

VALUE AT RISK (VAR)

What is the most I can lose on this investment? This is a question that almost every investor who has invested or is considering investing in a risky asset asks at some point in time. Value at Risk tries to provide an answer, at least within a reasonable bound. In fact, it is misleading to consider Value at Risk, or VaR as it is widely known, to be an alternative to risk adjusted value and probabilistic approaches. After all, it borrows liberally from both. However, the wide use of VaR as a tool for risk assessment, especially in financial service firms, and the extensive literature that has developed around it, push us to dedicate this chapter to its examination.

We begin the chapter with a general description of VaR and the view of risk that underlies its measurement, and examine the history of its development and applications. We then consider the various estimation issues and questions that have come up in the context of measuring VAR and how analysts and researchers have tried to deal with them. Next, we evaluate variations that have been developed on the common measure, in some cases to deal with different types of risk and in other cases, as a response to the limitations of VaR. In the final section, we evaluate how VaR fits into and contrasts with the other risk assessment measures we developed in the last two chapters.

What is Value at Risk?

In its most general form, the Value at Risk measures the potential loss in value of a risky asset or portfolio over a defined period for a given confidence interval. Thus, if the VaR on an asset is \$ 100 million at a one-week, 95% confidence level, there is a only a 5% chance that the value of the asset will drop more than \$ 100 million over any given week. In its adapted form, the measure is sometimes defined more narrowly as the possible loss in value from “normal market risk” as opposed to all risk, requiring that we draw distinctions between normal and abnormal risk as well as between market and non-market risk.

While Value at Risk can be used by any entity to measure its risk exposure, it is used most often by commercial and investment banks to capture the potential loss in value of their traded portfolios from adverse market movements over a specified period;

this can then be compared to their available capital and cash reserves to ensure that the losses can be covered without putting the firms at risk.

Taking a closer look at Value at Risk, there are clearly key aspects that mirror our discussion of simulations in the last chapter:

1. To estimate the probability of the loss, with a confidence interval, we need to define the probability distributions of individual risks, the correlation across these risks and the effect of such risks on value. In fact, simulations are widely used to measure the VaR for asset portfolio.
2. The focus in VaR is clearly on downside risk and potential losses. Its use in banks reflects their fear of a liquidity crisis, where a low-probability catastrophic occurrence creates a loss that wipes out the capital and creates a client exodus. The demise of Long Term Capital Management, the investment fund with top pedigree Wall Street traders and Nobel Prize winners, was a trigger in the widespread acceptance of VaR.
3. There are three key elements of VaR – a specified level of loss in value, a fixed time period over which risk is assessed and a confidence interval. The VaR can be specified for an individual asset, a portfolio of assets or for an entire firm.
4. While the VaR at investment banks is specified in terms of market risks – interest rate changes, equity market volatility and economic growth – there is no reason why the risks cannot be defined more broadly or narrowly in specific contexts. Thus, we could compute the VaR for a large investment project for a firm in terms of competitive and firm-specific risks and the VaR for a gold mining company in terms of gold price risk.

In the sections that follow, we will begin by looking at the history of the development of this measure, ways in which the VaR can be computed, limitations of and variations on the basic measures and how VaR fits into the broader spectrum of risk assessment approaches.

A Short History of VaR

While the term “Value at Risk” was not widely used prior to the mid 1990s, the origins of the measure lie further back in time. The mathematics that underlie VaR were largely developed in the context of portfolio theory by Harry Markowitz and others,

though their efforts were directed towards a different end – devising optimal portfolios for equity investors. In particular, the focus on market risks and the effects of the co-movements in these risks are central to how VaR is computed.

The impetus for the use of VaR measures, though, came from the crises that beset financial service firms over time and the regulatory responses to these crises. The first regulatory capital requirements for banks were enacted in the aftermath of the Great Depression and the bank failures of the era, when the Securities Exchange Act established the Securities Exchange Commission (SEC) and required banks to keep their borrowings below 2000% of their equity capital. In the decades thereafter, banks devised risk measures and control devices to ensure that they met these capital requirements. With the increased risk created by the advent of derivative markets and floating exchange rates in the early 1970s, capital requirements were refined and expanded in the SEC's Uniform Net Capital Rule (UNCR) that was promulgated in 1975, which categorized the financial assets that banks held into twelve classes, based upon risk, and required different capital requirements for each, ranging from 0% for short term treasuries to 30% for equities. Banks were required to report on their capital calculations in quarterly statements that were titled Financial and Operating Combined Uniform Single (FOCUS) reports.

The first regulatory measures that evoke Value at Risk, though, were initiated in 1980, when the SEC tied the capital requirements of financial service firms to the losses that would be incurred, with 95% confidence over a thirty-day interval, in different security classes; historical returns were used to compute these potential losses. Although the measures were described as haircuts and not as Value or Capital at Risk, it was clear the SEC was requiring financial service firms to embark on the process of estimating one-month 95% VaRs and hold enough capital to cover the potential losses.

At about the same time, the trading portfolios of investment and commercial banks were becoming larger and more volatile, creating a need for more sophisticated and timely risk control measures. Ken Garbade at Banker's Trust, in internal documents, presented sophisticated measures of Value at Risk in 1986 for the firm's fixed income portfolios, based upon the covariance in yields on bonds of different maturities. By the early 1990s, many financial service firms had developed rudimentary measures of Value

at Risk, with wide variations on how it was measured. In the aftermath of numerous disastrous losses associated with the use of derivatives and leverage between 1993 and 1995, culminating with the failure of Barings, the British investment bank, as a result of unauthorized trading in Nikkei futures and options by Nick Leeson, a young trader in Singapore, firms were ready for more comprehensive risk measures. In 1995, J.P. Morgan provided public access to data on the variances of and covariances across various security and asset classes, that it had used internally for almost a decade to manage risk, and allowed software makers to develop software to measure risk. It titled the service “RiskMetrics” and used the term Value at Risk to describe the risk measure that emerged from the data. The measure found a ready audience with commercial and investment banks, and the regulatory authorities overseeing them, who warmed to its intuitive appeal. In the last decade, VaR has become the established measure of risk exposure in financial service firms and has even begun to find acceptance in non-financial service firms.

Measuring Value at Risk

There are three basic approaches that are used to compute Value at Risk, though there are numerous variations within each approach. The measure can be computed analytically by making assumptions about return distributions for market risks, and by using the variances in and covariances across these risks. It can also be estimated by running hypothetical portfolios through historical data or from Monte Carlo simulations. In this section, we describe and compare the approaches.¹

Variance-Covariance Method

Since Value at Risk measures the probability that the value of an asset or portfolio will drop below a specified value in a particular time period, it should be relatively simple to compute if we can derive a probability distribution of potential values. That is basically what we do in the variance-covariance method, an approach that has the benefit

¹ For a comprehensive overview of Value at Risk and its measures, look at the Jorion, P., 2001, Value at Risk: The New Benchmark for Managing Financial Risk, McGraw Hill. For a listing of every possible reference to the measure, try www.GloriaMundi.org.


of simplicity but is limited by the difficulties associated with deriving probability distributions.

General Description

Consider a very simple example. Assume that you are assessing the VaR for a single asset, where the potential values are normally distributed with a mean of \$ 120 million and an annual standard deviation of \$ 10 million. With 95% confidence, you can assess that the value of this asset will not drop below \$ 80 million (two standard deviations below from the mean) or rise about \$120 million (two standard deviations above the mean) over the next year.² When working with portfolios of assets, the same reasoning will apply but the process of estimating the parameters is complicated by the fact that the assets in the portfolio often move together. As we noted in our discussion of portfolio theory in chapter 4, the central inputs to estimating the variance of a portfolio are the covariances of the pairs of assets in the portfolio; in a portfolio of 100 assets, there will be 49,500 covariances that need to be estimated, in addition to the 100 individual asset variances. Clearly, this is not practical for large portfolios with shifting asset positions.

It is to simplify this process that we map the risk in the individual investments in the portfolio to more general market risks, when we compute Value at Risk, and then estimate the measure based on these market risk exposures. There are generally four steps involved in this process:

- The first step requires us to take each of the assets in a portfolio and map that asset on to simpler, standardized instruments. For instance, a ten-year coupon bond with annual coupons C , for instance, can be broken down into ten zero coupon bonds, with matching cash flows:



The first coupon matches up to a one-year zero coupon bond with a face value of C , the second coupon with a two-year zero coupon bond with a face value of C and so

² The 95% confidence intervals translate into 1.96 standard deviations on either side of the mean. With a 90% confidence interval, we would use 1.65 standard deviations and a 99% confidence interval would require 2.33 standard deviations.

until the tenth cash flow which is matched up with a 10-year zero coupon bond with a face value of FV (corresponding to the face value of the 10-year bond) plus C . The mapping process is more complicated for more complex assets such as stocks and options, but the basic intuition does not change. We try to map every financial asset into a set of instruments representing the underlying market risks. Why bother with mapping? Instead of having to estimate the variances and covariances of thousands of individual assets, we estimate those statistics for the common market risk instruments that these assets are exposed to; there are far fewer of the latter than the former. The resulting matrix can be used to measure the Value at Risk of any asset that is exposed to a combination of these market risks.

- In the second step, each financial asset is stated as a set of positions in the standardized market instruments. This is simple for the 10-year coupon bond, where the intermediate zero coupon bonds have face values that match the coupons and the final zero coupon bond has the face value, in addition to the coupon in that period. As with the mapping, this process is more complicated when working with convertible bonds, stocks or derivatives.
- Once the standardized instruments that affect the asset or assets in a portfolio been identified, we have to estimate the variances in each of these instruments and the covariances across the instruments in the next step. In practice, these variance and covariance estimates are obtained by looking at historical data. They are key to estimating the VaR.
- In the final step, the Value at Risk for the portfolio is computed using the weights on the standardized instruments computed in step 2 and the variances and covariances in these instruments computed in step 3.

Appendix 7.1 provides an illustration of the VaR computation for a six-month dollar/euro forward contract. The standardized instruments that underlie the contract are identified as the six month riskfree securities in the dollar and the euro and the spot dollar/euro exchange rate, the dollar values of the instruments computed and the VaR is estimated based upon the covariances between the three instruments.

Implicit in the computation of the VaR in step 4 are assumptions about how returns on the standardized risk measures are distributed. The most convenient assumption both from a computational standpoint and in terms of estimating probabilities is normality and it should come as no surprise that many VaR measures are based upon some variant of that assumption. If, for instance, we assume that each market risk factor has normally distributed returns, we ensure that the returns on any portfolio that is exposed to multiple market risk factors will also have a normal distribution. Even those VaR approaches that allow for non-normal return distributions for individual risk factors find ways of ending up with normal distributions for final portfolio values.

The RiskMetrics Contribution

As we noted in an earlier section, the term Value at Risk and the usage of the measure can be traced back to the RiskMetrics service offered by J.P. Morgan in 1995. The key contribution of the service was that it made the variances in and covariances across asset classes freely available to anyone who wanted to access them, thus easing the task for anyone who wanted to compute the Value at Risk analytically for a portfolio. Publications by J.P. Morgan in 1996 describe the assumptions underlying their computation of VaR:³

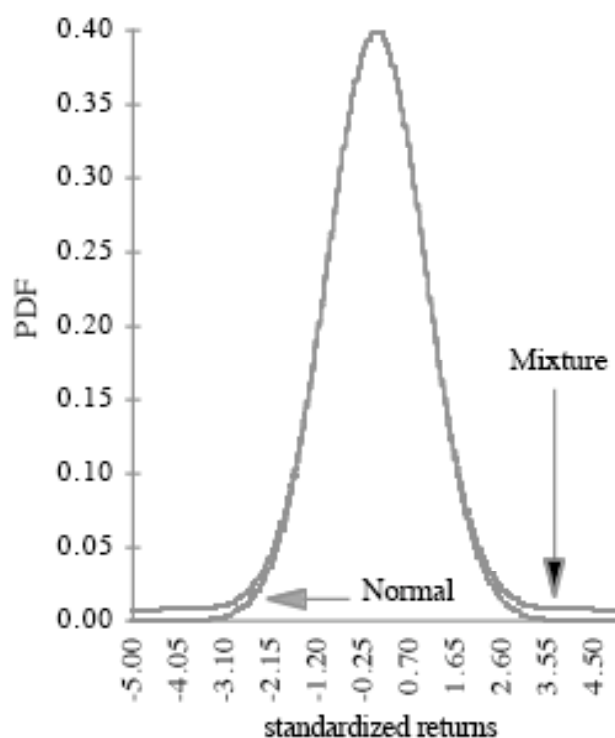
- Returns on individual risk factors are assumed to follow conditional normal distributions. While returns themselves may not be normally distributed and large outliers are far too common (i.e., the distributions have fat tails), the assumption is that the standardized return (computed as the return divided by the forecasted standard deviation) is normally distributed.
- The focus on standardized returns implies that it is not the size of the return per se that we should focus on but its size relative to the standard deviation. In other words, a large return (positive or negative) in a period of high volatility may result in a low standardized return, whereas the same return following a period of low volatility will yield an abnormally high standardized return.

³ RiskMetrics – Technical Document, J.P. Morgan, December 17, 1996; Zangari, P., 1996, An Improved Methodology for Computing VaR, J.P. Morgan RiskMetrics Monitor, Second Quarter 1996.

The focus on normalized standardized returns exposed the VaR computation to the risk of more frequent large outliers than would be expected with a normal distribution. In a subsequent variation, the RiskMetrics approach was extended to cover normal mixture distributions, which allow for the assignment of higher probabilities for outliers. Figure 7.1 contrasts the two distributions:

Figure 7.1

Standard normal and normal mixture probability density functions (PDF)



In effect, these distributions require estimates of the probabilities of outsized returns occurring and the expected size and standard deviations of such returns, in addition to the standard normal distribution parameters. Even proponents of these models concede that estimating the parameters for jump processes, given how infrequently jumps occur, is difficult to do.

Assessment

The strength of the Variance-Covariance approach is that the Value at Risk is simple to compute, once you have made an assumption about the distribution of returns and inputted the means, variances and covariances of returns. In the estimation process, though, lie the three key weaknesses of the approach:

- Wrong distributional assumption: If conditional returns are not normally distributed, the computed VaR will understate the true VaR. In other words, if there are far more outliers in the actual return distribution than would be expected given the normality assumption, the actual Value at Risk will be much higher than the computed Value at Risk.
- Input error: Even if the standardized return distribution assumption holds up, the VaR can still be wrong if the variances and covariances that are used to estimate it are incorrect. To the extent that these numbers are estimated using historical data, there is a standard error associated with each of the estimates. In other words, the variance-covariance matrix that is input to the VaR measure is a collection of estimates, some of which have very large error terms.
- Non-stationary variables: A related problem occurs when the variances and covariances across assets change over time. This nonstationarity in values is not uncommon because the fundamentals driving these numbers do change over time. Thus, the correlation between the U.S. dollar and the Japanese yen may change if oil prices increase by 15%. This, in turn, can lead to a breakdown in the computed VaR.

Not surprisingly, much of the work that has been done to revitalize the approach has been directed at dealing with these critiques.

First, a host of researchers have examined how best to compute VaR with assumptions other than the standardized normal; we mentioned the normal mixture model in the RiskMetrics section.⁴ Hull and White suggest ways of estimating Value at Risk when variables are not normally distributed; they allow users to specify any probability distribution for variables but require that transformations of the distribution

⁴ Duffie, D. and J. Pan, 1997, An Overview of Value at Risk, Working Paper, Stanford University. The authors provide a comprehensive examination of different distributions and the parameters that have to be estimated for each one.

still fall a multivariate normal distribution.⁵ These and other papers like it develop interesting variations but have to overcome two practical problems. Estimating inputs for non-normal models can be very difficult to do, especially when working with historical data, and the probabilities of losses and Value at Risk are simplest to compute with the normal distribution and get progressively more difficult with asymmetric and fat-tailed distributions.

Second, other research has been directed at bettering the estimation techniques to yield more reliable variance and covariance values to use in the VaR calculations. Some suggest refinements on sampling methods and data innovations that allow for better estimates of variances and covariances looking forward. Others posit that statistical innovations can yield better estimates from existing data. For instance, conventional estimates of VaR are based upon the assumption that the standard deviation in returns does not change over time (homoskedasticity), Engle argues that we get much better estimates by using models that explicitly allow the standard deviation to change of time (heteroskedasticity).⁶ In fact, he suggests two variants – Autoregressive Conditional Heteroskedasticity (ARCH) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) – that provide better forecasts of variance and, by extension, better measures of Value at Risk.⁷

One final critique that can be leveled against the variance-covariance estimate of VaR is that it is designed for portfolios where there is a linear relationship between risk and portfolio positions. Consequently, it can break down when the portfolio includes options, since the payoffs on an option are not linear. In an attempt to deal with options and other non-linear instruments in portfolios, researchers have developed *Quadratic Value at Risk* measures.⁸ These quadratic measures, sometimes categorized as delta-gamma models (to

⁵ Hull, J. and A. White, 1998, Value at Risk when daily changes are not normally distributed, *Journal of Derivatives*, v5, 9-19.

⁶ Engle, R., 2001, Garch 101: The Use of ARCH and GARCH models in Applied Econometrics, *Journal of Economic Perspectives*, v15, 157-168.

⁷ He uses the example of a \$1,000,000 portfolio composed of 50% NASDAQ stocks, 30% Dow Jones stocks and 20% long bonds, with statistics computed from March 23, 1990 to March 23, 2000. Using the conventional measure of daily standard deviation of 0.83% computed over a 10-year period, he estimates the value at risk in a day to be \$22,477. Using an ARCH model, the forecast standard deviation is 1.46%, leading to VaR of \$33,977. Allowing for the fat tails in the distribution increases the VaR to \$39,996.

⁸ Britten-Jones, M. and Schaefer, S.M., 1999, Non-linear value-at-risk, *European Finance Review*, v2, 161-

contrast with the more conventional linear models which are called delta-normal), allow researchers to estimate the Value at Risk for complicated portfolios that include options and option-like securities such as convertible bonds. The cost, though, is that the mathematics associated with deriving the VaR becomes much complicated and that some of the intuition will be lost along the way.

Historical Simulation

Historical simulations represent the simplest way of estimating the Value at Risk for many portfolios. In this approach, the VaR for a portfolio is estimated by creating a hypothetical time series of returns on that portfolio, obtained by running the portfolio through actual historical data and computing the changes that would have occurred in each period.

General Approach

To run a historical simulation, we begin with time series data on each market risk factor, just as we would for the variance-covariance approach. However, we do not use the data to estimate variances and covariances looking forward, since the changes in the portfolio over time yield all the information you need to compute the Value at Risk.

Cabedo and Moya provide a simple example of the application of historical simulation to measure the Value at Risk in oil prices.⁹ Using historical data from 1992 to 1998, they obtained the daily prices in Brent Crude Oil and graphed out the prices in Figure 7.2:

187; Rouvinez, C. , 1997, Going Greek with VAR, Risk, v10, 57-65.

⁹ J.D. Cabedo and I. Moya, 2003, Estimating oil price Value at Risk using the historical simulation approach, Energy Economics, v25, 239-253.

Figure 7.2: Price/barrel for Brent Crude Oil – 1992-99



They separated the daily price changes into positive and negative numbers, and analyzed each group. With a 99% confidence interval, the positive VaR was defined as the price change in the 99th percentile of the positive price changes and the negative VaR as the price change at the 99th percentile of the negative price changes.¹⁰ For the period they studied, the daily Value at Risk at the 99th percentile was about 1% in both directions.

The implicit assumptions of the historical simulation approach are visible in this simple example. The first is that the approach is agnostic when it comes to distributional assumptions, and the VaR is determined by the actual price movements. In other words, there are no underlying assumptions of normality driving the conclusion. The second is that each day in the time series carries an equal weight when it comes to measuring the VaR, a potential problem if there is a trend in the variability – lower in the earlier periods and higher in the later periods, for instance. The third is that the approach is based on the assumption of history repeating itself, with the period used providing a full and complete snapshot of the risks that the oil market is exposed to in other periods.

¹⁰ By separating the price changes into positive and negative changes, they allow for asymmetry in the return process where large negative changes are more common than large positive changes, or vice versa.

Assessment

While historical simulations are popular and relatively easy to run, they do come with baggage. In particular, the underlying assumptions of the model generate give rise to its weaknesses.

- a. Past is not prologue: While all three approaches to estimating VaR use historical data, historical simulations are much more reliant on them than the other two approaches for the simple reason that the Value at Risk is computed entirely from historical price changes. There is little room to overlay distributional assumptions (as we do with the Variance-covariance approach) or to bring in subjective information (as we can with Monte Carlo simulations). The example provided in the last section with oil prices provides a classic example. A portfolio manager or corporation that determined its oil price VaR, based upon 1992 to 1998 data, would have been exposed to much larger losses than expected over the 1999 to 2004 period as a long period of oil price stability came to an end and price volatility increased.
- b. Trends in the data: A related argument can be made about the way in which we compute Value at Risk, using historical data, where all data points are weighted equally. In other words, the price changes from trading days in 1992 affect the VaR in exactly the same proportion as price changes from trading days in 1998. To the extent that there is a trend of increasing volatility even within the historical time period, we will understate the Value at Risk.
- c. New assets or market risks: While this could be a critique of any of the three approaches for estimating VaR, the historical simulation approach has the most difficulty dealing with new risks and assets for an obvious reason: there is no historic data available to compute the Value at Risk. Assessing the Value at Risk to a firm from developments in online commerce in the late 1990s would have been difficult to do, since the online business was in its nascent stage.

The trade off that we mentioned earlier is therefore at the heart of the historic simulation debate. The approach saves us the trouble and related problems of having to make specific assumptions about distributions of returns but it implicitly assumes that the distribution of past returns is a good and complete representation of expected future

returns. In a market where risks are volatile and structural shifts occur at regular intervals, this assumption is difficult to sustain.

Modifications

As with the other approaches to computing VaR, there have been modifications suggested to the approach, largely directed at taking into account some of the criticisms mentioned in the last section.

a. Weighting the recent past more: A reasonable argument can be made that returns in the recent past are better predictors of the immediate future than are returns from the distant past. Boudoukh, Richardson and Whitelaw present a variant on historical simulations, where recent data is weighted more, using a decay factor as their time weighting mechanism.¹¹ In simple terms, each return, rather than being weighted equally, is assigned a probability weight based on its recency. In other words, if the decay factor is .90, the most recent observation has the probability weight p , the observation prior to it will be weighted $0.9p$, the one before that is weighted $0.81p$ and so on. In fact, the conventional historical simulation approach is a special case of this approach, where the decay factor is set to 1. Boudoukh et al. illustrate the use of this technique by computing the VaR for a stock portfolio, using 250 days of returns, immediately before and after the market crash on October 19, 1987.¹² With historical simulation, the Value at Risk for this portfolio is for all practical purposes unchanged the day after the crash because it weights each day (including October 19) equally. With decay factors, the Value at Risk very quickly adjusts to reflect the size of the crash.¹³

b. Combining historical simulation with time series models: Earlier in this section, we referred to a Value at Risk computation by Cabado and Moya for oil prices using a historical simulation. In the same paper, they suggested that better estimates of VaR could be obtained by fitting a time series model through the historical data and using the parameters of that model to forecast the Value at Risk. In particular, they fit an


¹¹ Boudoukh, J., M. Richardson and R. Whitelaw, 1998. "The Best of Both Worlds," *Risk*, v11, 64-67.

¹² The Dow dropped 508 points on October 19, 1987, approximately 22%.

¹³ With a decay factor of 0.99, the most recent day will be weighted about 1% (instead of 1/250). With a decay factor of 0.97, the most recent day will be weighted about 3%.

autoregressive moving average (ARMA) model to the oil price data from 1992 to 1998 and use this model to forecast returns with a 99% confidence interval for the holdout period of 1999. The actual oil price returns in 1999 fall within the predicted bounds 98.8% of the time, in contrast to the 97.7% of the time that they do with the unadjusted historical simulation. One big reason for the improvement is that the measured VaR is much more sensitive to changes in the variance of oil prices with time series models, than with the historical simulation, as can be seen in figure 7.3:

Figure 7.3: Value at Risk Estimates (99%) from Time Series Models



Note that the range widens in the later part of the year in response to the increasing volatility in oil prices, as the time series model is updated to incorporate more recent data.

3. Volatility Updating: Hull and White suggest a different way of updating historical data for shifts in volatility. For assets where the recent volatility is higher than historical volatility, they recommend that the historical data be adjusted to reflect the change. Assume, for illustrative purposes, that the updated standard deviation in prices is 0.8% and that it was only 0.6% when estimated with data from 20 days ago. Rather than use the price change from 20 days ago, they recommend scaling that number to reflect the change in volatility; a 1% return on that day would be converted into a 1.33% return

($\frac{0.8}{0.6} * 1\%$). Their approach requires day-specific estimates of variance that change over the historical time period, which they obtain by using GARCH models.¹⁴

Note that all of these variations are designed to capture shifts that have occurred in the recent past but are underweighted by the conventional approach. None of them are designed to bring in the risks that are out of the sampled historical period (but are still relevant risks) or to capture structural shifts in the market and the economy. In a paper comparing the different historical simulation approaches, Pritsker notes the limitations of the variants.¹⁵

Monte Carlo Simulation

In the last chapter, we examined the use of Monte Carlo simulations as a risk assessment tool. These simulations also happen to be useful in assessing Value at Risk, with the focus on the probabilities of losses exceeding a specified value rather than on the entire distribution.

General Description

The first two steps in a Monte Carlo simulation mirror the first two steps in the Variance-covariance method where we identify the markets risks that affect the asset or assets in a portfolio and convert individual assets into positions in standardized instruments. It is in the third step that the differences emerge. Rather than compute the variances and covariances across the market risk factors, we take the simulation route, where we specify probability distributions for each of the market risk factors and specify how these market risk factors move together. Thus, in the example of the six-month Dollar/Euro forward contract that we used earlier, the probability distributions for the 6-month zero coupon \$ bond, the 6-month zero coupon euro bond and the dollar/euro spot rate will have to be specified, as will the correlation across these instruments.

While the estimation of parameters is easier if you assume normal distributions for all variables, the power of Monte Carlo simulations comes from the freedom you have

¹⁴ Hull, J. and A. White, 1998, Incorporating Volatility Updating into the Historical Simulation Method for Value at Risk, *Journal of Risk*, v1, 5-19.

¹⁵ Pritsker, M., 2001, The Hidden Dangers of Historical Simulation, Working paper, SSRN.

to pick alternate distributions for the variables. In addition, you can bring in subjective judgments to modify these distributions.

Once the distributions are specified, the simulation process starts. In each run, the market risk variables take on different outcomes and the value of the portfolio reflects the outcomes. After a repeated series of runs, numbering usually in the thousands, you will have a distribution of portfolio values that can be used to assess Value at Risk. For instance, assume that you run a series of 10,000 simulations and derive corresponding values for the portfolio. These values can be ranked from highest to lowest, and the 95% percentile Value at Risk will correspond to the 500th lowest value and the 99th percentile to the 100th lowest value.

Assessment

Much of what was said about the strengths and weaknesses of the simulation approach in the last chapter apply to its use in computing Value at Risk. Quickly reviewing the criticism, a simulation is only as good as the probability distribution for the inputs that are fed into it. While Monte Carlo simulations are often touted as more sophisticated than historical simulations, many users directly draw on historical data to make their distributional assumptions.

In addition, as the number of market risk factors increases and their co-movements become more complex, Monte Carlo simulations become more difficult to run for two reasons. First, you now have to estimate the probability distributions for hundreds of market risk variables rather than just the handful that we talked about in the context of analyzing a single project or asset. Second, the number of simulations that you need to run to obtain reasonable estimate of Value at Risk will have to increase substantially (to the tens of thousands from the thousands).

The strengths of Monte Carlo simulations can be seen when compared to the other two approaches for computing Value at Risk. Unlike the variance-covariance approach, we do not have to make unrealistic assumptions about normality in returns. In contrast to the historical simulation approach, we begin with historical data but are free to bring in both subjective judgments and other information to improve forecasted probability

distributions. Finally, Monte Carlo simulations can be used to assess the Value at Risk for any type of portfolio and are flexible enough to cover options and option-like securities.

Modifications

As with the other approaches, the modifications to the Monte Carlo simulation are directed at its biggest weakness, which is its computational bulk. To provide a simple illustration, a yield curve model with 15 key rates and four possible values for each will require 1,073,741,824 simulations (4^{15}) to be complete. The modified versions narrow the focus, using different techniques, and reduce the required number of simulations.

- a. Scenario Simulation: One way to reduce the computation burden of running Monte Carlo simulations is to do the analysis over a number of discrete scenarios. Frye suggests an approach that can be used to develop these scenarios by applying a small set of pre-specified shocks to the system.¹⁶ Jamshidan and Zhu (1997) suggest what they called scenario simulations where they use principal component analysis as a first step to narrow the number of factors. Rather than allow each risk variable to take on all of the potential values, they look at likely combinations of these variables to arrive at scenarios. The values are computed across these scenarios to arrive at the simulation results.¹⁷
- b. Monte Carlo Simulations with Variance-Covariance method modification: The strength of the Variance-covariance method is its speed. If you are willing to make the required distributional assumption about normality in returns and have the variance-covariance matrix in hand, you can compute the Value at Risk for any portfolio in minutes. The strength of the Monte Carlo simulation approach is the flexibility it offers users to make different distributional assumptions and deal with various types of risk, but it can be painfully slow to run. Glasserman, Heidelberger and Shahabuddin use approximations from the variance-covariance approach to guide

¹⁶ Frye, J., 1997, "Principals of Risk: Finding Value-at-Risk Through Factor-Based Interest Rate Scenarios." NationsBanc-CRT.

¹⁷ Jamshidian, Farshid and Yu Zhu, 1997, "Scenario Simulation: Theory and Methodology." Finance and Stochastics, v1, 43-67. In principal component analysis, you look for common factors affecting returns in historical data.

the sampling process in Monte Carlo simulations and report a substantial savings in time and resources, without any appreciable loss of precision.¹⁸

The trade off in each of these modifications is simple. You give some of the power and precision of the Monte Carlo approach but gain in terms of estimation requirements and computational time.

Comparing Approaches

Each of the three approaches to estimating Value at Risk has advantages and comes with baggage. The variance-covariance approach, with its delta normal and delta gamma variations, requires us to make strong assumptions about the return distributions of standardized assets, but is simple to compute, once those assumptions have been made. The historical simulation approach requires no assumptions about the nature of return distributions but implicitly assumes that the data used in the simulation is a representative sample of the risks looking forward. The Monte Carlo simulation approach allows for the most flexibility in terms of choosing distributions for returns and bringing in subjective judgments and external data, but is the most demanding from a computational standpoint.

Since the end product of all three approaches is the Value at Risk, it is worth asking two questions.

1. How different are the estimates of Value at Risk that emerge from the three approaches?
2. If they are different, which approach yields the most reliable estimate of VaR?

To answer the first question, we have to recognize that the answers we obtain with all three approaches are a function of the inputs. For instance, the historical simulation and variance-covariance methods will yield the same Value at Risk if the historical returns data is normally distributed and is used to estimate the variance-covariance matrix. Similarly, the variance-covariance approach and Monte Carlo simulations will yield roughly the same values if all of the inputs in the latter are assumed to be normally distributed with consistent means and variances. As the assumptions diverge, so will the

¹⁸ Glasserman, P., P. Heidelberger and P. Shahabuddin, 2000, Efficient Monte Carlo Methods for Value at Risk, Working Paper, Columbia University.

answers. Finally, the historical and Monte Carlo simulation approaches will converge if the distributions we use in the latter are entirely based upon historical data.

As for the second, the answer seems to depend both upon what risks are being assessed and how the competing approaches are used. As we noted at the end of each approach, there are variants that have developed within each approach, aimed at improving performance. Many of the comparisons across approaches are skewed by the fact that the researchers doing the comparison are testing variants of an approach that they have developed against alternatives. Not surprisingly, they find that their approaches work better than the alternatives. Looking at the unbiased (relatively) studies of the alternative approaches, the evidence is mixed. Hendricks compared the VaR estimates obtained using the variance-covariance and historical simulation approaches on 1000 randomly selected foreign exchange portfolios.¹⁹ He used nine measurement criteria, including the mean squared error (of the actual loss against the forecasted loss) and the percentage of the outcomes covered and concluded that the different approaches yield risk measures that are roughly comparable and that they all cover the risk that they are intended to cover, at least up to the 95 percent confidence interval. He did conclude that all of the measures have trouble capturing extreme outcomes and shifts in underlying risk. Lambadrais, Papadopoulou, Skiadopoulus and Zoulis computed the Value at Risk in the Greek stock and bond market with historical with Monte Carlo simulations, and found that while historical simulation overstated the VaR for linear stock portfolios, the results were less clear cut with non-linear bond portfolios.²⁰

In short, the question of which VaR approach is best is best answered by looking at the task at hand? If you are assessing the Value at Risk for portfolios, that do not include options, over very short time periods (a day or a week), the variance-covariance approach does a reasonably good job, notwithstanding its heroic assumptions of normality. If the Value at Risk is being computed for a risk source that is stable and where there is substantial historical data (commodity prices, for instance), historical

¹⁹ Hendricks, D., 1996, Evaluation of value-at-risk models using historical data, Federal Reserve Bank of New York, Economic Policy Review, v2,, 39–70.

²⁰Lambadiaris, G. , L. Papadopoulou, G. Skiadopoulos and Y. Zoulis, 2000, VAR: Hisory or Simulation?, www.risk.net.

simulations provide good estimates. In the most general case of computing VaR for non-linear portfolios (which include options) over longer time periods, where the historical data is volatile and non-stationary and the normality assumption is questionable, Monte Carlo simulations do best.

Limitations of VaR

While Value at Risk has acquired a strong following in the risk management community, there is reason to be skeptical of both its accuracy as a risk management tool and its use in decision making. There are many dimensions on which researcher have taken issue with VaR and we will categorize the criticism into those dimensions.

VaR can be wrong

There is no precise measure of Value at Risk, and each measure comes with its own limitations. The end-result is that the Value at Risk that we compute for an asset, portfolio or a firm can be wrong, and sometimes, the errors can be large enough to make VaR a misleading measure of risk exposure. The reasons for the errors can vary across firms and for different measures and include the following.

a. Return distributions: Every VaR measure makes assumptions about return distributions, which, if violated, result in incorrect estimates of the Value at Risk. With delta-normal estimates of VaR, we are assuming that the multivariate return distribution is the normal distribution, since the Value at Risk is based entirely on the standard deviation in returns. With Monte Carlo simulations, we get more freedom to specify different types of return distributions, but we can still be wrong when we make those judgments. Finally, with historical simulations, we are assuming that the historical return distribution (based upon past data) is representative of the distribution of returns looking forward.

There is substantial evidence that returns are not normally distributed and that not only are outliers more common in reality but that they are much larger than expected, given the normal distribution. In chapter 4, we noted Mandelbrot's critique of the mean-variance framework and his argument that returns followed power law distributions. His critique extended to the use of Value at Risk as the risk measure of choice at financial

service firms. Firms that use VaR to measure their risk exposure, he argued, would be under prepared for large and potentially catastrophic events that are extremely unlikely in a normal distribution but seem to occur at regular intervals in the real world.

b. History may not a good predictor: All measures of Value at Risk use historical data to some degree or the other. In the variance-covariance method, historical data is used to compute the variance-covariance matrix that is the basis for the computation of VaR. In historical simulations, the VaR is entirely based upon the historical data with the likelihood of value losses computed from the time series of returns. In Monte Carlo simulations, the distributions don't have to be based upon historical data but it is difficult to see how else they can be derived. In short, any Value at Risk measure will be a function of the time period over which the historical data is collected. If that time period was a relatively stable one, the computed Value at Risk will be a low number and will understate the risk looking forward. Conversely, if the time period examined was volatile, the Value at Risk will be set too high. Earlier in this chapter, we provided the example of VaR for oil price movements and concluded that VaR measures based upon the 1992-98 period, where oil prices were stable, would have been too low for the 1999-2004 period, when volatility returned to the market.

c. Non-stationary Correlations: Measures of Value at Risk are conditioned on explicit estimates of correlation across risk sources (in the variance-covariance and Monte Carlo simulations) or implicit assumptions about correlation (in historical simulations). These correlation estimates are usually based upon historical data and are extremely volatile. One measure of how much they move can be obtained by tracking the correlations between widely following asset classes over time. Figure 7.4 graphs the correlation between the S&P 500 and the ten-year treasury bond returns, using daily returns for a year, every year from 1990 to 2005:

Figure 7.4: Time Series of Correlation between Stock and Bond Returns

Skintzi, Skiadoupoulous and Refenes show that the error in Var increases as the correlation error increases and that the effect is magnified in Monte Carlo simulations.²¹

²¹ Skintzi, V.D., G. Skiadoupoulous and A.P.N. Refenes, 2005, The Effect of Misestimating Correlation on Value at Risk, Journal of Alternative Investments, Spring 2005.

One indicator that Value at Risk is subject to judgment comes from the range of values that analysts often assign to the measure, when looking at the same risk for the same entity. Different assumptions about return distributions and different historical time periods can yield very different values for VaR.²² In fact, different measures of Value at Risk can be derived for a portfolio even when we start with the same underlying data and methodology.²³ A study of Value at Risk measures used at large bank holding companies to measure risk in their trading portfolios concluded that they were much too conservatively set and were slow to react to changing circumstances; in fact, simple time series models outperformed sophisticated VaR models in predictions. In fact, the study concluded that the computed Value at Risk was more a precautionary number for capital at risk than a measure of portfolio risk.²⁴ In defense of Value at Risk, it should be pointed out that there the reported Values at Risk at banks are correlated with the volatility in trading revenues at these banks and can be used as a proxy for risk (at least from the trading component).²⁵

Narrow Focus

While many analysts like Value at Risk because of its simplicity and intuitive appeal, relative to other risk measures, its simplicity emanates from its narrow definition of risk. Firms that depend upon VaR as the only measure of risk can not only be lulled into a false sense of complacency about the risks they face but also make decisions that are not in their best interests.

a. Type of risk: Value at Risk measures the likelihood of losses to an asset or portfolio due to market risk. Implicit in this definition is the narrow definition of risk, at least in conventional VaR models. First, risk is almost always considered to be a negative in VaR. While there is no technical reason why one cannot estimate potential profits

²² Beder, T.S., 1995, VAR: Seductive but Dangerous, *Financial Analysts Journal*, September-October 1995.

²³ Marshall, Chris, and Michael Siegel, "Value at Risk: Implementing a Risk Measurement Standard." *Journal of Derivatives* 4, No. 3 (1997), pp. 91-111. Different measures of Value at Risk are estimated using different software packages on the J.P. Morgan RiskMetrics data and methodology.

²⁴ Berkowitz, J. and J. O'Brien, 2002, How accurate are Value at Risk Models at Commercial Banks, *Journal of Finance*, v57, 1093-1111.

that one can earn with 99% probability, VaR is measured in terms of potential losses and not gains. Second, most VaR measures are built around market risk effects. Again, while there is no reason why we cannot look at the Value at Risk, relative to all risks, practicality forces us to focus on just market risks and their effects on value. In other words, the true Value at Risk can be much greater than the computed Value at Risk if one considers political risk, liquidity risk and regulatory risks that are not built into the VaR.

- b. Short term: Value at Risk can be computed over a quarter or a year, but it is usually computed over a day, a week or a few weeks. In most real world applications, therefore, the Value at Risk is computed over short time periods, rather than longer ones. There are three reasons for this short term focus. The first is that the financial service firms that use Value at Risk often are focused on hedging these risks on a day-to-day basis and are thus less concerned about long term risk exposures. The second is that the regulatory authorities, at least for financial service firms, demand to know the short term Value at Risk exposures at frequent intervals. The third is that the inputs into the VaR measure computation, whether it is measured using historical simulations or the variance-covariance approach, are easiest to estimate for short periods. In fact, as we noted in the last section, the quality of the VaR estimates quickly deteriorate as you go from daily to weekly to monthly to annual measures.
- c. Absolute Value: The output from a Value at Risk computation is not a standard deviation or an overall risk measure but is stated in terms of a probability that the losses will exceed a specified value. As an example, a VaR of \$ 100 million with 95% confidence implies that there is only a 5% chance of losing more than \$ 100 million. The focus on a fixed value makes it an attractive measure of risk to financial service firms that worry about their capital adequacy. By the same token, it is what makes VaR an inappropriate measure of risk for firms that are focused on comparing investments with very different scales and returns; for these firms, more conventional scaled measures of risk (such as standard deviation or betas) that focus on the entire risk distribution will work better.

²⁵ Jorion, P., 2002, How informative are Value-at-Risk Disclosures?, *The Accounting Review*, v77, 911-932.

In short, Value at Risk measures look at only a small slice of the risk that an asset is exposed to and a great deal of valuable information in the distribution is ignored. Even if the VaR assessment that the probability of losing more than \$ 100 million is less than 5% is correct, would it not make sense to know what the most you can lose in that catastrophic range (with less than 5% probability) would be? It should, after all, make a difference whether your worst possible loss was \$ 1 billion or \$ 150 million. Looking back at chapter 6 on probabilistic risk assessment approaches, Value at Risk is closer to the worst case assessment in scenario analysis than it is to the fuller risk assessment approaches.

Sub-optimal Decisions

Even if Value at Risk is correctly measured, it is not clear that using it as the measure of risk leads to more reasoned and sensible decisions on the part of managers and investors. In fact, there are two strands of criticism against the use of Value at Risk in decision making. The first is that making investment decisions based upon Value at Risk can lead to over exposure to risk, even when the decision makers are rational and Value at Risk is estimated precisely. The other is that managers who understand how VaR is computed, can game the measure to report superior performance, while exposing the firm to substantial risks.

a. Overexposure to Risk: Assume that managers are asked to make investment decisions, while having their risk exposures measured using Value at Risk. Basak and Shapiro note that such managers will often invest in more risky portfolios than managers who do not use Value at Risk as a risk assessment tool. They explain this counter intuitive result by noting that managers evaluated based upon VaR will be much more focused on avoiding the intermediate risks (under the probability threshold), but that their portfolios are likely to lose far more under the most adverse circumstances. Put another way, by not bringing in the magnitude of the losses once

you exceed the VaR cutoff probability (90% or 95%), you are opening ourselves to the possibility of very large losses in the worse case scenarios.²⁶

b. Agency problems: Like any risk measure, Value at Risk can be gamed by managers who have decided to make an investment and want to meet the VaR risk constraint. Ju and Pearson note that since Value at Risk is generally measured using past data, traders and managers who are evaluated using the measure will have a reasonable understanding of its errors and can take advantage of them. Consider the example of the VaR from oil price volatility that we estimated using historical simulation earlier in the chapter; the VaR was understated because it did not capture the trending up in volatility in oil prices towards the end of the time period. A canny manager who knows that this can take on far more oil price risk than is prudent while reporting a Value at Risk that looks like it is under the limit.²⁷ It is true that all risk measures are open to this critique but by focusing on an absolute value and a single probability, VaR is more open to this game playing than other measures.

Extensions of VaR

The popularity of Value at Risk has given rise to numerous variants of it, some designed to mitigate problems associated with the original measure and some directed towards extending the use of the measure from financial service firms to the rest of the market.

There are modifications of VaR that adapt the original measure to new uses but remain true to its focus on overall value. Hallerback and Menkveld modify the conventional VaR measure to accommodate multiple market factors and computed what they call a Component Value at Risk, breaking down a firm's risk exposure to different market risks. They argue that managers at multinational firms can use this risk measure to not only determine where their risk is coming from but to manage it better in the interests of maximizing shareholder wealth.²⁸ In an attempt to bring in the possible losses in the

²⁶ Basak, S. and A. Shapiro, 2001, Value-at-Risk Based Management: Optimal Policies and Asset Prices, *Review of Financial Studies*, v14, 371-405.

²⁷ Ju, X. and N.D. Pearson, 1998, Using Value-at-Risk to Control Risk Taking: How wrong can you be?, Working Paper, University of Illinois at Urbana-Champaign.

²⁸ Hallerback, W.G. and A.J. Menkveld, 2002, Analyzing Perceived Downside Risk: the Component

tail of the distribution (beyond the VaR probability), Larsen, Mausser and Uryasev estimate what they call a Conditional Value at Risk, which they define as a weighted average of the VaR and losses exceeding the VaR.²⁹ This conditional measure can be considered an upper bound on the Value at Risk and may reduce the problems associated with excessive risk taking by managers. Finally, there are some who note that Value at Risk is just one aspect of an area of mathematics called Extreme Value Theory, and that there may be better and more comprehensive ways of measuring exposure to catastrophic risks.³⁰

The other direction that researchers have taken is to extend the measure to cover metrics other than value. The most widely used of these is Cashflow at Risk (CFaR). While Value at Risk focuses on changes in the overall value of an asset or portfolio as market risks vary, Cash Flow at Risk is more focused on the operating cash flow during a period and market induced variations in it. Consequently, with Cash flow at Risk, we assess the likelihood that operating cash flows will drop below a pre-specified level; an annual CFaR of \$ 100 million with 90% confidence can be read to mean that there is only a 10% probability that cash flows will drop by more than \$ 100 million, during the next year. Herein lies the second practical difference between Value at Risk and Cashflow at Risk. While Value at Risk is usually computed for very short time intervals – days or weeks – Cashflow at Risk is computed over much longer periods – quarters or years.

Why focus on cash flows rather than value? First, for a firm that has to make contractual payments (interest payments, debt repayments and lease expenses) during a particular period, it is cash flow that matters; after all, the value can remain relatively stable while cash flows plummet, putting the firm at risk of default. Second, unlike financial service firms where the value measured is the value of marketable securities which can be converted into cash at short notice, value at a non-financial service firm takes the form of real investments in plant, equipment and other fixed assets which are far more difficult to monetize. Finally, assessing the market risks embedded in value, while

Value-at-Risk Framework, Working Paper.

²⁹ Larsen, N., H. Mausser and S. Uryasev, 2001, Algorithms for Optimization of Value-at-Risk, Research Report, University of Florida.

³⁰ Embrechts, P., 2001, Extreme Value Theory: Potential and Limitations as an Integrated Risk Management Tool, Working Paper (listed on GloriaMundi.org).

relatively straight forward for a portfolio of financial assets, can be much more difficult to do for a manufacturing or technology firm.

How do we measure CFaR? While we can use any of the three approaches described for measuring VaR – variance-covariance matrices, historical simulations and Monte Carlo simulations – the process becomes more complicated if we consider all risks and not just market risks. Stein, Usher, LaGattuta and Youngen develop a template for estimating Cash Flow at Risk, using data on comparable firms, where comparable is defined in terms of market capitalization, riskiness, profitability and stock-price performance, and use it to measure the risk embedded in the earnings before interest, taxes and depreciation (EBITDA) at Coca Cola, Dell and Cignus (a small pharmaceutical firm).³¹ Using regressions of EBITDA as a percent of assets across the comparable firms over time, for a five-percent worst case, they estimate that EBITDA would drop by \$5.23 per \$ 100 of assets at Coca Cola, \$28.50 for Dell and \$47.31 for Cygnus. They concede that while the results look reasonable, the approach is sensitive to both the definition of comparable firms and is likely to yield estimates with error.

There are less common adaptations that extend the measure to cover earnings (Earnings at Risk) and to stock prices (SPaR). These variations are designed by what the researchers view as the constraining variable in decision making. For firms that are focused on earnings per share and ensuring that it does not drop below some pre-specified floor, it makes sense to focus on Earnings at Risk. For other firms, where a drop in the stock price below a given level will give risk to constraints or delisting, it is SPaR that is the relevant risk control measure.

VaR as a Risk Assessment Tool

In the last three chapters, we have considered a range of risk assessment tools. In chapter 5, we introduced risk and return models that attempted to either increase the discount rate or reduce the cash flows (certainty equivalents) used to value risky assets, leading to risk adjusted values. In chapter 6, we considered probabilistic approaches to risk assessment including scenario analysis, simulations and decision trees, where we

considered most or all possible outcomes from a risky investment and used that information in valuation and investment decisions. In this chapter, we introduced Value at Risk, touted by its adherents as a more intuitive, if not better, way of assessing risk.

From our perspective, and it may very well be biased, Value at Risk seems to be a throwback and not an advance in thinking about risk. Of all the risk assessment tools that we have examined so far, it is the most focused on downside risk, and even within that downside risk, at a very small slice of it. It seems foolhardy to believe that optimal investment decisions can flow out of such a cramped view of risk. Value at Risk seems to take a subset of the information that comes out of scenario analysis (the close to worst case scenario) or simulations (the fifth percentile or tenth percentile of the distribution) and throw the rest of it out. There are some who would argue that presenting decision makers with an entire probability distribution rather than just the loss that they will make with 5% probability will lead to confusion, but if that is the case, there is little hope that such individuals can be trusted to make good decisions in the first place with any risk assessment measure.

How then can we account for the popularity of Value at Risk? A cynic would attribute it to an accident of history where a variance-covariance matrix, with a dubious history of forecasting accuracy, was made available to panicked bankers, reeling from a series of financial disasters wrought by rogue traders. Consultants and software firms then filled in the gaps and sold the measure as the magic bullet to stop runaway risk taking. The usage of Value at Risk has also been fed into by three factors specific to financial service firms. The first is that these firms have limited capital, relative to the huge nominal values of the leveraged portfolios that they hold; small changes in the latter can put the firm at risk. The second is that the assets held by financial service firms are primarily marketable securities, making it easier to break risks down into market risks and compute Value at Risk. Finally, the regulatory authorities have augmented the use of the measure by demanding regular reports on Value at Risk exposure. Thus, while Value at Risk may be a flawed and narrow measure of risk, it is a natural measure of short term risk for financial service firms and there is evidence that it does its job adequately.

³¹ Stein, J.C., S.E. Usher, D. LaGattuta and J. Youngen, 2000, A Comparables Approach to Measuring Cashflow-at-Risk for Non-Financial Firms, Working Paper, National Economic Research Associates.

For non-financial service firms, there is a place for Value at Risk and its variants in the risk toolbox, but more as a secondary measure of risk rather than a primary measure. Consider how payback (the number of years that it takes to make your money back in an investment) has been used in conventional capital budgeting. When picking between two projects with roughly equivalent net present value (or risk adjusted value), a cash strapped firm will pick the project with the speedier payback. By the same token, when picking between two investments that look equivalent on a risk adjusted basis, a firm should pick the investment with less Cashflow or Value at Risk. This is especially true if the firm has large amounts of debt outstanding and a drop in the cash flows or value may put the firm at risk of default.

Conclusion

Value at Risk has developed as a risk assessment tool at banks and other financial service firms in the last decade. Its usage in these firms has been driven by the failure of the risk tracking systems used until the early 1990s to detect dangerous risk taking on the part of traders and it offered a key benefit: a measure of capital at risk under extreme conditions in trading portfolios that could be updated on a regular basis.

While the notion of Value at Risk is simple- the maximum amount that you can lose on an investment over a particular period with a specified probability – there are three ways in which Value at Risk can be measured. In the first, we assume that the returns generated by exposure to multiple market risks are normally distributed. We use a variance-covariance matrix of all standardized instruments representing various market risks to estimate the standard deviation in portfolio returns and compute the Value at Risk from this standard deviation. In the second approach, we run a portfolio through historical data – a historical simulation – and estimate the probability that the losses exceed specified values. In the third approach, we assume return distributions for each of the individual market risks and run Monte Carlo simulations to arrive at the Value at Risk. Each measure comes with its own pluses and minuses: the Variance-covariance approach is simple to implement but the normality assumption can be tough to sustain, historical simulations assume that the past time periods used are representative of the

future and Monte Carlo simulations are time and computation intensive. All three yield Value at Risk measures that are estimates and subject to judgment.

We understand why Value at Risk is a popular risk assessment tool in financial service firms, where assets are primarily marketable securities, there is limited capital at play and a regulatory overlay that emphasizes short term exposure to extreme risks. We are hard pressed to see why Value at Risk is of particular use to non-financial service firms, unless they are highly levered and risk default if cash flows or value fall below a pre-specified level. Even in those cases, it would seem to us to be more prudent to use all of the information in the probability distribution rather than a small slice of it.

Appendix 1: Example of VaR Calculations: Variance – Covariance Approach

In this appendix, we will compute the VaR of a six-month forward contract, going through four steps – the mapping of the standardized market risks and instruments underlying this security, a determination of the positions that you would need to take in the standardized instruments, the estimation of the variances and covariances of these instruments and the computation of the VaR in the forward contract.

Step 1: The first step requires us to take each of the assets in a portfolio and map that asset on to simpler, standardized instruments. Consider the example of a six-month dollar/euro forward contract. The market factors affecting this instrument are the six-month riskfree rates in each currency and the spot exchange rate; the financial instruments that proxy for these risk factors are the six-month zero coupon dollar bond, the six-month zero coupon Euro bond and the spot \$/Euro.

Step 2: Each financial asset is stated as a set of positions in the standardized instruments. To make the computation for the forward contract, assume that the forward contract requires you to deliver \$12.7 million dollars in 180 days and receive 10 million euros in exchange. Assume, in addition, that the current spot rate is \$1.26/Euro and that the annualized interest rates are 4% on a six-month zero coupon dollar bond and 3% on a six-month zero coupon euro bond. The positions in the three standardized instruments can be computed as follows:

Value of short position in zero-coupon dollar bond

$$= \frac{\$12.7}{(1.04)^{180/360}} = -\$12.4534 \text{ million}$$

Value of long position in zero-coupon euro bond (in dollar terms) holding spot rate fixed

$$= \text{Spot } \$/\text{Eu} \frac{\text{Euro Forward}}{(1+r_{\text{Euro}})^t} = 1.26 * \frac{10 \text{ million}}{(1.03)^{180/360}} = \$12.4145 \text{ million}$$

Value of spot euro position (in dollar terms) holding euro rate fixed

$$= \text{Spot } \$/\text{Eu} \frac{\text{Euro Forward}}{(1+r_{\text{Euro}})^t} = 1.26 * \frac{10 \text{ million}}{(1.03)^{180/360}} = \$12.4145 \text{ million}$$

Note that the last two positions are equal because the forward asset exposes you to risk in the euro in two places – both the riskless euro rate and the spot exchange rate can change over time.

Step 3: Once the standardized instruments that affect the asset or assets in a portfolio been identified, we have to estimate the variances in each of these instruments and the covariances across the instruments. Considering again the six-month \$/Euro forward contract and the three standardized instruments we mapped that investment onto, assume that the variance/covariance matrix (in daily returns) across those instruments is as follows:³²

	<i>Six-month \$ bond</i>	<i>Six-month Eu bond</i>	<i>Spot \$/Euro</i>
<i>Six-month \$ bond</i>	0.0000314		
<i>Six-month Eu bond</i>	0.0000043	0.0000260	
<i>Spot \$/Euro</i>	0.0000012	0.0000013	0.0000032

In practice, these variance and covariance estimates are obtained by looking at historical data.

Step 4: The Value at Risk for the portfolio can now be computed using the weights on the standardized instruments computed in step 2 and the variances and covariances in these instruments computed in step 3. For instance, the daily variance of the 6-month \$/Euro forward contract can be computed as follows: (X_j is the position in standardized asset j and σ_{ij} is the covariance between assets i and j)

$$\begin{aligned} \text{Variance of forward contract} &= X_1^2\sigma_1^2 + X_2^2\sigma_2^2 + X_3^2\sigma_3^2 + 2X_1X_2\sigma_{12} + 2X_2X_3\sigma_{23} + 2X_1X_3\sigma_{13} \\ &= (-12.4534)^2(0.0000314) + (12.4145)^2(0.0000260) + (12.4145)^2(0.0000032) + 2(-12.4534)(12.4145) \\ &\quad (0.0000043) + 2(12.4145)(12.4145)(0.0000013) + 2(-12.4534)(12.4145)(0.0000012) \\ &= \$ 0.0111021 \text{ million} \end{aligned}$$

$$\text{Daily Standard deviation of forward contract} = 0.0111021^{1/2} = \$105,367$$

If we assume a normal distribution, we can now specify the potential value at risk at a 90% confidence interval on this forward contract to be \$ 173,855 for a day.

$$\text{VaR} = \$105,367 * 1.65 = \$173,855$$

³² The covariance of an asset with itself is the variance. Thus, the values on the diagonal represent the variances of these assets; the daily return variance in the six-month \$ bond is 0.0000314. The off-diagonal values are the covariances; the covariance between the spot \$/Euro rate and the six-month \$ bond is 0.0000012.