	422.87	422.87	
Price/S&P − 20%	96.12	0.00	0.00	106.91	422.87	
Position/S&P − 20%	$399,623	$0	$49	$1,241,941	$399,997	$2,041,610

expectation and between 1.07 percent to 1.30 percent with a 1 percent expectation. Over a two-week horizon, VAR increases significantly and ranges between 2.31 percent to 3.93 percent with a 5 percent expectation. The two-week, 1 percent expectation VARs are fairly consistent across methodologies, ranging between 3.48 percent and 3.95 percent as illustrated by the bar charts.

Figure 6 compares the results of the historical simulation for Portfolio 2 with the results of the Monte Carlo simulations for one-day and two-week returns, respectively. Of note is the significant change in return patterns based on increasing the holding period from one day to two weeks. In the case of one-day returns, all simulations display low-probability high-return/large-loss expecta-

Figure 5. VAR Calculations: Portfolio 2

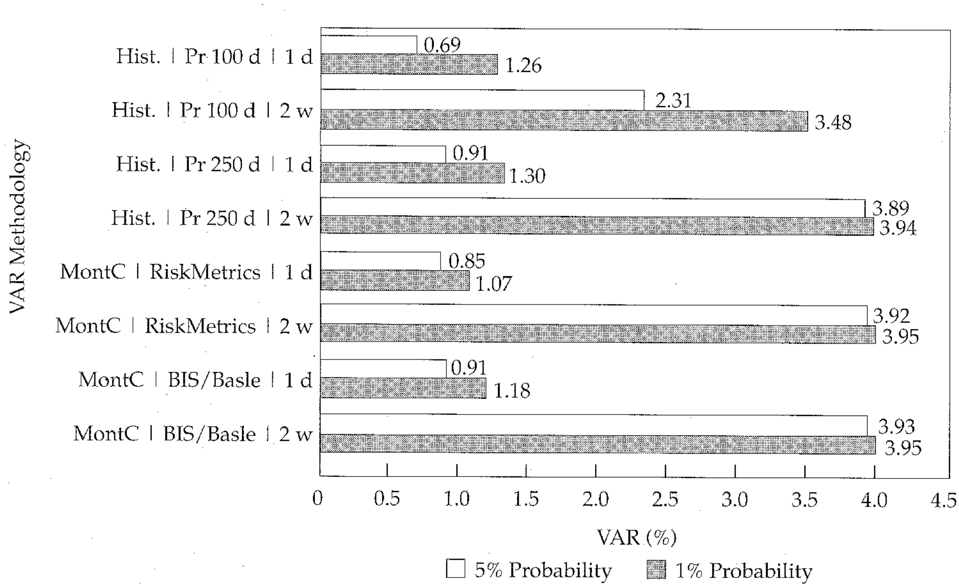

Figure 6. Distributions of One-Day and Two-Week Returns: Portfolio 2

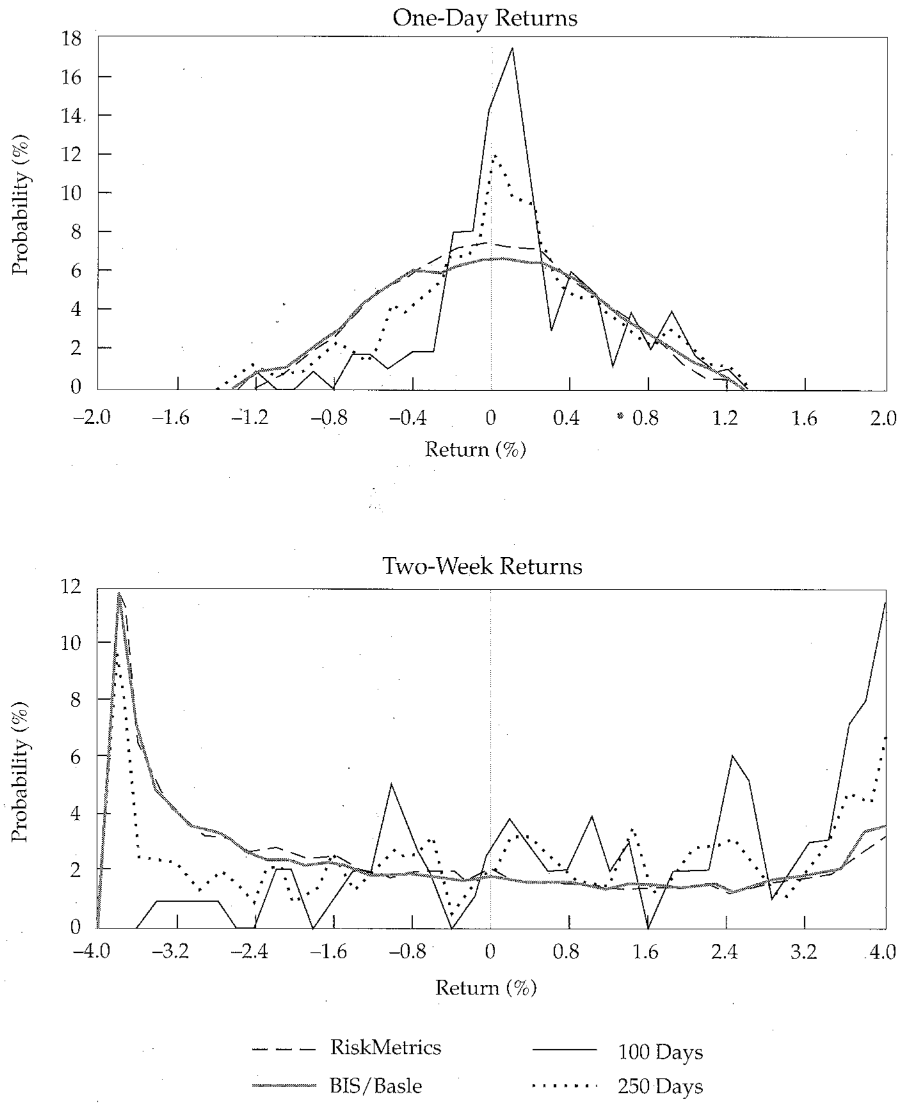

tions. In the case of two-week returns, the distribution changes to display bimodal behavior. This behavior is apparent in the lines that appear to be upside down "normal" distributions. A further observation is that the historical simulations produce high-probability high-return expectations and low-probability large-loss expectations relative to the Monte Carlo simulations.

The VAR calculations for Portfolio 2 expose several weaknesses of VAR, which can be managed with the addition of stress testing and limit policies. These weaknesses are illustrated by viewing the differences in the portfolio's return over various time horizons. The top panel of Figure 7 shows the return on Portfolio 2 on the day of construction as a function of the underlying asset price (the S&P 500). Given the starting S&P 500 level of approximately 529, the portfolio reflects a small positive return (less than 1 percent) if the S&P 500 rises to 544, but its greatest returns occur if the S&P 500 drops below 510. Small downward moves of the S&P 500 (between 510 and 529), and larger upward moves (above 544) produce losses.

The center panel of Figure 7 shows the return on Portfolio 2 at the end of the first two-week holding period, again as a function of the underly-

Figure 7. Portfolio Performance at One Day, at Two Weeks, and after June 545 Call Expires: Portfolio 2

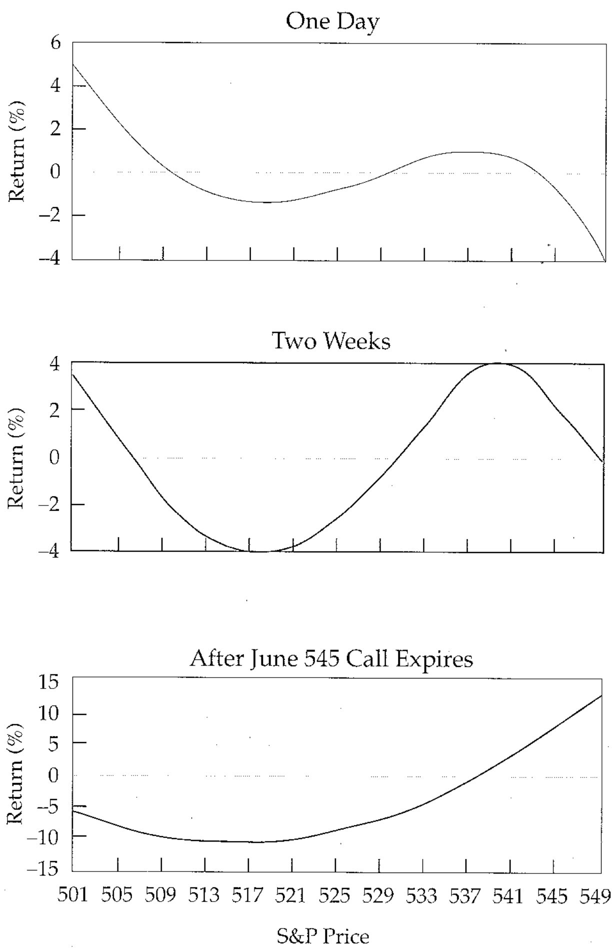

ing asset price (the S&P 500). At about the S&P value of 541, the magnitude of Portfolio 2's performance changes by a factor of four from the first one-day horizon to the first two-week horizon, as illustrated by the amplitude of the graphs (1 percent versus 4 percent, respectively). The price intervals under which Portfolio 2 loses and makes money change as well.

Portfolio 2 poses significantly different risks at different points in time. For example, as shown in the bottom panel of Figure 7, prior to the expiration of the short position in the June call (strike of 545), the portfolio presents the possibility of huge loss. After expiration of the June 545 call, however, Portfolio 2 no longer presents this possibility.

From management's perspective, VAR fails as a sole measure of risk because it does not reveal the true exposure a firm faces. VAR produces a small, finite number in all eight cases (less than 4 percent). Another common risk control measure, a prohibition on portfolios with negative gamma, may also fail. Portfolio 2 displays slightly positive gamma at construction and as of the time horizons in all panels of Figure 7.

Most users combine VAR with stress testing to address questions such as, "How much do I expect to lose the other 1 percent of the time?" Note that in the case of Portfolio 2, stress tests may fail to reveal the true nature of the firm's risk, producing finite pictures of profits and losses that depend upon the level of the S&P 500 assumed. As with VAR, the quality of the answer depends on the inputs, including the financial engineer's ability to select appropriate scenarios.

As experienced during the European currency crisis, the Gulf War, and the Mexican peso crisis, not only are key factors such as "maximum" volatility difficult to predict but also correlation relationships often change substantially during extreme market moves. The increasing complexity and optionality of many derivatives makes relevant scenario selection even harder. Given these challenges, many firms design stress tests to analyze the impact of large historical market moves. In our experience, portfolios do not necessarily produce their greatest losses during extreme market moves. Whether asset based or asset plus liability based, portfolios often possess Achilles' heels that require only small moves or changes between instruments or markets to produce significant losses. Stress testing extreme market moves will do little to reveal the greatest risk of loss for such portfolios. Furthermore, a review of a portfolio's expected behavior over time often reveals that the same stress test that indicates a small impact today indicates embedded land mines with a large impact during future periods. This trait is particularly true of options-based portfolios that change characteristics because of time rather than because of changes in the components of the portfolio. The need for other risk measures—for example, limits that restrict writing uncovered call options—is clear.

Portfolio 3

Portfolio 3 consists of the combination of Portfolio 1 and Portfolio 2. The portfolios are equally weighted, with the net investment in Portfolio 3 totaling $1 million.

Again, the VAR analyses, shown in Figure 8, reveal a wide range of risk profiles for the portfolio.

Figure 8. VAR Calculations: Portfolio 3

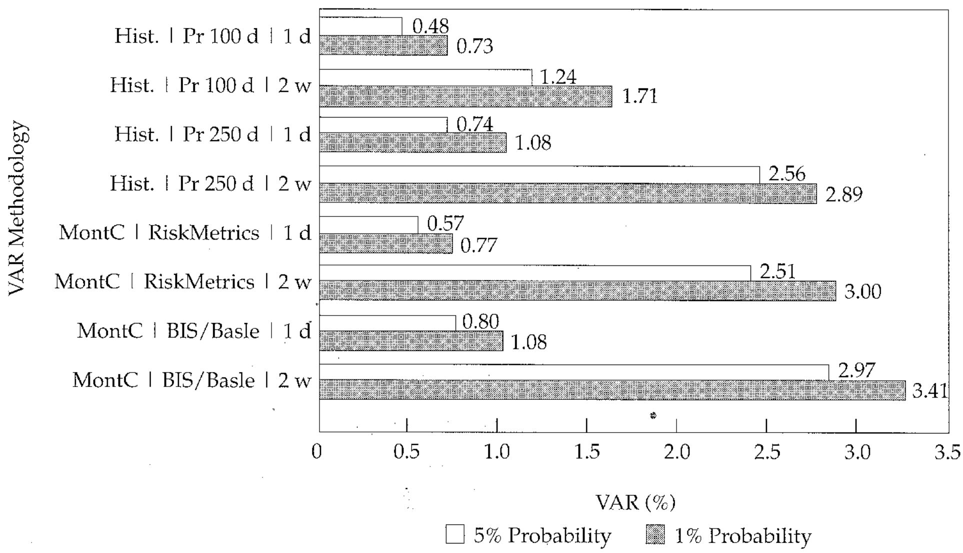

As in the case of Portfolios 1 and 2, historical simulations present a different view of risk than do the Monte Carlo simulations, and Portfolio 3's VAR differences are magnified over the two-week horizon. One-day returns for this multiasset class portfolio display more consistency under the VAR methodologies than one-day returns for the single-asset class portfolios that compose it.

The sensitivity to correlation assumptions is demonstrated by the difference in results between the Monte Carlo simulations under RiskMetrics and BIS/Basle factors. Under the RiskMetrics model, positive correlation is assumed between the Treasury strips in Portfolio 1 and the S&P 500 equity index positions in Portfolio 2. Under the BIS/Basle method, the correlation is assumed to be 1 between long positions and –1 between long and short positions. Not surprisingly, the BIS/Basle VAR factors are higher than RiskMetrics factors in all cases.

Figure 9 compares the results of the historical simulation for Portfolio 3 with the results of the Monte Carlo simulations for one-day and two-week returns. As with Portfolio 2, return patterns change significantly when the holding period is increased from one day to two weeks, but high-probability extreme events no longer occur.

Correlation assumptions are an important aspect of VAR. Firms select quite different answers to which exposures are allowed to offset each other and by how much. For example, is the Japanese yen correlated with movements in the Italian lira or the Mexican peso? Is the price of Saudi Light correlated with movements in the price of natural gas? If so, by how much? VAR requires that the user determine correlations not only within markets (for example, currency underlyings or commodity underlyings) but also across markets (for example, how do changes in the bond market in the United States affect the Australian equity market?). Given a portfolio with multiple instruments within and across markets, VAR varies significantly under alternate correlation assumptions. Pension funds have addressed the issues of correlation for decades in studying strategic versus tactical allocation of assets. Correlation issues are also a crucial component of performance measurement across asset classes. A single approach to assessing correlation does not exist, and opposite views are common. For example, what happens when a market breaks through its historical or implied trading pattern and violates the correlation assumption in place?

Recently, many currencies that previously had displayed little or no correlation with the Mexican peso made sympathy moves during the devaluation of the Peso. In some cases, the increased volume of barrier options on spreads (also known as knock-out or knock-in options) has been blamed for unexpected high correlations during periods when market levels approach strike levels, with both the writers and the buyers of the barriers

Figure 9. Distributions of One-Day and Two-Week Returns: Portfolio 3

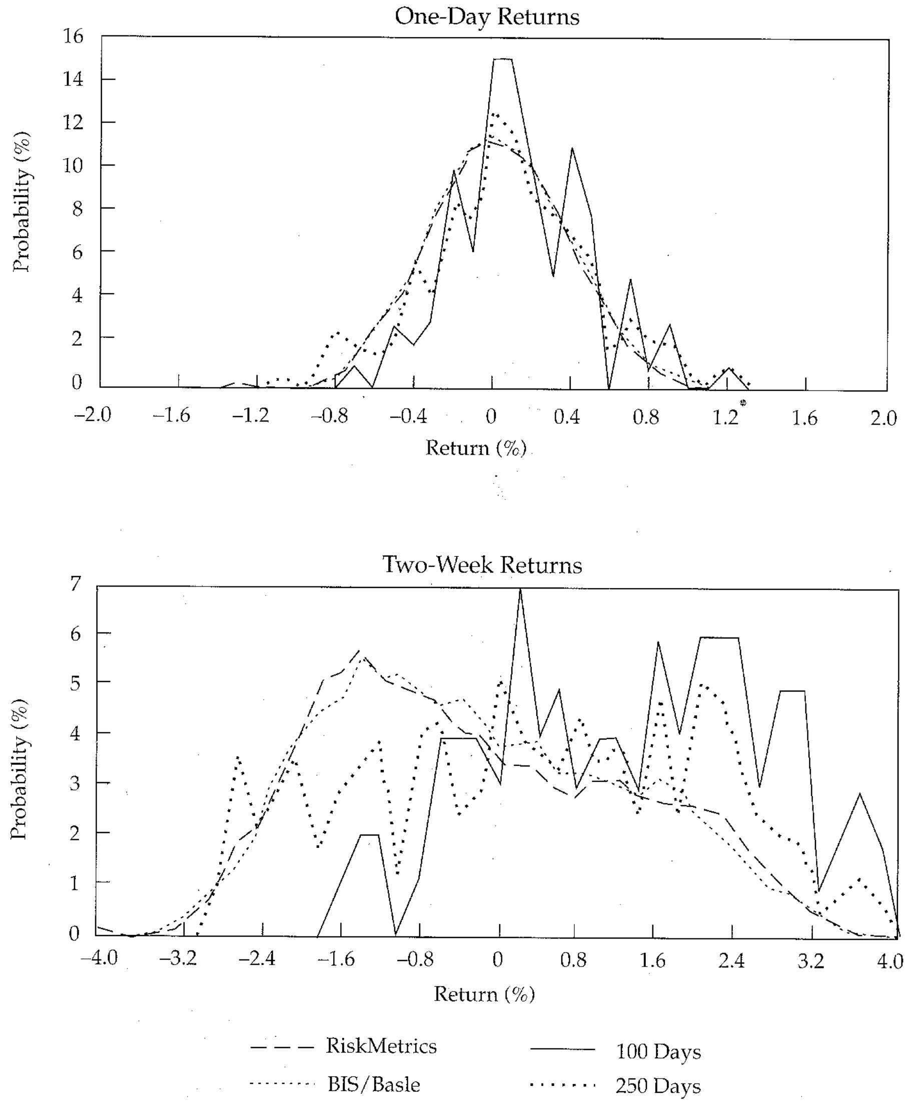

suspected of trading in large volume to influence the outcome.

Correlation assumptions also can mask risks that may be significant for many firms. For example, many portfolios display embedded rollover risk created through timing mismatches. An example is the common strategy that funds use to hedge long-dated foreign currency positions by rolling over short-dated forward foreign exchange contracts. Under common time horizons for VAR and its correlation assumptions, a flat currency risk position often appears. This pattern can mask the long-term rollover risk intrinsic in hedging the currency risk of 10- or 20-year securities with one- to three-month currency contracts.

In our review of different approaches to VAR, some firms assumed that all cash flows were correlated across all markets and others assumed a lower degree of correlation. Sophisticated mean–variance models—for example, the one used to compute the RiskMetrics data set—allow correlation for all instruments across all markets that are covered. At the other extreme are models that allow correlation only within asset classes (e.g., fixed income, foreign exchange, equity) and require perfect positive correlation across risk-factor groups. An example of this approach is the proposed amendment to the 1988 Basle Capital Accord on market risks.

VAR requires the use of mathematical models to value individual instruments, as well as to value the aggregate portfolio. Variance in the valuations produced by widely accepted models (termed "mark-to-model" risk) are well-documented and the subject of many research articles.[4] For example, the Black–Scholes versus Hull and White options models can produce differences of 5 percent or more in pricing, even when all input data are identical. In addition, the selection of probability distribution(s) (an assumption of anticipated or experienced market behavior) in one VAR model versus another is a topic of great debate among theoreticians and practitioners.

CONCLUSION

New studies of the differences in VAR are contemplated. Such studies will use additional computational techniques, alternate assumptions, and a broader range of portfolios in terms of number of positions, type of positions, and number of asset classes. As highlighted by the three portfolios, the picture of expected capital at risk is wildly dependent upon the VAR methodology and the assumptions behind the specific calculation. Not only do the eight VAR results for the individual portfolios differ significantly, but the magnitude of the difference does not follow a clear pattern with increasing complexity of the portfolio. Thus, dealers and end-users are in a precarious position. The dependence on technology and skilled professionals is greater than ever before. Although this dependence has produced invaluable advances in financial engineering and risk management, some firms have been lulled into a false sense of security. Often, firms forget the degree to which the output of models depends upon the modeler's perspective and assumptions. A firm's senior management and directors or trustees are shocked to learn that their firm's risk reports can change dramatically under alternate assumptions. This fact is even more surprising when the "alternate" assumptions are those they consider to be likely or reasonable.

Some firms make the mistake of equating VAR under a 99 percent expectation to the certainty or confidence that the firm will not lose more than the stated amount more than 1 percent of the time (i.e., fewer than three business days a year). As demonstrated by the sample portfolios, the 99 percent VAR changes significantly based on the time horizon, data base, correlation assumptions, mathematical models, and quantitative techniques that are used. Accordingly, VAR does not provide certainty or confidence of outcomes, but rather an expectation of outcomes based on a specific set of assumptions. Furthermore, many risk variables such as political risk, liquidity risk, personnel risk, regulatory risk, phantom liquidity risk, and others cannot be captured through quantitative techniques. Yet, as demonstrated by recent, well-publicized losses, such variables can cause significant risk. For this reason, VAR must be supplemented not only with stress testing but also with prudent checks and balances, procedures, policies, controls, limits, random audits, and appropriate reserves.

The BIS, the Group of Thirty, the Derivatives Product Group, the International Swaps and Derivatives Association, and many national regulators have declared VAR fundamental to current best practices in risk management. But models and math, although necessary to manage risk, are not sufficient to do so. The inability to capture many qualitative factors and exogenous risk variables points to the need to combine VAR with stress tests, checks and balances, procedures, policies, controls, limits, and reserves. Perhaps the limitations of quantitative techniques explain the recent announcement by Moody's that although 25 percent of its volatility rating for funds will be based on VAR, the remaining 75 percent will be based on qualitative factors. In sum, mathematics is integral to finance, but finance does not always follow mathematics.

Some regulators propose to allow firms to use their own internal VAR models plus assumptions, but others do not. Although the Basle amendment allows banks to select their own internal models to calculate VAR (subject to the proviso that certain assumptions are required for VAR's key factors and to oversight by national regulators), the National Association of Insurance Commissioners' proposed risk-based capital standard (RBC) requires insurers to use a single, rigid approach for their VAR-type calculations.[5] The use of a single, rigid approach may have a deleterious effect. Namely, because of a rigid correlation assumption, RBC may penalize the insurer for a successful asset allocation strategy. From a regulatory standpoint, the choice of standardized versus internally selected VAR approaches presents difficult trade-offs. The use of a single, rigid approach may stymie the development of new and improved risk measurement, creating expense for regulators, shareholders, employees, and taxpayers alike. A proper understanding of the assumptions behind VAR, as well as its limitations and pitfalls, will help all to benefit from this powerful technique.[6]

NOTES

1. Note that selection of key statistical parameters such as the mean and variance can significantly affect distributions and, therefore, the results of simulations.
2. Duration is defined here as the change in price with respect to yield.
3. The Monte Carlo simulations are based on the assumptions that the Treasury strip yields are lognormally distributed and that the average change, or "drift," in each yield is zero.
4. See Tanya Styblo Beder, "The Realities of Marking to Model," *Bank Accounting & Finance*, vol. 7, no. 4 (Summer 1994):4–12.
5. Note that a single RBC approach exists for life insurance companies, and a single alternate approach exists for property/casualty companies.
6. The author wishes to thank Frank Iacono, Maarten Nederlof, Tom Riesing, Anil Suri, and Charles Taylor for their invaluable assistance and input during the preparation of this article.