

Chapter 5: Measuring Risk—Introduction

5.1 Measures of Risk

Variance & Standard Deviation

If we model a factor as a random variable with a specified probability distribution, then the variance of the factor is the expectation, or mean, of the squared deviation of the factor from its expected value or mean. Let X be the random variable. Let μ be the mean: $\mu = E[X]$, where $E[X]$ denotes the expected value of X . We write the variance of X as $Var[X] = E[(X - \mu)^2]$. While the mean is a measure of the central tendency of the distribution, the variance measures the spread's distribution, i.e. how far the different realizations of X lie from the center.

The standard deviation of a random variable is the square root of the variance. One generally sees the standard deviation of a random variable denoted as σ . The variance is therefore σ^2 .

We often say that a risk factor with a greater variance has greater risk. We shall see that this is not the complete story.

In finance there are two different ways to estimate the volatility of a variable. One way is to look backwards and measure the historical volatility. We will see the formulas for estimating some historical volatilities in the next chapter. A second way exploits the fact that the volatility of some variable often plays a major role in setting the prices of certain financial securities. Therefore, one can use currently observed prices of these securities to back out the implied volatility on the variable. We say that implied volatilities are forward looking since the current security prices are determined by investors' forecasts of the variable's volatility for the horizons of the securities' cash flows. We will see the formulas for estimating implied volatilities in the later chapters on packaging risks.

The Normal Distribution

The normal distribution plays an important role in the practice of risk management. There are many reasons for this. It is a relatively simple and tractable model that seems to capture adequately important aspects of many random variables. Of course, it has its limitations, which we will discuss at various points in these lecture notes. For the moment, we will focus on its foundational use as a model of stock returns.

A Model of Stock Returns and Stock Prices

Suppose that we are analyzing a stock's possible movement through the horizon T , e.g., $T=1$ year. For the moment we will treat this horizon as a single investment period. The current or initial stock price is S_0 , and we model the stock price at T , S_T , as a random variable. As an example, we'll start at $S_0=\$100$. We assume that the stock pays no dividends. Define R_T , as the stock's continuously compounded return through T , meaning:

$$S_T = S_0 e^{R_T}, \quad (5.1)$$

or, equivalently,

$$R_T = \ln\left(\frac{S_T}{S_0}\right) = \ln(S_T) - \ln(S_0). \quad (5.2)$$

The return is obviously also a random variable.

We assume that this annual return is normally distributed with mean $\mu=10\%$ and standard deviation $\sigma=22\%$, writing

$$R_T \sim \text{Normal}(\mu, \sigma^2). \quad (5.3)$$

Since the stock's return is normally distributed, the mean return and the median return are the same:

$$\text{Median}(R_T) = \mu. \quad (5.4)$$

With the normal distribution, it is straightforward to construct confidence bounds around the median return. For example, the 1-standard deviation confidence bounds, corresponding to the 68% confidence interval are given by:

$$UR_T = \mu + \sigma, \quad (5.5)$$

$$LR_T = \mu - \sigma. \quad (5.6)$$

For our example, $UR_T=32\%$ and $LR_T=-12\%$. The top panel of Figure 5.1 shows the probability distribution of the returns with $\mu=10\%$ and $\sigma=22\%$, and marks these confidence bounds.

The probability distribution for the stock price is different from the distribution of returns in important ways. Rewriting the relationship between the stock price and return shown in equation (5.2) we have,

$$\ln(S_T) = \ln(S_0) + R_T. \quad (5.7)$$

Since the return is a normally distributed random variable, the equation above implies that the log of the price is normally distributed,

$$\ln(S_T) \sim \text{Normal}(\ln(S_0) + \mu, \sigma^2). \quad (5.8)$$

In that case, the price itself is log-normally distributed,

$$S_T \sim \text{Log-Normal}(\ln(S_0) + \mu, \sigma^2). \quad (5.9)$$

The bottom panel of Figure 5.1 shows the log-normal distribution of the stock price. Contrast the normal distribution of returns shown in the top panel with the log-normal distribution of prices shown in the bottom panel. The normal distribution of returns has tails that go out in both directions indefinitely. Although the probabilities may be small, every extremely negative or extremely positive return shows some positive probability. A log-normally distributed random variable can never go below zero, which is an appropriate feature for a distribution describing stock prices. Therefore, we say that the log-normal distribution is skewed, with upper tail on the right side of the distribution being much longer than the left tail.

Although the stock price distribution is skewed, there is still a one-to-one, monotonic correspondence between returns and stock prices, and the median stock price can be calculated from the median return, and since, we have:

$$\text{Median}(S_T) = S_0 e^{\mu}. \quad (5.10)$$

This median stock price at $T=1$ is \$110.52. Similarly, we can construct confidence bounds for the return and for the stock price. These are calculated by from the upper and lower bounds in equations (5.5) and (5.6):

$$US_T = S_0 e^{UR_T}, \quad (5.11)$$

$$LS_T = S_0 e^{LR_T}. \quad (5.12)$$

The upper confidence bound for our example is \$137.71, and the lower confidence bound is \$88.69. The bottom panel of Figure 5.1 marks these confidence bounds.

The skewness of the log-normal distribution of stock prices means that the mean and the median will not be equal. The mean of the lognormal distribution lies to the right of the median (i.e. above the median). The mean stock price reflects the variance, and this is what raises it above the median:

$$E(S_T) = S_0 e^{\mu + \frac{1}{2}\sigma^2} > S_0 e^{\mu} = \text{Median}(S_T). \quad (5.13)$$

In our example, the expected or mean stock price is \$113.22. This is also marked in the bottom panel of Figure 5.1. The greater the variance of the return, the more skewed is the lognormal distribution and the greater is the amount by which the mean stock price exceeds the median.

Slicing the Distribution

A key element of the revolution in finance that is risk management is the ability to carve risk into ever finer and finer components. Therefore, when we ask about a risk factor such as the price of a stock, we will want to know about more than just the expected return or the variance. We want to be able to ask questions about subsets of the distribution. For example, two classic questions are (i) “what is the probability that the stock price will be greater than X?”, and (ii) “what is the expected price of the stock, given that the price will be greater than X?” Since the normal distribution has been so well studied, it is straightforward to answer questions like this, and the answers will be important in later sections of these lecture notes, so we provide them here.

To answer the first question, we start by noting that $S_T > X$ exactly when $\ln(S_T) > \ln(X)$, so that:

$$\Pr(S_T > X) = \Pr(\ln(S_T) > \ln(X)).$$

Since $\ln(S_T)$ is normally distributed, if we subtract the mean and divide by the standard deviation, we will transform it to a standard normal random variable for which the relevant probabilities are readily to hand. Doing this to both sides of the expression inside the probability function gives us:

$$= \Pr\left(\frac{\ln(S_T) - \ln(S_0) - \mu}{\sigma} > \frac{\ln(X) - \ln(S_0) - \mu}{\sigma}\right).$$

So the expression on the left-hand-side of the inequality sign is a standard normally distributed random variable, which we write as z :

$$= \Pr\left(z > \frac{\ln(X) - \ln(S_0) - \mu}{\sigma}\right).$$

Taking advantage of symmetry around zero in the standard normal distribution, we can rewrite this as

$$= \Pr\left(z < -\frac{\ln(X) - \ln(S_0) - \mu}{\sigma}\right)$$

Rearranging the numerator on the right-hand-side of the inequality sign gives us:

$$\begin{aligned} &= \Pr\left(z < \frac{\ln(S_0) - \ln(X) + \mu}{\sigma}\right) \\ &= N\left(\frac{\ln(S_0) - \ln(X) + \mu}{\sigma}\right), \end{aligned}$$

where $N(\bullet)$ is the cumulative standard normal distribution function. Collapsing the intermediate lines of the derivation above, we have

$$\Pr(S_T) > X = N(\hat{d}_2), \tag{5.14}$$

Where

$$\hat{d}_2 = \frac{\ln(S_0) - \ln(X) + \mu}{\sigma}. \tag{5.15}$$

We can evaluate equation (5.15) in Excel using the NormSDist function. For our assumed parameters of $T=1$ years, $\mu=10\%$ and $\sigma=22\%$, we have $\hat{d}_2=0.93346$ and we arrive at the solution that $\Pr(S(t_i)>X)= 82\%$.

For the second question, we solve for the answer as follows:

$$E[S_T | S_T > X] = \frac{\int_X^\infty S_T f(S_T) dS_T}{\int_X^\infty f(S_T) dS_T}$$

$$= S_0 e^{\mu} N\left(\frac{\ln(S_0) - \ln(X) + (\mu + \sigma^2)}{\sigma}\right) / N(\hat{d}_2)$$

$$= S_0 e^{\mu} N(\hat{d}_1) / N(\hat{d}_2) \quad (5.16)$$

where,

$$\hat{d}_1 = \frac{\ln(S_0) - \ln(X) + (\sigma + \sigma^2)}{\sigma} = \hat{d}_2 + \sigma. \quad (5.17)$$

Using equations (5.16) and (5.17), in our numerical example, we have $\hat{d}_1 = 1.15346$ and the expected stock price at $T=1$ given that it is greater than \$90 is \$117.34. This should be contrasted with the value of XXXX from our Monte Carlo simulation.

Monte Carlo Simulation

Another way to analyze the random variables R_T and S_T is through Monte Carlo simulation. In a Monte Carlo simulation we essentially create the distribution through brute force, generating a large sample of the random variable. Once we have the large sample, if we want to ask questions about the properties of the distribution like those above, we can simply evaluate the sample and determine its properties. Of course, our answers will only be approximately correct, because the sample will not exactly produce the complete, true distribution. But if the sample is large enough, it is likely to be approximately close to the distribution. Moreover, as we shall see later, there are specialized techniques for improving the efficiency of the sample.

For this particular problem, a simulation does not seem to be necessary, since we obtain explicit formulas for the distribution and explicit solutions to the questions posed. There will be problems later for which no explicit calculation is available, and a Monte Carlo solution seems especially well suited for extracting a good answer. Since the practice of Monte Carlo simulation is so important, it is instructive to implement it even in this simple context. By implementing it in a context where we already know the answers, we can see how well the methodology works and appreciate the extent to which it only approximates the right answer.

Central to the operation of the simulation is a standard normal random variable, ε . We need to construct our two random variables R_T and S_T from ε . This is easily done. They are given by:

$$R_T = \mu + \sigma \varepsilon,$$

$$S_T = S_0 e^{R_T} = S_0 e^{\mu + \sigma \varepsilon}.$$

To simulate draws of the two random variables, we first produce a set of draws of the standard normal random variable, $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_k$. These can be readily generated in a standard Excel spreadsheet or with any of a number of other mathematical programs.¹ We then calculate the two variables for each of the separate draws, $i=1,2,\dots,k$, which gives us a sample of returns, $R_{T,1}, R_{T,2}, \dots, R_{T,k}$, and a corresponding sample of stock prices, $S_{T,1}, S_{T,2}, \dots, S_{T,k}$. We should find that the returns appear approximately normally distributed, with mean μ and variance σ^2 , and that the stock prices are log-normally distributed with mean parameter $\ln(S_0)+\mu$ and variance parameter, σ^2 .

Table 5.1 shows the first 10 draws of the standard normal random variable and the calculation of the corresponding returns and stock prices. The top panel of Figure 5.2 shows a histogram of returns for a simulation of 100 draws. The bottom panel of Figure 5.2 shows the corresponding histogram of stock prices. For a large enough sample of draws, this histogram should approximate the true probability distribution for the process we are simulating, and so can be used to estimate an answer for certain standard types of probability questions. For example,

- What is the expected cumulative return to $T=5$? In this small sample, the mean cumulative return at $T=5$ is 45.6%. This sample mean is our Monte Carlo estimate of the expected cumulative return. Expressed as an annual return this is 9.1%, which we can see is close to, but not exactly the same as the true model expected annual return of $M=10\%$.
- What is the expected stock price at $T=5$? In this sample, the mean stock price at $T=5$ is \$181.47. This corresponds to a cumulative return of 59.6% or 11.9% per annum.
- What is the median stock price at $T=5$? In this sample, the median stock price at $T=5$ is \$158.64. This corresponds to a cumulative return of 46.1% or 9.2% per annum.
- What is the volatility of cumulative returns at $T=5$? The sample standard deviation of cumulative returns at $T=5$ is 53.0%. When annualized, this is 23.7%. This is our Monte Carlo estimate of the model volatility. The true model volatility is 22%.
- What is the probability that the stock price at $T=5$ is greater than \$90? In our small sample, 84% of the paths end with a stock price greater than \$90. This is our Monte Carlo estimate of the requested probability.

¹ There are many different techniques for generating a draw of a standard normal random variable, ε . The main issues in choosing among them is the degree of precision with which the repeated use of the technique approximates the standard normal distribution. One method that can be used in Excel is the Inverse Transform Method. It uses Excel's approximation of the cumulative distribution function for a normally distributed random variable, the NORMINV function, and Excel's random number generator, the RAND function, which generates uniformly distributed numbers between 0 and 1. Since the cumulative distribution function ranges from 0 to 1, we can write $\text{Norminv}(\text{Rand}(),0,1)$, and it returns numbers with approximately the standard normal distribution. This technique has limitations set by Excel's approximation of the cumulative distribution function, among others.

- What is the expected stock price at $T=5$, given that it is greater than \$90? In our sample, the average stock price at $T=5$ among those paths for which the price is greater than \$90, is \$202.15. This is our Monte Carlo estimate of the requested conditional expected value.

As tedious as these types of calculations are, they are nevertheless readily doable with a computer.

Obviously, the accuracy of our estimated answers to these questions is limited, in part, by the size of the sample we take. A sample of 100 is useful for getting an initial feel for a problem, but is far too small for reliable results on any interesting questions. It is common to see results presented using a sample size of 10,000 runs, but there is nothing sacrosanct about this number. The right sample needed depends upon the degree of accuracy required and the particular function being estimated. The accuracy also depends on other elements of the simulation. For example, the formula and procedure used to generate the random number can affect the accuracy. Also, it should be clear that simply reproducing the distribution in the way we have described—simple sampling—is a sort of brute force technique. A number of techniques have been developed to deliberately select a sample that most efficiently reflects the properties of the underlying distribution, i.e., using the smallest sample size. See, for example, Latin hypercube sampling or orthogonal sampling. These techniques will not be explored in any more detail here.

More important than the size or technique of sampling, of course, is the question of whether the mathematical model we are using is the right one and whether the parameter values we have selected are right. As always, we are subject to the dictum ‘garbage in, garbage out.’

Extending the Model to Other Risk Factors

The model of stock returns and stock prices presented above embodies a simple, but important, technical trick. The normal distribution is a convenient modeling device, and it would be nice, from the perspective of the modeler, if all variables could be well described by the normal distribution. Unfortunately, they can't. That's just not the way the world is. Stock prices, for example, cannot fit the normal distribution, since the price of the stock can never go negative, while the normal distribution has tails that are unbounded in both directions. There is no getting around that fact. But while we cannot get around that fact, it turns out that a careful reworking of the problem allows us to still employ the normal distribution. It is not stock prices that are normally distributed, but stock returns. Stock prices are then related to returns by a simple function—exponential growth. Modeling stock prices remains a slight bit more complicated than would be the case if they themselves were normally distributed. However, by drilling down to the

underlying determinant of stock prices that can be modeled with the normal distribution, the extra complication is manageable.

The trick works for a few other variables. Take, for example, the price of some commodity like gold. As with stock prices, the price of gold cannot be negative, and so we cannot sensibly model the future gold price as a normally distributed random variable. However, we can drill down, and model the growth rate in the price of gold as a normally distributed random variable. In that case, the future gold price is log-normally distributed, and all of our earlier analysis of stock returns and stock prices can be re-employed to analyze growth rates in gold prices and the gold price itself.

Let the initial price of gold be denote P_0 , and suppose we are analyzing the evolution of the price of gold through the horizon $T=1$ year. The price of gold at T is P_T , a random variable. Define G_T , as the continuously compounded rate of growth in the price of gold through T , meaning:

$$P_T = P_0 e^{G_T}, \quad (5.1)$$

or, equivalently,

$$G_T = \ln\left(\frac{P_T}{P_0}\right) = \ln(P_T) - \ln(P_0). \quad (5.2)$$

The growth rate is also a random variable. This structure exactly mimics the structure above for the price of a stock and the rate of return on the stock that pays no dividends.

We assume that the annual growth rate in the price of gold is normally distributed with mean μ and standard deviation σ , writing

$$G_T \sim Normal(\mu, \sigma^2). \quad (5.3)$$

Since the growth rate is normally distributed, the mean growth rate and the median growth rate are the same:

$$\text{Median}(G_T) = \mu. \quad (5.4)$$

The 1-standard deviation confidence bounds, corresponding to the 68% confidence interval are given by:

$$UG_T = \mu + \sigma, \quad (5.5)$$

$$LG_T = \mu - \sigma. \quad (5.6)$$

The probability distribution for the price of gold is log-normal since

$$\ln(P_T) = \ln(P_0) + G_T. \quad (5.7)$$

and therefore,

$$\ln(P_T) \sim \text{Normal}(\ln(P_0) + \mu, \sigma^2), \quad (5.8)$$

$$P_T \sim \text{Log-Normal}(\ln(P_0) + \mu, \sigma^2), \quad (5.9)$$

$$\text{Median}(P_T) = P_0 e^\mu, \quad (5.10)$$

$$UP_T = P_0 e^{UG_T}, \quad (5.11)$$

$$LP_T = P_0 e^{LG_T}. \quad (5.12)$$

$$E(P_T) = P_0 e^{\mu + \frac{1}{2}\sigma^2} > P_0 e^\mu = \text{Median}(P_T). \quad (5.13)$$

With a quick inspection we can see that this is all just a relabeling of the equations used to examine the price of the non-dividend paying stock.

The model of stock returns and stock prices employs only the simplest version of the trick: embedding the normal distribution as describing the random return or growth rate. The trick has other incarnations. Any way that a risk factor can be related via a formula to some underlying combination of normally distributed random variables serves the same purpose. If you work in the risk management and analyze enough risk factors, you will come upon many other versions of this modeling trick.

Of course, the trick is not universally applicable. There will be many risk factors for which the normal distribution is simply not relevant, neither as a direct description of the factor itself, nor as an indirect descriptor. We will touch on some of these in these lecture notes with the objective of capturing the important lessons on measuring and mis-measuring risk. However, a full appreciation of the range of technical choices available in modeling various risk factors and a thorough sense of the tradeoffs involved is beyond these lecture notes.

Multiple Random Variables

Correlation

The correlation between two variables is a measure of the degree to which they tend to move together. The correlation coefficient lies between 1 and -1. When the correlation equals 1, we say that the two variables are perfectly correlated. When the correlation equals -1, we say they are perfectly negatively correlated. When the correlation equals 0, we say they are uncorrelated. In the case that the two variables are joint normally distributed, then being uncorrelated also means the two are independent of one another.

Suppose we have a pair of stocks, A and B , and the returns on the pair of stocks, $R_{A,t}$ and $R_{B,t}$, and each return is normally distributed. Then the correlation coefficient is:

$$\rho_{AB} = \frac{\sigma_{AB}}{\sigma_A \sigma_B} = \frac{E[(R_A - \mu_A)(R_B - \mu_B)]}{\sigma_A \sigma_B}. \quad (5.XX)$$

The Multivariate Normal Distribution

Just as the normal distribution is a convenient tool for analyzing certain random variables in isolation, so too is the multivariate normal distribution useful for analyzing certain sets of random variables. When there are only two random variables, the multivariate reduces to the bivariate normal distribution. The multivariate normal distribution is a natural extension of the single variable normal distribution. The key addition is that we must specify whether the variables move together or not. This requires specifying the correlations.

For the two variable problem—the bivariate normal—we only need to specify a single correlation coefficient.

For the multivariate problem, we need to specify the correlation matrix—the correlation coefficient between every pair of variables. Suppose we have N random variables. Then the correlation matrix has N rows and N columns:

$$\begin{bmatrix} \rho_{11} & \rho_{12} & \cdots & \rho_{1N} \\ \rho_{21} & \rho_{22} & & \\ \vdots & & \ddots & \vdots \\ \rho_{N1} & & \cdots & \rho_{NN} \end{bmatrix}$$

The entry ρ_{12} denotes the correlation of variable 1 with variable 2. This means that the diagonal values must be equal to 1: $\rho_{11}=1, \rho_{22}=1, \dots, \rho_{NN}=1$. It also means that the matrix is mirrored across the diagonal so that $\rho_{12}=\rho_{21}$. Consequently, there are $(N^2-N)/2$ correlations to be specified. In doing so, there are some restrictions on the combinations of correlations which we will not go into here.

The correlation matrix is related to the covariance matrix by the appropriate scaling by the respective volatilities as given in equation (5.XX).

Monte Carlo Simulation

When we want to simulate a pair of correlated risk factors—call them A and B—we need to be able to generate a pair of random shocks each period, $\varepsilon_{A,t}$ and $\varepsilon_{B,t}$, that are correlated with one another. If we have two factors, the task is straightforward. First, we generate two independent random variables, $x_{A,t}$ and $x_{B,t}$, each from the standard normal distribution. Second, we derive the pair of correlated shocks, $\varepsilon_{A,t}$ and $\varepsilon_{B,t}$, from these two independent shocks using a pair of formulas that assures each is from a standard normal distribution and that the two are correlated. The formulas are:

$$\begin{aligned}\varepsilon_{A,t} &= x_{A,t} \\ \varepsilon_{B,t} &= \rho x_{A,t} + \sqrt{1-\rho^2} x_{B,t}.\end{aligned}$$

For a larger number of correlated factors, the task is more complicated. One has to construct the variables using the Cholesky decomposition. That is one of the many details of more complete Monte Carlo simulation that we leave to other texts.

Tail Risk

The assumption of normality in returns and lognormality in stock prices leads to a very tractable and convenient model. This is a valuable property. But does the model fit the data? Are stock returns normally distributed? Are stock prices lognormal?

The lognormal model has proven very valuable for stock prices. At a certain level of precision, it appears to work very well. However, there appear to be a number of ways in which stock prices don't fit the model. One of the most important ones is the observation of fat tails, i.e. more large returns – especially large negative returns – than predicted by the normal distribution.

For example, the normal distribution implies that a 4 standard deviation event – an observation of a stock return above the mean plus 4 standard deviations or below the mean minus 4 standard deviations – should happen only 1 out of 15,800 times. However, looking at returns on the S&P500 index, we find that this happens 1 out of 293 times! This fact is a very general one for financial asset returns: returns appear to have fat tails, with more risk of extreme observations than the normal distribution implies.

If we look at a simple graph of the fit of the normal distribution with actual returns on the stock market, the problem of fat tails doesn't stand out as significant. And if we try to quantify numerically the significance of the poor fit, we find that overall the normal distribution doesn't do too badly. Overall, the discrepancy of the fit is not very large. However, looking at the picture "overall" is very misleading. The revolution in finance that is risk management is all about being able to carve up the distribution into many different pieces. No one has to buy the whole distribution any more. The distribution comes in component packages, and it is important to measure and evaluate these components accurately. It isn't satisfactory to have a good fit overall. Now it is important to be able to "fit" the tails of the distribution with some acceptable level of accuracy. In risk management, sometimes what one is "buying" or "selling" is just the tail of the risk distribution. When that's true, it's obviously important to get the tail right, and how successful we are in modeling the larger, more central portion of the distribution isn't relevant at all.

So the issue of the fat tails in the actual distribution of returns is a very important one for the practical implementation of risk management. A significant amount of research work has gone into finding better fitting models than the normal distribution. Fixing the fat tail problem is one of the criteria for a better fit, but there are also other ways in which the normal distribution fall short of our needs and which are guiding the work. More research work is still needed.

In the remainder of these lecture notes we will continue to primarily employ the normal distribution. It is perhaps the simplest foundational model on which we can build and develop the key insights and ideas encompassed by these lecture notes. If we were to substitute everywhere a different foundational model of risk—or a suite of such models, the overall set of insights would remain the same, although the implementation would be greatly complicated. The added complexity would only obscure the deep strategic issues. That said, it is important for the reader to appreciate the limitations of the models and formulas used in these notes. For some corporate problems the compromises involved in employing models built up from the normal distribution

will be small. There are so many other shortcomings in our assessments and models of many corporate problems that the errors springing from this detail are quite small. But for some other problems, the error will be economically meaningful.

Value-at-Risk

Variance is only one measure of risk. There are many others. Why?

One reason is because the probability distribution of a risky variable is a complex thing. When we try to summarize it with a single parameter, such as variance, we oversimplify it. Mathematically speaking, variance is one “moment” of a probability distribution. Most distributions have other, higher “moments” that are also relevant to accurately describing the risk of the distribution. For example, we have already seen that the log-normal distribution is skewed. Skewness is another “moment” of a probability distribution, and an important aspect of the risk of many probability distributions. Kurtosis is another, and there are more. An interesting property of the normal distribution is that variance does fully characterize the risk of the distribution. The normal distribution has no skewness and no excess kurtosis. In fact, in mathematics-speak, it’s cumulants above order two are zero, and the first two order cumulants are the mean and the variance. “Cumulants” are related to “moments.” So, for the normal distribution, specifying the variance is a complete specification of the risk. But this is not true generally for all probability distributions. Therefore, we can’t simply rely on variance nor any other single number as our only metric for risk. Depending upon the situation, other features of the probability distribution must be taken into account.

Although risk is very complicated, it is still important at times to have a simple metric that summarizes the risks we are taking. In certain instances, all we care about is the risk of loss, and not the shape of the entire distribution. Value-at-Risk (VaR) is a measure of the risk of loss, and it is useful in certain instances. But like variance, VaR is just one measure that cannot accurately summarize all aspects of risk relevant in all circumstances.

VaR is defined as a cutoff value for which the probability of a loss greater than the cutoff value equals a specified percentage. Commonly used percentages are 5% and 1%, and these calculations are called 5% VaR and 1% VaR, respectively. Sometimes these calculations are described by the probability that the losses will not be greater than the cutoff value, and so they are called 95% VaR and 99% VaR. The VaR is therefore essentially a one-tailed confidence bound.

VaR is a calculation developed for banks and portfolio managers owning a collection of securities and other tradeable assets. The risks of each of the individual securities and assets are evaluated together, taking into account their correlations one with another.

The calculation is careful to specify a time horizon over which a loss may occur before the security can be sold. Therefore, we speak of a 1-day VaR, in which the calculations measure the probable extreme price movements over the course of a full day, or a 1-week VaR in which the calculations must measure the probable extreme price movements over the course of a full week, or a quarterly VaR. The 1-week VaR will be much larger—i.e. show larger potential losses—than the 1-day VaR, since prices can move farther with greater probability over the longer time window. The quarterly VaR will be larger still. The choice of the appropriate time horizon for a VaR calculation is generally determined by the liquidity of the market for the securities in the portfolio. A VaR for a portfolio of emerging market stocks should use a longer time horizon than a VaR for a portfolio of US Treasury securities. How to handle a portfolio that contains both is an example of the complications that arise as one attempts to implement the methodology.²

For corporate risk management, VaR is a tricky risk measure to use. A non-financial corporation's assets are not generally composed significantly of tradeable securities. While plant and equipment, trademarks, product lines and intellectual property can be sold, the issues involved are very different from those relevant to a manager of a portfolio of securities. It is difficult to specify an appropriate time horizon, and just as difficult to determine if the appropriate metric is "value". Some analysts have attempted to develop similar measures, such as cash-flow-at-risk. We will discuss these issues later in these lecture notes when we try to model the non-financial corporation's hedging and liabilities management problems.

The origins of VaR as described in an excerpt from "Risk Management" by Joe Nocera, originally published in The New York Times Magazine, January 4, 2009

The late 1980s and the early 1990s were a time when many firms were trying to devise more sophisticated risk models because the world was changing around them. Banks, whose primary risk had long been credit risk — the risk that a loan might not be paid back — were starting to meld with investment banks, which traded stocks and bonds. Derivatives and securitizations — those pools of mortgages or credit-card loans that were bundled by investment firms and sold to investors — were becoming an

² A good textbook providing an in-depth study of VaR is Philippe Jorion's *Value at Risk—the New Benchmark for Managing Financial Risk*, McGraw Hill.

increasingly important component of Wall Street. But they were devilishly complicated to value. For one thing, many of the more arcane instruments didn't trade very often, so you had to try to value them by finding a comparable security that did trade. And they were sliced into different pieces — tranches they're called — each of which had a different risk component. In addition every desk had its own way of measuring risk that was largely incompatible with every other desk.

JPMorgan's chairman at the time VaR took off was a man named Dennis Weatherstone. Weatherstone, who died in 2008 at the age of 77, was a working-class Englishman who acquired the bearing of a patrician during his long career at the bank. He was soft-spoken, polite, self-effacing. At the point at which he took over JPMorgan, it had moved from being purely a commercial bank into one of these new hybrids. Within the bank, Weatherstone had long been known as an expert on risk, especially when he was running the foreign-exchange trading desk. But as chairman, he quickly realized that he understood far less about the firm's overall risk than he needed to. Did the risk in JPMorgan's stock portfolio cancel out the risk being taken by its bond portfolio — or did it heighten those risks? How could you compare different kinds of derivative risks? What happened to the portfolio when volatility increased or interest rates rose? How did currency fluctuations affect the fixed-income instruments? Weatherstone had no idea what the answers were. He needed a way to compare the risks of those various assets and to understand what his companywide risk was.

The answer the bank's quants had come up with was Value at Risk. To phrase it that way is to make it sound as if a handful of math whizzes locked themselves in a room one day, cranked out some formulas, and — presto! — they had a risk-management system. In fact, it took around seven years, according to Till Guldemann, an elegant, Swiss-born, former JPMorgan banker who ran the team that devised VaR and who is now vice chairman of SunGard Data Systems. "VaR is not just one invention," he said. "You solved one problem and another cropped up. At first it seemed unmanageable. But as we refined it, the methodologies got better."

Early on, the group decided that it wanted to come up with a number it could use to gauge the possibility that any kind of portfolio could lose a certain amount of money over the next 24 hours, within a 95 percent probability. (Many firms still use the 95 percent VaR, though others prefer 99 percent.) That became the core concept. When the portfolio changed, as traders bought and sold securities the next day, the VaR was then recalculated, allowing everyone to see whether the new trades had added to, or lessened, the firm's risk.

"There was a lot of suspicion internally," recalls Guldemann, because traders and executives — nonquants — didn't believe that such a thing could be quantified mathematically. But they were wrong. Over time, as VaR was proved more correct than not day after day, quarter after quarter, the top executives came not only to believe in it but also to rely on it.

For instance, during his early years as a risk manager, pre-VaR, Guldemann often confronted the problem of what to do when a trader had reached his trading limit but believed he should be given more capital to play out his hand. "How would I know if he should get the increase?" Guldemann says. "All I could do is ask around. Is he a good guy? Does he know what he's doing? It was ridiculous. Once we converted all the limits to VaR limits, we could compare. You could look at the profits the guy made and compare it to his VaR. If the guy who asked for a higher limit was making more money

with lower VaR” — that is, with less risk — “it was a good basis to give him the money.”

By the early 1990s, VaR had become such a fixture at JPMorgan that Weatherstone instituted what became known as the 415 report because it was handed out every day at 4:15, just after the market closed. It allowed him to see what every desk’s estimated profit and loss was, as compared to its risk, and how it all added up for the entire firm. True, it didn’t take into account Taleb’s fat tails, but nobody really expected it to do that. Weatherstone had been a trader himself; he understood both the limits and the value of VaR. It told him things he hadn’t known before. He could use it to help him make judgments about whether the firm should take on additional risk or pull back. And that’s what he did.

What caused VaR to catapult above the risk systems being developed by JPMorgan competitors was what the firm did next: it gave VaR away. In 1993, Guldemann made risk the theme of the firm’s annual client conference. Many of the clients were so impressed with the JPMorgan approach that they asked if they could purchase the underlying system. JPMorgan decided it didn’t want to get into that business, but proceeded instead to form a small group, RiskMetrics, that would teach the concept to anyone who wanted to learn it, while also posting it on the Internet so that other risk experts could make suggestions to improve it. As Guldemann wrote years later, “Many wondered what the bank was trying to accomplish by giving away ‘proprietary’ methodologies and lots of data, but not selling any products or services.” He continued, “It popularized a methodology and made it a market standard, and it enhanced the image of JPMorgan.”

JPMorgan later spun RiskMetrics off into its own consulting company. By then, VaR had become so popular that it was considered the risk-model gold standard.

5.2 Risk Factors and Exposure

Exposure is the variation in one variable— such as the cashflow from a project, the cash payoff on a security, the value of a project or asset or share of stock—that is associated with variation in another variable which we call the underlying risk factor. If changes in the factor are associated with changes in the value of an asset, then the asset value is said to be exposed to the factor.

The value of an asset can be exposed to risk from many different sources. Each of these is called a risk factor. If changes in the factor cause changes in the value of an asset, then the asset value is said to be exposed to the factor. The asset that is exposed could be, for example, the total market value of a company, the market value of a division or a project, the value of a supply contract, or the value of a security such as a debt liability, a futures or options contract, or the stock. In the US, companies with publicly traded stock must file a 10-K report annually with the

Securities and Exchanges Commission. This report must include a listing of key risk factors. A quick overview of any number of these filings will make clear the wide, wide array of variables that are risk factors.

A number of the risk factors are called “market risks”. These include exchange rates, stock market indexes, the rate of inflation, interest rates, and the prices of various commodities. We have data on these variables. Each of them is quoted on widely watched markets, and tracked carefully over time. We can get a handle on the history of these risk factors using standard statistical tools, as we shall see below.

Only a very few risk factors have easily observed market indexes associated with them. Take, for example, the demand for a single company’s specific products, which may wax or wane over time. Demand is an important determinant of the price the company can charge for its products and therefore of the company’s ultimate profitability and value. While the company may have methods for determining the demand for its products, there is no widely cited index and no readily usable data series we can assess with statistical tools. The underlying determinants of the demand for a single company’s specific products may be more nebulous still. Broad technological progress in society may make one company’s products, services and skills passé and another company’s suddenly relevant. Indexing something like broad technological progress is close to hopeless. There are many critical risk factors that are difficult to summarize with a simple statistic, including changing regulations, such as pollution standards for power plants, fuel economy standards for automakers, labor regulations and health insurance rules.

Despite the difficulty of reliably applying standard statistical tools to quantify the risk associated with certain factors, we will nevertheless try to think about these risks using the same mental concepts, and then wrestle with the adequacy or inadequacy of the effort. Confronting our limitations on these difficult to quantify risk factors may offer insights and cautions that apply just as well to those risk factors that appear to be easily indexed and quantified.

Risk factors need not be something exogenous to the company. A number of key risk factors have to do with the company’s own ability to execute on its strategy. The right way to manage risks will depend a lot on whether the risks are in any way under the control of the company or whether they are entirely exogenous.

Formula Exposures

Many contracts have explicit formulas defining the cash flows owed as a function of some market variable. For example, a floating rate bond may specify the interest rate used in determining the semi-annual coupon payments as a function of the LIBOR, the London Interbank Offered Rate, prevailing the period prior to the payment. The formula could set the rate in period t at $\text{LIBOR}_{t-1} + 4$ basis points. Another example would be an engineering contract for the construction of a nuclear power plant which escalates the wage rates used each month in accordance with one of the inflation indexes published by the government. The natural gas supply contracts between Russian and German companies typically index the delivered price of gas to the world market price of oil, usually with a lag of some months.

Each of these is an example of a clear and mechanically defined exposure. If we know the value of the underlying factor used in the contract—the LIBOR, the inflation index, the world market price of oil—we can calculate the correct cash flow. We know precisely how the cash flow will vary as the underlying risk factor varies: the exposure is defined by the contract formula and language.

Receivables denominated in a foreign currency are similar to these contract exposures. The contract specifies a fixed payment owed by a given date, but the recipient plans to convert the foreign currency paid into its home currency. The home currency cash flow is exposed to changes in the foreign exchange rate between the two currencies.

Linear and Non-linear Exposures

Each of the examples used at the opening of this section described a linear exposure. The cash flow owed varied by a constant factor with a change in the underlying risk factor. A few contractual exposures are linear like these. Many are not.

Suppose that the floating rate bond also had a cap setting the maximum interest rate to be paid in any semi-annual period. No matter how high the LIBOR, the contract interest rate would not exceed the cap. Suppose that the engineering contract for the nuclear power plant had a floor on the wage rates to be used. Many contracts for the delivery of LNG, liquefied natural gas, are indexed to oil, but with the index taking an ‘S’ shape, almost as if there were both a floor and a cap. If the price of oil escalates, the gas delivery price initially increases linearly with the price of oil. But if the escalation continues, then the rate of increase in the price of gas gradually falls. At

some point, there is a maximum or cap on the price of gas, independent of the price of oil. The same attenuation of the exposure occurs on the downside.

These are all examples of non-linear exposures.

Non-linear exposures are central to modern risk management.

Non-linear exposures are difficult to value properly using the risk-adjusted discount rate methodology. The revolution in finance that is risk management makes it possible to accurately value any well-defined exposure, no matter how non-linear. We will turn to the task of valuation in later chapters, starting with those on “pricing risk”.

The non-linearity of certain exposures creates major challenges in measuring, evaluating and controlling exposures. Why? Suppose I repeatedly take a bet with a very, very, very small positive payoff except for an event far, far, far off in the extreme tail of a normal probability distribution. It is entirely possible that over a long series of time the performance of my sequence of bets won't reflect the risk at all. It is easy to become complacent and imagine there is no major downside. Until one day, when, BANG. The bet explodes.

When a company's assets are just linear exposures, then we can evaluate them pretty well by looking at any sufficiently lengthy sample of performance over any window of time. Indeed, for a linear exposure, to accurately measure the risk, it is important not to be confused by one or two outliers or a particularly unrepresentative string of events. We try to look at a sufficiently large or lengthy sample so as to avoid an excessive focus on unrepresentative tail events. The reverse can be true when companies invest in non-linear exposures. Long strings of performance may reveal nothing about the important risk at hand. The true risk exposure only reveals itself occasionally, and possibly in exceptional times. To accurately assess risks out in the tail requires a concerted focus on extreme tail events or on whatever well defined subset of events create the large exposure.

Cash Flow Exposure and Value Exposure

Exposure is a concept about the relationship between two variables. One variable is the underlying risk factor. The other is the variable that is exposed. In these lecture notes we will focus primarily on two types of exposure: cash flow exposure and value exposure. The examples above are all about cash flow exposure. Each example involves a contractually defined exposure. A contract can define the contingent cash flow obligations between two parties, but it cannot define the contingent value of the obligations. The value depends upon factors outside of the

contract, first among them the price of risk. It is easy to define a contract or security. It's another thing entirely to value it.

For certain purposes a company may need to know about the exposure of certain cash flows to an underlying risk factor, while for other purposes the company may need to know about the exposure of the value. Of course, cash flow is what generates value, so there must be some close relationship between the cumulative cash flow exposures and the value exposure. Still, the two are distinct.

Cash flow exposure is critical, for example, to the comptroller of a firm needs to manage the firm's accounts payable and accounts receivable. One of the comptroller's objectives is to minimize the amount of the firm's capital that has to be held in liquid, low yielding assets—like cash. Knowing the volatility of the firm's cash flow is critical to determining the size of the balance required on the cash account. Even if every fluctuation in the weekly cash flow implies some future offsetting fluctuation, so that the value is always constant, the comptroller needs to understand the volatility in the timing mis-match so as to efficiently manage the payments and minimize the working capital costs.

Financial strategies designed to hedge value—i.e. to reduce value exposures—can be undermined if they create a large short-term cash flow exposure. Suppose a firm attempts to hedge a distant cash flow using a futures contract. Suppose the underlying risk factor moves, giving the firm an expected gain on its distant future cash flow. The firm will not earn any extra cash inflow today. The expected gain on the distant future cash flow is exactly offset by a loss on the futures contract. A futures contract typically requires a more or less immediate settling up of gains and losses, so the loss on the futures contract means an immediate cash outflow. If the size of the exposure is too large, then the firm may have difficulty making its payment on the hedge and could go bankrupt. This would undermine the so-called “perfect hedge”. In order for the value hedge to truly be perfect, the firm has to live long enough to enjoy the cash flow gain on its original position, and that requires being able to properly manage the near-term cash flow exposures. This topic is taken up in more detail in later chapters of the lecture notes.

A Comparables Approach to Measuring Cashflow-at-Risk for Non-Financial Firms

Stein, Jeremy, Stephen Usher, Daniel LaGattuta and Jeff Youngen, 2001, Journal of Applied Corporate Finance 13(4), 100-109.

The authors constructed cashflow-at-risk distributions for companies using quarterly EBITDA data for approximately 4,000 firms over the years 1991-1995. First,

they constructed a model of cash flow forecasts for each firm so as to be able to center the observed cash flows and identify the unexpected variation quarter to quarter. For this they used a simple autoregressive time series model based on the last four quarters of cash flow. Second, they sorted the companies into 81 different buckets, a high, medium and low bucket in each of four categories: market capitalization, profitability, industry risk and stock price volatility. Each bucket represents a set of comparable companies. Third, they calculated the probability distribution of forecast errors within each bucket, and then calculated the 5% lower tail confidence bound for each bucket.

Figure 5.X reproduces their Figure 1. It shows the year ahead C-FAR, or cashflow-at-risk, for three companies, Coca-Cola, Dell and Cygnus. The cashflow probability distribution for Coca-Cola has the lowest variance and therefore the smallest C-FAR or 5% loss figure. Cygnus has the highest variance and therefore the largest C-FAR.

[check on the number of companies in each bucket]

Table 5.X reproduces their Table 2, Panel A. The table shows all 81 buckets and the corresponding C-FAR figure, i.e. the 5% lower confidence bound. For a firm with \$100 in assets, the number in the cell shows how big a negative shock to one-quarter ahead EBITDA occurs with 5% probability.

These figures are helpful in giving us a sense of how the size of the cash flow exposure of many companies. However, it is difficult to use them for management, since they don't tell us how these total exposures are related to any underlying variables, and they don't allow us to identify how the exposures will respond to any management actions.

The accuracy of these exposure figures depends upon how representative is the time period over which the data was collected. Almost certainly a calculation that included the recent Global Financial Crisis and Great Recession would report larger cashflow exposures. What is the right figure as we go forward from today?

Payoff Diagrams

Payoff diagrams are an important tool in evaluating exposures. Figure 5.X displays three payoff diagrams corresponding to the three cash flow exposure examples used at the top of this section: the semi-annual coupon payment on a floating rate bond, the wage bill on an engineering construction contract for a nuclear power plant, and the invoice charge on a natural gas pipeline delivery contract. The variable on the horizontal axis of each diagram is the underlying risk factor—the LIBOR, the inflation index, the world oil price. The variable on the vertical axis is the cash flow exposed to the risk factor—the coupon payment, the wage bill, the natural gas invoice charge. The lines in each diagram show the exposures. Panel A of Figure 5.X, showing the interest exposure, includes a payoff line for the linear, uncapped exposure, as well as a payoff line for the non-linear, capped exposure. Panel B of Figure 5.X, showing the wage bill, includes a payoff line for the linear exposure, as well as a payoff line for the non-linear exposure which

reflects the floor wage rate. Panel C of Figure 5.X, showing the natural gas invoice exposure, includes a payoff line for the linear exposure, as well as a payoff line for the “S”-shaped exposure.

A payoff diagram is just a “what-if” calculation showing the “shape” of the exposure. It does not necessarily show the probability of any of the outcomes for the risk factor along the horizontal axis. Of course, one can combine the payoff diagram with information about the probability distribution. Figure 5.X shows the linear exposure on exchange rate movements faced by a manufacturer with a receivable outstanding. Overlaid on the horizontal axis is information about the probability distribution of the exchange rate.

Project or Asset Exposure

Most projects are a bundle of exposures. For example, a company that is launching a new product line is uncertain about both the demand for the product, the prices of multiple inputs, the efficiencies of its production process, and so on. Having a complete picture of exposure is complicated. The complications quickly grow as we move from examining an individual product line to examining the exposure for an entire business unit, and they become more complicated still as we move to examine the exposure of an entire firm.

The complications include the fact that most exposures do not come packaged as an explicit formula. Instead, they are embedded in the structure of business. For example, one business may expect to be able to raise its prices in-line with actual inflation, while another business cannot. A third business may be in-between. The three businesses have different exposures to inflation, and none of the exposures is written down in a contract. Each manager understands the competitive pressures and the characteristics of the demand for its products, and therefore has learned to what extent its prices can adjust in-line with inflation. This is how most exposures are identified.

Pro-forma Exposure

A pro-forma calculation of exposure is a straightforward exercise. It starts with a specific transaction, asset or security for which we want to measure the exposure and some underlying risk factor. We identify the mathematical relationship linking the two so that, for example, the value of the asset is expressed as a function of the risk factor, among other things. Finally, we

calculate how the value of the asset changes as the value of the risk factor changes, using the identified mathematical relationship.

For example, for a gold mining firm we know the quantity of gold produced each year, the costs of production, its tax rate and so on. We can construct a discounted cash flow model of the firm's value. The price of gold will be one input to this discounted cash flow valuation. Suppose we vary the price of gold and examine how the discounted value of the firm varies. This produces a measure of the exposure of the gold mining firm's value to fluctuations in the price of gold.

Tufano (1998) did just this exercise for his sample of gold mining firms. His simple model assumes a firm has a known and fixed quantity of reserves of gold in the ground, R , which it plans to extract at a constant rate over N years. Therefore the annual production is $Q=R/N$. He assumes a fixed annual cost of F and a marginal annual cost of C per unit. The production from the mine is sold at a price P . Therefore, the market value of the gold mine firm, V , is given as:

$$V = \sum_{t=1}^N \frac{[Q(P - C) - F](1 - \tau)}{(1 + r)^t}, \quad (5.XX)$$

where τ is the corporate income tax rate and r is the risk-adjusted discount rate for the mining company. We can clearly use this formula to examine how the value of the firm changes as we change any of the variables, e.g., Q , P , C , and F . This is a classic example of a pro forma exposure calculation.

The same exposure can be expressed a number of different ways. Equation (5.XX) can be used to construct a payoff diagram with P as the risk factor and V as the exposure measured. The slope of the exposure is:

$$\frac{\partial V}{\partial P} = Q(1 - \tau) \sum_{t=1}^N \frac{1}{(1 + r)^t}. \quad (5.XX)$$

We can also measure the exposure as a ratio of the percent change in value divided by the percent change in price. Expressed this way, we would have a gold price beta:

$$\beta \equiv \frac{\partial V/V}{\partial P/P} = \frac{PQ(1 - \tau) \sum_{t=1}^N \frac{1}{(1 + r)^t}}{V} = \frac{PQ}{Q(P - C) - F}. \quad (5.XX)$$

Tufano (1998) calculates this gold price beta for his sample of North American gold mining firms and arrives at an average gold price beta of approximately 2.

Note that Equation (5.XX) shows that the gold beta fluctuates with the level of the gold price. For example, as the price falls, the denominator approaches zero sooner than the numerator, so that the beta increases. When the price falls very close to the marginal operating cost, the firm has very high “operating leverage” and therefore a very high sensitivity to fluctuations in the price of gold. So a company’s exposure need not be fixed, but can change with the level of the gold price as well as with changes in other parameters.

The advantage of a pro forma exposure calculation is its simplicity. It is a useful starting point for thinking through a problem. However, its simplicity is also the main disadvantage. The exercise is purely mechanical in the sense that the mathematical relationship employed usually embodies a very superficial representation of the true links between the underlying risk factor and the exposure. In particular, while a host of other variables may appear in the mathematical relationship—or be implicitly assumed in the structure—the pro-forma calculation makes no effort to evaluate how these other variables might change together with the chosen risk factor, and makes no effort to scrutinize whether the mathematical relationship holds over the range of variation in the risk factor. When we dig deeper and try to identify the full structure of the exposure, we are moving to calculate the economic exposure.

One of the restrictive assumptions made in equation (5.XX) links movements in today’s price of gold—the price at the time when the value is being calculated—and changes in the expected price throughout the horizon of production. Note that the calculation in (5.XX) assumes that a \$1 change in the gold price today implies an equal change in the price for the full life of the mine. Is this likely to be the case? We could have made a different implicit assumption. The point is that in order to implement a pro forma calculation of a mine’s exposure we are forced to determine a relationship between variations in today’s price of gold and the expectation about future prices. For most practical problems it is difficult to determine what this relationship ought to be. Sometimes a simplistic relationship is used just because the beta can then be easily calculated. But this, of course, doesn’t give us any confidence in the reliability of our calculation as a measure of the true exposure. In the next chapter we will carefully scrutinize the assumption about the relationship between changes in the current price of gold and the changes in our expectation about the future price of gold at different horizons.

Economic Exposure

How good are the calculations of exposures of gold mining firms made using equation (5.XX)? Tufano (1998) compared the gold betas calculated “analytically” using equation (5.XX) against the empirically observed gold betas calculated from data on the company’s stock returns. Figure 5.X reproduces Tufano’s Figure 2 which graphs the analytically predicted betas against the empirically observed betas. Each point represents the observation for one firm in one year. If for each firm in each year the two calculated betas were equal, then the points would graph along the 45° line. For betas below 2, this seems to be approximately correct. However, for betas above 2, the empirically observed betas are generally lower than the analytically predicted betas. Why?

One possible explanation has to do with how firms adjust their production in the face of changes in the price. Recall that the analytically predicted betas should be high when gold prices are low. Equation (5.XX) assumes the quantity produced stays constant, no matter how low the price falls. But this assumption probably doesn’t hold. As the price falls and approaches the marginal cost of production, there are many actions the company may take that have been assumed away in equation (5.XX). The company may temporarily shut down production, and save its reserves for a time when the price is once again high enough to generate a good margin to cost. If the company has this operating flexibility, then the value will not drop as quickly in the face of a drop in price, and the gold beta will not be as high.

This comparison of empirically derived with analytically derived gold betas illustrates a typical problem with pro-forma exposure calculations. They assume a given and fixed operating strategy and this mis-represents a company’s true exposure to fluctuations in key variables.

An economic exposure calculation attempts to take into account the interrelationships across risk factors, so that if changing one risk factor also implies changing another, then this should be taken into account. In our gold mining example, the pro forma exposure measured by equations (5.XX) and (5.XX) assumes that the cost of production, C and F , do not change when the selling price of gold, P , changes. Whether or not this is true needs to be scrutinized. What if the underlying cause for a change in the price of gold is the general rate of inflation in the economy? Then we might expect the cost of production to increase in tandem with the price of gold. The pro forma exposure measured by equations (5.XX) and (5.XX) also assumes that the rate of production, Q , does not change when the selling price of gold changes. It is easy to imagine cases in which it does not, and equally easy to imagine cases in which it does change. Measuring economic exposure is about thinking these implicit assumptions through more

completely and being explicit about the full constellation of interactions at the heart of the risk being measured.

An especially good example of how economic exposure differs from a pro forma exposure is the topic of foreign exchange rate exposure of two companies. One is selling its products abroad, while the other sells exclusively to its domestic market.

A typical pro forma exposure calculation for the company selling its products abroad will include a variable for the quantity of sales, a variable for the price measured in foreign currency, and a variable for the exchange rate, among other variables. The pro forma exposure to the exchange rate is measured by taking the derivative with respect to the exchange rate, while keeping the quantity and price constant. An economic exposure calculation questions whether or not the price and quantity will be constant as the exchange rate changes. A change in the exchange rate may imply some adaptation to the list price as denominated in a foreign currency. The full economic exposure must reflect the combined effect.

A typical pro forma exposure calculation for the company selling its products to its domestic market will show zero exposure since the exchange rate doesn't even appear in the formula for its cash flows. An economic exposure calculation questions whether or not the price and quantity in the domestic market might not be impacted by the exchange rate. If the company is competing for sales against imports, the exchange rate change may make the imports cheaper, thereby either forcing the company to lower its own prices or to accept a smaller share of the market. The full economic exposure must reflect the combined effect.

An economic exposure calculation should also reflect how operating and investment decisions will naturally be adjusted in the face of changes in a risk factor. If the costs of various inputs change, the company may adjust the quantities of each that it uses in order to maximize its profits. The full economic exposure must reflect both the change to the input costs and the adaptation to the quantities of each input purchased.

Calculating economic exposure “bottom up” is a challenging task. There are many variables and interactions that need to be taken into account. And the structure of the interactions may differ depending upon the size of the variation in the underlying risk factors one considers. Usually some compromises are made that help to yield a useful answer that is qualitatively better than a pro-forma calculation, even if not everything has been thoroughly incorporated.

Scenario Analysis

Oftentimes it is simply impossible to fully specify the mathematical structure of the problem, incorporating all of the subtle interactions between various risk factors? And, in certain cases, trying to simplify the problem too much will be misleading. Sometimes it is foolish to attempt to specify the mathematical structure completely, because the imposing nature of the problem only invites us to delude ourselves into accepting oversimplified representations so that we are at least comforted with some sort of an answer.

One solution to this predicament is to conduct scenario analysis. Versions of scenario analysis are practiced in a variety of institutional contexts. Military strategists are familiar with war gaming and with developing key scenarios for which plans of action must be outlined. In corporate management, the “Shell scenarios” is the most well known example of a now widely practiced type of exercise.

Scenarios provide alternative views of the future. They identify some significant events, main actors and their motivations, and they convey how the world functions. We use scenarios to explore possible developments in the future and to test our strategies against those potential developments. Shell has been using scenarios for 30 years. (from the Shell website, August 2010)

Scenario analysis can serve many different purposes within an organization, and there are different versions implemented in different institutions.

In financial risk management, scenario analysis is often employed in order to explore for ‘worst case’ situations and to evaluate “what ifs” for the simultaneous occurrence of key risk factors. A version of this practice was recently performed on major banks in the US and the EU under the label “stress tests.” Scenario analysis involves detailing some constellation of events—a scenario—and then evaluating the consequences for the relevant project, asset or security’s value.

Factor Models of Exposure

Until now we have analyzed exposures mostly “bottom up”. Exposures are understood in terms of the building blocks creating the exposure. What does the contract formula say? What are the elements of the project’s cash flow? These have been modeled, or theoretical exposures. Another way to approach the problem is by empirical observation. How do the exposures reflect themselves in the data? Factor models are a tool for estimating exposures. Factor models are applied to time series data on stock prices and other market indexes. Because the stock price reflects the aggregate value of the many projects and assets of a firm, and is therefore a mix of

many different exposures, one can think of these as “top down” models of exposure. The model usually abstracts from the particular structure of the firm, and this is another way in which these are “top down” models.

The categorization of models into “bottom up” or “top down” is not as hard and fast as that. And models in both classifications contain theory and can be tested against data. Nevertheless, the generalizations are worth stating.

The One-Factor Model of Stock Returns

A classic exposure model for stocks decomposes each stock’s risk into two pieces. One piece is driven by a single underlying factor shared in common with other stocks. This is called the systematic risk. In principle, this common underlying factor could be some complicated synthesis of many macroeconomic variables. In practice, the common factor is proxied by the return on a broad, diversified portfolio of stocks and is called the market return. While the underlying factor is common to all stocks, the strength of the exposure differs among the stocks. Some stocks will have a high exposure to the factor, with the stock’s return increasing or decreasing sharply, on average, as the factor varies. Other stocks will have a low exposure to the factor, with the stock’s return increasing or decreasing only slightly, on average, as the factor varies. The second piece of a stock’s risk is the non-systematic or idiosyncratic risk particular to each individual stock. This is the residual fluctuations in the stock’s return that are uncorrelated to the market return.

A version of the one-factor model is written:

$$r_S = E[r_S] + \beta_S (r_M - E[r_M]) + \varepsilon_S, \quad (5.XX)$$

where r_S is the random return on a stock, $E[\bullet]$ is the expected value, β_S is a coefficient measuring the stock’s exposure to the market return, r_M , and ε_S is the idiosyncratic component of the stock’s return. The model assumes that the two random returns, r_S and r_M , are joint normally distributed, and, by construction, that r_M and ε_S are uncorrelated. The beta coefficient is the covariance of the stock’s excess return with the market’s excess return, scaled by the market variance

$$\beta_S = \frac{\text{Cov}(r_S, r_M)}{\text{Var}(r_M)}. \quad (5.XX)$$

The stock’s beta measures the stock’s exposure to movements in the return on the market. If the beta is 1, then a 1% increase in the return on the market implies a 1% increase in the return on the

stock, on average. If the beta is 2, then a 1% increase in the return on the market implies a 2% increase in the return on the stock.

In the one-factor model, since the two pieces of risk are uncorrelated with one another, the variance of the stock's excess return, σ_S^2 , is equal to the sum of the variance of the two respective factors:

$$\sigma_S^2 = \beta_S^2 \sigma_M^2 + \sigma_{\varepsilon_S}^2. \quad (5.XX)$$

In a regression used to estimate equation (5.XX) the standard deviation of the residuals is an estimate for σ_{ε} .

Systematic Exposure of Individual Stocks

A study by Andersen et al (2000) estimated the volatilities of the 30 individual stocks in the Dow Jones Industrial Average over the period 1993-1998. The median annual volatility of the stocks was 28%. This varied from a high of 42% for Walmart to a low of 22% for United Technologies. ...

With some reorganization and relabeling, equation (5.XX) can also be written in a more familiar regression equation:

$$R_S = \alpha_S + \beta_S R_M + \varepsilon_S, \quad (5.XX)$$

where α_S is a constant, R_S is the return on a stock in excess of the risk-free rate, r_f , so that $R_S = r_S - r_f$, and similarly, R_M is the return on the market in excess of the risk-free rate, so that $R_M = r_M - r_f$.

The Capital Asset Pricing Model (CAPM) is a version of this one-factor model, but with some additional assumptions generating conclusions about expected returns such that $\alpha_S = 0$. For the CAPM we have

$$R_S = \beta_S R_M + \varepsilon_S, \quad (5.XX)$$

and,

$$E[R_S] = \beta_S E[R_M]. \quad (5.XX)$$

Multifactor Models

The obvious drawback to the one-factor model is how it oversimplifies the systematic risk of stocks. There is no reason to assume only a single systematic source of risk. There may be

several. It's no use to insist that the single factor is really just a summary statistic for an array of factors. And one stock may be very sensitive to only one subset of factor and less sensitive to another subset, while a second stock may be very sensitive to all factors, and so on. No single factor model could accurately reflect this diversity of exposures.

A multi-factor model is given by:

$$r_S = E[r_S] + \beta_{S1} F_1 + \beta_{S2} F_2 + \dots + \beta_{SN} F_N + \varepsilon_S, \quad (5.XX)$$

where F_1, F_2, \dots, F_N are the N systematic factors chosen, and $\beta_{S1}, \beta_{S2}, \dots, \beta_{SN}$ are the N coefficients measuring the stock's exposure to each of the systematic factors. The one-factor model is just a special case in which $N=1$, and the market model sets $F_1=M$. The N systematic factors may be things such as the percent change in industrial production, the percent change in expected inflation, the term premium in interest rates, and the default premium on corporate bonds. In equation (5.XX), the factor variables F_1, F_2, \dots, F_N must be constructed so that each has mean zero. Note that equation (5.XX) assured this for the one factor model. The model assumes that the $N+1$ random variables, r_S and F_1, F_2, \dots, F_N , are joint normally distributed. Under certain important assumptions about the chosen factors and how they were constructed, each factor beta coefficient is the covariance of the stock's excess return with the respective factor, scaled by the factor variance

$$\beta_{Si} = \frac{\text{Cov}(r_S, F_i)}{\text{Var}(F_i)}. \quad (5.XX)$$

These factor betas are sometimes called factor sensitivities or factor loadings.

The variance of the stock's excess return, σ_S , can be decomposed into the variance of the N systematic factors plus the variance of the non-systematic factor:

$$\sigma_S^2 = \beta_{S1}^2 \sigma_{F1}^2 + \beta_{S2}^2 \sigma_{F2}^2 + \dots + \beta_{SN}^2 \sigma_{FN}^2 + \sigma_{\varepsilon_S}^2. \quad (5.XX)$$

Once again, the estimate for σ_{ε} is the the standard deviation of the residuals from the regression.

In order to determine the factor sensitivities, we can estimate the regression equation:

$$r_S = \alpha_S + \beta_{S1} G_1 + \beta_{S2} G_2 + \dots + \beta_{SN} G_N + \varepsilon_S, \quad (5.XX)$$

where G_1, G_2, \dots, G_N are the same N systematic factors chosen earlier, but without the requirement that each have mean zero. Where specifying F_1, F_2, \dots, F_N , requires some estimation of the mean

values for the data employed, specifying G_1, G_2, \dots, G_N , does not. Nor does the equation require that we specify the expected return on the stock. If all we want is to determine betas, then the form of equation (5.XX) is sufficient.

When analyzing the returns to many diverse stocks and attempting to identify the foundational risk factors for the economy as a whole, it is common to use macroeconomic variables like those mentioned earlier: the percent change in industrial production, the percent change in expected inflation, the term premium in interest rates, and the default premium on corporate bonds. However, in more targeted analyses of specific groups of companies, where one is not necessarily searching for the economy-wide underlying foundational risk factors, it is sometimes useful to determine each company's factor loading on certain industry-specific risk factors. This type of model was what Tufano (1998) employed to estimate the sensitivity of gold mining stocks to fluctuations in the price of gold:

$$r_{i,t} = \alpha_i + \beta_{i,g} r_{g,t} + \beta_{i,m} r_{m,t} + \varepsilon_{i,t}, \quad (5.XX)$$

where $r_{i,t}$ is the return on the stock of gold mining firm i in period t , $r_{g,t}$ is the total return on an investment in gold in period t , and $r_{m,t}$ is the return on a composite stock market index in period t . The regression estimates the two betas, which tells us the exposure of gold company stocks to changes in the price of gold and to fluctuations in the stock market. Tufano (1998) fit this regression equation for data on 48 gold mining companies in the United States and Canada during the early 1990s. On average, he found that the typical gold beta, $\beta_{i,g}$, was approximately 2, so that an increase in the return on gold of 1% yields an increase in the stock return of 2%.³ The average gold beta varied significantly over time, ranging between 1 and 4 over a 5-year period. There was also significant variation in the gold betas of the different firms. These are the empirically observed gold betas discussed earlier in this chapter.

Another version of the multi-factor model is known as the industry index model. The risky in a company's stock is divided into three pieces: the piece attributable to general marketwide movements, the piece attributable to the relative performance of the company's industry or some chosen set of comparable companies, and the idiosyncratic risk of the particular stock:

$$r_{i,t} = \alpha_i + \beta_{i,m} r_{m,t} + \beta_{i,comp} r_{comp,t} + \varepsilon_{i,t}, \quad (5.XX)$$

³ Peter Tufano, 1998, The Determinants of Stock Price Exposure: Financial Engineering and the Gold Mining Industry, *Journal of Finance* 53:1015-1052.

where $\beta_{i,comp}$ is the sensitivity of the stock return of company i with $r_{comp,t}$, the return on the portfolio of comparable companies in period t . This model can be used to decompose a company's stock performance over time into the returns attributable to each of the factors and the residual that reflects the company's relative performance above and beyond that due to market-wide or industry-wide factors. This decomposition may be helpful in performance evaluation and designing a compensation system that rewards management's contribution to relative performance and not performance due to market-wide or industry-wide factors.

[insert material on Ralston case]

5.3 The Many Determinants of Risk

Risk is not constant. If we ask a question like “How risky is an investment in oil?”, the right answer is, “It depends.” It depends upon whether the investment is going to be held for 1 day, for 1 year, or for 10 years. It depends upon whether oil prices are high or low. It depends upon the current state of the market. It depends upon how we are making the investment, and whether the instruments used are liquid or illiquid. And it depends upon the social context and whether the institutions have been constructed so that this risk is well defined or so that this risk must include a number of political and systemic factors.

Horizon

In finance, we are often evaluating risk factors over different time horizons, and we expect the variance to depend upon the time horizon. For example, most risk factors will exhibit a greater variance the further out in time one looks—i.e. the variance of the variable's realization two years from now is greater than the variable's realization one week from now. Therefore, it is necessary to be precise about the horizon over which one is measuring the variance or standard deviation.

Level

Some simple models of risk have the same volatility no matter the level of the variable. Look back at the simple model of stock returns at the top of this chapter. The volatility of the return is specified by a single parameter. No matter what is the initial stock price, S_0 , the volatility of the return is the same. Similarly, when we estimate a factor model, we are implicitly assuming

that the exposure is constant at all levels of the underlying variables in the data set that we employ, and so properly measured by a single coefficient, β .

However, the volatility may vary with the level of the parameters. We have already seen this in the case of the gold mining companies discussed above, where Equation (5.XX) can be used to understand how the beta might vary with the parameters of the problem, including with the level of the gold price.

Time Dependent Risk

The volatility of the stock market and of many other market variables often vary over time. Figure 5.XX is a graph of the standard deviation of monthly returns on a value weighted stock market index from 1926 to 1997. It is taken from a paper by Campbell, Lettau, Malkiel and Xu (2001). As one can see in the graph, the volatility has varied tremendously year-by-year.

Figure 5.XX is a graph of the average R^2 from estimating a market model for each of a set of stocks. The R^2 measures the portion of a stock's volatility that is accounted for by the market, i.e. systematic risk. The remaining portion, $1-R^2$, is the portion of a stock's volatility that is idiosyncratic risk. It appears from the graph as if the portion of risk that is systematic has been declining over time, while the portion that is idiosyncratic has been increasing.

Credit Risk

Many business exposures involve contractual relationships between two parties. Often times the contract specifies how an underlying exposure is to be shared. A contractual exposure can be deceptively straightforward. The formula may, on first glance, seem clear and unambiguous, but few things in business—as in life, in general—are exactly as they seem on the surface. While one clause of a contract may specify a clear and unambiguous formula, there may be another clause somewhere in the contract that specifies the circumstances under which the formula does not apply. Contracts generally include a force majeure clause that excuses each party from performance under certain classes of exceptional events. And no matter how firm the contract language, there may be circumstances under which the party obligated to make a payment according to the formula simply cannot do so. There's no getting blood out of a stone.

The danger that one party may not be able to perform on a contract when it owes money to the other party is called credit risk. Credit risk needs to be acknowledged and factored into management assessments of its true exposure from a deal. The key point that needs to be

evaluated is how the counterparty's ability to pay correlates with its obligation to pay. If, when it owes money, it is always able to pay, then there is little credit risk. However, if, when it owes money, it is least likely to be able to pay, then there is great credit risk.

Once credit risk is acknowledged and evaluated, a company can choose how it will manage it. It may be useful to limit the size of its position with any single customer or supplier. It may change the terms of contracts so as to protect itself. Or it may find it necessary to accept the credit risk, but knowing that it is exposed, manage its operations and finances accordingly.

Market Risk & Force Majeure

Sometimes a party to a contract may play a role in bringing about an act of force majeure! In late summer 2010, a long lasting heat wave across Russia caused a major drop in supplies of Russian wheat sending prices skyrocketing. Glencore, a major international trader in commodities, had a number of contracts obligating it to deliver Russian wheat at fixed prices. The increasing cost of purchasing the wheat was exposing Glencore to significant losses. Glencore's deputy chief executive officer of its International Grain unit, Nikolai Demyanov, urged the Russian government to ban further exports of wheat. A government ban would count as a force majeure and excuse Glencore from having to take losses on the delivery contracts. The next day the Russian government imposed the ban.

Liquidity Risk

A key market risk that companies need to give attention to is liquidity risk. When we speak of the value of a corporate asset—e.g., an airline's aircraft, or a consumer product firm's trademarks or brands, a paper goods manufacturer's forest acreage, or a biotech company's patents—we are always implicitly making an assumption about liquidity. We make a distinction between the price one can get if one takes the normal amount of time to market the asset effectively to all potential buyers, on the one hand, and the 'firesale price' one may be forced to accept if one has to sell the asset hurriedly. Company's must take care to avoid getting themselves into positions in which they must dispose of assets in a firesale, thus losing value.

The liquidity of a market can vary markedly. Traders on the stock market know that there are certain hours of the day when the market is thick with both buyers and sellers, so that the market is at its most liquid, and other hours of the day when there is little liquidity. In addition to this very regular variability in liquidity, the liquidity varies across years and certain historical episodes. The same is true for the markets for corporate assets. Most importantly, a company needs to watch the liquidity of the markets it uses to finance its operations—the money market, the bond market and the market in its own stock. If these markets suddenly seize up and there is

no liquidity, then the company may be unable to access fresh capital for some period of time. So long as it doesn't need fresh capital, there is no problem. But if some of its major debt is coming due, or other exigencies require it to raise new money, and the market is closed, it can face a liquidity crisis.

Constellation Energy's 2008 Liquidity Crisis

Constellation Energy is a major utility based in Baltimore, Maryland. It owns several electric generating stations, including four nuclear reactors. It also has wholesale and retail electric power business, and owns the local electric and gas utility for Baltimore. It also ran a major trading operation focused on electricity, natural gas and coal. In mid-2008, a combination of factors, including its recent swift expansion of its coal business, record commodity prices and the revelation of errors in its risk management data systems, left the firm with the sudden need for \$X billion in capital. This discovery was announced to the market in the company's 2Q earnings release on July 30. The company immediately set about trying to line up the necessary emergency lines of credit. Unfortunately, raising this amount of capital takes some time. Before the company could complete a transaction, the Lehman Brothers insolvency was announced on September 18, precipitating the global financial crisis. Credit markets worldwide seized up. It was impossible for Constellation to close on its new credit lines. The company was forced instead to sell a major stake in the firm at less than ½ the price that the shares had been trading at after the 2Q release. On July 30 Constellation's stock price had been \$XX/share. By September xX, when the new sale of shares closed, the stock price had fallen to \$XX/share.

5.4 Model Risk

Past Performance is No Guarantee of Future Results

Historical data plays a critical role in our estimates of risk and exposure. After all, experience is a great teacher. But we need to be cautious when we extrapolate from past history into the future. The world economy changes quickly. This has important implications for the relevance of the data to our decisions about the future.

First, we need to be recognize how limited is our historical data. Many of the asset markets and lines of business that concern us are relatively new. The sample of observations we have from the past is usually quite short. It may not be representative.

+ black swans

Oddly, we sometimes actually handicap ourselves by focusing on an even smaller subset of data than is truly available. We do this in an attempt to exploit more advanced techniques that are only recently available and that rely on features of the data only included for more recent samples. For example, Shiller & mortgages, Hamilton & oil.

We need to be careful about consuming a given model and a given data set in isolation, as if God's truth lies fully expressed within one data series. A dash of wisdom developed from looking at the bigger picture over a longer historical window of time may be helpful in understanding what can be made out of any individual data set.

Second, the future is also going to be different from the past. The past is only prologue. The economy is changing, and that means that the structure between variables evidenced in past data may only have modest relevance to the structure they will evidence in the future.

Here be dragons

Psychological and Cognitive Biases in Assessing Risk

stuff

The Dynamics of Bubbles

stuff