9

Benchmarked Costs of Capital

As an investor, your problem is how to form good portfolios. As a corporate manager, your problem is how to get your own firm into third-party investors' portfolios. So you need to know the right discount rate (i.e., high enough an expected rate of return) at which enough investors will "bite." In earlier chapters, this discount rate was just time-based and all you had to do was to offer the same expected rate of return. In this chapter, we begin adding a risk component.

We will now assume that your investors simply benchmark all investment opportunities (including your stocks, bonds, projects, etc.) to other prominent asset classes in the economy. In particular, we assume that they will evaluate your firm based on two characteristics: (1) whether your project payoffs are more like short-term or long-term investments; and (2) whether your payoffs are more like safe debt or risky equity. Short-term and safe bond-like projects can get away with offering investors lower average rates of return; long-term and risky stock-like investments must offer higher expected rates of return. This means we need to take another look at bills, bonds, and stocks in the overall economy. What is the appropriate risk-free rate of return for projects of similar durations, and what is the (so-called) equity premium for the expected rate of return on stocks above bonds?

9.1 What You Already Know

Let's take stock (pun intended). Unless otherwise noted, in this and the next chapter, you are primarily taking the stance of a manager who wants to attract and retain external investors and invest their money into projects that they like. You already know the right train of thought for capital budgeting purposes: As a corporate manager, your task is to determine whether you should accept or reject a project. You make this decision with the NPV formula. To determine the discount factor in the NPV formula, you need to estimate the appropriate cost of capital — or, more precisely, the *opportunity* cost of capital for your investors. This means that you need to judge what a fair expected rate of return, $E(r)$, is for your project, given your project's characteristics. When compared to "similar" projects elsewhere, if your project offers a lower expected return, then you should not put your investors' money into your project but instead return their money to them. If your project offers a higher expected return, then you should go ahead and invest *their* money into *your* project. Put differently, your goal now is to learn what your investors, if asked, would have wanted you to invest in on their behalves. Of course, it still remains difficult to determine what "similar" is, but this is a devil in the details.

Unfortunately, the perfect market assumptions of financial Utopia are no longer enough to proceed. You must begin to speculate more about your investors' preferences. What do investors like and dislike? You already know two relevant project attributes:

Far-Off vs. Nearby Payments: Long-term Treasury bonds have (usually) been offering higher yields *per-annum* than short-term Treasury bills. Presumably, this is because investors are more reluctant to part with their money when payment is farther down the line. In this sense, you can think of long-term as "undesirable" relative to short-term. Investors (usually) seem to like getting money sooner.

Equities vs. Bonds: The stock market has offered higher average rates of return than the bond market. Presumably, this is because investors are more reluctant to part with their money when all they get is a fuzzy, risky claim, like equity, whose repayment depends more on the success of the project. In this sense, you can think of equity as "undesirable" relative to bonds. Investors like getting money with less variance.

(A quick clarification: high expected rates of return usually mean that investors dislike an asset's attributes — this asset could *not* be sold for a high price because investors needed to be compensated extra for something.)

As an executive, you should assume that if investors dislike an attribute in the wider financial markets, they will also dislike it in your own projects. If you offer them a project that pays off more like stock-market equity, it has to offer the higher expected rate of return of stocks. If you offer them a project that pays off more like bond-market principal and interest, it can offer the lower expected rate of return of bonds. And if you compare two projects, one with payoffs farther in the future than another, the former should offer higher per-period expected returns — just as long-term bonds usually offer higher expected returns than short-term bills. The focus of this chapter is therefore to assess what rates of return you can expect in these different types of investment.

In a perfect market, these rules must surely be correct for the most simple of all investment projects: Imagine firms that do nothing but invest in Treasury bonds. Such fixed income investment funds are actually quite popular. They should offer about the same expected rates of return as their Treasury bond holdings. If they offer lower expected returns, investors can buy the bonds themselves (perhaps with some extra transaction costs in the real world). If they offer more, investors will quickly bid up the price of the fund until the expected returns become about the same. The same is true for equity. You already know about firms that basically do nothing but hold S&P 500 stocks — funds like VFIAX (S&P 500). They should offer about the same expected rates of return as the stock market. And funds that invest 50-50 (which exist, too!) should offer 50-50.

Although we are still remaining in the perfect-market framework, we do not have to lean too heavily on it. For example, assume that investors like bitcoin-like stocks for no good reason. If your stock is very bitcoin-like, it should still be priced similarly to other bitcoin-like stocks. The deeper reason why investors like bitcoin-like stocks are secondary. All you need to learn as a corporate manager is to judge how much your own projects are just like what investors can buy and sell elsewhere in the financial markets. It is a little more difficult, though. Without a perfect market, what stocks should you consider similar? This judgment can border on philosophy: relevant characteristics can even be defined as whatever aspects make projects have the same pricing.

The big question of this chapter is: how can you assess the appropriate expected rate of return on the standard benchmarks, i.e., on risk-free investments and on stocks? In the next chapter, you will learn methods to judge how similar your given project is to each of these benchmarks.

9.2 The Risk-Free Rate — Term Compensation

How do you assess the risk-free rate of return? As notation, we use r_F, but although this is very common, it is not the definitive standard. (Others may be using rf or just F.) We are also not pedantic enough to write the subscripts for the specific time-frame we are working with.

Most corporations want to discount nominal cash flows, so they use the standard nominal rate from U.S. Treasuries. In the rare case that a corporation needs to discount real cash flows, the U.S. Treasury also offers quotes on inflation-adjusted real bonds (TIPS).

There is one small issue, though — which Treasury? What if the yield curve is upward-sloping (as it usually is)? For example, in December 2021, Treasuries yielded 0.4% per annum over one year, 1.3% over five years, and 1.9% over thirty years.

So think about the basics of your own project. You want to match your projects' cash flows to the most similar risk-free bond benchmark. You should choose the risk-free bond yield that most closely mirrors the specific expected cash flows. For example, to value a safe project that operates for three years, use the 1-year Treasury yield to discount the expected cash flow for the first year's NPV term, the 2-year Treasury for the second year's NPV term, and the 3-year Treasury for the third year's NPV term. If you had to use just one risk-free rate for multiple cash flows (because your Dilbertian boss says so), choose an average of the three rates or simply the 2-year bond. There are better duration-matching ways to do this, but unless you are a bond trader, the extra precision is rarely worth it.

Matching cash flows to similar maturity bonds is not a law of nature but a reasonable (and loose) approach. Think about the opportunity cost of capital for a small investment that does not vary systematically with anything else. If your corporation's investors are willing to commit their money for ten years, they could earn the yield on a ten-year risk-free bond instead. It is this ten-year rate that would then be more indicative of the opportunity cost of capital on your own project cash flow that will materialize in ten years than would, say, the rate on a one-year or thirty-year bond. If your project's cash flow will occur in three months, your investors could alternatively only earn the lower rate of return on the three-month bill.

Of course, to your investors, your project's cash flows are not likely to be exactly like the analogous U.S. Treasury payments. Thus, you can consider some refinements. It may be more appropriate to use an opportunity cost more similar to corporate than to Treasury bonds. For example, you may deem it to be better to use the interest from short-term corporate bonds issued by investment-grade companies. Fortunately, after you take into account that quoted yields have to be reduced by the expected default premium, the average historical rate of return is probably not that different. Or you may deem long-term non-investment-grade (i.e., high-yield) bonds (except perhaps mortgages) to be more similar and appropriate. Their cost of capital (expected not quoted rate of return) seems to be quite a bit higher.

(Omitted solvenow)

9.3 The Equity Premium — Risk Compensation

Appropriate compensation for a risk-free investment over a given time frame is the easy part. This is the cost of risk-free capital. Now comes the hard part: appropriate compensation for taking risk. This is the cost of risky capital. Although most corporate projects are not risk-free, you can think of them as some combination of a safe part (a debt-financed claim) and a risky part (an equity-financed claim). Indeed, you have already learned that you can always split a medium-risky project into claims that have safer and riskier payoffs. Therefore, you usually need to know the appropriate cost of capital on the risky part, too — the task at hand now.

Unfortunately, the expected rate of return on risky assets is much more difficult to estimate than the risk-free rate. First, what is a good benchmark for corporate risk? Hmmm — what is the most canonical risky asset in the economy? The stock market! Thus corporate financiers usually rely on the equity premium as a benchmark:

(Omitted eq)

which is the extra expected rate of return that risky equity projects have to offer above and beyond what risk-free bonds are offering. (It is a difference of two rates, so you can use either two nominal or two real rates.) Later, when you want to determine the expected rate of return on a project that consists only of one asset that is the stock market, say the VFIAX (S&P 500) fund, you would add back the interest rate you just subtracted out here. It is easier to think about the "extra" of the risk premium above the term premium (in the risk-free rate) rate. The equity premium $\left[\mathsf{E}(r_M) - r_F\right]$ is also sometimes called the market risk premium. In common use, the terms can refer either to realized rates of return or expected rates of return, although the latter is more common and we will use it mostly in this sense in this chapter. (This ambiguity is not my fault.)

This equity premium is a number of first-order importance for everybody. It is not just the corporations that want to know it for their cost-of-capital estimation. *You* also want to know it as an investor when you decide how much of your money you should invest in stocks rather than bonds. Unfortunately, in real life, the equity premium is not posted anywhere — and *no one really knows the correct number*. Worse: Not only is it difficult to estimate, but the estimate often has a large influence over all financial decision-making. *C'est la vie!*

(Omitted fig)

There are a number of methods to guesstimate the equity premium. Unfortunately, for many decades now, these methods have disagreed with one another. It should thus come as no surprise to you that practitioners, instructors, finance textbook authors, and everyone else have been confused and confusing. For example, each finance textbook seems to have its own little estimate, as you can see in Figure **??**. There was a trend, though: both the disagreement and the average recommended estimate were slowly declining over the decades.

So we finance-textbook authors have two choices:

1. We can throw you one estimate, pretend it is the correct one, and hope that you won't ask questions. It would be a happy fairy tale ending. Unfortunately, it would also be a lie.
2. We can confess to the truth. We can tell you how different methods can lead to different estimates — and how we are really all in the same boat. Worse, we are not sure where the boat has holes.

In this book, I am going to take the second route. I will explain to you what each method suggests and actually means. You can then make up your own mind as to what you deem to be best. (I will tell you my own personal estimate at the end.) This approach also has an important advantage: you won't be surprised if your boss uses some other equity premium to come to different conclusions. At least you will understand why.

Let's discuss one-by-one — and in order of prominence — the six most prominent methods that form the bases of common equity-premium estimates.

Method 1. Historical Averages I

The first and most common guesstimation method is to ignore the fact described in Chapter that average returns — unlike standard deviations and market-betas — tend not to be very reliably predictable. This makes them difficult or outright hazardous to extrapolate. Nevertheless, Method 1 presumes that whatever the average equity premium was in the past will also be the case in the future.

(Omitted fig)

Figure **??** plots the average *geometric* performance of the stock market (with dividends) over the last T years. You choose point T on the x-axis based on how interested you are in shorter (younger) or longer (older) historical data. The graph shows the historical geometric average rates of return on stocks, (long-term) investment-grade corporate bonds, long-term Treasury bonds, short-term Treasury bills, and inflation. The difference between the black stock-market line at the top and the faint (red) Treasury bills line is what most people call the equity premium. Over the last 50-100 years, stocks have outperformed short-term bonds by about 7-8%/year.

Let's discuss this estimates and its interpretation in more detail. In particular, we want to be clear about how to deal with this benchmarks for assessing your own short-term and long-term project opportunities. Most interesting corporate projects, like factories, IT infrastructure, research, or brand names, deliver cash flows over many years.

➤ The Stock Market is a Long-Term Asset

A natural way to think of stock returns is that they have two components:

A Term Premium: One part is related to the fact that stocks are long-term investments, earning and paying dividends on average for 10-100 years. Like long-term Treasuries, they therefore have to offer a higher expected rate of return than Treasury bills.

A Residual Pure-Risk Premium: Another part is the remainder, presumably related to the fact that stocks are also riskier investments than long-term Treasuries and thus have to pay more (the pure risk premium).

The empirical evidence in Figure **??** shows that over the last 30-60 years, all but about 4% of stocks' 12% rate of return has come from the term premium (the fact that investors earned a high return for payments — dividends or interest — occurring on average far in the future). The extra 4% was the extra compensation that stock investors earned above long-term Treasury investors. For 1980–2021,

(Omitted eq)

However, be warned: many investors call the total equity premium the risk premium, thus including in it both a term-premium and a pure residual risk premium.

➤ Disagreeing About Risk and Equity Premium Estimates

Let's reconcile different historical equity premium figures that you may encounter. The basic facts for historical annual rates of return in the United States are

Note: inflation hits all classes the same.

	Arithmetic			Geometric		
	Stocks	Bonds	Bills	Stocks	Bonds	Bills
1926-2020	12.0%	6.1%	3.4%	10.1%	5.7%	3.3%
1970-2020	12.3%	9.0%	4.5%	10.8%	8.4%	4.4%

If you use 95 years of historical data, arithmetic rates of return, and a spread over short-term T-bills,2]subsect:geoarireturnsGeometric vs. Arithmetic Returns and Extrapolation you can settle on an equity premium estimate as high as $12.0\% - 3.4\% = 8.6\%$. Versions of this 8.6% number were etched into the minds of generations of students, practitioners, and finance professors by many textbooks from the 1980s. (This is especially in the context of the CAPM, which will be explained in the next chapter.) The fact that 8.6% seems too big to be a reasonable a-priori expectation of what an investor should have gotten (from taking on what seems to be just a modest amount of risk) is called the equity premium puzzle. It is important to understand what it means and what the alternatives are.

Don't jump too quickly to conclusions:

1. Arithmetic means are often misleading, because they ignore the (relative) effects of volatility on compounding. Stocks really didn't earn a buy-and-hold investor 8.6% per annum above short-term Treasuries.
2. Figure **??** shows that for the last 30-60 years, the equity premium puzzle seems to have been more of a term-premium puzzle than a risk-premium puzzle: Long-term bonds outperformed short-term bills by about 4.6%. Stocks — themselves more long-term assets — outperformed long-term bonds by "only" 3.5%.

   However, the longer 1930-2021 sample was relatively worse for long-term Treasury bonds. Here, long-term bonds outperformed short-term bills by "only" 2.2% but stocks outperformed long-term bonds by 4.4%.

Let's reconcile the roughly 100-year arithmetic equity premium over T-bills (which measures both a term premium and a risk premium) with a roughly 50-year geometric equity premium over T-bonds (which measures only a risk premium):

Arithmetic Equity Premium 1926 to 2020 over Short-Term T-Bills		8.6%
Instead use later Sample Period 1970 to 2020	–0.8% ≈	7.8%
Instead use Long-Term T-Bonds	–4.5% ≈	3.3%
Instead use Geometric Returns	–0.9% ≈	2.4%

Both the high 8.6% estimate and the much lower 2.4% alternative estimate in Figure **??** can thus follow from historical U.S. data. To be clear, the 8.6% premium is not wrong. It is one valid measure of the superior performance of U.S. stocks over U.S. Treasury bills, but it has only its very specific meaning. (And being arithmetic, it was also not compoundable.) So depending on context, 2.4% may be the more appropriate measure.

Let's discuss the differences one by one:

1. **Sample Period?:** You have to judge what historical sample is appropriate. You probably want to end the sample recently (last year). But it is not clear whether you should start, say, in 1926 (which is when the most reliable of our common finance databases, commonly called just CRSP (www.crsp.org," begins) or in 1970 (about half-way). Although your estimate can seem statistically more reliable if you use more years, using the long sample means that you are then leaning more heavily on the heroic assumption that the world has not changed. Are the world and its expected risk and reward choices really still the same today as they were in, say, 1830, 1871, 1926, or 1970? (And is the United States really the right country to consider alone? Did it just happen to have had an unusually lucky streak during [the first half of] the "American Century," which is unlikely to repeat? In this case, the average country's experience may be a better forecast for today's United States, too.) No one knows the best sample choice. As for me, I intuitively prefer a shorter sample of half a century. But smart people can disagree.

   Equity performance itself does not seem to have deteriorated much. As Figure **??** showed, the equity premium was lower in this 50-year sample not because (more volatile) stocks performed worse (they did not), but because (less volatile) Treasury bonds performed better — and bonds continue to have higher yields than bills even as of 2022.

2. **Long-Term or Short-Term Bonds?:** You have to judge whether short-term or long-term bonds are the appropriate benchmark. From the perspective of a manager who needs to decide about a short-term project, using a similar-term short-term interest rates as the benchmark also makes sense.

   However, from the perspective of a corporate manager who needs to commit funds to a long-term project with cash flows over decades, it does not. It is not possible for corporations to quickly move in and out of decisions to build, say, power plants. Building a plant is a long-term decision. If all investors can earn higher yields in Treasuries by committing their money for 20 years, and if your own plant requires them to commit their money for 20 years, too, then your plant should also be benchmarked to this long-term expected rate of return. Conveniently, the term spread between 1-year and 20-year risk-free rates (though not the rate of return on rolling over 1-year bills over 20 years) can be easily looked up on the web every day. There is little uncertainty.

3. **Geometric or Arithmetic?:** You have to judge whether you should use geometric or arithmetic rates of return in your benchmark cost of capital in the NPV formula. The answer is not clear, as you may recall from Section . Many corporations incorrectly compound the annual arithmetic average stock return or equity premium without much thought. However, doing so means that they expect multi-year performance of stocks relative to bonds to be better in the future than it was in the past.

   But there is a simpler argument based on the rule of comparing apples to apples. How do you think about your own expected cash flows? I bet you do so in geometric terms. If you think in terms of arithmetic expected cash flows compounded over many periods — i.e., if you consider the expected cash flow on a project that first earns +200% and then –100% [for a complete overall loss] to be a success with a positive average rate of return, then you should use the arithmetic average. Hardly anyone thinks this way.

We will return to compounding concerns in Section **??**.

In sum, one good perspective is that the equity premium puzzle has been more of a term puzzle than a risk puzzle at least for half a decade.

(Omitted anecdote)

➤ Uncertainty About Historical Estimates

Forgive me, but I have not even mentioned another big problem: the large margin of error. The standard deviation of stock returns of 20%/year translates into a standard error of about $20\%/\sqrt{100} \approx 2\%$ if you use a 100-year sample. If you are willing to assume that the stock-market process has not changed over the last 100 years, and that stock returns are roughly normally distributed, then you can use some additional statistical artillery: You can then be about 95% sure (a confidence range popular in statistics) that the true mean geometric stock return over long bonds was between 0% and 8% from 1926 to 2020. Frankly, this large a range on the appropriate cost of capital for equity would not be the kind of accuracy you like when you are a manager who has to decide whether to invest money. You already knew — or at least should have reasonably believed — that the equity premium should not have been negative.

➤ "Pennies in Front of Steamrollers" and the "Peso Problem"

To make matters even more complex, some economists believe that even the observed historical data are not telling the full story, either. Let me explain this by analogy with a little detour into placing asymmetric bets.

The odds in roulette are always in favor of the casino and against you. For example, when you bet on a single number (out of 36+1 numbers), your return when winning is a payout of 35-to-1. How good is a bet of $1 each on the first 35 roulette numbers? Well, each one of your number bets will win 35/37≈95% of the time and lose 2/37≈5% of the time. When one of your 35 numbers wins, the other 34 will have lost (thus, losing you $34), but you receive $35 for your single correct number. Thus, you would win a net of $1 in 95% of your gambles. Congratulations — you have just created a strategy that wins most of the time. However, in the remaining 5% of your gambles, you lose all $35. Your net profit is therefore about 95% · $1 + 5% · (–$35) ≈ –$0.80 per roll of the roulette ball.

Playing this game is not only not in your interest (it has a negative expected rate of return), but also incredibly risky. In fact, over long enough a time period, you are guaranteed to lose all your money. Economists also sometimes call strategies like this (where you win small amounts most of the time and lose large amounts very rarely) "picking up pennies in front of a steam-roller."

However, let's say you were unaware of the physics of roulette and tried to infer your chances from your historical performance alone. After eight gambles, there is about a 50-50 chance that all of them won (and you earned $1 each). What if this turned out to have been the case? You might even conclude that roulette is like free money, never losing. Of course, because we know the physics of the Roulette wheel, we know that this reasoning is delusional.

Similarly, maybe we just happen to live in world in which the stock market has never "rolled" the rare worst outcomes. The true expected rate of return could have been zero or even negative. Some economists indeed believe that worse disasters have been possible, but their probabilities have been tiny (say, 1-in-100 years or lower) — and they just "happened not to have happened" in the last 100-200 years. The super-volcano did not blow; the asteroid did not hit; climate change did not make earth uninhabitable for humans. For a hypothetical example,

nge:
e.world/

	Asteroid	Normal
Probability	0.0001	0.9999
Stock Return	–99.99%	+0.01%

and thus

True Average Expected Rate of Return:	0%
Average Rate of Return Inferred Given Luck of No Asteroid:	0.01%

Presumably, you would not like a zero expected rate of return for a risky investment if you can invest in Treasuries instead. Trust me that if the above probabilities were correct, then there would have been about a $0.9999^{100} \approx 99\%$ chance that over the last 100 years, not even one such asteroid would have hit. If you happen to have lived in such a lucky world — called "the U.S. of the last 100 years" — you would have calculated a historical average rate of return of 0.01%. Alas, it would be too optimistic an estimate of the true expected rate of return.

This is sometimes called the Peso problem, based on an otherwise obscure academic paper about the currency spread of the Mexican Peso. When you say "Peso problem," financial economists will know exactly what you mean!

There is some empirical evidence that investors behave exactly as if they fear such a Peso crash — but we do not know whether such a fear is (or was) rational and we are not sure how much of the historical equity premium it can explain. A reasonable order of magnitude is that extra compensation for crash risk could account for at most a 1-2%/year equity premium — and perhaps for nothing.

➤ In Conclusion

If your estimate of the forward-looking equity premium is based on the "historical averages I" method, then you can defend a choice of 1% (for long-term cash flows). If you are aggressive, you can defend even a choice of 8% (for short-term cash flows), and equity premium ranges from 0% to beyond 10% if need be (or, more cynically, if you are an expert witness paid to so opine). Are you in awe or disgust about our uncertainty and the wide possible range of estimates here? For me, it's both.

Method 2. Historical averages II

The second method for estimating the equity premium is to look at historical realizations in the opposite light, as in Figure **??**. Maybe stocks have become more desirable — perhaps because more investors have become less risk-averse. They would have competed to own more stocks, and thus have driven up the prices. This would imply lower expected rates of return in the future! High past rates of return would then be indicative of low future expected rates of return.

An even more extreme version of this argument suggests that high past equity returns could have been due not just to high a-priori equity premiums, but also to historical "bubbles" in the stock market. (We will explain this further in Chapter .) The proponents of the bubble view usually cannot quantify the appropriate equity premium, but they do argue that it is lower after recent market run-ups — exactly the opposite of what proponents of the *Historical Averages I* guesstimation method argue. However, you should be aware that not everyone believes that there were *any* bubbles in the stock market, and very few credible economists believe that the U.S. stock market over the entire century was one big bubble. (They do believe there was at least a small Peso problem, however.)

Method 3. Current predictive ratios

The third method for estimating the equity premium is to try to predict the stock-market rate of return actively with historical dividend yields (i.e., the dividend payments received by stockholders). Higher dividend yields should make stocks more attractive and therefore predict higher future equity premiums. This equity premium estimation is usually obtained in two steps:

1. Estimate a statistical regression that predicts next year's equity premium with this year's dividend yield
2. Substitute the currently prevailing dividend yield into your estimated regression formula in order to predict.

In 2020, dividend yields were about 1.5% per annum. This is so low that the regression-predicted equity premium was negative — which makes no sense. Variations of this method have used interest rates or earnings yields, typically with similar results. In any case, the empirical evidence suggests that this method does not yield great predictions — for example, it predicted low equity premiums in the 1990s, which was a period of superb stock-market performance.

Academics disagree whether such methods work for short-term equity-premium predictions (say 1-5 years). But all agree that we do not have the data to test whether the dividend-yield can predict 10-50 year equity premiums. And it is estimates for the very-far-away expected cash flows where corporate finance managers are most in need of equity premium estimates. Therefore, most managers can neglect these regressions.

Method 4. Philosophy

The fourth method is to wonder how much rate of return is required to entice reasonable investors to be indifferent between stocks and bonds. Even with a geometric equity premium as low as 3%, over 25 years, an equity investor would end up with more than twice the money of a bond investor. Naturally, in a perfect market, nothing should come for free, and the reward for risk-taking should be just about fair. Therefore, equity premiums of 6-8% just seem too high for the amount of risk observed in the stock market. This philosophical method generally suggests reasonable equity premiums of about **1%** to **3%**.

Method 5. Surveys: Ask the Experts

What to choose? Welcome to the club! No one knows the true equity premium. So, the fifth method is to ask the experts — or anyone else who may *or may not* know. It's the blind leading the blind. The ranges of estimates have varied widely (and they are often also conveniently tilted in the interest of those giving them):

- Duke professors John Graham and Campbell Harvey have been surveying between 200 and 500 CFOs every quarter for decades now. Their pre-Covid median and mean one-year equity premium estimates (above the 1-year Treasury) have been meandering between about 3% and 5% per annum, with the most recent 2019 estimates in about the middle of this range. About half of their respondents' mean estimates were inside the 1% to 7% range; the other half were outside. In 2019, 10-20% of the respondents even predicted a negative equity premium. (A negative expected equity premium seems in itself somewhat odd from a theory perspective, because if the *ex-ante* equity premium had been negative, presumably few people would have kept their money in the stock market. However, a minority of executives could well have held this view, anyway. In hindsight, i.e. *ex-post*, even the average executive was proven too pessimistic. Despite Covid, the 2019-2021 years were quite good for stock-market investors.)
- The U.S. CFO estimates seem in line with those of CFOs from other countries. For example, in 2012, Pablo Fernandez reported that analysts and companies in the United States, Spain, Germany, and the United Kingdom all used average estimates of between 5% and 6% — just like finance professors, and with the same typical range from about 3% to 8%. And this estimate further increased by another 1% over the following 3 years.
- For decades, the consulting firm McKinsey has used a standard of around 5%.
- The Social Security Administration sometimes uses an estimate of around 4%.
- In 2021, The California Public Employee Retirement System (CalPERS), the largest pension fund in the world, assumed that it will earn about 6.8% on its approximately $500 billion of holdings (compared to realized returns of 6.9% over 10 years and 8.4% over 30 years). With almost all its money invested in equities, this comes to an equity premium estimate of about 5% over long-term Treasuries.

  If this seems somewhat high to you, it probably is. But lowering this estimate would mean that California's politicians would have to set aside more money for their unfunded pension obligations *today*. Obviously, they would prefer to leave the optimistic estimate as is, and kick the can down the road to their successors.
- Compared to professional investors, retail investors (like those on Robinhood) tend to believe more in recent history. They turn more bullish after a good stock market and more bearish after a bad stock market. They are probably not the best predictors to follow.

In sum, finance professionals nowadays seem to work with equity-premium estimates of about 3-4% per year. Of course, these estimates were themselves likely based on the first four methods, and they occur in echo chambers — they are what analysts, companies, consultants, students, and professors have been reading in corporate finance textbooks (like this one) for many years now. (Hmm — maybe I should try claiming 3.14159% and then see how many surveyees will repeat this pi-oneering estimate back in ten years.)

6. Internal Cost of Capital (ICC) and Accounting Models

A hybrid method combining survey methods and analysis is the "Internal Cost of Capital." Basically, this method uses analysts' consensus projections about S&P 500 earnings (over the next few years) with a perpetuity model to back out a cost of capital that makes the S&P 500 price equal to the analysts' discounted future earnings. Because analyst estimates vary over the business cycle, researchers usually use the average of many ICCs over many years.

Until the mid-1980s, this geometric average was generally lower than the historical average performance, consistent with the view that the 20th century was the lucky American Century. However, more recently, it has agreed more with the historical expected rate of return in suggesting much higher expected stock-market rates of return for the future. (And, as with historical estimates, different variants can give estimates with a much larger range, say, from 0% all the way to 7%.)

There are some accounting-based models that are based on similar principles and are often claimed by their proponents as panaceas — or at least as better alternatives. Alas, when I looked at some of these models with a more skeptical eye, I could not share their enthusiasm for three reasons. First, these models are too "cute": each has been tweaked just a little here and there to make it look good on their data. Second, these models tended to work well in the first parts of their samples and not so well in the second parts. Third, if they really worked half as reliably as they are claimed to work, then investment funds should have flocked to them like flies. Many looked at these models and they did not. This is not to say that no such model works — just that those that I checked up on in more detail did not seem to hold up.

9.4 Forward-Looking Benchmarks?

The risk-free rate and the equity premium are the two most important numbers in economics and finance. If the risk-free rate is high, you should save more and consume less. If the equity premium is high, you should allocate more of your savings into diversified risky stocks and less into bonds. The previous section has taught you about how to view the historical data.

But you are probably not interested in *historical* performance for its own sake. You are probably interested in the *future* expected performance instead. (When you want to judge whether the road goes uphill or downhill, looking at the rearview mirror may be better than nothing, but it is not ideal.)

So what is the appropriate forward-looking expected equity premium *today*? Sadly, no one can tell you the authoritative estimate. Such an authority does not exist. Everyone is guessing. Unfortunately, unless your project has no market-risk type of exposure, you usually have to take a stance. (I will explain in Section how you can finesse this, but doing so will have its own drawbacks.1]subsect:neutralizingbetaNeutralizing Market Exposure) I failed to shield you from the estimation dilemma. I can only give you the considerations that you can contemplate when you are picking *your* estimates.

If you are hoping I will rescue you in future chapters, by either giving you the correct numbers or telling you that you do not really need them to make decisions, I can't. Even more involved financial models, in particular the CAPM in the next chapter, ask *you* to provide the very same equity premium estimate. They just try to inform you about the expected rate of return for projects *relative* to Treasuries and the stock market. Given *your* estimate of how much risky average stock-market projects should earn relative to safe projects, plus the market-beta, the CAPM then tells you the benchmark cost of capital for *your* projects. But unless your projects have zero exposure to stock-market-type risk, the models themselves require *you* to input your equity-premium estimate.

The need for good alternatives (benchmarks) is important to capital budgeting in corporations. They measure the opportunity cost of capital. But you also need them if you are an investor, i.e. on the buying side. Like everybody else, you cannot let your limited knowledge stop you from making investment decisions. You do need to be your own judge: what are your prevailing (economy-wide) opportunities? Where do you want to place your money?

Term and Risk

I admit that I could not tell you the correct equity premium estimate. No one can. But I am not altogether useless, either. I can teach you at least how to avoid some basic errors. You have already read about one important aspect, albeit in the context of historical averages. Short-term and long-term projects should have different benchmarks. This insight is very important and you can easily get this right. So let's discuss it in more detail.

Risk-free Projects: The correct approach is obvious for risk-free projects. If your project is short-term, the correct benchmark is the rate of return on short-term bills, not long-term bonds. If your project is long-term, the correct benchmark is the rate of return on long-term bonds, not short-term bills.

Risky Projects: The correct approach is less obvious for risky projects. Remember that stocks are themselves assets based on many long-term cash-flows (even though you can always sell the shares instantly, just as you can sell Treasury bonds at any moment).

- If you have a project with a payoff that is as risky as the stock market and with a similarly long horizon, the stock market is your correct benchmark. The stock market's expected rate of return reflects both the term and the risk premium. If you think that the last 30 years (1990-2020) are representative for the purposes of estimating the future, Figure **??** tells you that you should expect to deliver a 12% average geometric rate of return on stocks. (If you want to decompose this further, about 5% is the short-term benchmark [the premium for saving money], 2% is the premium for the long-term nature of payoffs, 4% is the premium for taking risk, and the remaining 1% is combinations.)
- If you have a project with a payoff that is as risky as the stock-market payouts or earnings, but lasts for only a year or so, then the equity premium without the term premium is a better benchmark. Thus, a discount rate of about 2% less — i.e., about 8% — would seem better.

Some finance professors believe that you should use a higher risk premium (higher than 3%) for long-term cash flows — that is, more term premium in stocks than in Treasuries. But only the Treasury term premium is easy to measure. The jury is still out, and this extra "kicker" would likely be small.

➤ Geometric or Arithmetic Cash Flows and Benchmarks?

How does the NPV formula work under uncertainty? Over one time period, a geometric average rate of return is the same as the arithmetic average rate of return. This arithmetic average rate of return was itself calculated as the compounded (geometric) average over many smaller time intervals. Now, commonly-published benchmark rates of return are usually quoted as *annual* or even shorter-term rates of return, not as, say, 30-year rates of return. You have to translate the shorter-term rate of return statistics you are given into the expected longer rate of return statistics you need.

Does it make sense to compare arithmetic average returns across long-term project cash flows with different volatilities? Would you rather invest for T years in a ($100) project A with an average annual rate of return of 5% and a standard deviation of 40% (twice that of the stock market), or in a project B with an average rate of return of 2% and zero variance? Would you take project A if the financial markets also offered you project B? If you use the NPV formula on the arithmetic averages as

(Omitted eq)

you would conclude that you should take A. But this would be wrong.

The reason is that the expected rate of return over T periods $E((1+r)^T)$ is not $[1+E(r)]^T$. Geometric rates of return are smaller than arithmetic rates of return. (Remember: a rate of return of 50% followed by one of –50% leaves you with a –25% compounded rate of return.) As you already know, if the distribution follows a normal bell curve, with modest variance, then the geometric rate of return is about half the variance squared less than the arithmetic rate of return. With A's mean of 5% and standard deviation of 40%, you should really expect to earn $5\% - 40\%^2/2 \approx -3\%$ per T in your project, whereas the financial-market benchmark B offers +2% per T. For a long-term project, you would be better off turning down A and going with B.

For long-term cash flows, NPV really makes sense only if you use the appropriately compounded, i.e., geometric, expected rates of returns. Fortunately, many investors think of the expected cash flows (not returns) in the NPV numerator, which are more akin to the outcome of a geometric compounding mental process. If they use a –$100 flow today and a $150 flow in 10 years, they implicitly mean that they expect a compound rate of return of 50%, which they want to compare to geometric opportunity rates of return in the financial market elsewhere.

➤ Term and Averaging

What do you expect as a rate of return on the stock-market benchmark? If you expect the stock market to deliver 12% over the next year, with a 20% standard deviation, you should expect it to deliver about $12\% - 20\%^2/2 \approx 10\%$ over the very long run. The 2% difference is roughly the historical difference between arithmetic and geometric rates of return on the U.S. stock market over the last 50 years.

Now put together your knowledge of the term premium and risk premium when you want to benchmark your own either short-term or long-term risky cash flows. For a long-term project, you could invest either in the stock market or in Treasury bonds. As an investor, how much would you expect to earn above the stock market?

Begin Important

Whatever your base estimate is of the short-term market-risk premium EQPST ("equity premium, short term estimate"), the following rough adjustment is required to keep your estimate of the long-term market-risk premium consistent with your short-term market-risk premium estimate (assuming that the risk-reward tradeoffs will remain similar over the next few decades):

	Arithmetic	Geometric
Relative to Short-Term Bills	EQPST	≈ EQPST – 2%
Relative to Long-Term Bonds	≈ EQPST – 2%	≈ EQPST – 4%

End Important

For example, if you believe that the stock market will outperform Treasury bills by 6% over a typical one year, then you should expect the stock market to outperform Treasury bonds by a (compound) ≈2% over the next 30 years. One can quibble whether these adjustment recommendations are off by up to 1% (could be 3% per year), but they are in the right ballpark.

When you evaluate short-term market-risk-level projects, you can use your EQPST base estimate in the top left corner of the table in the box as a reasonable benchmark. When you evaluate long-term projects, you should use the estimate in the bottom right corner of the table. Whatever else you do, do not make the mistake of thinking they should be the same.

The decomposition of the stock-market return into a term premium and an equity premium matters for investments that are not 100% like stocks. For investing 100% in stocks, whatever term premium you add on one end is subtracted back from the other ($TP + (MRP - TP) = MRP$, where TP means term premium and MRP means marginal risk premium). For short-term investments, you can expect a high equity premium but a low term premium. For long-term investments, you can expect a low equity premium but a high term premium. But if you have other types of investment, e.g. one that is more like 50% stock and 50% bond, it matters ($TP + 0.5 \cdot (MRP - TP) \neq MRP$). This will become even clearer in the next chapter.

Investors need to think about the same kind of adjustments. When evaluating stock investments, fund managers should add the equity premium estimate and the term premium estimate, too, to arrive at what they can expect. Expecting to earn 6% above short-term Treasuries over the next year is consistent with expecting to earn 2-3% above long-term Treasuries over the long run.

Do not take the rules too literally. It is not unusual for managers to be more conservative for long-term projects and assess higher hurdle rates on them. This is more likely related to their uncertainty about their cash flows and to imperfect market premia than the proper assessment of long-term average rates of return of stock and bond investments. For example, a tax-exempt pension fund that believes in using historical estimates for its expectation should not expect an investment in the U.S. stock market to outperform an investment in long-term Treasuries by more than 2% per annum over the decades, even if it has the perspective that the stock market will outperform Treasury bills by 6% over the next year.

Incidentally, do you remember Figure ??? Some of the disagreements over estimates stem from the fact that textbooks can mean different things by “equity premium.” The most common estimate is probably the highest estimate, the EQPST.

My Personal Opinions

The choice of geometric vs. arithmetic and Treasury bills vs. bonds is determined by application and not by opinion. Many earlier textbooks fail to explain the difference, resulting in miscalculated costs of capital. However, the choice of a relevant historical sample to assess the future is, in the end, opinion. For me, I tend to believe that the last 50 years are more relevant than the last 100 years. Thus, I recommend an equity premium of about 2-3% for long-term cash flows — which is much lower than the 5% that would be touted in other books. *Yet*, I also emphasize that I then use the 10-year term premium, which is 2-4% higher than the 1-year term premium. In Chapter , we will also discuss imperfect market premiums which can often further increase my long-term cost-of-capital estimates.

I also emphasize that it is important to *be consistent*. Do not use 3% for investing in one project and 8% for investing in another similar project. Being consistent can reduce your relative mistakes in choosing one project over another.

Finally, be aware that managers often care less about the scientific merits of costs of capital estimates than they care about whether they want to take or not take a project — whether they want to overstate or understate its value. “Expert” witnesses often cherry-pick estimates as low as 0% or as high as 8%, depending on the paying clients’ desires. I often find these estimates less believable the further away they are from my own assessment and the further they violate the spirit of the correct term adjustment. And I find anything outside a 1% to 8% range just progressively more difficult to swallow.

In our perfect capital market, companies should also use their cost-of-capital estimates as their hurdle rates. They should undertake every project whose rate of return exceeds its cost of capital and reject all others. However, in the real and imperfect world, companies do not trust that their own estimates are unbiased and accurate. Thus, it is more common for companies to have hurdle rates that are quite a bit higher (on average by about 4% in recent surveys) than their cost-of-capital estimates. We will return to these issues when we discuss more realistic world scenarios. Note that these scenarios still build conceptually on the perfect capital market scenario, so you are not wasting your time learning them in this perfect-market scenario first.

(Omitted solvenow)

9.5 Asset vs. Equity Costs of Capital

It is important that you always distinguish between the asset cost of capital and equity cost of capital.]eq:assetbetaAsset and equity betas Debt is always safer than the underlying project and equity is always riskier. Thus, equity should have a higher cost of capital than the assets.

Let’s work a short example. Say that you can buy a retail mall at a price that suggests an expected rate of return of 6%. However, when you look at REITs (real estate investment trusts, which are stock-like equity investments) of retail malls at YAHOO!FINANCE, you see that those seem to offer much higher expected rates of return, say 12%. Drop the deal? Not necessarily.

To compare the two investments, you have to take into account that REITs are typically already highly levered. It is easy to obtain a 50% mortgage on a retail mall, and even 80% is possible. If an 80% mortgage has an expected rate of return of 4% per annum, then the asset cost of capital for the underlying REIT project is

(Omitted eq)

The 6% mall looks like a great deal. This calculation is called unlevering the cost of capital. Alternatively, you could have calculated a levered cost of capital for your proposed mall, assuming you could obtain the same mortgage terms,

(Omitted eq)

This suggests an expected rate of return of 14%.

9.6 Deconstructing Quoted Rates of Return

In this chapter, we assumed that stocks offer higher *expected* rates of return than bonds. This was surreptitious. We changed the scenario. Such higher expected returns make sense only if investors are less willing to hold stocks than bonds *given* equal expected rates of returns. A good reason for such preferences is that investors are risk-averse. This indeed seems to have been the case and for a long time. (In Figure , you saw that riskier asset classes had higher average rates of return from 2005-2021, too.)

Let's return to the decompositions from Section : You learned that in a perfect and risk-neutral world, stated rates of return consist of a term premium and a default premium. On average, the default premium is zero, and the expected rate of return would just be the term premium.0.5]subsect:deconstructdefaultpremiumTerm and default premiums All projects with payoffs at the same time offered the same expected rate of return. This is no longer true if investors are risk-averse.

There can be only one risk-free rate in a perfect market for any given time horizon. A risk-free corporate bond would not have to offer either a default or a risk premium, either. Yet with the potential of default comes risk. For example, let's say we start a mutual fund that invests $100 in stocks. To fund it, we issue a bond for $90. With a risk-free rate of 0% and a stock market that has a mean rate of return of 12% and risk of 16.7%, the bond has to promise to pay more than the risk-free rate. In fact, involved calculations can show that the bond has to promise to pay about 2.7%. Its expected rate of return will only be about 1.6%, however. (Its standard-deviation will be about 4%; its market-beta will be about 0.13.)

In this case, the expected rate of return on stocks offers an extra risk premium.

(Omitted eq)

You need to be careful in distinguishing between the default premium and the risk premium. The default premium is zero on average. It is positive only in terms of the promise (i.e., what you get when everything turns out well). If you are risk-averse, it is only the risk premium that increases your expected rate of return in the long run.

Unfortunately, the expected rate of return (or, equivalently, the risk premium) is never posted in the real world. It is always only the stated rate of return that is publicly posted.

In the real world, the premium that investment-grade corporate bonds quote above equivalent Treasuries is due both to the risk premium and the default premium (and perhaps some other imperfect premiums discussed in later chapters). It's simple: Corporate bonds just won't always pay as much as they promise. For corporate projects and equity shares, the risk premium is considerably larger.

Begin Important

Never forget:

- You should think of benchmarks in terms of expected rates of return. If you use historical average returns, you usually effectively assume that these averages are also representative of future expected rates of return.
- The expected rate of return in the denominator of the NPV formula is not a stated (i.e. promised or quoted) return, because it does not include a default premium.
- The variability of payments — including the possibility of default — must be handled in the NPV numerator (through the expected cash flows), and not in the NPV denominator (through the expected rate of return). (It is only covariation with the market or other assets that is handled in the denominator.)
- You cannot discount promised cash flows with promised rates of return elsewhere in the economy. One promise is not the same as another.

End Important

9.7 "The Benchmarking Method" and More Benchmarks

Stocks and Treasury bonds are not the only two benchmark assets that you can use. Depending on the project to be valued, managers often use other benchmarks, too. For example, instead of the risk-free Treasury, some corporate managers use bonds that are similar to what they can issue themselves — e.g., investment-grade or junk bonds, mortgage bonds, collateralized bonds, prime borrower bank financing, etc. In all these cases, it is important to consider that publicly quoted comparables always include default premia, and that your own firm will also have to offer default premia. This is so important that I will repeat the repeat: I beseech you never to confuse expected rates of return with promised rates of return. Just because a non-investment grade bond offers 2-5% above the risk-free rate does not mean that it expects to pay off 2-5% above the risk-free rate. Future defaults will erode the difference. Expected rates of return are much more alike.

Even within the small sector of the economy that are equity fund managers, many different stock benchmarks are in use. They tend to choose the most relevant to your own projects. This can, but need not be, the VFIAX (S&P 500). Depending on the goal of the fund, the benchmark can also be the value portfolio (VVIAX) for older firms, the growth portfolio (VWUAX) for startup firms, the small-cap portfolio (VSMAX) for smaller firms, the REIT portfolio (VGSLX) for real-estate portfolios, your industry portfolios, momentum portfolios, profitability portfolios, and so on. Private equity, venture capital, and hedge funds often have their own set of benchmarks, too.

Some corporate managers can benchmark their performance to some underlying commodities. For example, the expected rate of return on Exxon can be closely linked to the price of oil. If the appropriate expected rate of return on oil is, say, 20%, then Exxon's oil storage operations should similarly yield an expected rate of return of 20%. Working with benchmark portfolios also allows firms to hedge out their risk. For example, Phillips 66 (PSX) is a gas station operator. If it wants its returns to be based more on its capabilities running gas stations and less on the world price of oil, it can use oil futures contracts to reduce its net exposure to oil price variability. If Phillips hedges out its entire relevant risk, it would no longer even need to estimate the expected rate of return on its benchmark any longer. It has become irrelevant.

In principle, for firms with unhedged risk exposures, it always works the same way: as a corporate manager, first you assess the expected rate of return on some underlying benchmark portfolios. Then you assess the expected rates of return on your own internal investment opportunities. How similar are your projects and to which benchmark? Can your projects be viewed as combinations of your benchmarks? If your opportunities beat the publicly available alternatives in risk-reward, you should invest. Otherwise, you should return the funds to your investors.

Our valuation method is essentially just comparing your opportunities to the price at which your investors can buy similar ones elsewhere. This is also why such a method is called an asset-pricing model, even though the model is then phrased in terms of expected returns. Expected returns are *never* posted. Only prices are. But all the economic insights just amount to this: "opportunities with similar characteristics — and in particular, risk characteristics — should offer similar expected rates of return."

Again, let it sink in: as a corporate manager with exposure to a particular risk, you need an appropriate expected rate of return — an opportunity cost of capital — as the denominator in the NPV formula. If your project offers a lower expected return than what your investors can earn elsewhere in projects with similar risk, then you should not put your investors' money into your project but instead return their money to them. They can then buy those better projects instead. If, on the other hand, your project offers more expected return, then you should go ahead and invest their money into your project.

Summary

This chapter covered the following major points:

- For each project cash flow, you need to estimate the expected rate of return on equivalent benchmark investments. This is the "opportunity cost of capital" that corporations can use as their costs of capital in the terms of the NPV formula.
- The most important benchmarks are the expected rate of return to low-risk fixed income assets (such as Treasury bonds) and to high-risk equity assets (such as the S&P 500).
- For r_F, you should use bonds that match the timing of your project's cash flows. Thus, cash flows farther in the future usually have higher opportunity costs of capital.
- It is difficult to estimate the equity premium. There is no clear consensus on what it should be or how to estimate it best. Reasonable estimates for the equity premium $(E(r_M) - r_F)$ can range from about 1%/year

(geometric) for long-term payoffs to 8%/year (arithmetic) for short-term payoffs. Estimates of about 1-3% seem common for most long-term project cash flows.

- Investors care about geometric rates of return, not arithmetic rates of return. When projects have different risk, the two averages can be very different.
- The correct benchmarks adjust properly for term and risk, but when based on historical estimates require judgment about what historical sample period is most representative of the future.
- Both bond and stock benchmarks have expected rates of return that are due to a number of factors — first and foremost, term and risk. So do other benchmark portfolios and assets. It does not have to be bonds and stocks. By choosing better benchmarks that are more similar to their own projects, managers can often obtain better estimates for their costs of capital.
