

## Part 5. Thinking About Risk

Most financial decisions involve payments made in the future, and these payments are usually uncertain. For example, no one knows for sure what the price of a stock will be next year, or how a new project will turn out. Because people do not like risk, they demand a higher return from their investments that are uncertain. In this section, we will review the basic statistics involved in measuring uncertainty and discuss how individuals and businesses can reduce the amount of risk they face.


### Distributions of Random Variables

One way to describe a random variable is by its distribution. The distribution of a random variable is a summary of the possible outcomes along with the probabilities of getting each of the outcomes. To take a simple example, say that we roll two dice and add the values together. The possible outcomes range from 2 (if both dice come up 1) to 12 (if both dice come up 6). Assuming that the dice are not loaded, so that each side has the same probability of occurring, we can calculate the probability of getting each outcome.

Table 5.1.

Outcome	Probability
2	1/36
3	2/36
4	3/36
5	4/36
6	5/36
7	6/36
8	5/36
9	4/36
10	3/36
11	2/36
12	1/36

Another way of representing the distribution is with a histogram. A histogram shows graphically the probability of each outcome. The advantage of this is that it is easy to see the “shape” of the distribution; where the middle is and how spread out the outcomes are.



If there are a lot of possible outcomes, writing down the distribution can be bothersome, so it would be helpful if we could summarize the distribution with a few numbers. One important characteristic of a distribution is the “middle” or “average” outcome. A good way to measure this is with what is called the “expected value”. Mathematically, the expected value is a weighted average of outcomes, with the weights equaling the probabilities of the outcomes. The formula is:

$$\text{Expected Value} = E(X) = \sum_{i=1}^n P_i X_i$$

where  $X_i$  is outcome  $i$ ,  $P_i$  is the probability of outcome  $i$ , and there are  $n$  different outcomes. The expected value works like a regular average, but takes into account the fact that not all outcomes are equally likely. For example, take a distribution where the outcome will be 100 with a probability of 90% and 0 with a probability of 10%. A simple average of the outcomes would give us 50 (  $50 = (100+0)/2$  ). However, we are much more likely to get 100 than 0, so the outcome of 100 should get extra importance. The expected value is calculated as  $100*0.9 + 0*0.1 = 90$ , a result closer to 100. In the case of the two dice, the expected value is in the middle of the distribution and equal to 7 (as an exercise, see if you get this value).

It is important to remember that even though it is called an expected value, you may get something entirely different when you take a single draw from the distribution. Indeed, in the dice example, you are more likely to get something besides 7 than to get exactly 7. However, if you observe this random value again and again (i.e. keep rolling the dice) the expected value tells you what number you would expect to get if you calculated the average of the outcomes over time.

We also need a measure of how dispersed the possible outcomes are. We want to capture two things: how far away an outcome is from the expected value (farther away means more uncertainty), and how likely it is to get that outcome. One measure for this is the standard deviation. The formula is,

$$\text{Standard Deviation} = \sigma = \sqrt{\sum_{i=1}^n P_i (X_i - E(X))^2}$$

To see how the standard deviation works, we will do the calculation for our dice example (see Table 5.2). In the first step (column 3) we calculate the deviation of each outcome from the expected value ( $X_i - E(X)$ ). For example, the first outcome is 2, which is  $-5$  from the expected value of 7. Some outcomes in our example are larger than the expected value (a positive deviation) while others are smaller (a negative deviation). If we averaged them together as they are in column 3, the positive and negative values would cancel out, which is not what we want. To make them all positive, we square the deviations (column 4). This also has the effect of making outcomes far from the middle relatively more important in our measure. Next, we want to determine the weighted average of the squared deviations. We multiply each squared deviation by its probability (column 5) and then add them up. This sum is called the variance of the distribution. However, because we squared all the numbers, the variance is in terms of “units squared”. To correct for this, we take the square root of the variance to arrive at the standard deviation.

Table 5.2.

(1) Outcome	(2) Probability	(3) Deviation	(4) Squared Deviations	(5) (2) x (4)
2	1/36	-5	25	0.6944
3	2/36	-4	16	0.8889
4	3/36	-3	9	0.75
5	4/36	-2	4	0.4444
6	5/36	-1	1	0.1389
7	6/36	0	0	0
8	5/36	1	1	0.1389
9	4/36	2	4	0.4444
10	3/36	3	9	0.75
11	2/36	4	16	0.8888
12	1/36	5	25	0.6944

Variance = 5.83

Standard Deviation = 2.42

The standard deviation is a common measure of the uncertainty or risk of a random variable. The reason for this is that it captures much of our intuition about risk. The farther the possible outcomes are from the expected value, the greater the standard deviation. Also, the more likely it is to get an outcome away from the expected value, the greater the standard deviation. These characteristics, combined with the mathematical convenience of the standard deviation, make it a very popular measure of risk in finance.

For a practice problem, we will look at situation where stock returns depend on the state of the economy. When you are investing in stock, the return you get over the next year will be uncertain. It could be higher if the economy does well, it could be lower if the economy does

poorly. To be more concrete, we start with a situation with three possible situations in the economy and three possible outcomes for the stock.

Economy	Probability	Stock Return
Recession	15%	-4%
Steady Growth	65%	8%
Rapid Growth	20%	10%

To summarize this distribution, you would calculate the expected return and the standard deviation of the returns (the answers are 6.6% and 4.52 respectively). In this way you could compare the average return and the uncertainty of the return of this stock with other stocks to determine which one is the best to invest in (we will do this in a later section of the course).

#### *Estimating the expected value and the standard deviation using historical data*

In the previous two examples, we knew the outcomes and their probabilities with certainty. This is rarely the case in actual financial decision-making. For example, with any stock there are a great number of possible outcomes, and it is difficult to say what exactly the probability of any given outcome will be. However, even if we do not know the entire distribution of outcomes, or if it would be too inconvenient to write out, we can still summarize the distribution based on estimates of the expected value and standard deviation using historical data.

From a statistical point of view, we are using past values of the random variable to estimate the expected value and standard deviation of the distribution that produced the data. To estimate the expected value, we calculate the mean of our data,

$$\text{Mean} = \frac{1}{T} \sum_{i=1}^T X_i$$

We estimate the standard deviation of our distribution by calculating the standard deviation of our data (treating our data as a “sample”),

$$\text{Standard Deviation} = \sqrt{\frac{1}{T-1} \sum_{i=1}^T (X_i - \bar{X})^2}$$

These formulas are very similar to the previous formulas, except that we no longer assume that we know what the probabilities are. Instead, we let how often things happened in history be our measure of how likely they are to happen in general.

*A comment on some of the assumptions we have been making*

The view of risk developed here makes some strong assumptions. In particular, we are assuming that we can adequately summarize a distribution by the expected value and standard deviation, and that we can use historical data to make reliable estimates of their values. Are these reasonable assumptions? It depends. The calculations we have done here are a very good starting point and for some purposes may be enough. Other times, more sophisticated calculations may be needed. And in some situations, even our best estimates may not be that reliable.

**Example:** We are trying to determine the financial viability of some project. Because this project has a number of unique aspects, it is difficult to estimate an accurate standard deviation from historical data. On the other hand, from our experience in the business we can place some rough probabilities on particular events. For example, we might expect that the probability our construction costs exceed our estimates is no more than 20%. From these rough probabilities we can determine notions of the expected return to the project and the range of outcome. Even though we cannot come up with precise measures of uncertainty, the act of trying to quantify the risks we face will be very useful in the decision-making process. How we incorporate this uncertainty into our decision is a topic for the sections on capital budgeting.

**Example:** We are deciding whether to invest in stocks, bonds or some mix of both. Our decision will depend in part on how much we can expect to earn with our investments and how much risk we will face. Here, looking at historical data may be of some help. We will see that over history, stocks have offered a higher expected return, but also have a higher standard deviation of returns, in other words, more risk. Our investment decision will depend on how we value the extra return we get from stocks compared with how we feel about taking on the extra risk.

Of course, just because something happened in the past doesn't mean that it is going to happen in the future. However, we will see in a later section of the course that there are good reasons for stocks to have more risk and a higher expected return. We can draw the conclusion that it is likely that this pattern will continue. By looking at history we can get a sense of how big this effect is. While you wouldn't expect that future expected returns and standard deviations to be *exactly* the same as the values from history, they are likely to be close enough that this is very useful information.

**Risk Aversion**

When determining the effect of risk on our financial decisions we need to know both the amount of risk we face and what our attitude towards risk is. Generally it seems that people do not like risk. There are exceptions to this, Las Vegas and the lottery being two examples, but in financial circumstances, people seem to be willing to pay to avoid risk. Confusingly, two different terms are used to describe this behavior. The first, risk aversion, refers to how much an individual dislikes risk. Risk tolerance, refers to how much individuals are willing to put up with risk. In other words, risk aversion is just the opposite of risk tolerance. People in finance will use both terms when discussing attitudes towards risk.

As a minor point, it needs to be mentioned that the standard deviation, our measure of risk, only describes risk well for symmetrical distributions. Many risks we face do not have symmetrical outcomes (in statistical terms they are skewed) and this may affect how individuals treat these risks. Take the risk of your house burning down. In any given year, the chance of that happening is very small, less than 1%. There is an extremely large probability of no loss, but a very small probability of a huge loss. People generally do not like risks with the possibility (if ever so small) of extreme losses (although if the risk gets small enough, sometimes they will ignore it altogether). This leads to a very active insurance market to protect individuals from the costs of these events. On the other hand, the lottery represents the opposite kind of distribution. Chances are that you won't win anything and you lose the cost of your ticket. There is a very (very) small probability of winning, but if you do win, you win big. The expected value of a lottery is negative (that is why governments use them to raise revenue) and so risk averse individuals logically shouldn't participate. However, people seem to like risks with no chance of a big loss and at least some chance at a big gain. Because of this, lotteries are very popular despite the fact that they are not good investments.

### *The risk-return tradeoff*

Risk averse individuals, when offered a choice between an investment with an expected return of 5% and a standard deviation of 7%, and an investment with the same expected return and a standard deviation of 9%, will choose the first investment. If everyone in the market was risk averse, no one would want to hold the second investment. In order to get people to hold it, it would have to offer a higher return; enough to compensate the investors for having to hold the extra risk. The relationship that investments with more risk will tend to offer higher returns is called the **risk-return tradeoff**. The risk-return tradeoff is one of the central ideas of finance.

### *Measuring risk aversion*

Measuring the amount of risk is relatively easy – we already know how to calculate standard deviations. And in a later section, we will look at other ways of measuring risk. Measuring risk aversion is a different matter. For this, we need to know how an individual feels about risk, and this can be difficult to do. Most investment books will have simple quizzes you can use to develop a sense of your own risk averseness. However, these are just guides; they don't come up with numbers as precise as a standard deviation. However, there is a different way to approach this problem.

We need to distinguish between the risk averseness of a specific individual and the risk averseness of the market as a whole. The market risk averseness is just the average risk averseness of all the individuals who participate in financial markets, by borrowing, lending or investing. This average risk averseness will show up as the difference in expected returns across securities of different risks. Say that a security with no risk offers a return of 3%. Another security with a standard deviation of 2% offers an expected return of 4%. Because the second security has more risk, it has to offer a risk premium to get people to hold it. Because we can observe actual numbers in the marketplace, we can quantify this relationship. Increasing the standard deviation by 2% requires an additional 1% of expected return. This is a measure of the market risk averseness. This is sometimes also called the “market price” of risk since people can buy and sell securities of different risk at this premium. (In practice, it's not quite this simple since there are other measures of risk.)

The market price of risk shows up in both capital-budgeting and investment decisions. For a company that is trying to raise money in the market by issuing stocks and bonds, it tells the

company the return it will have to pay the investors given the amount of risk in the company's securities. For an individual investor, the market risk-return tradeoff describes the options they face. If they are more risk averse than the market (that is, than the average investor) they should hold assets with lower risk than average. If they are less risk averse they should hold a riskier selection of assets to take advantage of their higher expected returns.

### Strategies for reducing risk - I

While you can try to avoid risk in life, it is not always the best thing to do. Investing in stocks, while risky, offer the promise of higher returns. Driving a car involves the risk of accidents, but in most places it is necessary to do. However, while you may be in situations where you cannot avoid risk completely, there are things you can do to reduce risk.

#### *Diversification*

One of the important principles of investing is that you should not “put all your eggs in one basket”. In financial language, you should diversify your investments. If you only invested in one asset, say in shares of Microsoft, your entire investment outcome depends on the fortunes of that company. If Microsoft goes on to dominate other software markets you could get a great return. However, if it loses its near monopoly on operating systems and business software you could find that much of your wealth is gone. Having a diversified portfolio means spreading your wealth across a number of different companies besides Microsoft. If Microsoft does poorly, it is likely that some other company will do well. While this does not completely eliminate risk, the chance that all your investments will lose money is small. Of course, the reverse is also true. Even if Microsoft does well, it is likely that some other company will not do as well, bringing down your average return. Diversification cuts down on risk by reducing both the upside and the downside.

For example, take two stocks that have different responses to changes in the economy. When the economy is doing well (the state of the economy is “good”) Stock A has a low return and Stock B has a high return. When the economy is doing poorly, the returns are reversed. If each state of the economy is equally likely, each stock has an expected return of 5% with a standard deviation of 5%.

	Economy is Good	Economy is Bad	Expected Return	Standard Deviation
Stock A	0%	10%	5%	5%
Stock B	10%	0%	5%	5%
50/50	5%	5%	5%	0

If we construct a new portfolio that is 50% Stock A and 50% Stock B, the expected return is unchanged, but we have eliminated all uncertainty, since the higher-than-average return of one stock will cancel the lower-than-average return of the other stock in each state of the economy. Given the choice of a portfolio of All Stock A, All Stock B, and a 50/50 mix, we would prefer the mix since it offers the same expected return with lower risk.

This is the benefit of diversification. By mixing together stocks with different patterns (distributions) of returns, we can put together a portfolio with the same return, but less risk, than the stocks taken individually.

### *Pooling Risk and Insurance*

If you drive a car, you know that there are risks. While the chance of an accident is very small, the cost if it happens can be very large. However, if you could get together with 10,000 of your closest friends, you could all agree to contribute some money to a fund that will reimburse anyone who has an accident. By doing this, you are changing the distribution of the risk you face. Without this agreement, you face a large probability of no cost and a very small probability of a very large cost. With this agreement, you face a guarantee of a small loss (the money you contribute) but you also avoid the probability of a large loss. Since individuals are risk averse, they are willing to make this trade. Of course, insurance companies are the financial intermediaries that handle this transaction.

### **Correlations and Covariances**

Once we start looking at more than one security, we need to take into account how the risks of each security are related. In the auto insurance example, we implicitly made the assumption that the risks were unrelated, and so if you had a large enough group of cars, some will be in accidents while most will not. The chance that every driver has an accident at the same time is unlikely. In other situations, this may not be a good assumption. When the economy is doing well, most stocks will do well - they will have higher-than-average returns *at the same time*. We have two measures of how closely the outcomes of two random variables are related. The **covariance** is similar to the variance in that it measures an average deviation from the expected value. The formula is,

$$Covariance(X, Y) = \sum_{i=1}^n P_i (X_i - E(X))(Y_i - E(Y))$$

If the two variables tend to move in the same direction, that is, *A* tends to be above its average when *B* is above its average, the variables have a positive covariance. If they tend to move in opposite directions, *A* is below its average when *B* is above, they have a negative covariance.

Covariance captures two different aspects of the relationship. It measures the closeness of the connection between the two variables. It also measures how large the changes in the variables are. The larger the standard deviations of *A* and *B* are, the bigger the covariance between *A* and *B*. Sometimes we only want to know the direction of the connection. We can remove the magnitude of the variation by dividing by the standard deviations of *X* and *Y*. This measure is called the **correlation** between *X* and *Y*:

$$Correlation(X, Y) = \frac{Covariance(X, Y)}{\sigma_X \sigma_Y}$$



Correlation runs on a scale from -1 to 1. A value of 1 is perfectly positive correlation, the variables always move in the same direction. A value of -1 is perfectly negative correlation, the variables always move in opposite directions. A value of 0 means the variables are uncorrelated, there is no connection in their movements.

### **Strategies for reducing risk - II**

We can use information about correlation to improve our risk-reduction strategies.

#### *Hedging*

Hedging is when you simultaneously take two opposite financial positions, with the idea that a poor performance by one will be made up for by a good result by the other. For example, you are a technology company that has issued a number of stock options. If your stock price increases, you will have to pay your employees the value of the option. To hedge this risk, you can buy a contract in the market that pays you if the stock price increases. The costs and benefits of the stock-price increase cancel out and your risk is gone. Of course, this is not free. Buying the option contracts cost money. However, paying a small amount to reduce risk may be a worthwhile tradeoff.

#### *Modern Portfolio Theory*

Correlation can also play an important role in determining which stocks you want to hold. This is because correlation tells you how much diversification benefit you get from an investment. Say that you hold stock in the Alpha Corporation. You decide to become more diversified by holding stock in Beta Industries. If the returns to Alpha and Beta are perfectly correlated, that is, they always move in the same direction, there will not be any diversification benefit. Every time the returns to Alpha fall, the returns to Beta will also fall, and you haven't reduced your risk. If the returns to the two companies are independent, then there will be some diversification. Some of the time when the returns to Alpha fall, the returns to Beta will be higher than average. Finally, if the returns to the two companies are perfectly negatively correlated, then whenever one company's return is lower than average, the return to the other company will be higher than average. This will have a substantial effect on reducing risk and will provide maximum diversification.

We will examine "Modern Portfolio Theory" and the role of correlation in a later part of the course.