

Uncertainty, Default, and Risk

Risk-neutral Promised vs. Expected Returns; and Debt vs. Equity

You are now entering the world of uncertainty and abandoning the pleasant idea that you have perfect foresight. We shall still pretend, however, that you live in a perfect market with no taxes, no transaction costs, no differences of opinion, and infinitely many investors and firms. But you will learn in this chapter that the presence of uncertainty adds quite a bit of additional complexity and realism.

Net present value still rules supreme, but you will now have to face the sad fact that it is no longer easy to use. It is *not* the NPV concept that is difficult. Instead, it is the *inputs* that are difficult — the expected cash flows and appropriate costs of capital that you now have to guesstimate.

In a world of uncertainty, there are scenarios in which you will get more cash than you expected and scenarios in which you will get less. The single most important insight under uncertainty is that you must always draw a sharp distinction between *promised* (or *quoted* or *stated*) returns and *expected* returns. Because firms can default on payments or go bankrupt in the future, expected returns are lower than promised returns.

After some necessary statistical background, this chapter will cover two important finance topics: First, you must learn how much lenders should charge borrowers if there is the possibility of default. Second, you must learn how to work with the two building blocks of financing — namely, debt and equity.

6.1 An Introduction to Statistics

Statistics has the reputation of being the most painful of the foundation sciences for finance — but you absolutely need to understand it to describe an uncertain future. Yes, it can be a difficult subject, but if you have ever placed a bet in the past, chances are that you already have a good intuitive grasp of what you need. In fact, I had already sneaked the term “expected” into previous chapters, even though it is only now that this book covers what this precisely means.

Random Variables and Expected Values

The most important statistical concept is the expected value, which is the probability-weighted average of all possible outcomes. It is very similar to a mean or average. The difference is that the latter two names are used if you work with *past* outcomes, while the expected value applies if you work with *future* outcomes. For example, say you toss a coin, which can come up either heads or tails with equal probability. You receive \$1 if the coin comes up heads and \$2 if the coin comes up tails. Because you know that there is a 50% chance of \$1 and a 50% chance of \$2, the expected value of each coin toss is \$1.50. If you repeated this infinitely often, and if you recorded the series of realization[realizations] (actual outcomes), the mean would converge to exactly \$1.50. Of course, in any one throw, \$1.50 can never come up — the expected value does not need to be a possible realization of a single coin toss.

Begin Important

The expected value is just the mean (a fancy word for average) if you could repeat an experiment (the random draws) infinitely often. **End Important**

Economists also like to call the expected value an ex-ante (“before the fact”) value to distinguish it from the realized ex-post value (“after the fact”).

To make it easier to work with uncertainty, statisticians have invented the concept of the random variable. It is a variable whose outcome has not yet been determined. In the coin toss example, you can define a random variable named c (for “coin toss outcome”) that takes the value \$1 with 50% probability and the value \$2 with 50% probability. (Random variables are often written with tildes over them, such as \tilde{c} , but we will dispense with this formality in our book.) The expected value of c is \$1.50. To denote the expected value, we use the notation E . In this bet,

(Omitted eq)

After the coin has been tossed, the actual outcome c could, for example, be $c = \$2$. After the toss, this c is no longer a random variable. Also, if you are certain about the future outcome — for example as in cases in which there is only one possible outcome — then the actual realization and the expected value are the same. The random variable is then the same as an ordinary nonrandom variable. Is the expected outcome of the coin toss a random variable? No: You know the expected outcome is \$1.50 even before the toss of the coin. The expected value is known; the uncertain outcome is not. The expected value is an ordinary nonrandom variable; the possible outcome is a random variable. Is the outcome of the coin throw *after* it has come down heads a random variable? No: It is an actual outcome and you know what it is (heads), so it is no longer a random variable.

(Omitted fig)

A random variable is defined by the probability distribution of its possible outcomes. The coin throw distribution is simple: the value \$1 with 50% probability and the value \$2 with 50% probability. This is sometimes graphed in a histogram, which is a graph that has the possible outcomes on the X-axis and the frequency (or probability) on the Y-axis. Figure ?? shows the histogram for the coin throw.

Begin Important

The best way to think of a random variable is as a placeholder for a histogram with a name. **End Important**

One final reminder: In this chapter, we are eliminating our certainty assumption. But we are *not* (yet) eliminating our perfect market assumption. The assumption of no-disagreement means that we all must agree on the probabilities of all possible outcomes. An example of an imperfect market would be if you believed that there was a 60% probability of an outcome of \$1, and I believed there was a 50% probability of \$1. (And even after the actual draw, we will still not know who was ultimately right!)

► Fair Bets

A fair bet is a bet that costs its expected value. If repeated infinitely often, both the person offering the bet and the person taking the bet would expect to end up breaking even. For example, call D your payoff based on the following structure:

- There is a $1/4$ chance that you will be paid \$2;
- a $1/4$ chance that you will be paid \$10;
- and a $2/4$ chance that you will be paid \$8.

You can simulate this payoff structure by drawing a card from a complete deck. If it is ♣, you get a value V of \$2; if it is ♦, you get \$10, and if it is ♥ or ♠, you get \$8. What would be a fair price for this card bet? The uncertain payoff is a random variable, which we called D . First, you must determine $E(D)$. It is

(Omitted eq)

If you repeat this bet a zillion times, you would expect to earn \$7 zillion. On average, each bet would earn \$7, although some sampling variation in actual trial draws would make this a little more or a little less. If it costs \$7 to buy each single bet, it would be fair.

Generally, the procedure to compute expected values is always the same: Multiply each outcome by its probability and add up all these products. For a random variable X ,

(Omitted eq)

This is the formula that you used above,

(Omitted eq)

Note that the formula is general. It works even with outcomes that are impossible. You would just assign probabilities of zero to them.

Begin Important

You must understand the following:

1. The difference between an ordinary variable and a random variable
2. The difference between a realization and an expectation
3. How to compute an expected value, given probabilities and outcomes
4. What a fair bet is

End Important

(Omitted solvenow)

Variance and Standard Deviation

In finance, we often need to measure the (average) reward that you expect to receive from making an investment. Ordinarily, we use the expected value of the investment as our measure of reward. We also often need to measure a second characteristic of an investment, its risk. Thus, we also need summary measures of how spread out the possible outcomes are. These two concepts will play starring roles in the next few chapters, where you will explore them in great detail. For now, if you are curious, think of risk as a measure of the variability of outcomes around your expected mean outcome. The most common measure of risk is the standard deviation, which takes the square root of the sum of squared deviations from the mean — a mouthful. Let's just do it once for our card-draw problem. Recall our formula: the expectation is probability-weighted values. First, work out each squared deviation from the mean:

The first outcome is \$2. The mean is \$7, so the deviation from the mean is $\$2 - \$7 = -\$5$. You need the squared deviation from the mean, which is $(-\$5)^2 = +\25 . The units are strange — dollars squared — and impossible to interpret intuitively. Don't even try.

The second outcome is \$10, so the deviation from the mean is $\$10 - \$7 = +\$3$. You need the squared deviation from the mean, which is $(+\$3)^2 = +\9 .

The third outcome is \$8, so the deviation from the mean is $\$8 - \$7 = +\$1$. You need the squared deviation from the mean, which is $(\$1)^2 = +\1 .

Together, in one table, this is

Probability	1/4	1/4	2/4
Outcome	\$2	\$10	\$8
Net of the Mean (\$7)	-\$5	+\$3	\$1
Squared Net of Mean	\$25	\$9	\$1

Now compute the expected value of these squared deviations, which is called the variance:

(Omitted eq)

The standard deviation (SD, sometimes Std.Dev., Sdv, or Sdev) is therefore

(Omitted eq)

There you have it — our mouthful: The standard deviation is the square root of the average squared deviation from the mean. Unlike the variance, the standard deviation has sensible units. Together, the mean and standard deviation allow you to characterize your bet. It is common phrasing, though a bit loose, to state that you expect to earn \$7 (the expected value) from a single card draw, plus or minus \$3 (the standard deviation).

(Omitted solvenow)

Risk Neutrality (and Preview of Risk Aversion)

Fortunately, the expected value is all you need to learn about statistics for this chapter. This is because we are assuming — only for learning purposes — that everyone is risk-neutral. Essentially, this means that investors are willing to write or take any fair bet. For example, if you are risk-neutral, you would be indifferent between getting \$1 for sure and getting either \$0 or \$2, each with 50% probability. And you would be indifferent between earning 10% from a risk-free bond and earning either 0% or 20%, again with fifty-fifty probability, from a risky bond. You have no preference between investments with equal expected values, no matter how safe or uncertain these investments may be.

If, instead, you are risk-averse, you would not want to invest in the more risky alternative if both the risky and safe alternatives offered the same expected rate of return. And you would prefer the safe \$1 to the unsafe \$0 or \$2 investment. You would prefer a 10% risk-free bond to the unsafe corporate bond that would pay either 0% or 20%. In this case, if I wanted to sell you a risky project or a risky bond, I would have to offer you a higher expected rate of return as risk compensation. I might have to pay you, say, 5 cents to get you to be willing to accept the project that pays off \$0 or \$2 if you can instead earn \$1 elsewhere. And, I would have to lower the price of my corporate bond so that it offers you an expected rate of return higher than 10% — say, “either 1% or 21%” instead of “either 0% or 20%.”

It is true that if you are risk-averse, you should not accept fair bets. (You can think of this as the definition of risk aversion.) But would you really worry about a bet for either +\$1 or −\$1? Probably not. For small bets, you are probably close to risk-neutral — I may not have to pay you even 1 cent extra to induce you to take this bet. But what about a bet for plus or minus \$100? Or for plus or minus \$10,000? My guess is that you would be fairly reluctant to accept the latter bet without getting extra compensation for risk bearing. If you are like most investors, you are more risk-averse when the bet is larger. To take the plus or minus \$10,000 bet, I would probably have to offer you several hundred dollars extra.

However, your own personal risk aversion is not what matters in financial markets. Instead, the financial markets set investments prices in line with the market’s aggregate risk aversion. The reason is risk sharing. For example, if you could share the \$10,000 bet with 10,000 other students in your class, your own part of the bet would be only plus or minus \$1. And some of your colleagues may be willing to accept even more risk for relatively less extra risk compensation — they may have healthier bank accounts or wealthier parents. Therefore, when you can lay bets across many investors, the effective risk aversion of the group will be lower than that of any of its members. And this is exactly how financial markets work: Their aggregate risk absorption capabilities are considerably higher than those of their individual investors. In effect, the financial markets are less risk-averse than individual investors.

You will study risk aversion in the next chapters. In this chapter, we will focus on pricing under risk neutrality. But, as always, all tools you learn in this simpler scenario will remain applicable in the more complex scenario

in which investors are risk-averse. Moreover, in the real world, the differences between promised and expected returns that are discussed in this chapter are often more important (in terms of value) than the extra compensation for risk aversion that is ignored in this chapter.

(Omitted solvenow)

6.2 Interest Rates and Credit Risk (Default Risk)

Most loans in the real world are not risk-free, because the borrower may not fully pay back what was promised. We will assume that there is one exception, which is that U.S. Treasuries are risk-free loans in nominal terms. In principle, the United States can always tax more and print more dollars to satisfy all promised bond payments. Therefore, it is reasonable to assume the United States cannot default. (Intelligent people can disagree. Washington politics is so dysfunctional that the U.S. may actually default not for lack of dollars, but by choice.) So, how do you compute appropriate expected rates of return for risky bonds?

(Omitted anecdote)

Risk-Neutral Investors Demand Higher Promised Rates

Now, put yourself into the position of a banker. Assume that a 1-year Treasury note offers a safe annual rate of return of 10%. Your immediate problem is that you are contemplating making a 1-year loan of \$1 million to me. What interest rate should you charge me on the loan? If you are 100% certain that I will fully pay the agreed-upon amount, you can just charge me 10%. You earn just as much from me as from the Treasury note. Both will pay back \$1,100,000.

However, in the real world, there are few borrowers for whom you can be 100% certain that they will fully repay a loan. For example, assume you believe there is only a 50% chance that I will pay back the principal plus interest. (If I do have assets in excess of my liabilities, I will be called solvent and presumably pay it back). There is also a 50% chance that I will default (fail to pay all that I have promised). This is often informally called bankruptcy. In this case, I may only be able to pay back \$750,000 — all that I have left. If, as the bank, you were to charge me a 10% interest rate, your expected payout would be

(Omitted eq)

Your *expected* return would not be \$1,100,000, but only \$925,000. Your *expected* rate of return would not be +10%, but only $\$925,000 / \$1,000,000 - 1 = -7.5\%$. Extending such a loan would not be — pardon the pun — in your best interest: You can do better by investing your \$1,000,000 into government Treasury notes.

(Omitted anecdote)

You should conclude that you must demand a higher interest rate from risky borrowers as a banker, even if you just want to “break even” (i.e., expect to earn the same \$1,100,000 that you could earn in Treasury notes). If you solve

(Omitted eq)

for the desired promised repayment, you will find that you must ask me for \$1,450,000. The promised interest rate is therefore $\$1,450,000/\$1,000,000 - 1 = 45\%$. Of this 45%, 10% is the time premium that the Treasury pays. Therefore, you can call the remaining 35% the default premium — the difference between the promised rate and the expected rate that you, the lender, would have to demand just to break even. It is very important that you realize that the default premium is not extra compensation for your taking on more risk, say, relative to holding Treasuries. You don't receive any such extra compensation in a risk-neutral world. The default premium just fills the gap between the expected return and the promised return.

You rarely observe expected rates of return directly. Newspaper and financial documents almost always provide only the promised interest rate, which is therefore also called the quoted interest rate or the stated interest rate. When you read a published yield-to-maturity, it is also usually only a promised rate, not an expected rate — that is, the published yield is an internal rate of return that is calculated from promised payments, not from expected payments. Of course, you should never make capital budgeting decisions based on promised IRRs. You almost always want to use an expected IRR (YTM). But you usually have easy access only to the promised rate, not the expected rate. On Wall Street, the default premium is often called the credit premium, and default risk is often called credit risk.

(Omitted solvenow)

A More Elaborate Example With Probability Ranges

This distinction between expected and promised rates is so important that it is worthwhile to work another more involved example. Assume again that I ask you to lend me money. You believe that I will pay you what I promise with 98% probability; that I will repay half of what I borrowed with 1% probability; and that I will repay nothing with 1% probability. I want to borrow \$200 from you, which you could alternatively invest into a government bond promising \$210 (i.e., a 5% interest rate). What interest rate would you ask of me?

If you ask me for a 5% interest rate, next year (time 1), your \$200 investment today (time 0) will produce the following:

Payoff (C_1)	Rate of Return (r)	Frequency (Prob)
\$210	+5.0%	98% of the time
\$100	−50.0%	1% of the time
\$0	−100.0%	1% of the time

Therefore, your expected payoff is

(Omitted eq)

Your expected return of \$206.80 is less than the \$210 that the government promises. Put differently, if I *promise* you a rate of return of 5%,

(Omitted eq)

then your expected rate of return would be only

(Omitted eq)

This is less than the 5% interest rate that Uncle Sam promises — and surely delivers.

You need to determine how much I have to promise you just to break even. You want to expect to end up with the same \$210 that you could receive from Uncle Sam. The expected loan payoff is the probability-weighted average payoff. You want this payoff to be not \$206.80 but the \$210 that you can earn if you invest your \$200 into government bonds. You need to solve for an amount x that you receive if I have money,

(Omitted eq)

The solution is that if I promise you $x \approx \$213.27$, you will expect to earn the same 5% interest rate that you can earn in Treasury notes. This \$213.27 for a cash investment of \$200 is a *promised* interest rate of

(Omitted eq)

Such a promise provides the following:

Payoff (C_1)	Rate of Return (r)	Frequency (Prob)
\$213.27	+6.63%	98% of the time
\$100.00	−50.00%	1% of the time
\$0.00	−100.00%	1% of the time

This comes to an *expected* interest rate of

(Omitted eq)

(Omitted solvenow)

Deconstructing Quoted Rates of Return — Term and Default Premiums

The difference of 1.63% between the promised (or quoted) interest rate of 6.63% and the expected interest rate of 5% is the default premium — it is the extra interest rate that is caused by the default risk. Of course, you only receive this 6.63% *if* everything goes perfectly. In our perfect market with risk-neutral investors,

(Omitted eq)

Begin Important

Except for 100%-safe bonds, the promised (or quoted) rate of return is higher than the expected rate of return.

Never confuse the promised rate with the (lower) expected rate. If you only remember one thing from this book, this should be it! **End Important**

Begin Important

Financial securities and information providers rarely, if ever, provide information about expected rates of return. They almost always provide only quoted rates of return, which are higher than expected rates of return. **End Important**

On average, under risk-neutrality, the expected rate of return is the expected term premium plus the expected default premium. Because the *expected* default premium is zero *on average*,

(Omitted eq)

If you want to work this out, you can compute the expected realized default premium as follows: You will receive $6.63\% - 5\% = 1.63\%$ in 98% of all cases; $-50\% - 5\% = -55\%$ in 1% of all cases (note that you lose the term premium); and $-100\% - 5\% = -105\%$ in the remaining 1% of all cases (i.e., you lose not only all your money, but also the term premium). Therefore,

(Omitted eq)

In addition to the 5% term premium and the 1.63% default premium, in the real world, there are also other premiums that we have not yet covered:

Risk premiums that compensate you with (even) *higher* expected rates of return for your willingness to take on risk. They will be the subject of Chapter .

Imperfect market premiums (e.g., liquidity premiums) that compensate you for future difficulties in finding buyers for your bonds. They will be the subject of Chapter .

In normal times, for corporate bonds, the risk (and often also the imperfect market) premiums are typically considerably lower than the term and default premiums.

(Omitted solvenow)

Credit Ratings and Default Rates

To make it easier for lenders to judge the probability of default, a number of data vendors for credit ratings have appeared. For individuals, lenders can obtain credit ratings from Experian, Transunion, and Equifax — you should request a free credit report for yourself from the [Federal Trade Commission](#) if you have never seen one. For small companies, Dun & Bradstreet provides similar credit scores. For larger corporations, the [U.S. Securities Exchange Commission](#) recognizes nine agencies: A.M. Best, DBRS [Morningstar], Egan-Jones, Fitch, HR Mexico, Japan JCR, Kroll [Duff&Phelps], Moody's, and S&P. The biggest are Moody's and Standard&Poor's (S&P). For a fee, these agencies

rate the probability that the issuer's bonds will default. This fee depends on a number of factors, such as the identity of the issuer, the desired detail in the agencies' investigations and descriptions, and the features of the bond (e.g., a bond that will pay off within one year is usually less likely to default before maturity than a bond that will pay off in thirty years; thus, the former is easier to grade).

(Omitted tbl)

The credit rating agencies ultimately do not provide a whole set of default probabilities (e.g., 1% chance of 100% loss, 1.2% chance of 99% loss, etc.), but just an overall rating grade. Table ?? shows the categories for Moody's and Standard & Poor's. It is then up to the lender to translate the rating into an appropriate compensation for default risk. The top rating grades are called investment grade, while the bottom grades are called speculative grade (or junk grade).

Ratings have limited usefulness:

1. They do not consider common risk, wherein many bonds would default at the same time. This will be an important concept in the next few chapters. See, most bond buyers should care more about the (small) risk of all their bonds blowing up at the same time and care less about one small individual bond in their many-bond portfolios defaulting. But common risk assessments are *not* what rating agencies provide.
2. Unlike most other financial market experts, rating agencies are not liable for their ratings or perspectives even if they deliberately deceive investors. (The [2010 Dodd-Frank Act](#) repealed this exemption, but the SEC has granted indefinite “no-action” relief for most ratings.)
3. The strangest aspect, however, is how the rating agencies are paid. They collect fees for rating securities by the investment banks — how critical would you be of their bond products in this case? Not surprisingly, although they need to maintain some independence and reputation, the agencies have also often been complicit when it comes to being manipulated — some would even call it bribed. A good part of the Great Recession (of 2008-9), falls squarely on the shoulders of the rating agencies, which earned billions providing optimistic ratings for issues explicitly engineered by investment banks to have high ratings. And although they have taken some steps to improve the situation, the basic conflicts of interest are still there. When public attention will have moved on to another “issue of the day,” chances are that the ratings will return back to business as usual.
4. Ratings change over time: after their issue, bonds move up or down with probabilities of about 3-10% each per year.

Nevertheless, despite all their flaws, they are a useful source of information for potential bond buyers.

► Empirical Evidence on Default

Ed Altman and his coauthors (from New York University) collected corporate bond statistics from 1971 to 2019. Figure ?? gives you a sketch of how likely default was. (It is standard to define default as missing at least one coupon payment. It is *not* complete ultimate non-payment.) The average default rate was about 3.5% per year — but the left plot shows that it was much higher in recessions, where defaults typically shot up to over 10%. For

example, in the Great Recession, about 11% of bonds failed to pay. By 2010, the worst seemed to have passed. In retrospect, the financial crisis of the Great Recession ended up “not so bad” for most public corporations. (Lucky!) The right plot shows that corporate bonds originally rated A or better rarely defaulted, even 10 years after issue. However, about half of all CCC junk bonds would fail to pay at least one coupon within the first five years of issue.1]ch:corpclaimsSenior and Junior Bonds Interestingly, despite Covid, 2021 had the lowest default rate on record.

(Omitted fig)

Moody's monthly *Default Report* lists recovery rates after bonds default. In 2005, they reported that from 1982 on, recovery rates were about 60% for senior secured bonds (the bonds promised first repayment, with specific collateral), 45% for senior unsecured bonds (without specific collateral), and 30% for junior bonds (to be paid only after senior bonds). The typical recovery in a default was about 30-40 cents on the dollar, with 25 cents in recessions and 50 cents in booms. Low-rated bonds would pay less. These numbers seem to have remained similar in the 2010-2016 time period, too, but there is a lot of idiosyncratic variation across individual bonds.

► Bond Contract Option Features

Before I show you how bonds are priced, I need to let you know that bonds in the real world differ from one another not just in credit risk. Most bonds have additional contract features that may also influence their quoted rates of return. For example, many corporate bonds allow the issuer to repay the loan early. (The same applies to almost all domestic mortgages.) If the interest rates in the future fall, this can be a good thing for the borrower and a bad thing for the lender. The borrower would pay off the loan and borrow more cheaply elsewhere. If the interest rates in the future rise, the borrower gets to pay just the earlier low interest rate. For example, assume that the interest rate is 10% today and you are lending me \$90,909 in exchange for my promise to pay you \$100,000 next year. One second after you extend the loan, one of two scenarios can happen:

1. The interest may fall to 5%. I would then simply repay your \$90,909 loan and refinance at this lower interest rate elsewhere — thereby saving myself a lot of money..
2. The interest rate may rise to 15%. In this case, I keep my \$100,000 promise to pay next year — I received \$90,909 for a loan that should have given me only $\$100,000/1.15 \approx \$86,957$.

This would not be a good arrangement for you — unless you are appropriately compensated for giving me this option to prepay. Borrowers who want the right to repay without penalty therefore have to pay higher interest rates when they issue such bonds. Virtually all mortgage bonds in the United States allow prepayment and therefore carry higher interest rates than they would if they did not have a prepayment feature. Loosely speaking, you can classify these contract option features as default premiums, too, because on average they tend not to add or subtract from your expected rate of return. Sometimes they increase the amount paid, and sometimes they decrease the amount paid by the lender — just as a bond from a solvent borrower would pay more to the lender and an insolvent borrower.

(Omitted solvenow)

Differences in Quoted Bond Returns in 2021

So how do real-world credit risk and reflecting bond credit ratings translate into differences in promised (quoted) bond yields? Table ?? lists the borrowing rates of various issuers on December 2021.

(Omitted tbl)

The data looks broadly consistent with the theory — bonds that have higher default risk have to offer higher promised rates of return, especially over longer horizons where they have more opportunities to get into trouble. Bonds with higher (better) credit ratings can find lenders at lower quoted interest rates (higher bond prices).

Do lenders who extend loans to riskier creditors end up with the same average rates of return? This would be the case in a perfect market in which lenders and borrowers are risk-neutral. The evidence suggests that this is not exactly true, but it is also not as far from reality as one might think. Most of the investment-grade bond spread above the Treasury that you see in Table ?? simply made corporate borrowers and lenders come out about even. That is, the *expected* rates of return were much more similar to one another than the *promised* rates of return in the table suggest. If you put a gun to my head and asked me for an opinion, I would guess that about 80% of the promised spreads over Treasury are credit risk (sans some tax considerations); about 10% are due to measurement or contract features (e.g., the timing of coupons or some option contingencies); and only about 10% are extra compensation that the creditors earn on average above and beyond what they would earn in term-equivalent U.S. Treasuries. (And, of course, all interest income is taxable.)

You already saw in Figure ?? that actual defaults by reasonably large corporations were cyclical and increased only briefly in the Great Recession. Similarly but even more extreme, the interest rates that lower-rated corporations paid (which are *promised* rates that reflect expectations of non-payment over the full life of the bond) were cyclical and spiked in 2009. Figure ?? plots the historical yields of the 20-year Treasury bond, of Moody's-rated Aaa and Baa investment-grade bond portfolios, and non-investment grade (high-yield) bond portfolios. The typical investment-grade bond promised about 100-200 basis points above the Treasury, while the typical junk bond promised about 200-600 basis points above the Treasury. (Junk is relative — non-publicly traded corporations usually pay even higher rates on non-collateralized obligations.) Any non-investment grade corporation that had to borrow during the Great Recession of 2007-9 was in serious trouble. Promised interest rates were spiking at 23% per annum, more than 2,000 bps higher than Treasuries, as investors were fleeing to safety. After 2010, things got back to normal (albeit with short spooked outliers in 2011 and 2020 for high-yield corporate bonds).

We could discuss risk and reward for many other types of credit. Credit risks are not always similar. Mortgage investors in Arizona can face different interest rates than London banks. Incidentally, the latter is called the London Interbank Offered Rate (Libor). Libor plays an important role, because it is also a commonly used benchmark for about \$350 trillion of derivatives. Even your mortgage interest rate may be tied to Libor. However, the industry is now changing over to the Secured Overnight Financing Rate ([SOFR](#)).

(Omitted anecdote)

Nerdnote: Joseph Stiglitz won a Nobel Prize in economics for the insight that high interest rates can themselves induce firms to take much riskier projects — so much that no interest rate can end up inducing lenders to provide capital to a firm. This is called credit rationing.

(Omitted fig)

Figure ?? shows the promised yields on about 7,500 individual corporate bonds based on credit-ratings on a particular day. (The numbers are not exact, because bonds can have many different contract clauses, and I am just graphing them by maturity and credit rating.) There are very few AAA bonds, so the fitted line here is least reliable. Similarly, going beyond 10-20 years, the fitted lines become less and less reliable. The figure shows clearly how longer-term corporate bonds tend to have not just higher but *increasingly* higher *promised* spreads over Treasuries.

(Omitted fig)

Credit Default Swaps (CDS)

The financial world is always changing and innovating. The components of bond returns described above used to be primarily a conceptual curiosity — firms would borrow money from their lenders, paying one interest rate that just contained all premiums. But then, with the introduction of credit default swap[credit default swaps] (often abbreviated credit swap[credit swaps] or CDS), some premium components suddenly became themselves tradable.

Here is an example of a CDS: A large pension fund that owned a \$15 million bond issued by *Hospital Corporation of America* (HCA) may have wanted to purchase a \$10 million credit swap from a hedge fund that in turn wanted to bet that HCA would *not* go bankrupt. Upon HCA bankruptcy, the hedge fund would owe the pension fund \$10 million. The *Wall Street Journal* reported that this CDS contract cost about \$130,000, but rose to over \$400,000 one month later, because of a potential buyout deal that increased the risk of future default. And, in this case, purchasing the CDS in June would have been a lucky deal for the pension fund and an unlucky deal for the hedge fund — because HCA did indeed go bankrupt.

The best way to think of such credit swaps is as insurance contracts, in which the swap sellers (e.g., hedge funds) are the insurance providers. The seller of the swap thus takes on the credit risk from the buyer — just like an insurance company takes on some risk from the insured party — in exchange for a premium payment upfront. The insurer then pays out in case of a credit event (e.g., a failed payment or bankruptcy) — typically for one particular bond within a given number of years. The payment itself can be formula-determined, or it can be a guarantee by the CDS seller to buy the bond at a predetermined price. One way of thinking of the upfront cost (the \$130,000 that increased to \$400,000) is that it contains the bond's default premium.

Credit swaps make it easier for different funds to hold different premiums of a bond. In our example, the pension fund decided to earn primarily the term premium component of HCA's bonds, divesting itself of the credit risk. The hedge fund took over the credit premium. It decided to speculate that HCA would not go bankrupt, and it could do so without having to take a large cash position in HCA's bonds. Of course, other hedge funds could have speculated with CDS that HCA would go bankrupt.

Credit swaps are typically traded in lots of \$5 million and last for 5 years (but 3 to 10 years is not unusual, either). This market used to be over-the-counter (OTC) — that is, negotiated one-to-one between two parties. However, after the financial crisis, the market for standardized CDS has largely migrated to clearing houses (like ICE Clear Credit and the London Common Exchange). These houses guarantee the contract in case one of the two parties

goes bankrupt. This market is also very big: in 2020, there was about [\\$10 trillion outstanding in single-name swaps](#), i.e., swaps that relate to a specific company. However, credit swaps on indexes are even more important and popular than single-name credit swaps — and, frankly, index CDS are better insurance for most investors than single-name swaps.

In ancient times — in the first edition of this book in 2007 — I wrote that no one knew any longer who was really holding most of the credit risk in the economy. I gave as an example the German bank IKB, which had collapsed to everyone's surprise, because it had owned too many financial securities that were tied to U.S. mortgages. I was either clairvoyant or just lucky. This is still the case. Financial markets are global and there is no global regulator with a full perspective of what large trading desks are doing simultaneously in many different venues. Are their bets hedged or doubled up?

CDS were one of the dominoes that caused the financial crisis of 2008. In 2008, before clearing houses, the CDS market collapsed temporarily altogether. Investors no longer wanted to trust even solid corporations and banks, simply because they did not know what their actual exposures were — and neither did the Federal Reserve or the U.S. Treasury department. (It's as if everyone suddenly awoke to the fact that they thought someone else was driving the bus — and no one was.) But everyone had just learned that an important insurer, American International Group (AIG), had a financial arm that had sold too many CDS and the traders had booked them as outright profit. This had worked for a long while — until it no longer did in 2008. The U.S. Treasury (i.e., we, the taxpayers) had to bail out AIG, because it feared that too many bond buyers were relying on AIG's insurance, and would themselves have to default if the AIG insurance became worthless.

In the believe-it-or-not category, it is still possible to manipulate CDS. There have been cases in which some enterprising creditors purchased a lot of CDS positions and then pressured the underlying company to declare bankruptcy in order to trigger the CDS payout. Also in the same category, it is up to a committee of major banks that can decide whether a bond is really in default or not when the creditor offers an exchange. These banks may themselves be highly conflicted, invested on one or the other side of the CDS.

However, CDSs are not intrinsically evil. Like most financial instruments, they can be used to reduce or increase risk. The social problem, even today, is that (1) traders have incentives to speculate too much with them, because their risk is hard to measure, and banks and their shareholders intrinsically like risk (but in cases of a crisis, it is taxpayers that will be on the hook again!); and (2) no one really knows what is going on in parts of this (and many other) financial markets. If large players speculate too much with them (and pay non-recoverable bonuses based on estimated profitability) and then fail, we the people may have no choice but to bail them out *again*.

6.3 Uncertainty in Capital Budgeting

Let's now return to the basic tasks of capital budgeting: selecting projects under uncertainty. Your task is to compute present values with imperfect knowledge about future outcomes. The principal tool in this task will be the payoff table (or state table), which assigns probabilities to the project value in each possible future-value-relevant scenario. For example, the value of a factory producing computer hard disk drives may depend on computer sales (say, low, medium, or high), whether hard disks have become obsolete (yes or no), whether the economy is in a recession or expansion, and what the oil price (a major transportation cost factor) turns out to be. It is the manager's task to

create the appropriate “state” table, which specifies what variables and scenarios are most value-relevant and how the business will perform in each of them. Clearly, it is not an easy task even to understand what the key factors are, much less to determine the probabilities under which these factors will take on one or another value. Assessing how your own project will respond to them is an even harder task — but it is an inevitable one. If you want to understand the value of your project, you must understand what your project’s key value drivers are and how your project will respond to these value drivers. Fortunately, for many projects, it is usually not necessary to describe all possible outcomes in the most minute detail — just a dozen or so scenarios are often enough to cover the most important possibilities.

Present Value with Outcome-Contingent Payoff Tables

We begin with the hypothetical purchase of a building for which the future value is uncertain. Next year, this investment will be worth either \$60 thousand (with 1 / 4 probability) or \$100 thousand (with 3 / 4 probability). (In case you are worried that real firms last longer than one year, you can think of these values as themselves reflecting further future outcomes for the firm.) To help you remember the two possible states, let’s just call the bad outcome “Hurricane” and the good outcome “Sun.” Don’t take this naming literally — it’s just a memory aid.

► **The Building’s Expected Value**

If you own the full building, your payoff table, omitting thousands henceforth, is as follows:

Event	Probability	Value
Hurricane	1 / 4	\$60
Sun	3 / 4	\$100
⇒ Expected Future Value		\$90

The expected future building value of \$90 (thousand) was computed as

(Omitted eq)

This is *not* yet discounted. It is only your expectation of the future outcome.

Now assume that the appropriate expected rate of return for a project of type “building” with this type of riskiness and with 1-year maturity is 20%. (This 20% discount rate is provided by demand and supply in the financial markets, and it is assumed to be known by you, the manager.) Your goal is to determine the present value — the appropriate price — for the building *today*.

There are two methods to arrive at the present value of the building — and they are almost identical to what you have done earlier. You only need to replace the known value with the expected value, and the known future rate of return with an expected rate of return. The first PV method is to compute the expected value of the building next period and to discount it at the cost of capital, here 20%:

(Omitted eq)

The second method is to compute the discounted state-contingent value of the building, and then take expected values. To do this, augment the earlier table:

Event	Probability	Value	Discount Factor	⇒	PV
Hurricane	1/4	\$60	1/(1+20%)	⇒	\$50
Sun	3/4	\$100	1/(1+20%)	⇒	\$83.33

If it “hurricanes,” the present value is \$50. If the sun shines, the present value is \$83.33. Thus, the expected value of the building can also be computed as

(Omitted eq)

Both methods lead to the same result: You can either first compute the expected value of the investment next year ($1/4 \cdot \$60 + 3/4 \cdot \$100 = \$90$) and then discount this expected value of \$90 to \$75; or you can first discount all possible future outcomes (\$60 to \$50, and \$100 to \$83.33) and then compute the expected value of the discounted values ($1/4 \cdot \$50 + 3/4 \cdot \$83.33 = \$75$.)

Begin Important

Under uncertainty, in the NPV formula,

- known future cash flows are replaced by expected cash flows, and
- known appropriate rates of return are replaced by appropriate expected rates of return.

You can first do the discounting and then take expectations, or vice-versa. The order does not matter. **End Important**

► The State-Contingent Rates of Return

What would the rates of return be in the two states, and what would your overall expected rate of return be? If you have bought the building for \$75 and it will be sunny, your actual rate of return will be

(Omitted eq)

If it's rainy, your rate of return will be

(Omitted eq)

Therefore, your expected rate of return is

(Omitted eq)

The probability state-weighted rates of return add up to the expected overall rate of return. This is as it should be: After all, you derived the proper price of the building today using a 20% expected rate of return.

(Omitted solvenow)

6.4 Splitting Uncertain Project Payoffs into Debt and Equity

State payoff tables help you understand cash flow rights. They make it easy to understand one of the most important concepts in finance: the difference between a loan (also called debt or leverage) and levered ownership (also called levered equity or simply equity or stock). Almost all companies and projects are financed with both debt and levered equity. You already know in principle what debt is. Levered equity is simply what accrues to the business owner *after* the debt is paid off. We leave it to later chapters to make a distinction between financial debt and other obligations — for example, tax obligations — and to cover the control rights that flow from securities — for example, how debt can force borrowers to pay up and how equity can replace poorly performing managers.

You probably already have an intuitive understanding about the distinction between debt and equity. If you own a house with a mortgage, you really own the house only after you have made all debt payments. If you have student loans, you *yourself* are the levered owner of your future income stream. That is, you get to consume “your” residual income only *after* your liabilities (including your nonfinancial debt) are paid back. But what will the levered owner and the lender get if the company’s projects fail, if the house collapses, or if your career takes a turn toward Rikers Island? What is the appropriate compensation for the lender and the levered owner? The split of net present value streams into loans (debt) and levered equity lies at the heart of finance.

You now know how to compute the present value of state-contingent payoffs — your building paid off differently in two different states of the world. Thus, your building was a state-contingent claim — its payoff depended on the outcome. But it is just one of many possible state-contingent claims. Another might promise to pay \$1 if the sun shines and \$25 if the hurricane hits. Using payoff tables, you can work out the value of *any* state-contingent claim and, in particular, the value of our two most important state-contingent claims, debt and equity.

The Loan

Let’s assume you want to finance the building purchase of \$75 (thousands) with a mortgage of \$70. In effect, the single project “building” is being turned into two different projects, each of which can be owned by a different party. The first project is “Mortgage Lending.” The second project is “Residual Building Ownership,” that is, ownership of the building but bundled with the obligation to repay the mortgage. The “Residual Building Ownership” investor will not receive a dime until *after* the debt has been satisfied. As already explained, such residual ownership is called levered equity, or just equity (or even stock) in the building. This avoids calling it “what’s-left-over-after-the-loans-have-been-paid-off.”

What sort of interest rate would the creditor demand? To answer this question, you need to know what will happen if the building were to be worth less than the mortgage promise. Let’s say that the value of the building will be \$60 next year in the hurricaney state of the world. We are assuming that the owner could walk away, and the creditor could repossess the building but not any of the borrower’s other assets. Such a mortgage loan is called a no-recourse loan. There is no recourse other than taking possession of the asset itself. This arrangement is called limited liability. The building owner cannot lose more than the money that he originally puts in. Limited liability is

the mainstay of many financial securities: For example, if you buy stock in a company in the stock market, you cannot be held liable for more than your investment, regardless of how badly the company performs.

To compute the present value for the project “Mortgage Lending,” return to the problem of setting an appropriate interest rate, given credit risk (from Section ??)sect:creditriskCredit Risk). Start with the following payoff table:

Event	Prob	Value	Discount Factor
Hurricane	1/4	\$60	1/1.20
Sun	3/4	Promised	1/1.20

(Omitted anecdote)

The creditor receives the property worth \$60 if it hurricanes, or the full promised amount (to be determined) if the sun shines. To break even, the creditor must solve for the payoff to be received if the sun shines in exchange for lending \$70 today. This is the “quoted” or “promised” payoff:

(Omitted eq)

You can solve this for the necessary promise, which is

(Omitted eq)

in repayment, paid by the borrower only if the sun shines.

With this promised payoff of \$92 (if the sun shines), the lender’s rate of return will be the promised rate of return:

(Omitted eq)

The lender would not provide the mortgage at any lower promised interest rate. If it hurricanes, the owner walks away, and the lender’s rate of return will be

(Omitted eq)

Therefore, the lender’s *expected* rate of return is

(Omitted eq)

The stated rate of return is 31.4% (and it is not an exorbitant rate!), but the expected rate of return is 20%. After all, in our risk-neutral perfect market, anyone investing for one year expects to earn an expected rate of return of 20%.

Nerdnote: Special liability and tax rules apply to private residences. Mortgages can have limited liability (“non recourse”) or unlimited liability (“full recourse”). The latter can also have further nasty tax consequences, where a capital loss in the home can create a large ordinary income tax obligation, adding insult to injury. (If interested, google for “cancellation-of-debt income.”) Moreover, as a home owner, you can deduct interest only on the first \$1 million in mortgage; and capital losses on the home do not create a tax credit, but large capital gains can create a tax obligation.

In financial economics, we usually assume that we are dealing with sophisticated borrowers and lenders. However, this ignores questions of ethics regarding limited liability. Is it fair to lenders to repudiate loans if it is legal to do so? Is it different if the lenders were friends and family? Did they consider your limited liability when they extended the loan?

The Levered Equity

As the residual building owner, what rate of return would you expect as proper compensation? You already know the building is worth \$75 today. Thus, after the loan of \$70, you need to pay in \$5 — presumably from your personal savings. Of course, you must compensate your lender: To contribute the \$70 to the building purchase today, you must promise to pay the lender \$92 next year. If it hurricanes, the lender will confiscate your house, and all your invested personal savings will be lost. However, if the sun shines, the building will be worth \$100 minus the promised \$92, or \$8. Your payoff table as the levered equity building owner is as follows:

Event	Prob	Value	Discount Factor
Hurricane	1 / 4	\$0	1 / 1.20
Sun	3 / 4	\$8	1 / 1.20

This allows you to determine that the *expected* future levered building ownership payoff is $1 / 4 \cdot \$0 + 3 / 4 \cdot \$8 = \$6$. Therefore, the present value of levered building ownership is

(Omitted eq)

Your rates of return are

(Omitted eq)

The expected rate of return of levered equity ownership, that is, the building with the bundled mortgage obligation, is thus still

(Omitted eq)

Reflections On The Example: Payoff Tables

Payoff tables are fundamental tools to help you think about projects and financial claims. Admittedly, they can sometimes be tedious, especially if there are many different possible states. (There may even be infinitely many states, as in a bell-shaped, normally-distributed project outcome — but you can usually approximate even the most continuous and complex outcomes fairly well with no more than a dozen discrete possible outcomes.)

Table ?? shows how elegant such a table can be. It describes everything you need in a very concise manner: the state-contingent payoffs, expected payoffs, net present value, and expected rates of return for your house scenario. Because owning the mortgage and the levered equity is the same as owning the full building, the last two columns must add up to the values in the “Building Value” column. You could decide to be any kind of investor: a creditor (bank) who is loaning money in exchange for promised payment; a levered building owner who is taking a “piece left over after a loan”; or an unlevered building owner who is investing money into an unlevered project (i.e., the whole piece). All three investments are just state-contingent claims.

(Omitted tbl)

Begin Important

Whenever possible, in the presence of uncertainty, write down a payoff table to describe the probabilities and payoffs in each important possible “state of the world.” **End Important**

(Omitted solvenow)

Reflections on The Example: Debt and Equity Risk

We have only briefly mentioned risk. It was just not necessary to illustrate the main insights. In a risk-neutral world, all that matters is the expected rate of return, not the uncertainty about what you will receive. Of course, you can assess the risk even in our risk-neutral world where risk earns no extra compensation (a risk premium). So, which investment is riskiest: full ownership, loan ownership, or levered ownership?

(Omitted fig)

Figure ?? plots the histograms of the rates of return for each of the three types of investments. The equity loses everything (–100%) with a $1/4$ probability but earns 60% with $3/4$ probability. The debt loses about 14.3% with $1/4$ probability and gains 31.4% with $3/4$ probability. The full ownership loses about –20% with $1/4$ and gains 33.3% with $3/4$ probabilities. As the visuals show, the loan is least risky, followed by the full ownership, followed by the levered ownership. There is an interesting intuition here. By taking the mortgage, the medium-risk project “building” has been split into one more risky project (“levered building”) and one less risky project (“mortgage”). The combined “full building ownership” project therefore has an average risk.

Of course, regardless of leverage, all investment projects in our risk-neutral world expect to earn a 20% rate of return. After all, 20% is the universal time premium for investing money in our hypothetical world. (The default premium is a component only of promised interest rates, not of expected interest rates; see Section ??.) By assuming that investors are risk-neutral, we have assumed that the risk premium is zero. Investors are willing to take any investment that offers an expected rate of return of 20%, regardless of risk. (If investors were risk-averse, debt would offer a lower expected rate of return than the project, which would offer a lower expected rate of return than equity.)

Although our example was a little sterile because we assumed away risk preferences, it is nevertheless very useful. Almost all projects in the real world are financed with loans extended by one party and levered ownership held by another party. Understanding debt and equity is as important to corporations as it is to building owners. After all, stocks in corporations are basically levered ownership claims that provide money only *after* the corporation has paid back its liabilities. The building example has given you the skills to compute state-contingent, promised, and expected payoffs, as well as state-contingent, promised, and expected rates of return. These are the necessary tools to work with debt, equity, or any other state-contingent claim. And really, all that will happen later when we introduce risk aversion is that you will add a few extra basis points of required compensation — more to equity (the riskiest claim), fewer to the project (the medium-risk claim), and still fewer to debt (the safest claim).

(Omitted solvenow)

What “Leverage” Really Means — Financial and Operational Leverage

I have already mentioned that debt is often called leverage and equity is called “levered equity.” Let me now explain why. A lever is a mechanical device that can amplify effects. Just like mechanical leverage lets you take on a bigger weight in exchange for a proportionately greater range of mechanical movement, so financial leverage lets you take on a bigger investment in exchange for a proportionately greater range of financial returns. In finance, a lever is something that allows a smaller equity investment to still control the firm and be more exposed to the underlying firm’s gain or loss than unlevered ownership. That is, with leverage, a small change in the underlying project value translates into a larger change in value for levered equity, both up and down. You have seen this leverage mechanism in our house example above, and specifically in Figure ?? . Ordinary ownership would have cost you \$75. But with leverage, you could take control of the house with cash of only \$5. In addition, it also meant that if the sun had shone, you would have earned $(\$8 - \$5)/\$5 = 60\%$, not just $(\$100 - \$75)/\$75 = 33.3\%$; but if it had hurricaned, you would have earned -100% (lost *everything*), not just $(\$60 - \$75)/\$75 = -20\%$. Leverage had amplified your stake.

(Omitted tbl)

Financial debt is a lever, but it is not the only one. Leverage can also be calculated using all corporate liabilities (which may include, e.g., accounts payable and pension obligations). More importantly, because leverage is a general concept rather than an accounting term, you should think of it in even broader terms. The idea of leverage is always that a smaller equity investment can control the firm and is more sensitive to firm value changes. Table ?? illustrates some different types of levers. In this table, you can pay \$475 for machine and labor, and receive either \$200 or \$1,000 in product revenues, plus \$150 as resale value for the machine. In the first line of the table in the exhibit, you can see that in the bad state, you lose 26%; and in the good state, you earn 142%. The second line shows that financial leverage can magnify these rates of return into -100% or $+540\%$. But instead of taking on financial debt, you could also lease the machine, which costs you \$250 in leasing fees (with no residual ownership of the machine at the end), and pay for labor of \$75. In this case, you have effectively levered up, increasing your risk to -38% and $+208\%$ but without taking on any financial leverage. It is the lease that has now become your leverage! And you can also combine real and financial leverage. Finally, there can even be differences in the degree to which the production technologies themselves are levered. The final example in the fourth line shows a different method of production, which is intrinsically more levered (and which happens to be roughly equally good in this example).

Working with More Than Two Possible Outcomes

In the real world, possible outcomes can often range from 0 to infinity. Can you use the same method if you have more than two scenarios? For example, assume that the building could be worth \$60, \$70, \$80, \$90, or \$100 with equal probability (for an expected value of \$80) and that the appropriate expected interest rate is 20%. It follows that the building has a PV of $\$80/1.20 \approx \67 . If a loan promised to pay \$60 at time 1, how much would the lender expect to receive? The full \$60, of course, because the building is always worth at least this much:

(Omitted eq)

But if a loan promised \$61, how much would it expect to receive then? It would expect \$60 for sure, plus the extra “marginal” \$1 with 80% probability (because there is an 80% chance that \$61 is covered; only if the outcome is \$60, which happens 20% of the time, would it not receive the full \$61). Thus, for the \$61 loan promise, it would expect to receive \$60.80. In fact, it would expect only 80 cents for each dollar promised between \$60 and \$70. So, if a loan promised x between \$60 and \$70, it would expect to receive

(Omitted eq)

If a loan promised \$71, how much would it expect to receive? It would expect \$60 for sure, plus \$8 for the next promised \$10, plus 60 cents on the dollar for anything above \$70, i.e., \$68.60,

(Omitted eq)

(Omitted fig)

Figure ?? plots these expected payoffs as a function of the promised payoffs. With this figure, mortgage valuation becomes easy. For example, how much would the loan have to promise to provide \$50 today? The expected payoff would have to be $(1 + 20\%) \cdot \$50 = \60 . This is on the linear segment, so you would have to promise \$60. Of course, you cannot offer an expected payoff of more than \$80, so forget about borrowing more than $\$80/1.2 \approx \66.67 today.

The same approach would also work when possible value outcomes are “normally distributed” (i.e., following a bell-curve). The math is a more complex, but the method remains the same.

(Omitted solvenow)

An Error: Discounting Promised Cash Flows with the Promised Cost of Capital

A common mistake is the attempt to avoid the need to estimate expected values by discounting promised cash flows with promised discount rates. After all, both numbers reflect default risk. The two default issues might cancel out one another, and you might end up with the correct inference. *Or they might not cancel out, in which case you will end up with an incorrect decision!*

To illustrate, say the appropriate expected rate of return is 20%. A newly available bond investment promises \$25 for a \$100 investment with fully insured principal but a 50% probability of default on the interest payment. Say that other risky bonds in the economy offer 11.4% more (for a total interest rate of 31.4%). If you discounted the promised interest payment of \$25 with the quoted interest rate on your benchmark bonds, you would get

(Omitted eq)

Wrong! Instead, you must work with expected values:

(Omitted eq)

Summary

This chapter covered the following major points:

- Uncertainty means that a project may not return its promised amount.
- A random variable is one whose outcome has not yet been determined. It is characterized by its distribution of possible future outcomes.
- The “expected value” is the probability-weighted sum of all possible outcomes. It is the “average” or “mean,” but it is applied to the future instead of to a historical data series. It is the standard measure called “reward” when financiers talk about risk vs. reward.
- Risk neutrality means indifference between a safe bet and a risky bet if their expected rates of return are the same.
- The possibility of future default causes promised (quoted) interest rates to be higher than expected interest rates. Default risk is also often called credit risk.
- Most of the difference between promised and expected interest rates is due to default. For bonds, other premiums (like extra compensation for bearing more risk — the “risk premium” — though not the term premium) are typically smaller than the default premium.
- Credit ratings can help judge the probability of potential losses in default. Moody’s and S&P are the two most prominent vendors of ratings for corporate bonds.
- The key tool for thinking about uncertainty is the payoff table. Each row represents one possible outcome, which contains the probability that the state will come about, the total project value that can be distributed, and the allocation of this total project value to different state-contingent claims. The state-contingent claims “carve up” the possible project payoffs.
- Most real-world projects are financed with the two most common state-contingent claims — debt and equity. Their payoff rights are best thought of in terms of payoff tables.
- Debt and equity are methods to parcel out total firm risk into one component that is safer than the overall firm (debt) and one that is riskier than the overall firm (equity).
- The presence of debt “levers up” equity investments. That is, a smaller upfront cash investment becomes more exposed to swings in the value of the underlying firm. However, there are also other leverage mechanisms that firms can choose (e.g., leasing or technology).
- If debt promises to pay more than the project can deliver in the worst state of the world, then the debt is risky and requires a promised interest rate in excess of its expected interest rate.
- NPV is robust to modest errors in the expected interest rate (the discount rate) for near-term cash flows. However, NPV is not necessarily robust with respect to modest errors in either expected cash flows or discount rates for distant cash flows.
- NPV is about discounting *expected* cash flows with *expected* rates of return. You must not discount *promised* cash flows with *promised* rates of return.