CHAPTER ONE

INTRODUCTION TO VALUE AT RISK (VaR)

CHAPTER OUTLINE

Risk measurement has preoccupied financial market participants since the dawn of financial history. However, many past attempts have proven to be impractically complex. For example, upon its introduction, Harry Markowitz's Nobel prize-winning theory of portfolio risk measurement was not adopted in practice because of its onerous data requirements.[1] Indeed, it was Bill Sharpe who, along with others,[2] made portfolio theory the standard of financial risk measurement in real world applications through the adoption of the simplifying assumption that all risk could be decomposed into two parts: systematic, market risk and the residual, company-specific or idiosyncratic risk. The resulting Capital Asset Pricing Model theorized that since only undiversifiable market risk is relevant for securities pricing, only the market risk measurement β is necessary, thereby considerably reducing the required data inputs. This model yielded a readily measurable estimate of risk that could be practically applied in a real time market environment. The only problem was that β proved to have only a tenuous connection to actual security returns, thereby casting doubts on β's designation as the true risk measure.[3]

With β questioned, and with asset prcing in general being at a bit of a disarray with respect to whether the notion of "priced risk" is really relevant, market practitioners searched for a replacement risk measure that was both accurate and relatively inexpensive to estimate. Despite the consideration of many other measures and models, Value at Risk (VaR) has been widely adopted. Part of the reason leading to the widespread adoption of VaR was the decision of JP Morgan to create a transparent VaR measurement model, called RiskMetrics.™ RiskMetrics™ was supported by a publicly available database containing the critical inputs required to estimate the model.[4]

Another reason behind the widespread adoption of VaR was the introduction in 1998[5] by the Bank for International Settlements (BIS) of international bank capital requirements that allowed relatively sophisticated banks to calculate their capital requirements based on their own internal modes such as VaR. In this chapter, we introduce the basic concept of VaR as a measurement tool for market risk. In later chapters, we apply the VaR concept to the measurement of credit risk and operational risk exposures.

1.1 ECONOMICS UNDERLYING VaR MEASUREMENT

Financial institutions are specialists in risk management. Indeed, their primary expertise stems from their ability to both measure and manage risk exposure on their own behalf and on behalf of their clients – either through the evolution of financial market products to shift risks or through the absorption of their clients' risk onto their own balance sheets. Because financial institutions are risk intermediaries, they maintain an inventory of risk that must be measured carefully so as to ensure that the risk exposure does not threaten the intermediary's solvency. Thus, accurate measurement of risk is an essential first step for proper risk management, and financial intermediaries, because of the nature of their business, tend to be leading developers of new risk measurement techniques. In the past, many of these models were internal models, developed in-house by financial institutions. Internal models were used for risk management in its truest sense.

Indeed, the VaR tool is complementary to many other internal risk measures – such as RAROC developed by Bankers Trust in the 1970s.[6] However, market forces during the late 1990s created conditions that led to the evolution of VaR as a dominant risk measurement tool for financial firms.

The US financial environment during the 1990s was characterized by the de jure separation of commercial banking and investment banking that dated back to the Glass Steagall Act of 1933.[7] However, these restrictions were undermined in practice by Section 20 affiliates (that permitted commercial bank holding companies to engage in investment banking activities up to certain limits), mergers between investment and commercial banks, and commercial bank sales of some "insurance" products, especially annuities. Thus, commercial banks competed with investment banks and insurance companies to offer financial services to clients in an environment characterized by globalization, enhanced risk exposure, and rapidly evolving securities and market procedures. Concerned about the impact of the increasing risk environment on the safety and soundness of the banking system, bank regulators instituted (in 1992) risk-adjusted bank capital requirements that levied a capital charge for both on- and off-balance sheet credit risk exposures.[8]

Risk-adjusted capital requirements initially applied only to commercial banks, although insurance companies[9] and securities firms had to comply with their own reserve and haircut regulations as well as with market forces that demanded capital cushions against insolvency based on economic model-based measures of exposure – so called economic capital. Among other shortcomings of the BIS capital requirements were their neglect of diversification benefits, in measuring a bank's risk exposure. Thus, regulatory capital requirements tended to be higher than economically necessary, thereby undermining commercial banks' competitive position vis-à-vis largely unregulated investment banks. To compete with other financial institutions, commercial banks had the incentive to track economic capital requirements more closely notwithstanding their need to meet regulatory capital requirements. The more competitive the commercial bank was in providing investment banking activities, for example, the greater its incentive to increase its potential profitability by increasing leverage and reducing its capital reserves.

JP Morgan (now JP Morgan Chase) was one of a handful of globally diversified commercial banks that were in a special position relative to the commercial banking sector on the one hand and the investment banking sector on the other. These banks were caught in between, in a way. On the one hand, from an economic perspective, these banks could be thought of more as investment banks than as commercial banks, with large market risks due to trading activities, as well as advisory and other corporate finance activities. On the other hand this group of globally diversified commercial banks were holding a commercial banking license, and, hence, were subject to commercial bank capital adequacy requirements. This special position gave these banks, JP Morgan being a particular example, a strong incentive to come out with an initiative to remedy the capital adequacy problems that they faced. Specifically, the capital requirements for market risk in place were not representative of true economic risk, due to their limited account of the diversification effect. At the same time competing financial institutions, in particular, investment banks such as Merrill Lynch, Goldman Sachs, and Salomon Brothers, were not subject to bank capital adequacy requirements. As such, the capital they held for market risk was determined more by economic and investor considerations than by regulatory requirements. This allowed these institutions to bolster significantly more impressive ratios such as return on equity (ROE) and return on assets (ROA) compared with banks with a banking charter.

In response to the above pressures, JP Morgan took the initiative to develop an open architecture (rather than in-house) methodology, called RiskMetrics. RiskMetrics quickly became the industry benchmark in risk measurement. The publication of RiskMetrics was a pivotal step moving regulators toward adopting economic capital-based models in measuring a bank's capital adequacy. Indeed, bank regulators worldwide allowed (sophisticated) commercial banks to measure their market risk exposures using internal models that were often VaR-based. The market risk amendments to the Basel accord made in-house risk measurement models a mainstay in the financial sector. Financial institutions worldwide moved forward with this new approach and never looked back.

1.1.1 What is VaR?

It was Dennis Weatherstone, at the time the Chairman of JP Morgan, who clearly stated the basic question that is the basis for VaR as we know it today – "how much can we lose on our trading portfolio by tomorrow's close?" Note that this is a risk measurement, not a risk management question. Also, it is not concerned with obtaining a portfolio position to maximize the profitability of the bank's traded portfolio subject to a risk constraint, or any other optimization question. Instead, this is a pure question of risk measurement.

There are two approaches to answering Weatherstone's question. The first is a probabilistic/statistical approach that is the focus of the VaR measure. To put the VaR approach into perspective, we briefly consider the alternative approach – an event-driven, non-quantitative, subjective approach, which calculates the impact on the portfolio value of a scenario or a set of scenarios that reflect what is considered "adverse circumstances."[10]

As an example of the scenario approach, consider a specific example. Suppose you hold a $1 million portfolio of stocks tracking the S&P 500 index. For the purpose of our discussion we may assume that the tracking is perfect, i.e., there is no issue of tracking error. To address the question of how much this portfolio could lose on a "bad day," one could specify a particular bad day in history – say the October 1987 stock market crash during which the market declined 22 percent in one day. This would result in a $220,000 daily amount at risk for the portfolio if such an adverse scenario were to recur.

This risk measure raises as many questions as it answers. For instance, how likely is an October 1987-level risk event to recur? Is the October 1987 risk event the most appropriate risk scenario to use? Is it possible that other historical "bad days" should instead be used as the appropriate risk scenario? Moreover, have fundamental changes in global trading activity in the wake of October 1987 made the magnitude of a recurrence of the crash even larger, or, instead, has the installation of various circuit-breaker systems made the possibility of the recurrence of such a rare adverse event even smaller? In chapter 3, we discuss how these questions may be answered in implementing scenario analysis to perform stress testing of VaR-based risk measurement systems.

In contrast to the scenario approach, VaR takes a statistical or probabilistic approach to answering Mr. Weatherstone's question of how much could be lost on a "bad day." That is, we define a "bad day" in a statistical sense, such that there is only an x percent probability that daily losses will exceed this amount given a distribution of all possible daily returns over some recent past period. That is, we define a "bad day" so that there is only an x percent probability of an even worse day.

In order to more formally derive VaR, we must first define some notation. Since VaR is a probabilistic value the 1 percent VaR (or

VaR calculated on the basis of the worst day in 100 days) will yield a different answer than the 5 percent VaR (calculated on the basis of the worst day in 20 days). We denote a 1 percent VaR as $VaR_{1\%}$, a 5 percent VaR as $VaR_{5\%}$, etc. $VaR_{1\%}$ denotes a daily loss that will be equaled or exceeded only 1 percent of the time. Putting it slightly differently, there is a 99 percent chance that tomorrow's daily portfolio value will exceed today's value less the $VaR_{1\%}$. Similarly, $VaR_{5\%}$ denotes the minimum daily loss that will be equaled or exceeded only 5 percent of the time, such that tomorrow's daily losses will be less than $VaR_{5\%}$ with a 95 percent probability. The important practical question is how do we calculate these VaR measures?

1.1.2 Calculating VaR

Consider again the example used in the previous section of a $1 million equity portfolio that tracks the S&P 500 index. Suppose that daily returns on the S&P 500 index are normally distributed with a mean of 0 percent per day and a 100 basis point per day standard deviation. Weatherstone's question is how risky is this position, or, more specifically, how much can we lose on this position by tomorrow's market close?

To answer the question, recall first the basic properties of the normal distribution. The normal distribution is fully defined by two parameters: μ (the mean) and σ (the standard deviation). Figure 1.1 shows the shape of the normal probability density function. The cumulative distribution tells us the area under the standard normal density between various points on the *X*-axis. For example, there is

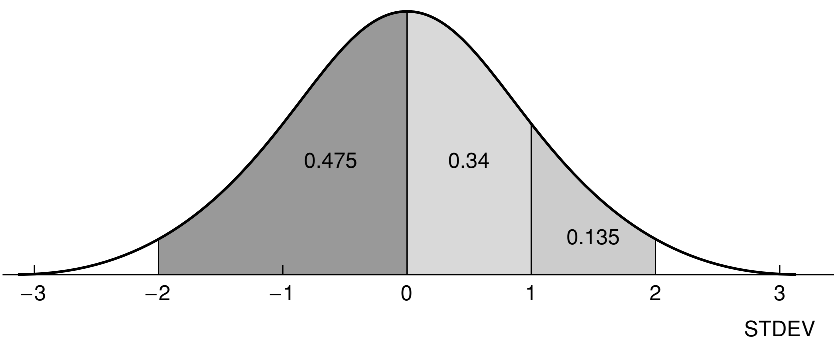

Figure 1.1 The normal probability distribution

Table 1.1 Normal distribution cumulative probabilities for commonly used VaR percentiles

$Prob(X < z)$	0.1%	0.5%	1.0%	2.5%	5.0%	10%
z	−3.090	−2.576	−2.326	−1.960	−1.645	−1.282
VaR	\$30,900	\$25,760	\$23,260	\$19,600	\$16,450	\$12,820

a 47.5 percent probability that an observation drawn from the normal distribution will lie between the mean and two standard deviations below the mean. Table 1.1 shows the probability cutoffs for the normal distribution using commonly used VaR percentiles.[11]

Reading table 1.1 is simple. Given that X is a standard normal random variable (with mean zero and standard deviation one) then, for example, Prob($X < -1.645$) = 5.0 percent. Stated more generally, for any normally distributed random variable, there is a 5 percent chance that an observation will be less than 1.645 standard deviations below the mean. Returning to our equity portfolio example, the daily fluctuations in the S&P 500 index are assumed to be normally distributed with a zero mean and a standard deviation of 100 bp. Using the properties of the normal distribution shown in table 1.1, there is a 5 percent chance that the S&P 500 will decline tomorrow by more than 1.645 × 100 bp = 1.645 percent. Based on the \$1 million equity portfolio in the example, this represents a minimum daily loss of \$16,450 (0.01645 × \$1 million), which will be exceeded only 5 percent of the time. Thus, the equity portfolio's $VaR_{5\%}$ = \$16,450. That is, there is a 5 percent chance that daily losses on the S&P 500-linked equity portfolio will equal or exceed \$16,450. Alternatively, we could say that our portfolio has a 95 percent chance of being worth \$983,550 or more (\$1,000,000 − \$16,450) tomorrow. Using table 1.1, we can compute other VaR measures. For example, $VaR_{1\%}$ = \$23,260 (2.326 × 0.01 × \$1 million), and so on, as shown in table 1.1. We can define VaR for whatever risk level (or confidence level) is deemed appropriate.

We have thus far considered only daily VaR measures. However, we might want to calculate the VaR over a period of time – say a week, a month or a year. This can be done using the daily VaR model and the "square root rule."[12] The rule states that the J-day VaR is $\sqrt{J} \times$ (*daily VaR*). Thus, the one week (5 business days) $VaR_{5\%}$ for the equity portfolio example is $\sqrt{5}$ × \$16,450 = \$36,783. Similarly, the annual (using

250 days as the number of trading days in a year) $VaR_{5\%}$ for the equity portfolio example is $\sqrt{250} - \$16{,}450 = \$260{,}097$; that is, there is a 5 percent probability that the equity portfolio will lose $260,097 or more (or a 95 percent likelihood that the portfolio will be worth $739,903 or more) by the end of one year.

VaR can be calculated on either a dollar or a percentage basis. Up until this point, we have calculated the dollar VaR directly by examining the probability distribution of dollar losses. Alternatively, we could have calculated the percentage VaR by examining the probability distribution of percentage losses as represented by the distribution's standard deviation. For example, consider the weekly $VaR_{5\%}$ computed as $36,783 for the equity portfolio example. If instead of calculating the 5 day dollar $VaR_{5\%}$, the 5 day standard deviation of S&P 500 index returns were instead computed, we would obtain $100 \text{ bp} \times \sqrt{5} = 2.23607$ percent. Calculating the 5 day percentage $VaR_{5\%}$ we obtain $1.645 \times 2.23607 = 3.6783$ percent. This states that there is a 5 percent probability that the S&P 500-linked equity portfolio's value will decline by 3.6783 percent or more over the next week. Given a $1 million portfolio value, this translates into a $36,783 ($1m × 0.036783) dollar $VaR_{5\%}$.

To be widely adopted as a risk measure, VaR certainly appears to satisfy the condition that it be easy to estimate. However, does it satisfy the other condition – that VaR is an accurate risk measure? The answer to that question hinges on the accuracy of the many assumptions that allow the easy calculation of VaR. Unfortunately, it is often the case that the simplicity of the VaR measures used to analyze the risk of the equity portfolio, for example, is in large part obtained with assumptions not supported by empirical evidence. The most important (and most problematic) of these assumptions is that daily equity returns are normally distributed. As we examine these (and other) assumptions in greater depth, we will find a tradeoff between the accuracy of assumptions and ease of calculation, such that greater accuracy is often accompanied by greater complexity.

1.1.3 The assumptions behind VaR calculations

There are several statistical assumptions that must be made in order to make VaR calculations tractable. First, we consider the *stationarity requirement*. That is, a 1 percent fluctuation in returns is equally likely to occur at any point in time. Stationarity is a common assumption in financial economics, because it simplifies computations considerably.

A related assumption is the *random walk* assumption of intertemporal unpredictability. That is, day-to-day fluctuations in returns are independent; thus, a decline in the S&P 500 index on one day of x percent has no predictive power regarding returns on the S&P 500 index on the next day. Equivalently, the random walk assumption can be represented as the assumption of an expected rate of return equal to zero, as in the equity portfolio example. That is, if the mean daily return is zero, then the best guess estimate of tomorrow's price level (e.g., the level of the S&P 500 index) is today's level. There is no relevant information available at time t that could help forecast prices at time $t + 1$.

Another straightforward assumption is the *non-negativity requirement*, which stipulates that financial assets with limited liability cannot attain negative values.[13] However, derivatives (e.g., forwards, futures, and swaps) can violate this assumption. The *time consistency requirement* states that all single period assumptions hold over the multiperiod time horizon.

The most important assumption is the *distributional assumption*. In the simple equity portfolio example, we assumed that daily return fluctuations in the S&P 500 index follow a normal distribution with a mean of zero and a standard deviation of 100 bp. We should examine the accuracy of each of these three assumptions. First, the assumption of a zero mean is clearly debatable, since at the very least we know that equity prices, in the particular case of the S&P 500, have a positive *expected* return – the risk free rate plus a market risk premium.[14] To calibrate the numbers for this non-zero mean return case, let us assume a mean risk free rate of 4 percent p.a. and a risk premium of 6 percent p.a. A total expected return, hence, of 10 percent p.a. translates into a mean return of approximately four basis points per day (i.e., 1000 bp/250 days = 4 bp/day). Hence, an alternative assumption could have been that asset returns are normally distributed with a mean return of four basis points per day rather than zero basis points per day. As we shall see later, this is not a critical assumption materially impacting overall VaR calculations.

Similarly, the assumption of a 100 bp daily standard deviation can be questioned. Linking daily volatility to annual volatility using the "square root rule" we can see that this is equivalent to assuming an annualized standard deviation of 15.8 percent p.a. for the S&P 500 index. The "square root rule" states that under standard assumptions,[15] the J-period volatility is equal to the one period volatility inflated by the square root of J. Here for example, the daily volatility is assumed to be 1 percent per day. Assuming 250 trading days in a year gives us an annual volatility of 1 percent/day $\times \sqrt{250}$ = 15.8 percent p.a. Historically, this is approximately the observed order of magnitude for the volatility of well-diversified equity portfolios or wide-coverage indices in well-developed countries.[16]

The most questionable assumption, however, is that of normality because evidence shows that most securities prices are not normally distributed.[17] Despite this, the assumption that continuously compounded returns are normally distributed is, in fact, a standard assumption in finance. Recall that the very basic assumption of the Black–Scholes option pricing model is that asset returns follow a log-normal diffusion. This assumption is the key to the elegance and simplicity of the Black–Scholes option pricing formula. The instantaneous volatility of asset returns is always multiplied by the square root of time in the Black–Scholes formula. Under the model's normality assumption, returns at any horizon are always independent and identically normally distributed; the scale is just the square root of the volatility. All that matters is the "volatility-to-maturity." Similarly, this is also the case (as shown earlier in section 1.1.2) for VaR at various horizons.

1.1.4 Inputs into VaR calculations

VaR calculations require assumptions about the possible future values of the portfolio some time in the future. There are at least three ways to calculate a rate of return from period t to $t + 1$:

- absolute change $\quad \Delta S_{t,t+1} = s_{t+1} - s_t \quad$ (1.1)
- simple return $\quad R_{t,t+1} = (s_{t+1} - s_t)/s_t \quad$ (1.2)
  (or $1 + R_{t,t+1} = s_{t+1}/s_t$)
- continuously compounded return $\quad r_{t,t+1} = \ln(s_{t+1}/s_t). \quad$ (1.3)

Which computational method is the right one to choose? Let us examine which of these methods conforms to the assumptions discussed in section 1.1.3.

Calculating returns using the absolute change method violates the stationarity requirement. Consider, for example, using historical exchange rate data for the dollar–yen exchange rate through periods when this rate was as high as ¥200/$ or as low as ¥80/$. Do we believe that a change of ¥2 is as likely to occur at ¥200/$ as it is at ¥80/$?

Probably not. A more accurate description of exchange rate changes would be that a 1 percent change is about as likely at all times than a ¥1 change.[18]

The simple return as a measure of the change in the underlying factor, while satisfying the stationarity requirement, does not comply with the time consistency requirement. In contrast, however, using continuously compounded returns does satisfy the time consistency requirement. To see this, consider first the two period return defined as simple return, expressed as follows:

$$1 + R_{t,t+2} = (1 + R_{t,t+1})(1 + R_{t+1,t+2}).$$

Assume that the single period returns, $1 + R_{t,t+1}$ and $1 + R_{t+1,t+2}$, are normally distributed. What is the distribution of the two period return $1 + R_{t,t+2}$? There is little we can say analytically in closed form on the distribution of a product of two normal random variables.

The opposite is true for the case of the two period continuously compounded return. The two period return is just the sum of the two single period returns:

$$r_{t,t+2} = r_{t,t+1} + r_{t+1,t+2}.$$

Assume again that the single period returns, $r_{t,t+1}$ and $r_{t+1,t+2}$, are normally distributed. What is the distribution of the two period return? This distribution, the *sum* of two normals, does have a closed form solution. The sum of two random variables that are jointly normally distributed is itself normally distributed, and the mean and standard deviation of the sum can be derived easily. Thus, in general, throughout this book we will utilize the continuously compounded rate of return to represent financial market price fluctuations.

The mathematics of continuously compounded rates of return can be used to understand the "square root rule" utilized in section 1.1.2 to calculate multiperiod, long horizon VaR. Suppose that the compounded rate of return is normally distributed as follows:

$$r_{t,t+1},\ r_{t+1,t+2} \sim N(\mu, \sigma^2).$$

For simplicity, assume a zero mean ($\mu = 0$) and constant volatility over time. In addition, assume that the two returns have zero correlation; that is, returns are not autocorrelated. The importance of these assumptions will be discussed in detail in chapter 2. The long horizon (here, for simplicity, two period) rate of return is $r_{t,t+2}$ (the sum of $r_{t,t+1}$ and $r_{t+1,t+2}$) is normally distributed with a mean of zero (the sum of the two zero mean returns) and a variance which is just the sum of the variances,[19] which is $2\sigma^2$. Hence, the two period continuously compunded return has a standard deviation which is $\sqrt{(2\sigma^2)} = \sigma\sqrt{2}$. More generally, the J-period return is normal, with zero mean, and a variance which is J times the single period variance:

$$r_{t,t+J} = r_{t,t+1} + r_{t+1,t+2} + \ldots + r_{t+J-1,t+J} \sim N(0, J\sigma^2). \tag{1.4}$$

This provides us with a direct link between the single period distribution and the multiperiod distribution. If continuously compounded returns are assumed normal with zero mean and constant volatility, then the J-period return is also normal, with zero mean, and a standard deviation which is the square root of J times the single period standard deviation. To obtain the probability of long horizon tail events all we need to do is precisely what we did before – look up the percentiles of the standard normal distribution. Thus, using the above result, the VaR of the J-period return is just $\sqrt{J}$ times the single period VaR.

There is one exception to the generalization that we should use continuously compounded rates of return rather than absolute changes in the level of a given index to measure VaR. The exception is with all interest rate related variables, such as zero coupon rates, yields to maturity, swap spreads, Brady strip spreads, credit spreads, etc. When we measure the rate of change of various interest rates for VaR calculations, we measure the rate of absolute change in the underlying variable as follows:

$$\Delta i_{t,t+1} = i_{t+1} - i_t.$$

That is, we usually measure the change in terms of absolute basis point change. For example, if the spread between the yield on a portfolio of corporates of a certain grade and US Treasuries of similar maturity (or duration) widened from 200 to 210 basis points, we measure a 10 basis point change in what is called the "quality spread." A decline in three month zero rates from 5.25 percent annualized to 5.10 percent p.a., would be measured as a change of $\Delta i_{t,t+1} = -15$ bp.

Calculating VaR from unanticipated fluctuations in interest rates adds an additional complication to the analysis. Standard VaR calculations must be adjusted to account for the effect of duration (denoted D), i.e. the fact that a 1 percent move in the risk factor (interest rates) does not necessarily mean a 1 percent move in the position's value, but rather a –D percent fluctuation in value. That is:

$$\textit{1 bp move in rates} \rightarrow -D \textit{ bp move in bond value.} \tag{1.5}$$

To illustrate this, consider a \$1 million corporate bond portfolio with a duration (D) of 10 years and a daily standard deviation of returns equal to 9 basis points. This implies a $VaR_{5\%}$ = 1.645 × .0009 × 10 × \$1 million = \$14,805. However, in general, simply incorporating duration into the VaR calculation as either a magnification or shrinkage factor raises some non-trivial issues. For example, VaR calculations must take into account the convexity effect – that is, duration is not constant at different interest rate levels. This and other issues will be discussed in Chapter 3 when we consider the VaR of nonlinear derivatives.

1.2 DIVERSIFICATION AND VaR

It is well known that risks can be reduced by diversifying across assets that are imperfectly correlated. Indeed, it was bank regulators' neglect of the benefits of diversification in setting capital requirements that motivated much of the innovation that led to the widespread adoption of VaR measurement techniques. We first illustrate the impact of diversification on VaR using a simple example and then proceed to the general specification.

Consider a position in two assets:

- long \$100 million worth of British pound sterling (GBPs);
- short \$100 million worth of Euros.

This position could be thought of as a "spread trade" or a "relative value" position[20] that represents a bet on a rise in the British pound (GBP) relative to the Euro. In order to determine the risk of this position, we must make some assumptions. First, assume that returns are normally distributed with a mean of zero and a daily standard deviation of 80 basis points for the Euro and 70 basis points for the GBP. The percentage $VaR_{5\%}$ of each position can be calculated easily. For the Euro position a 1.645 standard deviation move is equivalent to a move of 1.645 × 80 = 132 bp, and for the GBP a 1.645 standard devation move is equivalent to a move of 1.645 × 70 = 115 bp. Thus, the dollar $VaR_{5\%}$ of the positions are, \$1.32 million for the Euro position (\$100m $\times$ 0.0132) and \$1.15 million for the GBP position (\$100m $\times$ 0.0115).

What is the risk of the entire portfolio, however? Total risk, without accounting for the effect of diversification, could be thought of as the summation of each position's VaR: \$1.32m + \$1.15m = \$2.41 million. However, this summation does not represent an economic measure because risks are not additive. Intuitively, the likelihood of losing money on both parts of this position are slim. This is because the correlation between the \$/Euro rate and the \$/GBP rate is likely to be fairly high and because the two opposite positions (one long and one short) act as a hedge for one another. With a relatively high correlation between the two risk factors, namely, the \$/Euro rate and the \$/GBP rate, the most statistically likely event is to see gains on one part of the trade being offset by losses on the other. If the long GBP position is making money, for example, then it is likely that the short position in the Euro is losing money. This is, in fact, precisely the nature of spread trades.

For the purpose of this example, we shall assume a correlation of 0.8 between the \$/GBP and the \$/Euro rates. This correlation is consistent with evidence obtained by examining historical correlations in the exchange rates over time. What is the $VaR_{5\%}$ for the entire foreign currency portfolio in this example?

To derive the formula for calculation of the VaR of a portfolio, we use results from standard portfolio theory. The continuous return on a two-asset portfolio can be written as follows:

$$r_p = wr_1 + (1 - w)r_2 \tag{1.6}$$

where w represents the weight of the first asset and $(1 - w)$ is the fraction of the portfolio invested in the second asset.[21] The variance of the portfolio is:

$$\sigma_p^2 = w^2\sigma_1^2 + (1 - w)^2\sigma_2^2 + 2w(1 - w)\sigma_{1,2}, \tag{1.7}$$

where σ_p^2, σ_1^2 and σ_2^2 are the variances on the portfolio, asset 1 and asset 2, respectively and $\sigma_{1,2}$ is the covariance between asset 1 and 2 returns. Restating equation (1.7) in terms of standard deviation (recall that $\sigma_{1,2} = \rho_{1,2}\sigma_1\sigma_2$) results in:

$$\sigma_p = \sqrt{\{w^2\sigma_1^2 + (1 - w)^2\sigma_2^2 + 2w(1 - w)\rho_{1,2}\sigma_1\sigma_2\}}, \tag{1.8}$$

where $\rho_{1,2}$ is the correlation between assets 1 and 2. However, the percentage $VaR_{5\%}$ can be stated as $1.645\sigma_p$. Moreover, the 5 percent percentage VaR for asset 1 (asset 2) can be denoted as $\%VaR_1$ ($\%VaR_2$) and can be expressed as $1.645\sigma_1$ ($1.645\sigma_2$). Substituting the expressions for $\%VaR_p$, $\%VaR_1$ and $\%VaR_2$ into equation (1.8) and multiplying both sides by 1.645 yields the portfolio's percentage VaR as follows:

$$\%VaR_p = \sqrt{\{w^2\%VaR_1^2 + (1-w)^2\%VaR_2^2 + 2w(1-w)\rho_{1,2}\%VaR_1\%VaR_2\}}. \quad (1.9)$$

Equation (1.9) represents the formula for the percentage VaR for a portfolio consisting of two assets.[22] However, equation (1.9) is not directly applicable to our spread trade example because it is a zero investment strategy and therefore the weights are undefined. Thus, we can restate equation (1.9) in terms of the dollar VaR. To do that, note that the dollar VaR is simply the percentage VaR multiplied by the size of the position. Thus, the weights drop out as follows:

$$\$VaR_p = \sqrt{\{\$VaR_1^2 + \$VaR_2^2 + 2\rho_{1,2}\$VaR_1\$VaR_2\}}. \quad (1.10)$$

Note that in the equation (1.10) version of the VaR formula, the weights disappeared since they were already incorporated into the dollar VaR values.

Applying equation (1.10) to our spread trade example, we obtain the portfolio *VaR* as follows:

$$\begin{aligned}\$VaR_p &= \sqrt{\{\$1.32^2 + (-\$1.15)^2 + 2 \times 0.80 \times \$1.32 \times (-\$1.15)\}} \\ &= \$0.64\text{MM}.\end{aligned}$$

In the example, the British pound position is long and therefore the $VaR = \$100\text{m} \times 0.0132 = \$1.35$ million. However, the Euro position is short and therefore the $VaR = -\$100\text{m} \times 0.0115 = -\$1.15$ million. These values are input into equation (1.10) to obtain the VaR estimate of $640,000, suggesting that there is a 5 percent probability that the portfolio will lose at least $640,000 in a trading day. This number is considerably lower than the sum of the two VaRs ($2.41 million). The risk reduction is entirely due to the diversification effect. The risk reduction is particularly strong here due to the negative value for the last term in the equation.[23]

There is a large economic difference between the undiversified risk measure, \$2.41 million, and the diversified risk VaR measure \$0.64 million. This difference is an extreme characterization of the economic impact of bank capital adequacy requirements prior to the enactment of the market risk amendment to the Basel Accord which recognized correlations among assets in internal models calculating capital requirements for market risk as part of overall capital requirements. Use of the undiversified risk measure in setting capital requirements (i.e. simply adding exposures) is tantamount to assuming perfect positive correlations between all exposures. This assumption is particularly inappropriate for well-diversified globalized financial instititutions.

1.2.1 Factors affecting portfolio diversification

Diversification may be viewed as one of the most important risk management measures undertaken by a financial institution. Just how risk sensitive the diversified portfolio is depends on the parameter values. To examine the factors impacting potential diversification benefits, we reproduce equation (1.8) representing the portfolio's standard deviation:

$$\sigma_p = \sqrt{\{w^2\sigma_1^2 + (1 - w)^2\sigma_2^2 + 2w(1 - w)\rho_{1,2}\sigma_1\sigma_2\}}.$$

Assuming $\sigma_1^2 = \sigma_2^2 \equiv \sigma^2$ the standard deviation can be rewritten as:

$$\sigma_p = \sigma\sqrt{\{1 - 2w(1 - w)(1 - \rho)\}}. \tag{1.11}$$

Minimizing risk could be viewed as minimizing the portfolio's standard deviation. Using equation (1.11), we can examine the parameter values that minimize σ_p.

Considering the impact of the position weights, w, we can solve for the value of that minimizes σ_p. For simplicity, assume that the position weights take on values between zero and 1 (i.e., there are no short positions allowed). The product of the weights, $w(1 - w)$, rises as w rises from zero to 0.5, and then falls as w rises further to 1. Since $(1 - \rho)$ is always positive (or zero), maximizing $w(1 - w)$ results in maximal risk reduction. Thus, the portfolio with $w = 0.50$ is the one with the lowest possible volatility. For $w = 0.50$, $w(1 - w) = 0.25$. In contrast, if $w = 0.90$, the risk reduction potential is much lower, since $w(1 - w) = 0.09$. This implies that risk diversification is reduced by asset concentration (i.e. 90 percent of the portfolio invested in a single position). This illustrates the diversification effect – risk is reduced when investments are evenly spread across several assets and not concentrated in a single asset.

Equation (1.11) also illustrates the power of correlations in obtaining risk diversification benefits. The correlation effect is maximized when the correlation coefficient (denoted ρ) achieves its lower bound of -1. If the correlation between the two porfolio components is perfectly negative and the portfolio is equally weighted (i.e., $w = 0.50$ and $\rho = -1$), then the portfolio's standard deviation is zero. This illustrates how two risky assets can be combined to create a riskless portfolio, such that for each movement in one of the assets there is a perfectly offsetting movement in the other asset – i.e., the portfolio is perfectly hedged.[24] Finally, equation (1.11) shows that the greater the asset volatility, σ, the greater the portfolio risk exposure – the so-called volatility effect.

1.2.2 Decomposing volatility into systematic and idiosyncratic risk

Total volatility can be decomposed into asset-specific (or idiosyncratic) volatility and systematic volatility. This is an important decomposition for large, well-diversified portfolios. The total volatility of an asset within the framework of a well-diversified portfolio is less important. The important component, in measuring an asset's marginal risk contribution, is that asset's systematic volatility since in a well-diversified portfolio asset-specific risk is diversified away.

To see the role of idiosyncratic and systematic risk, consider a large portfolio of N assets. As before, suppose that all assets have the same standard deviation σ and that the correlation across all assets is ρ. Assume further that the portfolio is equally weighted (i.e., all weights are equal to $1/N$). The portfolio variance is the sum of N terms of own-asset volatilities adjusted by the weight, and $N(N-1)/2$ covariance terms:

$$\sigma_p = \sqrt{\{N(1/N)^2\sigma^2 + 2[N(N-1)/2](1/N)(1/N)\rho\sigma^2\}}. \qquad (1.12)$$

And, hence, we obtain:

$$\sigma_p = \sqrt{\{\sigma^2/N + [(N-1)/N]\rho\sigma^2\}}, \qquad (1.13)$$

or simplifying terms:

$$\sigma_p = \sigma\surd\{1/N + \rho(N - 1)/N\}. \quad (1.14)$$

As N gets larger, i.e., the portfolio becomes better diversified, the first term, $1/N$, approaches zero. That is, the role of the asset's own volatility diminishes. For a large portfolio of uncorrelated assets (i.e., $\rho = 0$) we obtain:

$$\lim_{N\to\infty}\sigma_p = \lim_{N\to\infty}\surd\{\sigma^2/N\} = 0. \quad (1.15)$$

In words, the limit of the portfolio's standard deviation, as N goes to infinity, is zero – a riskfree portfolio. This results from the assumption of a portfolio with an infinite number of uncorrelated assets. Fluctuations in each asset in the portfolio would have a diminishingly small impact on the portfolio's overall volatility, to the point where the effect is zero so that all risk is essentially idiosyncratic. Of course, practically speaking, it would be impossible to find a large number of uncorrelated assets to construct this hypothetical portfolio. However, this is a limiting case for the more realistic case of a portfolio with both idiosyncratic and systematic risk exposures.

To summarise:

- High variance increases porfolio volatility.
- Asset concentration increases portfolio volatility.
- Well-balanced (equally weighted) portfolios benefit the most from the diversification effect.
- Lower correlation reduces portfolio volatility.
- Systematic risk is the most important determinant of the volatility of well-diversified portfolios.
- Assets' idiosyncratic volatility gets diversified away.

1.2.3 Diversification: Words of caution – the case of long-term capital management (LTCM)

Risk diversification is very powerful and motivates much financial activity. As an example, consider the economic rationale of what is known as a fund of hedge funds (FoHF). A FoHF invests in a number of different hedge funds. For example, suppose that a FoHF distributes a total of \$900 million equally among nine hedge funds. Suppose the annualized standard deviation of each of these funds is 15 percent p.a. This is a fairly realistic assumption as such funds are in the habit of levering up their positions to the point that their total volatility is in the vicinity of overall market volatility. The undiversified standard deviation of this investment in the annual horizon is 15 percent p.a. or \$135 million (0.15 × \$900 million).

Suppose that the FoHF managers have two important selection criteria. First, they try to choose fund managers with excess performance ability as measured by their fund's Sharpe ratio – the ratio of expected excess return over and above the risk free rate to the volatility of the fund's assets. Suppose the standard that is applied is that managers are exected to provide a Sharpe ratio of at least 2 – that is, an expected excess return equal to twice the fund's volatility. Thus, the fund's target expected return is equal to the riskfree rate, say 5 percent, plus 2 × 15 percent, for a total expected return of 35 percent p.a.

The second criterion that the FoHF managers apply in choosing funds is to choose funds with low levels of cross correlation in order to better diversify across the investments. They pick one fund which is a macro fund (betting on macroeconomic events), another fund that is in the risk arbitrage business (betting on the results of mergers and acquisitions), another fund in the business of fixed income arbitrage, and so on. Suppose the FoHF managers are successful in obtaining a portfolio of nine such funds with distinct investment strategies such that they believe that a zero correlation across all strategies is a reasonable assumption.

As constructed above, the FoHF portfolio will achieve a strong diversification benefit. Using equation (1.15), the diversified portfolio's standard deviation is $\sqrt{9} \times \$15$ million = \$45 million, much lower than the undiversified standard deviation of \$135 million. This represents a standard deviation for the entire FoHF equal to 5 percent (\$45m/\$900m), or about one-third of the single investment standard deviation of 15 percent. Moreover, since the expected excess return is still 35 percent per fund, the FoHF can achieve a much higher Sharpe ratio. The FoHF's expected excess return is 30 percent with a standard deviation of only 5 percent, thereby yielding a Sharpe ratio of 6 (= 30/5); far in excess of the target Sharpe ratio of 2.

This example describes the structure of Long Term Capital Management's (LTCM) hedge fund. The highly distinguished proprietors of the hedge fund based in Greenwich, Connecticut, thought that they were investing in a well-designed FoHF consisting of a number of trading strategies that were assumed to be independent. This independence assumption across all strategies allowed the firm to lever up significantly. The strategies that LTCM invested in included such trades as: (i) trading on-the-run vs. off-the-run US Treasuries; (ii) trading mortgage backed securities hedged by US Treasury futures; (iii) trading Japanese warrants hedged by related equities; (iv) trading Swedish vs. Italian government bonds betting on Euro convergence; (v) trading global swap spreads (over Treasuries); and (vi) trading long positions in corporates hedged by government bonds, betting on declines in spreads. Theoretically as well as empirically these trades had little in common and it was reasonable to expect that their correlations were close to zero.

Of course, that assumption proved to be fatally false. When Russia defaulted on its sovereign debt obligations in August 1998, all trades turned south together, thereby raising correlations and eliminating diversification benefits just at the moment when they were most needed. The stellar reputations of the managers of LTCM should serve as a reminder to us that any model is only as good as the quality of its assumptions about the model inputs and parameter estimates. As we evaluate different VaR models, we will return to this theme repeatedly, particularly when we attempt to measure operational risk and credit risk VaR using the “incomplete” data sets that are currently available.