Exploring the Limitations of Value at Risk: How Good Is It in Practice?

ANDREAS KRAUSE

ANDREAS KRAUSE
is a lecturer in finance at the University of Bath School of Management in the UK.
mnsak@bath.ac.uk

Many persons might have attained to wisdom had they not assumed that they already possessed it.

— Seneca

In the early 1990s we witnessed a number of spectacular company failures attributed to the inappropriate use of derivatives and a lack of sufficient internal controls, e.g., Orange County (1994, losses of US$1.8bn), Metallgesellschaft (1994, US$1.3bn), Barings (1995, US$1.3bn), and Daiwa (1995, US $1.1bn). Although in the decades before these events entire financial systems collapsed, not only in developing countries like Argentina, Brazil, or Mexico but also in developed countries as the U.S. savings and loan crises in the 1980s and the Japanese banking crises of the 1990s show, these events made apparent the risks faced by companies and consequently financial systems. Losses could no longer be attributed only to bad government policies and a lack of stringent supervision. The need for an improved risk management, especially for financial organizations, became clear.

In response to these events, a new method of measuring risk had been developed that focused on the losses that companies can make with a reasonable probability: value at risk (VaR). Developed in 1993 and made widely available since 1994 through J.P. Morgan's RiskMetrics®, complemented by CreditMetrics® in 1997 and CorporateMetrics® and PensionMetrics® in 1999, VaR has been proven very popular with financial institutions and regulators alike—the Bank for International Settlements and the SEC, among others.

Originally VaR was intended to measure the risks in derivatives markets, but it became widely applied in financial institutions to measure all kinds of financial risks, primarily market and credit risks. It also evolved from an instrument to measure risks to a tool in active risk management. Recently it made a step toward becoming an important method in enterprise-wide risk management by integrating all risks into a single framework.

VaR has moved well beyond use in financial institutions. At first companies with a large exposure to financial markets, e.g., commodity trading houses, applied it to these kinds of activities, and from there it spread to other business areas. Nowadays VaR is embraced by an ever-increasing number of individual companies using it as their chosen method to develop an enterprise-wide risk management. This trend is clearly facilitated by the fact that VaR is easily understood by non-specialists, which cannot be said of many alternative approaches.

The dangers of the widespread use of VaR are an overreliance on the results it provides, misinterpretation, and even misuse. It is therefore imperative that anyone using VaR be aware of its problems and limitations. In this article we will explore in detail these limitations, which unfortunately do not feature prominently in most books on VaR or risk management in general. We assume that the reader has a basic understanding of risk management and, ideally, VaR.

THE BASIC IDEA OF VaR

> *The whole of science is nothing more than a refinement of everyday thinking.*
>
> — Albert Einstein

Ever since Markowitz [1952], it has become common to measure the risk of an investment by the standard deviation of outcomes. However, the usual perception of risk is that of a possible loss relative to a benchmark. Any risk measure therefore should concentrate on these losses rather than the variability of outcomes which also includes profits, usually not regarded as risk. Using the standard deviation as a risk measure therefore does not address the way risk is usually seen.

It would be more appropriate to concentrate on the lower tail of the distribution of outcomes where the losses are located. As being concerned with the size of these possible losses, a useful risk measure would be one which reflects the amount that can reasonably be lost. We therefore use the expected worst loss that occurs with a certain probability as our risk measure. For example, we could then report the risk by stating that with a probability of 99% (the confidence level) the losses do not exceed US$13m. Or, in other words, the losses exceed US$13m with a probability of only 1%. Obviously we have to specify the time horizon in which these losses may occur. At which confidence level the risk is reported, depends on the risk aversion of the individuals involved, besides regulatory requirements.

The above idea is the basis for the definition of value at risk (VaR). Let V be the value of an investment at the end of the time horizon. We then define a V^* such that

$$Prob(V \leq V^*) = \int_{-\infty}^{V^*} dF(V) = 1 - c \tag{1}$$

where c denotes the confidence level and $F(V)$ the cumulative distribution function of V. Hence the value of the investment is below V^* with a probability of $1 - c$. Based on this equation we can now define the VaR. We only have to define a benchmark such that any outcomes below this benchmark are regarded as losses. It is common to define losses either relative to the status quo, i.e., the current value of the investment V_0, or relative to the expected outcome $E[V]$, such that any opportunity costs are included. This gives rise to the following definitions of the VaR, depending on the benchmark:

$$\begin{aligned} VaR_{c,\Delta T}^{Zero} &= V_0 - V^* \\ VaR_{c,\Delta T}^{Mean} &= E[V] - V^* \end{aligned} \tag{2}$$

where ΔT denotes the time horizon of the possible losses. In most cases it is more convenient to express the VaR in terms of returns. Define R^* and μ such that

$$\begin{aligned} V^* &= (1 + R^*)V_0 \\ E[V] &= (1 + \mu)V_0 \end{aligned} \tag{3}$$

We then can rewrite Equation (2) as

$$\begin{aligned} VaR_{c,\Delta T}^{Zero} &= -V_0 R^* \\ VaR_{c,\Delta T}^{Mean} &= -V_0(R^* - \mu) \end{aligned} \tag{4}$$

Exhibit 1 illustrates this definition of the VaR. In order to apply VaR in risk management we now have to estimate the relevant parameters. While the determination of the current value of the investment, V_0, and the expected return, μ, usually do not impose difficulties, it is much more complicated to estimate R^* properly. There are generally two methods to estimate R^*, the non-parametric method and the parametric method.

The non-parametric method uses historical data and infers from them directly the $(1 - c)$-quantile of returns, which gives R^*. This estimate of R^* is then used to calculate the VaR according to Equation (4).

The alternative parametric method assumes that returns follow a certain distribution. We then use past data on returns to estimate the relevant parameters of the distribution. These parameter estimates are then used to calculate, in most cases analytically, the $(1 - c)$-quantile of this distribution, which as before gives an estimate of R^* that can be used to determine the VaR from Equation (4).

The most widely used distribution for this purpose is the normal distribution, although other distributions like the t-distribution, generalized error distribution, or the Lévy-distribution have been used.

EXHIBIT 1
Definition of Value at Risk

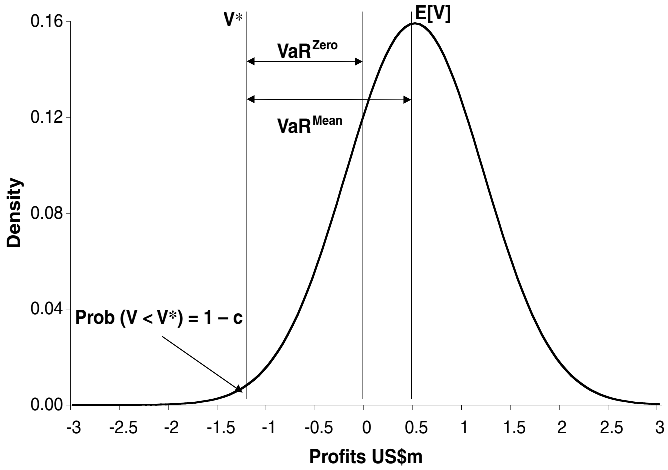

Let $\phi(z)$ denote the standard normal density function. We can now define an α such that

$$\int_{-\infty}^{-\alpha} \phi(z)dz = 1 - c \tag{5}$$

It is then straightforward to show that

$$R^* = \mu - \alpha\sigma \tag{6}$$

where σ denotes the standard deviation of returns. Inserting into Equation (4) gives

$$\begin{aligned} VaR_{c,\Delta T}^{Zero} &= V_0(\alpha\sigma - \mu) \\ VaR_{c,\Delta T}^{Mean} &= V_0\alpha\sigma \end{aligned} \tag{7}$$

Exhibit 2 illustrates the determination of the VaR under the assumption of normally distributed returns. As we can see from Exhibit 3, the results of the two methods do not have to be identical. This is due to the fact that in reality the distribution of returns is not exactly normal, but sometimes substantial deviations are observed, especially in the tails that are evaluated for the VaR. The example compares the distribution of daily Nasdaq Composite Index returns from March 1999 to March 2001 with and without the assumption of a normal distribution.

EXHIBIT 2
Determination of the VaR with Normally Distributed Returns

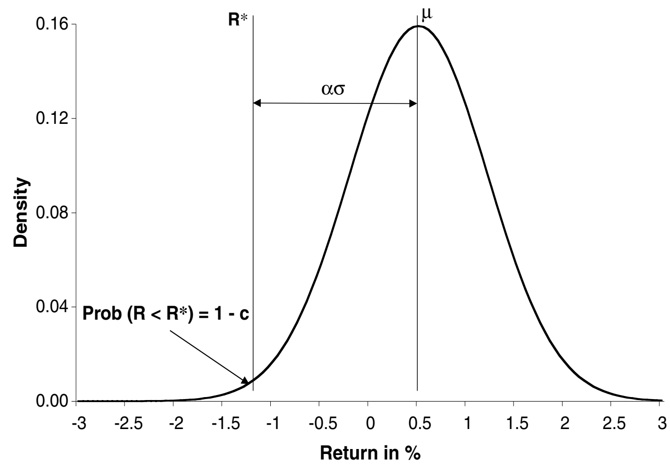

EXHIBIT 3
Comparison of the Parametric and Non-Parametric Methods in Estimating the VaR from the Nasdaq Composite Index March 1999-March 2001

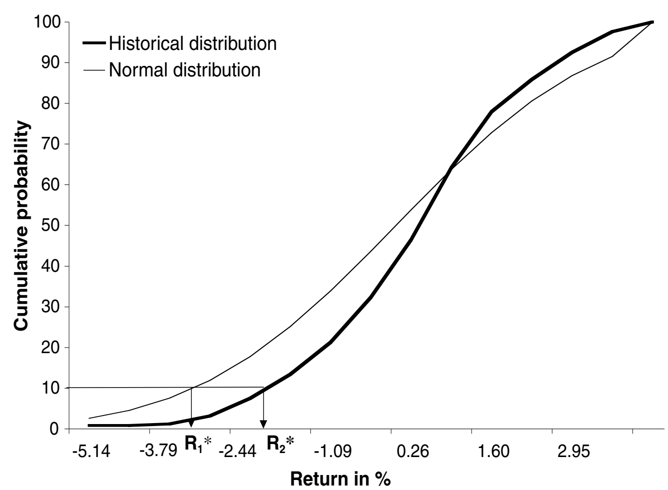

We see from Equation (7) that the standard deviation plays a crucial role in the determination of the VaR. By using historical data, the implicit assumption made is that the distribution does not change over time. Empirical evidence clearly shows that this assumption cannot be upheld as the variance changes over time (heteroscedasticity). In many applications it is now common to model the variance of returns as a stochastic process, usually of the

GARCH family, to provide a better prediction of future volatility, i.e., variance. Here once again parameters of the stochastic process have to be estimated based on historical data and the prediction for the standard deviation is then inserted into Equation (7) to estimate the VaR.

A final aspect that has to be considered is the time horizon ΔT. In most cases returns and their standard deviations are given on an annual basis, which rarely coincides with the time horizon used for the VaR (usually 10 working days for applications in financial markets as required by banking regulation). Hence we have to transform these annual returns into returns for the required time horizon. The simplest way to do this is to assume that returns are identically and independently distributed in each time period. With this assumption we can easily transform any reference expected return, μ^*, and standard deviation, σ^*, into the required time horizon used to determine the VaR:

$$\begin{aligned} \mu &= \mu^* \Delta T \\ \sigma &= \sigma^* \sqrt{\Delta T} \end{aligned} \tag{8}$$

where ΔT is measured relative to the time length of the reference period; e.g., for μ^* and σ^* given annually, a time horizon of one month corresponds to $\Delta T = 1/12$.

We will return to several problems associated with estimating R^* later, but will first explore whether VaR is such a good risk measure as claimed above.

THE QUALITIES OF VaR AS A RISK MEASURE

There is always a well-known solution to every human problem—neat, plausible, and wrong.

H.L. Mencken

Artzner et al. [1999] define desired properties of risk measures and call those fulfilling these properties *coherent*. They ensure that investments which are more risky are assigned a higher value of the risk measure.

As long as returns of an investment have an elliptically shaped distribution, e.g., are normally distributed, it can be shown that VaR is always a coherent risk measure. However, in general VaR does not have to be coherent as the following two examples show.

The VaR is defined by a cutoff return, R^*, beyond which the distribution of returns is not further evaluated. Hence any losses larger then R^* are not weighted according to their size and are thus treated alike. Suppose you are making one of the investments whose left tails of the return distributions are shown in Exhibit 4. The first investment has a higher probability of losses around US$4m, while the second investment has this higher probability close to US$5m. Above the 5%-quantile the two distributions are identical, hence the 95% VaR, $VaR^{Zero}_{0.95,\Delta T}$, is the same in both cases. As the losses beyond the VaR are larger in the second case, we would reasonably conclude that this is the riskier investment. The VaR, however, suggests that both investments have identical risk. Clearly VaR is not representing the risk of the two investments properly.

EXHIBIT 4
Example for the Incoherence of VaR

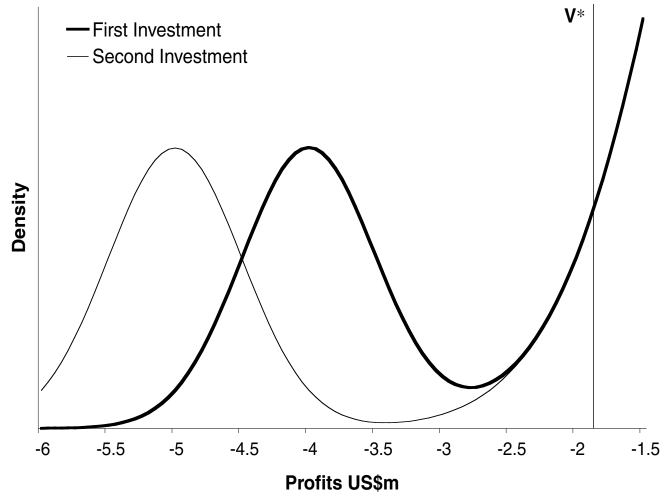

As a second example, suppose that you have 100 assets whose outcomes are independent of each other and the payoff profile at the end of the time horizon is identical for every $i = 1, \ldots, 100$:

$$V_i = \begin{cases} 5 & \text{with probability } 0.99 \\ -100 & \text{with probability } 0.01 \end{cases} \tag{9}$$

Now compare the following two portfolios, the first consisting of 100 units of asset k, $1 \leq k \leq 100$, and the second portfolio consisting of one unit of each of the 100 assets. The diversification of risks in the second portfolio has clearly reduced the risk as a complete loss of the investment is highly unlikely. When calculating the 95% VaR, $VaR^{Zero}_{0.95,\Delta T}$, we can easily derive that for the first portfolio it is –500 and for the second, diversified portfolio it is 200. Therefore, according to the VaR measure, the second portfolio should be riskier. This clearly contradicts our previous argumentation. Therefore the risk measure is also in this case not coherent.

In both of these examples the actual risk and the results from the risk measure are not consistent with each other, and hence VaR is not a coherent risk measure. It can be shown that the reason for this finding is that unless specific conditions on the distribution of outcomes are met, VaR is not subadditive, i.e., $VaR(V_1 + V_2) \leq VaR(V_1) + VaR(V_2)$ is generally not valid.

In many cases the conditions for the validity of subadditivity are (at least approximately) fulfilled and VaR is (nearly) coherent. But remarkably, the payoff distributions of options and similar derivatives mostly violate these conditions. And it was the use of such instruments that caused many of the large losses in the early 1990s and led to the development of VaR. Another example of payoffs that violate the conditions for VaR being coherent are catastrophic risks. Here the large losses that occur with a very small probability, with the payoffs otherwise having the desired properties, cause this violation.

We may summarize this section by stating that VaR is not a coherent risk measure, and thus the VaR may in certain situations not reflect the risk adequately. We show below how this might be exploited by companies seeking to manipulate their VaR.

ESTIMATING THE VAR

First get the facts, then you can distort 'em as much as you want.

— Mark Twain

The VaR estimate is based on past data, i.e., it uses the historical distribution of outcomes of the investment. However, to evaluate the risk of an investment, it is of no interest how large this risk has been in the past, but rather how much risk there is within the time horizon; hence the future distribution of outcomes would be the relevant to consider. As long as the distribution of outcomes remains stable, i.e., does not change over time, the VaR can easily be deducted from the historical distribution. In reality the distributions are not stable over time; most notably, the variance of outcomes and the correlations change. Relying solely on historical data can therefore give an inadequate risk measure.

Models to derive and estimate the future distribution, or certain parameters of this distribution like the variance or correlations, are often employed to address this problem. Forecasting techniques usually use time series models, for example the GARCH models to forecast the variance of outcomes, or calculate the implied volatility in option prices for this purpose. The predicted and the realized distribution of outcomes in many cases show large differences, despite these forecasting efforts. Exhibit 5 shows how much the distribution of outcomes, and hence the VaR, may differ between subsequent years. We used the daily returns of the Nasdaq Composite Index from March 1999 to March 2000 and from March 2000 to March 2001, respectively. Usually the distribution is estimated with only relatively few and recent data, e.g., from the past year, supposing that the change of the distribution is smaller the closer the data are to the present time. But as this example shows, even then large deviations may be observed.

Even if the problem of forecasting the future distribution of outcomes can be overcome with a reliable fore-

EXHIBIT 5

Example of the Changes in the Distribution in Subsequent Years for the Nasdaq Composite Index March 1999-March 2001

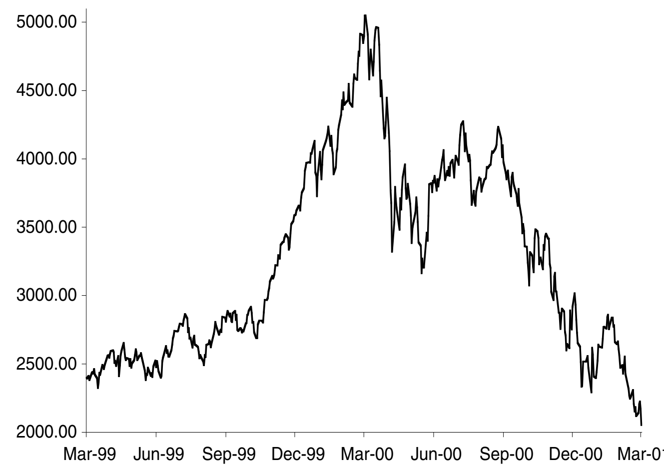

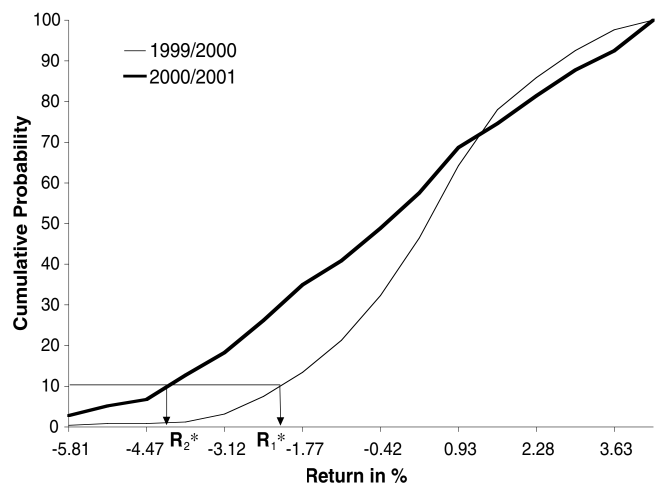

casting method, a problem arises from the small size of the data set used in its estimation. As extreme events causing very large losses are rare, they are in most cases not included in the data. Hence the VaR estimate gives only a risk assessment of the investment under "normal" market conditions. Extreme events like a stock market crash are not included. What is especially not included is that the correlations in most of such cases tend to increase significantly, increasing losses further by not allowing for the diversification effect to work as anticipated. To overcome this shortcoming, VaR should be complemented by stress testing to enable risk managers to include such events into their overall risk assessment.

Estimation Error

Unfortunately, even with this issue resolved, there remains a problem with the estimation of the VaR itself. As the true probability distribution is not known in general, it has to be estimated from the data. A good estimation of the lower tail of the distribution is of essential importance for the VaR estimation. The complication arises from the fact that by definition only few observations are made at the tails, and hence the estimation errors are large.

Kendall [1994] shows that the standard error of the *c*-quantile q is given by

$$SE(q) = \sqrt{\frac{c(1-c)}{Tf(q)}} \tag{10}$$

where T denotes the total number of observations and $f(\cdot)$ the density function evaluated at the quantile q. Exhibit 6 shows the confidence intervals of such an estimation. The graph shows the estimate of the VaR with a time horizon of one day for different confidence levels with its 95% confidence intervals for supposedly standard normally distributed returns, an investment of US$100m, and 250 observations, corresponding to one year of daily data. We can easily see that the estimation error is substantial and increases the higher the confidence level becomes.

When explicitly assuming the outcomes to be normally distributed, information on the entire distribution can be used in the estimation of the tails. Therefore the standard error reduces significantly:

$$SE(q) = \alpha\sigma\sqrt{\frac{1}{2T}} \tag{11}$$

EXHIBIT 6
Estimation Errors of the VaR

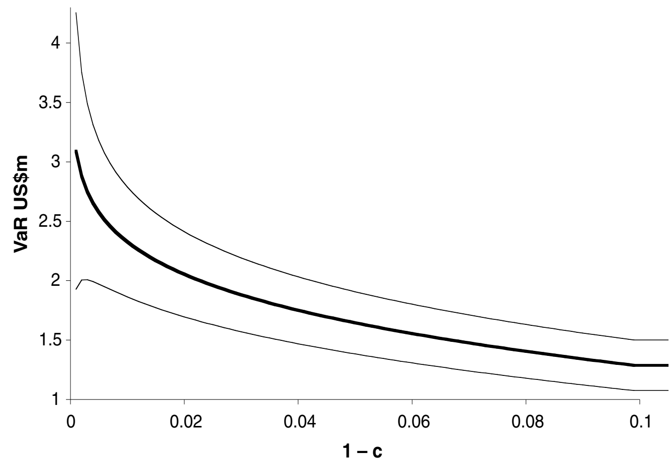

EXHIBIT 7
Standard Errors in Parametric and Non-Parametric Estimation

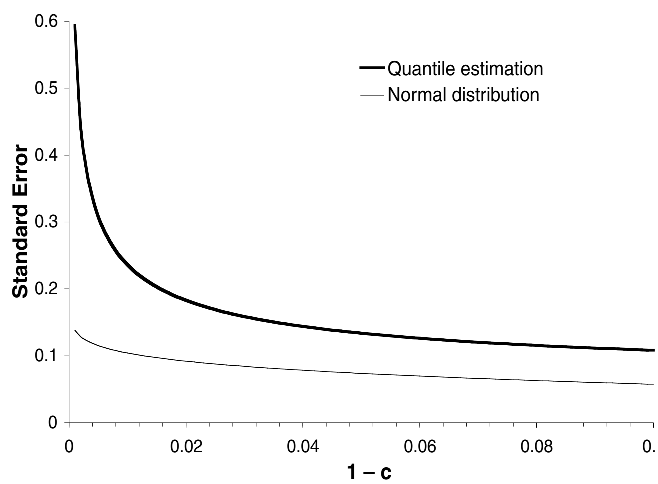

where α is defined as before. Exhibit 7 compares the two standard errors assuming returns to follow a standard normal distribution and having $T = 250$ observations. Although the estimation error is reduced, it has to be kept in mind that this has only been achieved with the additional assumption of normally distributed outcomes. This assumption, as stated above, is usually not fulfilled, and the estimation is therefore subject to errors arising from this misspecification. The overall error therefore may not be reduced at all when making this additional assumption.

A further assessment of the error can be derived by comparing the estimated VaR with the maximal VaR that is consistent with the data. This relationship can be derived using the theorem of Bienaymé-Chebyshev, which states that for each random variable x with a finite variance σ^2 and expected value μ it is

$$Prob(|x-\mu| \geq \varepsilon) \leq \frac{\sigma^2}{\varepsilon^2} \tag{12}$$

Assuming a symmetric distribution we get

$$Prob(x-\mu \leq \varepsilon) \leq \frac{1}{2}\frac{\sigma^2}{\varepsilon^2} \tag{13}$$

Using the definition of the VaR from Equation (2), we can rewrite Equation (13) as

$$c = Prob\left(V - E[V] \leq -VaR_{c,\Delta T}^{Mean}\right) \leq \frac{1}{2}\frac{\sigma^2}{\left(VaR_{c,\Delta T}^{Mean}\right)^2} \tag{14}$$

Hence we obtain that

$$VaR_{c,\Delta T}^{Mean} \leq \frac{\sigma}{\sqrt{2c}} \tag{15}$$

Therefore the maximal VaR is given by $\sigma/\sqrt{2c}$. Exhibit 8 compares the maximal VaR with the estimated VaR assuming returns to be standard normally distributed and the current investment to be US$100m. We see that the difference can be substantial, especially for large confidence levels.

Stahl [1997] uses this result to justify the rule set by the Basel Committee on Banking Supervision that to obtain the relevant exhibits for the capital adequacy rules, the estimated VaR has to be multiplied by a factor of three. This is exactly the factor by which the 99% VaR, the basis for the regulation, has to be multiplied to give the maximum VaR as above.

Estimation Bias

Companies are usually exposed to multiple risks, and the correlations between those risks have a crucial importance for the VaR estimate as they determine the strength of the diversification effect which reduces the total risk. Like the VaR estimate, the estimate for the correlation matrix is usually based on a relatively small number of past data, typically one year of daily data or about 250 data points. As the number of instruments traded is often quite large, the dimension of the correlation matrix in some cases may be close to that of the number of observations, or may even exceed them. It has been shown by Ju and Pearson [1999] that the VaR estimate under these circumstances shows a substantial downward bias. The risk is substantially underestimated; under realistic assumption it can easily be that the estimated VaR is only about half of the true VaR.

EXHIBIT 8
Estimated and Maximal VaR

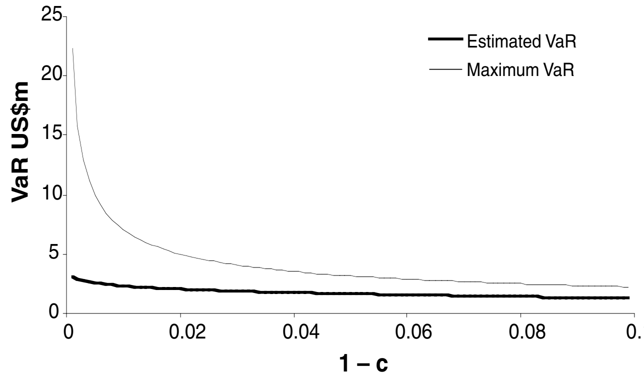

The reason for this finding is that when the dimensionality of the correlation matrix is sufficiently close to the dimensionality of the data vector, the estimated matrix becomes nearly non-singular (if the dimensionality of the correlation matrix exceeds the dimensionality of the data vector, it actually is non-singular). In these cases it is possible to construct portfolios with an estimated VaR close to zero, i.e., being virtually risk-free according to the estimated VaR, although they are inherently risky.

Given those econometric problems in estimating the VaR—changing parameters, large estimation errors, and downward bias—it is apparent that it is far from being a precise risk measure. It may only give a rough indication of the magnitude of risks in a company.

REGULATING RISK WITH VAR

Any statistical relationship will break down when used for policy purposes.

— Charles Goodhard

Exploiting the Estimation Bias

When using VaR to restrict the risk taking within a company or as the basis for performance-related pay, the behavior of those targeted by the rules will be affected in many cases not only in the desired direction. Suppose a trader faces a VaR constraint or performance-related pay based on his VaR. In this case traders have a clear incentive to understate the risk they actually take. They can easily do so by exploiting the bias in the VaR estimate as explained in the previous section. They choose a portfolio whose VaR estimate is lower than the actual risk taken. As the expected return is determined by the "true" risk, the trader will, at least on average, generate a large profit with a relatively low VaR, compensating him for the true risk. In this sense a trader who is active in multiple markets, e.g., by trading on the yield curve or in related stocks, can systematically exploit the way VaR is estimated.

This systematic exploitation of the estimation bias requires the trader to be able to coordinate his actions in several markets. As usually a single trader or a single trading desk only trades in a small number of assets, relative to the total number of assets traded in a company, it is very unlikely that the bias is exploited at the corporate level. The bias of uncoordinated individual exploitations is usually not easily aggregated at the corporate level as individual effects cancel each other out such that the total risk of the corporation is not affected to the same degree.

However, this example shows that individual traders can engage in investments whose risks are systematically underestimated. Although the effect on the risk at the corporate level may be negligible, it may cause the company to pay out bonuses to traders that are too high and thus reduce the profitability of the company.

Manipulating the VaR

It is also possible to use the properties of the VaR by manipulating its value without actually reducing risks. Making use of the observation that VaR is not coherent, a manager can reduce the VaR from a current value of VaR_0 to a desired level of $VaR_1 < VaR_0$ using the following strategy as shown in Danielsson [2001]: write V_0 call options on the investment with a strike price of R_0^*, corresponding to VaR_0, and buy V_0 put options with a strike price of R_1^*, corresponding to VaR_1. This transaction changes the probability distribution of outcomes as illustrated in Exhibit 9. The VaR is reduced to VaR_1, as required, but the downside risk has been increased and the expected return reduced.

Such strategies may be employed by individual traders who otherwise would violate their risk limits as defined by VaR. Furthermore, companies as a whole or individual divisions could easily coordinate their activities to manipulate their VaR in such a manner. The incentives are especially strong for financial institutions that are subject to capital adequacy requirements based on their VaR. To comply with the regulation they could easily and legally resort to such methods in order to avoid detection or punishment. Hence unlike the exploitation of the bias in estimating the VaR, this property of the VaR, based on its incoherence, can also be exploited at the corporate level.

EXHIBIT 9
Manipulating the VaR with Options

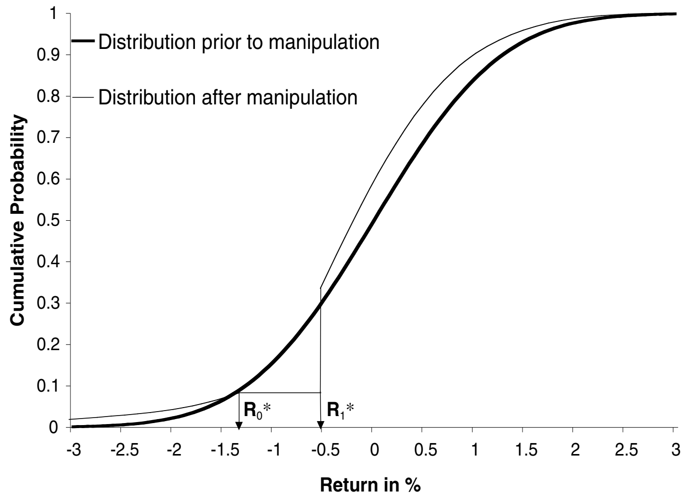

Systemic Risk

VaR has slowly become the standard method in risk management, especially in financial markets and banking. Consequently many market participants will apply similar or identical techniques in risk management, and reactions to events and strategies developed are likely to be similar, too. As the market data used to estimate the VaR are therefore endogenously determined by the behavior of market participants, we could face the problem that a crisis arises only because actions taken as the result of a VaR estimate become self-fulfilling, e.g., from the liquidation of an investment. In that sense the VaR estimate becomes obsolete as a risk measure because the distribution of outcomes changes significantly, what is not taken into account in the VaR estimation.

The reaction to an unforeseen event that violates the VaR constraint, e.g., through an increase in the risk, would be to sell this investment. Thus with many market participants acting in the same way, this may then cause the loss to be realized. A similar result was observed in the 1987 stock market crash at the NYSE, where portfolio insurance activities, designed to prevent large losses, actually aggravated them. Hence a regulation of risk via VaR, and probably any other strict mechanism, can under certain circumstances actually increase the systemic risk that it was intended to reduce.

THE BENEFITS OF USING VaR

> *You don't want the decimal point—you want the order of magnitude.*
>
> — Lev Borodovsky

The previous sections have shown that VaR is not unproblematic to use, it is not a coherent risk measure, its estimation is subject to large errors, the estimate is downward biased, and these shortcomings can be exploited by individuals within the company as well as the company as a whole. However, these shortcomings do not necessarily imply that VaR is not a useful tool in risk management.

The obvious benefit of VaR is that it is easily and intuitively understood by non-specialists. It can therefore be well communicated within a company as well as between the company and regulators, investors, or other stakeholders. Furthermore, it can address all types of risks in a single framework, which not only allows the aggregation of risks but also further facilitates communication.

With the help of VaR we can address most problems arising from risks. It can be used to set risk limits for individual traders, divisions, and the entire company; it facilitates performance measurement and can be used as the basis for performance-related pay; it facilitates decisions on the allocation of capital and gives an indication of the total capital requirement; finally, it can help to decide which risks to reduce, if necessary.

No other risk management system developed thus far addresses all these aspects in a single framework, while still being accessible to managers and a wide range of employees. Therefore VaR has proved to be a very useful tool that is readily accepted, despite its shortcomings. But it is exactly these shortcomings that limit the extent to which VaR can be used. The VaR estimate should not be taken as a precise number, but it provides an indication as to how much risk is involved. It also aids in detecting any trends in the behavior of individuals, divisions, or the company as a whole.

Properly used, VaR is a powerful but still simple tool in risk management. On the other hand, overreliance on its results and justifying important decisions solely on its basis are likely to be counterproductive. No risk management system can replace the sound judgment of managers, and those using it should be aware of its limits.

The benefits of the simplicity of VaR cannot be underestimated. A much more precise and improved method that is not understood by decision makers is of much less value than an easily understood method, even if it gives only a rough estimate of the risks involved, provided these limits are understood. Results based on systems that are not understood are either ignored or used without the necessary precautions. In both cases decisions are likely to be inferior.

CONCLUSIONS

> *After all, you can't expect to be paid a fat salary just for plugging numbers into a formula.*
>
> — Brealey and Myers [2000]

In this article we explored limitations in the use of value at risk (VaR). It was shown that under certain circumstances VaR does not give an appropriate risk measure, its estimation is subject to large estimation errors, and a downward bias in the estimation can easily be exploited by employees or the entire company to their own benefit. We also argued that despite these shortcomings VaR has its advantages that are rooted in the easy access non-specialists have to its basic concepts.

Currently research is under way to improve the estimation procedure for VaR so as to overcome some of the above-mentioned problems. Alternatives to VaR have been developed that employ a similar idea, but provide a better risk measure. One proposed measure is expected shortfall (ES), which uses the expected loss given that the VaR is exceeded. It thus includes more information on the distribution of losses than the VaR and can be shown to be a coherent risk measure. However, currently estimation procedures for ES are subject to even larger estimation errors than VaR, and it is less intuitively understood by non-specialists. Therefore, thus far ES is rarely applied by companies. Future research into how to use ES or other risk measures effectively in risk management may then improve its applicability. In the meantime, the appropriate use of VaR with full awareness of its limitations can improve decision making in companies.

REFERENCES

Acerbi, C., C. Nordio, and C. Sirtori. "Expected Shortfall as a Tool for Financial Risk Management." 2002.

Ahn, D.-H., J. Boudoukh, and M. Richardson. "Optimal Risk Management Using Options." *Journal of Finance*, 54 (1999), pp. 359-375.

Artzner, P., F. Delbaen, and J.-M. Eber. "Coherent Measures of Risk." *Mathematical Finance*, 9 (1999), pp. 203-228.

Brealey, R.A., and S.C. Myers. *Principles of Corporate Finance*. 6th ed. New York: Irwin McGraw-Hill, 2000.

Danielsson. J. "The Emperor Has No Clothes: Limits to Risk Modelling." *Journal of Banking and Finance*, 26 (2002), pp.1273-1296.

Hendricks, D. "Evaluation of Value-at-Risk Models Using Historical Data." *Federal Reserve Bank of New York Economic Policy Review,* 1996, pp. 39-69.

Jorion, P. "Risk^2: Measuring the Risk in Value at Risk." *Financial Analysts Journal*, 52 (1996), pp. 47-56.

Ju, X., and N. Pearson. "Using Value-at-Risk to Control Risk Taking: How Wrong Can You Be?" *Journal of Risk*, 1 (1999), pp. 5-36.

Kendall, M. *Kendall's Advanced Theory of Statistics*. New York: Halsted Press, 1994.

Markowitz, H. "Portfolio Selection." *Journal of Finance*, 7 (1952), pp. 77-91.

Stahl, G. "Three Cheers." *Risk*, 10 (1997), pp. 67-69.

Yamai, Y., and T. Yoshiba. "Comparative Analyses of Expected Shortfall and Value-at-Risk (2): Expected Utility Maximization and Tail Risk." In *Monetary and Economic Studies*, 2002, pp. 95-115.

——. "Comparative Analyses of Expected Shortfall and Value-at-Risk: Their Estimation Error, Decomposition, and Optimazation." In *Monetary and Economic Studies*, 2002, pp. 87-122.

——. "On the Validity of Value-at-Risk: Comparative Analyses of Expected Shortfall." In *Monetary and Economic Studies*, 2002, pp. 57-85.