10

The Capital Asset Pricing Model

What cost of capital does your project have to offer? The last chapter explained how you can determine the expected rate of return if your project is like other assets — such as Treasuries, the stock market, or some other traded financial assets.

But what about projects that are not easily decomposed into obvious combinations? How would you judge how much of other assets you would need to mimic your project? And which other assets should you then choose as your benchmarking portfolios?

This is the domain of the Capital Asset Pricing Model (CAPM). It hypothesizes that the only two important aspects are the rate of return on the risk-free rate and the stock-market. If you are willing to buy into its premise, the CAPM gives you a beautifully simple and intuitive formula that relates how much reward your investment project has to offer in order to compensate your investors for its risk. The risk is the market beta. By assumption, the formula works for *any* kind of project. You can then use its costs of capital in your NPV calculations.

We will first briefly review what you already know. Then you will learn all about the CAPM. And you will get to apply it.

And then I will have to deliver the bad news: although the CAPM remains the dominant model in practice (and in job interviews), not only is it often poorly applied, but it has also performed so poorly (and increasingly poorly over the last few decades) that finance researchers no longer believe that it is a good model for predicting which stocks have high and which stocks have low expected rates of return. Its empirical validity was miserable even under the most sympathetic assumptions. On average, you would have been better off simply ignoring it — or even assuming the opposite. It's our collective conundrum. Welcome to the club.

Wait a little — I will put it all in perspective for you.

Funny Youtube Video from your's truly: "CAPM crunched"

10.1 What You Already Know

We are still going after finding the answer to the same central question — what should we, as corporate managers, consider to be a good (opportunity) cost of capital for the projects we are considering deploying? We already know one good measure from the previous chapter: We can compare the price of our own project to those of benchmark projects that are already traded in competitive public markets. But there can be a problem. It is not always easy to find good obvious benchmarks for all our projects. We would love a model that is a bit less specific and a bit more generic in its prescriptions. It should tell us about the cost of capital even if we do not have a good benchmark. That is, we would love a model that gave us the equivalent of a universal benchmark.

To get such a model, we have to push on some assumptions. We will assume that investors are risk-averse. From the previous chapters, you know that if expected rates of return were the same, investors would prefer short-term low-risk project cash flows (like overnight Treasuries) to long-term unsafe project cash flows (like the stock market). We also assume that investors are smart, so they diversify their portfolios to reduce their risk exposures. The types of risk that investors consider noxious can only be the parts that they cannot wash out by diversification and that remain left over even when all assets are just tiny parts of their large overall portfolios.

We also continue to lean on our perfect-market assumptions just as we have in recent chapters. Thus, not being dummies, investors collectively snatch up the best projects — those that have low risk and high expected rates of return. In fact, anyone contemplating selling a project with more reward than it deserves would be needlessly giving away bargains that would attract a gazillion bidders. Anyone contemplating selling projects with too unfavorable risk contributions for their rewards would not receive a single offer. There is really only one correct choice of price. Consequently, what investors purchase in the real world at the correct prices must be subject to some trade-off: Projects that drive up investors' overall portfolio risk must offer higher expected rates of return.

These simplifications will leave you with a consistent conceptual framework: External investors dislike overall portfolio risk (standard deviation) and like reward (expected returns). They care about their overall financial investments portfolio. They are smart, so they are diversified. With a few more assumptions, this means most likely that they are holding something akin to the overall stock-market portfolio for risk and something like risk-free Treasuries for safety.

Given this CAPM conceptual framework about how external investors want to invest, we can skip back to the perspective of the corporate manager. We know that our assumed model says investors will judge our project in terms of how it is "more like stocks" rather than how it is "more like Treasuries." The measure that gives us the relative type is the "market beta." The CAPM puts this all together — the risk-free rate of return (like Treasuries), the expected rate of return on risky assets (like stocks), and the market-beta that measures how much of each.

In the "risk-reward game," we measure investors' reward are their portfolios' expected rate of return. Investors' risk is their *overall* portfolio risk, *not* your project's *own* standard-deviation risk. Your own project's contribution to investors' overall portfolio risk is then best measured by the market beta of your project.

A good way to think of beta is as a measure of your project's "noxiousness." A project that decreases in value when the market decreases in value (and increases when it increases) has a positive market beta. It's noxious — investors don't like it. A project that increases in value when the market decreases in value, and vice-versa, has

a negative market beta. It's less noxious — investors like it more. That is, projects with lower market betas help investors (who already otherwise hold market-like portfolios) suffer less overall portfolio risk.

Another way to think of non-synchronicity is that oranges in winter are more valuable than oranges in summer (when oranges are plentiful). Non-synchronicity is a desirable attribute, and beta is a measure of synchronicity.

10.2 The Capital Asset Pricing Model (CAPM)

The capital asset pricing model (CAPM) gives you a formula that relates the appropriate expected rate of return (cost of capital) for your project — if you feed it your project's single relevant risk characteristics (the market beta), the risk-free rate of return and the expected return on the stock-market (or equivalently the equity premium). Note that you also had to provide the expected rate of return on the stock-market in the previous benchmarking chapter. It is the most uncertain and thus difficult input into the CAPM.

The CAPM states that an investment's cost of capital is lower when it offers better diversification benefits for investors who hold the overall market portfolio — less required reward for less risk contribution. Market beta is the model's measure of risk contribution. Projects contributing more risk (market beta) require higher expected rates of return (for you to want them on behalf of your investors); projects contributing less risk require lower expected rates of return. According to the CAPM, nothing but the risk-free rate, the expected equity premium, and the market beta matters. No other financial assets need to be investigated to judge your project.

Begin Important

To estimate the required expected rate of return for a project or firm — that is, the cost of capital — according to the CAPM, you need three inputs:

1. The risk-free rate of return, r_F.
2. The expected rate of return on the overall market, $E(r_M)$ — or, alternatively, the equity premium $E(r_M) - r_F$.
3. The project's beta with respect to the market, β_i.

The CAPM formula is

(Omitted eq)

where i is the name of your project and $E(r_i)$ is your project's expected rate of return. All model inputs are forward-looking: the risk-free rate, the equity premium, and the market beta of the asset.

You need to memorize the CAPM formula. End Important

The CAPM specifically ignores the stand-alone risk (standard deviation) of your project. That is, investors do not care about your projects' variance, because they are smart enough to diversify away any idiosyncratic risk. Investors care only about your project's market betas, because it is betas that measure the component of risk that your project contributes and that investors holding the wide market portfolio would not have diversified away.

On a pragmatic level, the CAPM is seductive. It limits your attention to just two benchmark assets. It gives you a coherent universal measure of where projects lie on the spectrum between stocks and bonds. More market beta

means "more like stocks," and thus higher expected rates of return ("just like stocks"). Less market beta (down to zero) means "more like bonds," and thus lower expected rates of return ("just like bonds").

Without going into detail, economists also love a deep "economic equilibrium model" justification for the CAPM that I will largely spare you. In this view, financial markets are perfect, each and every investor faces the same tradeoffs and uses the model, and each and every asset is priced by it. When all the assumptions are satisfied, it implies mathematically that the CAPM must hold. Necessarily, there could then not be any benchmarks other than the risk-free rate and the stock market, and the only valid measure of risk would be the market beta. This CAPM justification, with its stringent assumptions, is too orthodox and simply not realistic.

Unfortunately, more important than philosophy, the empirical data soundly rejects the CAPM, as I will explain below in more detail. For now, let me just say that you must still study the CAPM not only because it is conceptually important but also because every finance dinosaur in the real world is using it — and, more than likely, it will growl questions about the CAPM in your job interview.

(Omitted solvenow)

The Security Market Line (SML)

Let's first use the CAPM formula as a recipe. If you believe that the risk-free rate is 3% and the expected rate of return on the market is 8%, then the CAPM states that

(Omitted eq)

Therefore, a project with a beta of 0.5 should have a cost of capital of $3\% + 5\% \cdot 0.5 = 5.5\%$, and a project with a beta of 2.0 should have a cost of capital of $3\% + 5\% \cdot 2.0 = 13\%$. The CAPM gives the opportunity cost for your investors' capital: If the project with the beta of 2.0 cannot earn this expected rate of return of 13%, you should not take this project and instead return the money to your investors. Your project would add too much risk for its reward. Your investors have better opportunities elsewhere.

(Omitted fig)

The CAPM formula is often graphed as the security market line (SML), which shows the relationship between the expected rate of return of a project and its beta. Figure **??** draws a model-perfect security market line for seven assets. Each investment asset (such as a stock or a project) is a point in this coordinate system. Because all assets in our example properly follow the CAPM formula, they must lie on a straight line. The SML is just the graphical representation of the CAPM formula. The slope of this line (the reward per unit of beta) is the equity premium, $\mathsf{E}(r_M) - r_F$, and the intercept is the risk-free rate, r_F.

(Omitted fig)

Alas, in the real world, even if the CAPM holds, you would not have the data to draw Figure **??**. The reason is that you do not know true expected returns and true expected market betas. Figure **??** plots a version where you have to rely only on what most investors have and rely on — observable historical data averages. Thus you can only fit an

"estimated security market line," not the "true security market line." And you have to hope that your historical data has provided good, unbiased estimates of the true forward-looking market beta and true forward-looking expected rates of return. (Both are big assumptions!) If the fitted line on the data looks straight, you would not immediately throw out the CAPM. In any case, any workable version of the CAPM in real life can only state that there should roughly be a linear relationship between the data-estimated market betas and the data-estimated expected rates of return, just as drawn in Figure **??**.

(Omitted solvenow)

The CAPM in the Present Value Formula

If you take the CAPM at face value, it gives you a good denominator for the NPV formula, the opportunity cost of capital, $E(r)$:

(Omitted eq)

(Subscripts here are time rather than firms.) Together, the CAPM and the NPV formulas tell you that cash flows that correlate more with the overall market are of less value to your investors and therefore require higher expected rates of return ($E(r)$) in order to pass muster (well, to pass the hurdle rate, which is determined by the alternative opportunities that your model presumes your investors have). 3]pg:asset-pricing-model-names-bAsset-Pricing Model

The CAPM is called an asset-pricing model, even though it is most often expressed in terms of a required expected rate of return rather than in terms of an appropriate asset price. Fortunately, the two are equivalent — you can always work with the CAPM return first, and then discount the expected cash flow into an appropriate price. A given expected rate of return implies a given price. (If you do not know the fair price, you will have to take two aspirin and work with a more difficult version of the CAPM formula. It is called certainty equivalence [CEV] and is explained in the companion online chapter.)

Equity and Asset Betas

As in Section , it is important that you always distinguish between asset costs of capital and equity costs of capital. And, fortunately, just as you can take weighted averages of expected rates of return, you can take weighted averages of betas. Thus, whatever worked in Section with the overall costs of capital also works here with market betas. Done. You can skip the rest of this section, or endure a few more examples.

For example, assume that the risk-free rate is 4% and the equity premium is 5%. You own a $100 million project with an asset beta of 2.0 that you can finance with $20 million of risk-free debt. Truly risk-free debt always has a beta of 0. To find your equity beta, write down the formula for your asset beta (firm beta):

(Omitted eq)

Solve this to find that your market beta of equity is 2.5. It is this market beta of equity that you would find reported on sites like YAHOO! FINANCE. You would not want to base your hurdle rate for your entire firm's typical average

project on your equity beta:1]sect:blendedcostsofcapitalTypical, average, and marginal betas Such a mistake would recommend you use a hurdle rate of $\mathsf{E}(r_i) = r_F + [\mathsf{E}(r_M) - r_F] \cdot \beta_i = 4\% + 5\% \cdot 2.5 = 16.5\%$. This would be too high. Instead, you should require your average projects to return $\mathsf{E}(r_i) = 4\% + 5\% \cdot 2.0 = 14\%$.

	20% Debt	80% Equity	100% Project
Beta	0.0	2.5	2.0
⇒ Cost of Capital	4%	16.5%	14.0%

In both cases, the capitalization-weighted average of debt and equity rate of return is always the overall project asset rate of return.

Conversely, if your project is private but the potential future owners are well-diversified, you may have to find its hurdle rate by looking at public comparables. Let's presume you find a similarly sized firm with a similar business that YAHOO!FINANCE lists with a beta of 2, or perhaps better yet, the firm's industry. Remember that financial websites always list only the equity beta. The CAPM tells you that the expected rate of return on the equity is $4\% + 5\% \cdot 2 = 14\%$.

However, even in a perfect capital market, this is not necessarily the hurdle rate for your project. When you look further on YAHOO!FINANCE, you may notice that your comparable is itself also financed with 90% debt and 10% equity. (Companies in financial distress and banks often have leverage that is this unusually high. If the comparable had very little debt, then a debt beta of 0 might have been a good assumption, but unfortunately, in this case it is not.) Corporate debt rarely has good historical return data that would allow you to estimate a debt beta. Consequently, practitioners often estimate the expected rate of return on debt via debt comparables based on the credit rating. Say your comparable's debt is rated BB and say that BB bonds have offered *expected* rates of return of 100 basis points above the Treasury. (This might be 200 basis points *quoted* above the Treasury). With the Treasury standing at 4%, you would estimate the comparable's cost of capital on debt to be 5%. The rest is easy. The expected rate of return on your project should be

(Omitted eq)

This would make a good hurdle rate estimate for your project.

1]sect:blendedcostsofcapitalTypical, average, and marginal betas

Does Risk Reduction Create Value?

In the 1960s and 1970s, many firms became conglomerate[conglomerates], that is, companies with widely diversified and often unrelated holdings. Can firms add value through such diversification? The answer is "usually no." Diversification indeed reduces the standard deviation of the company's rate of return (diversified companies are less risky). Yet, in a perfect market, your investors can just as well diversify risk for themselves. They don't need the firm to do it for them. This is a more general perfect-market insight than the details of what follows. Again: if investors can do it without the firm, the firm cannot add value by doing it for them.

As in the previous section, we can elaborate on this in the context of the CAPM. However, the basic idea should hold in any reasonable framework, e.g., if projects have different cash flow horizons and thus different term-based

costs of capital. Thus, you can consider it "done" and you can skip this section, too, if you already fully understand this. Otherwise, endure the example below.

If your $900 million firm ABC with a beta of 2 and a risk of 20% is planning to take over the $100 million firm DEF (e.g., with a beta of 1 and also a risk of 20%), the resulting firm is worth $1 billion. ABC + DEF indeed has an idiosyncratic risk lower than 20% if the two firms are not perfectly correlated, but your investors (or a mutual fund) could just have held 90% of their portfolios in ABC and 10% in DEF and thereby achieved the very same diversification benefits. If anything, a merger takes away your investors' freedom: They no longer have the ability to buy, say, 50% of their portfolios in ABC and 50% in DEF. (In a perfect CAPM world, this would not matter, because any investor could simply repurchase the firm and undo this.) The CAPM makes it explicit that the cost of capital does not change unduly. Say both firms follow the CAPM pricing formula, and say that the risk-free rate is 3% and the equity premium is 5%,

(Omitted eq)

and

(Omitted eq)

The newly formed company will have an expected rate of return (cost of capital) of

(Omitted eq)

(This is not a property of the CAPM but a property of expected rates of return. It would also work with the simpler benchmarking approach of the previous chapter.)

For the combined conglomerate, the market beta is now

(Omitted eq)

The merged company will still follow the CAPM,

(Omitted eq)

Its cost of capital has not unduly increased or declined. In an ideal [CAPM] world, no value has been added or destroyed — even though ABC + DEF will have a risk lower than the 20% per annum that its two constituents had.

(Diversification could add value if risk-averse investors are too foolish to diversify themselves — a possibility which the CAPM has assumed away. With foolish investors, the conglomeration could make these investors' portfolios more diversified and less risky, making them better off. However, although some such an investor may exist, it seems implausible that they are plentiful and important enough to determine pricing in the world.)

Deconstructing Quoted Rates of Return

As in Section , the CAPM provides just an underpinning to the expected rate of return, not the quoted rate of return. (In fact, the example I constructed in Section was surreptitiously based off the CAPM. I even told you the market-beta in the text as an aside. With a market-beta of 0.13, the expected rate of return on the risky bond was $0.13 \cdot 12\% \approx 1.58\%$.)

Short-Term and Long-Term Projects?

Although the CAPM formally recognizes only one SML in theory, we use different risk-free rates for different project horizons in practice. Thus, short-term projects would have lower costs of capital than long-term projects. For example, you might assess, say, a 3% equity premium and use prevailing rates for the Treasuries — say a one-year Treasury at 1% and a thirty-year Treasury of 2% — and thus assess

		Your Project β_i					
		–2	–1	0	1	2	3
Short-Term Projects	$\mathsf{E}(r_i) = 1\% + \beta \times 3\%$	–5%	–2%	1%	4%	7%	10%
Long-Term Projects	$\mathsf{E}(r_i) = 2\% + \beta \times 3\%$	–4%	–1%	2%	5%	8%	11%

Recall that we are not sure whether we should use the same equity premium (here 3%) for both near and far project cash flows. It is due to ignorance that we typically use the same equity-premium estimate regardless of an equity premium that depends on the term.

(Omitted solvenow)

10.3 Estimating the Extra Input: Market Beta

We already discussed estimating the risk-free rate and equity premium in the previous chapter. We also discussed beta estimation in the chapter before. Nevertheless, because beta is the only novel aspect relative to benchmarking, let's discuss it a little more. Beta tells you how the rate of return of your project fluctuates with that of the overall market. Unlike the previous two inputs, which are the same for every project in the economy, the beta input depends on your own specific project characteristics: Different projects have different betas.

Just as with the risk-free rate and the expected rate of return on the stock-market (or equivalently, the equity premium) in Chapter , investors are really interested in the *future* market betas of your projects and not in their historical market betas. No one really cares about the past for its own sake.0.5]subsect:history-repeatsWill history

Nerdnote: If the CAPM truly held, long-term bonds would have higher expected rates of return than short-term bonds, and this could be explained exactly by their positive market beta. Alas, researchers and text-book writers do not know why long-term bonds have had negative market betas for a few decades now. However, be aware that applying the CAPM to long-term bonds would so obviously contradict reality that few bond investors have been tempted to use it in this context. Instead, bond investors heavily use adjusted yield-curve estimates.

repeat itself? But as usual, you often have no choice other than to rely on estimates, and these are usually based largely on statistical analysis of historical data.

First, the good news. Betas can be forecast much more accurately than average rates of return (which we use as estimates of expected rates of return). Use daily rates of return and shrink them towards 1, and you are good, especially for more diversified portfolios. (Or look them up at https://www.ivo-welch.info/research/betas/.) Now the bad news. The further ahead you need to estimate, the less reliable betas become. See, betas drift themselves slowly over horizons of many years. It's like shooting at a moving target — not as good as shooting at a fixed target, but if the target is moving slowly enough, it shouldn't end up too bad. *C'est la vie.*

Market Beta Estimation Based on Historical Data

The basic mechanics of finding the historical market beta for a project with historical rates of return is easy. You run a market-model regression. The independent variable is the rate of return on the stock market. The dependent variable is the rate of return on your project. It is also good practice to subtract the risk-free T-bill rate from both your project's and the stock market's rates of return. Any statistical software package (and common computer spreadsheet programs like Excel or Openoffice) can readily calculate the coefficients a and b in the market-model regression:

(Omitted eq)

The slope b is the market beta. It's a good thing that we use b as a symbol instead of β, because the b that the regression spits out is only an estimate of a true beta (β), and not the true and unknowable beta itself.

This is only the basics. To get a better forward-looking market beta estimate, you should do the following:

1. Use daily stock returns, not monthly stock returns.
2. Use about two years' worth of data. Between one and five years of data will do.
3. For memory sake, I earlier recommended a "shrink" of 50% towards 1.0. A more exact shrink of your first-pass market beta would be 30-40%, depending on the timing of the cash flow that you intend to use it on:

$<$ 1 Year	$(1-0.3)\times b + 0.3\times 1$
$>$ 5 Years	$(1-0.4)\times b + 0.4\times 1$

   This 0.3 (or 0.4) factor is used partly because it reduces the impact of historical outliers, and partly because true market betas drift over long horizons. If you want, you can shrink beta by another 10% if your project and firm are small (to compensate for additional estimation "noise").

   For example, if your statistical software gives you a first-pass market-beta estimate for your project of 2.0, and you want to estimate a CAPM cost of capital for a project cash flow in 1 year, then use $(1–0.3)\times 2.0+0.3 = 1.7$. If you want to estimate it for a project cash flow in 10 years, use 1.6. If your first-pass estimate is -1.0, and the cash flow is in 1 year, use $(1-0.3)\times -1.0 + 0.3 = -0.4$.

Nerdnote: If you want to estimate future market-beta even better, then shrink not towards 1, but towards a smaller constant, like 0.6-0.8, if your firm and project are small.

It does not matter much which particular stock-market index you use as your independent variable. The S&P 500 with or without dividends is fine. There are also other more sophisticated methods, but the above three guidelines cover the most important basics. It is unlikely that you can improve much on them. These market betas are as good as simple estimates are going to get.

In practice, you may encounter two common estimation practices that dramatically worsen the quality of estimated market betas. So let me warn you:

1. If you have good daily data, do not estimate market beta with monthly return data. (And if you have no choice [as, for example, with hedge funds, which report rates of return only monthly], then shrink more — think 50-60%, not 30-40%.)
2. If you have your firm's own stock returns, do not use industry portfolio returns as stand-ins for your firm. Although industry betas move less than stock-specific betas and thus seem appealing, in reality industry betas (or peer betas) are much worse predictors for stocks than the stocks' own market betas.

If you see either practice, tell the dinosaur using them that the mammals are taking over and they'd better evolve and adapt!

Market Beta Estimation Based on Theoretical Consideration

As a corporate manager, you are rarely interested in the market beta of an industry or even a stock. Usually, you are interested in the market beta of a potential project. Sometimes, your firm may not even be publicly traded, so you would not have any historical price data to begin with. In this case, corporate CAPM users sometimes rely on economic intuition rather than historical statistics. To see the logic, rearrange the CAPM formula. Now, do you think your project cash flows and its future project values (which are influenced by changes in the economy) are likely to move more or less with the overall stock market (and possibly the overall economy)?

(Omitted eq)

The right side of this formula helps translate your intuition into a beta estimate. What rate of return (above the risk-free rate) will your project have if the market were to have +10% or –10% rate of return (above the risk-free rate)? Clearly, such guesswork is difficult and error-prone — but it can occasionally provide a market-beta estimate when no other is available. But be aware that such estimates are almost always poor.

If you do not believe me that your estimate is going to be so bad that you may as well just go back to the peer benchmarks from Chapter , then I dare you to try. Randomly pick five stocks from *YAHOO!*FINANCE. Do not peek at their market betas. Explain to me what they should be, and then check your claims against their actual market betas. If you can accurately assess which market betas are farther from 1.0, then you are a better intuitive economist than I am. In fact, I have almost no economic intuition as to why entire asset classes, such as long-term bonds, have had negative market betas over the last 20 years and positive market betas before then.

Moreover, *please* stand back and think for a moment what you are really doing here. If you are dealing with a new project that has never seen the light of day and that has no historical data, would you even want to use the CAPM? And are you a fully diversified unconflicted owner-investor who cares only about market-risk and not about

idiosyncratic project risk, and who has access to a perfectly competitive capital market? There are whole sections of the economy that cannot meet these criteria. In particular, if you are an entrepreneur who satisfies them, I would really like to meet you — I have never met such an entrepreneur. (Besides, are you already so convinced by the beauty and logic of the CAPM that you are willing to believe that it is a good description of the real world — even though I have shown you zero empirical evidence about the model itself up to this point?)

(Omitted solvenow)

10.4 Neutralizing Equity-Premium Uncertainty

Do you recall my claim that the risk-free rate and the equity premium were the two most important numbers in finance, regardless of whether you are using the CAPM or not? Well, you also want to know the market-beta for the same reason. It is an extremely useful number, too.

It is very easy to short .5]subsect:securities-microstruct:shortingShortingthe stock market (e.g., using an S&P 500 future or ETF). This allows you to "innoculate" or "hedge" your project against overall stock-market risk. Just short the right amount of stock, which is exactly the ratio that market beta gives you..5]sect:stocksbondscashExchange-Traded Fund For example, if you have \$100 million invested in an asset with a market beta of 3, you can short $3 \cdot \$100 = \$300$ million in the market and thereby reduce your market risk to zero. If the stock market happens to go down by 1%, you would expect (a) your project to go down by 3% but (b) your hedge to go up by the same $3 \cdot 1\%$. The CAPM formula even suggests that your equity-premium estimate is now irrelevant.

However, a short market position can also increase the variance of your project outcomes: You may end up in a scenario in which your own project underperforms and the stock market outperforms. You may even go bankrupt because of it. Your project's idiosyncratic-risk component and your errors in estimating betas now become more important. This is not a problem if your project owners are highly diversified, and your particular project is just a tiny fraction of their wealth that they don't care a great deal about. Yet, it is a problem if they are not; or if you, as the corporate manager, care about your one specific project a great deal (or if there are bankruptcy costs, as you will learn in Chapter).

Putting this together, from your perspective as the CEO of one small company in a large market, you can render a degenerate version of the CAPM formula to be nearly true by definition. If you are shorting the correct amount of stock market, it won't matter whether you are overestimating or underestimating the equity premium. The limits to this strategy are your estimation uncertainty about beta and your idiosyncratic risk tolerance. In the real world, a full short may neither be possible nor desirable. If you do not immunize your company against market risk, then it matters to you what the equity-risk premium is — and whether the CAPM is right in the first place.

You may object that you would not want to short the stock market — betting against the market was historically not a smart maneuver. But, as a CFO, do you really know better whether you should be long or short the stock market? If it is fairly priced, so be it. Leave this choice to your investors. If they want to bet on or against the overall stock market, they do not need you to do it for them. You are only "abusing" the insights of market-betas to avoid or at least reduce the consequences of your ignorance about your project's and the overall market's cost of capital.

Nerdnote: The strategy of neutralizing the market could work only for a small number of companies and investors. If it became too widespread, it would change the investment opportunity set.

10.5 Is the CAPM the Right Model?

The CAPM Assumptions Are Not Innocuous

Although the CAPM edifice is reasonable, it does not mean that this edifice "obviously" holds. The CAPM model leans a lot more on the perfect-market assumptions (and then some) than our earlier chapters did. Are most financial markets really so perfect? Do most investors really hold diversified stock-market portfolios? Do they really care *only* about risk and reward in their financial-asset portfolios and nothing else?

Stand back for a moment. How can the CAPM edifice crumble? Consider the following examples:

- If you own a house, chances are that much of your current wealth is invested in the equity portion of your house, and you are not as diversified as you should be. You should then try to find stocks that reduce your house risk exposure, not stocks that reduce your financial market risk exposure. You should like stocks that go up when your house value goes down.
- If you are under 40 years of age, chances are that much of your lifetime wealth is in your human capital. It is not diversified. And only *you* can invest easily in *your* education: No one else can. *You* need to hedge *your* career, not mine. *I* need to hedge *my* career, not your's. You should like stocks that go up when the value of your expertise goes down.
- If you are a tech engineer and work in Silicon Valley, you should short technology stocks as a hedge against their tanking. Conversely, you should not mind losing in your financial portfolio when technology stocks boom (and you end up rich from your employer's stock options, anyway). Yet many engineers in Silicon Valley are so irrationally overconfident, excited, and/or convinced of technology and their (stock-picking) abilities that they end up buying mostly technology stocks for their portfolios, instead. They double up rather than hedge. It may have worked for them so far, but just wait…
- Do firms really live in near-perfect capital markets? Entrepreneurs often need to scrape together whatever capital they can. If they cannot easily find many capital providers, they may have to pay much higher costs of capital than suggested by the CAPM. And they may be forced to invest most of their own wealth — to the point of bankrupting themselves if their projects fail.
- Entrepreneurs are notorious for staking their entire life's savings on their startups. They are hardly ever diversified and usually short of funds.

So, even though the theoretical CAPM assumptions are simple enough to be appealing, their applicability is actually quite narrow — the CAPM considers a scenario in which all investors do not care about anything but the risk and return in the public financial markets, and they all have (largely) the same investments and investment opportunities. Don't think the CAPM *has* to be true just because it seems reasonable.

What if every investor chose portfolios for personal reasons and not with the same perspective — some looking to hedge their houses, others their job, others their industry, others their product's failure? Then it may well be that some assets offer higher or lower expected rates of return than suggested by the CAPM — the CAPM would not hold. In this case, corporate managers really should not rely on the CAPM. Instead, they should stick to a more holistic approach (e.g., higher hurdle rates) or the less ambitious peer benchmarking approach from Chapter .

The Empirical Evidence Firmly Rejects the CAPM Among Stocks

Sadly (for finance professors), it indeed turns out to be true that the empirical evidence does not support the CAPM. Thus the best advice in real life is not to use the CAPM. It does not work. Use the benchmarking approach instead.

Huh? WTH? Did you really read me right?

Yes you did. Don't use the CAPM. The empirical evidence firmly contradicts it. This means that the common corporate use of the CAPM — to obtain hurdle rates for projects in the NPV formula — is based on no more than wishful thinking. It is contrary to empirical evidence.

(Omitted fig)

I could summarize decades worth of academic work, but understanding the principal problem does not require a higher degree. It only requires basic common sense and open eyes. At its most basic implication, the CAPM says that high-beta stocks should offer higher average rates of return than low-beta stocks.

To show that this was not so, I downloaded the market-beta portfolios (publicly posted by the famous researcher Ken French). Ken sorts stocks by market-betas into quintile portfolios (here value-weighted) and then tracks the return performance of these portfolios over the following year. Figure **??** plots them. The realized ex-post market-beta was 0.75 for the low-beta quintile and 1.47 for the high-beta quintile. Market-betas were stable. Because stocks delivered almost 12% per annum from 1964 to 2021, high-beta stocks should have delivered about $(1.47 - 0.75) \cdot 12\% \approx 10\%$ more per annum than low-beta stocks. Instead, they delivered *less*.

This is a robust finding. Within the asset class of stocks, the empirical evidence shows that higher market-beta stocks did not have higher average rates of return in the past than lower market-beta stocks. It's not just the more subtle CAPM problem that benchmarks other than the equity market *also* matter; it's the very unsubtle problem that beta itself does not seem to matter. Yikes 😬! Yikes 😬! Yikes 😬!

Ergo, using a market-beta of 1 on every project (the ultimate shrinkage) would not have done any harm predicting future average returns. Standing in the past, if you had wished to predict future expected rates of return, you would have done better assuming that every stock had a market-beta of 1. This would have been the same as benchmarking all stocks indiscriminately to the VFIAX (S&P 500) index. Standing here today in 2022, I think that this is still good advice about future average stock returns.

Begin Important

Although one can predict the market-beta of individual stocks quite well, you should use a good estimate of market-beta only for diversification and hedging purposes (neutralizing the equity premium).

For estimating differential expected rates of return (costs of capital) within the asset class of publicly-traded stocks, you should skip the beta estimation and use your estimate of the expected rate of return on the market instead. Another way of saying this is to use the CAPM formula with the (obviously incorrect) input as if every market beta input was 1.0. **End Important**

➤ In Defense of the Use of the CAPM

If the evidence is against the CAPM, then why do we finance professors torture you with it? We may indeed have sadistic streaks (as our captive PhD students can testify), but this is not why. This "why" is much easier to answer than how stocks are priced in the real world or what the best estimate of the appropriate hurdle rates for your project should be.

***Across* asset classes:** The CAPM was not rejected across asset classes. Stocks had higher average rates of return than bonds which had higher average rates of return than bills. In this sense, some high-beta assets offered higher average returns than some low-beta assets.

Caveat: Although this is evidence that investors have indeed been rewarded for taking on more beta-risk, almost any measure of risk would have also predicted this to be the case. Declaring victory for the CAPM based on this rough evidence is overreaching.

Impeccable intuition: The CAPM shines through its simplicity and focus on diversification. It gets executives away from the false notion that public investors care about the idiosyncratic risk of projects that they can diversify away. Thus, corporate diversification into a conglomerate for its own sake can reduce its own risk but not market risk. It cannot add value. Investors can diversify themselves. They don't need the firm to do it for them. The CAPM logic reinforces this important point.

Strong Belief: Many instructors and practitioners find the CAPM to be so plausible that they are willing to live with the "absence of CAPM evidence." They do not take this absence to mean "evidence of CAPM absence." Thus, they adopt the CAPM based on faith and not on evidence — actually, more like contrary to the evidence about the basic association between market-betas and average rates of return shown above. If you do this, you must be aware that this is what you are doing.

Stand-in for Expected Cash Flow Default: The CAPM often assigns higher costs of capital to projects that are more likely to fail. If you have not fully adjusted your expected cash flow estimates downwards to account for the possibility of for failure (a common human as well as managerial error), the CAPM cost of capital may help imposing a higher hurdle rate on riskier cash flows. It's a crutch.

Stand-in for Imperfect-Market Factors: The CAPM often assigns higher costs of capital to projects that do not satisfy the perfect-market assumptions and that face higher costs of capital. Again, this can accidentally result in better cost-of-capital estimates not *because* of the CAPM, but *despite* the CAPM. It's another crutch.

Such a Great Idea: The CAPM is so intuitive and appealing that it would be "rediscovered" again and again by those who were not forced to learn it. Those who cannot remember the past are condemned to repeat it

Everyone uses it: The CAPM is *the* standard. Exhibit **??** shows that over 80% of the CFOs of large companies report that they always or almost always use the CAPM. Multifactor models are sort of like the CAPM, but throw in a few more factors beyond the market factor (e.g., changes in oil prices). The problem is that every firm may have its own factors, so multifactor models can give very different results for different users. Beyond the CAPM, no model is in wide use (and this includes specific versions of its multifactor cousin known as the APT). Consequently, you have no choice but to understand the CAPM model well — *if you will work for a corporation, then the CAPM is the benchmark model that your future employer will likely use and will expect you to understand well.* Chances are that you will be interrogated about it in your job interview.

Again, the CAPM is simply *the* standard. The CAPM is also used as a benchmark by many investors (e.g., to rate their investment managers), government regulatory commissions, courts (in tort cases), and so on. It is literally the dominant, if not the only, widely used model to estimate the cost of capital. There is even a whole section on the CFA exam about the CAPM!

(Omitted fig)

Alternatives — please stand up: The famous social psychologist Kurt Lewin wrote that "there is nothing more practical than a good theory." If not the CAPM, then what else? There are alternatives, but none are universally accepted. (My own recommendation is to go with the benchmarking approach from the previous chapter or to use the [incorrect] market-beta of 1.0 in the CAPM formula for every stock. It's not a great theory, but it has predicted better.)

Market Hedging: Even if the market beta does not measure the average rate of return, it does guide managers about how much market risk they face — and, if they so desire, how to neutralize it and focus on their real expertise.

Be aware that my treatment of the CAPM in an introductory corporate finance textbook borders on heresy. Most textbooks still make the CAPM their centerpiece. They do this not because the authors believe in it, but because it is dogma that new finance students (and many old finance professors) are too fragile to deserve the hard truth. I am sorry — I wish I could have told you a happy bedtime story about how the world is nice and orderly, too. But it would have been a lie.

(Omitted anecdote)

➤ If you must use it…

If you still want to use the CAPM, here is my advice. As a corporate executive, you should always first think hard about why and when you want to use the CAPM. Think about whether it is useful for your own cost-of-capital estimates. Think about whether the CAPM errors seem too large to be useful for your particular needs. And understand what you are getting. Do simpler benchmarks first — do they agree with the CAPM estimate?

Accuracy: The CAPM is a poor model if you want precision. If you believe that CAPM expected rates of return should be taken seriously on any digits after the decimal point, then you need a shrink. Please realize that, at best, the CAPM could only ever offer expected rates of return that are in the ballpark, plus or minus a few percentage points, perhaps. And maybe not even that…

Actually, if accuracy is important, you are in trouble no matter what. Finance does not have *any* models that can offer physics-level precision for stocks. Fortunately, you may not have to be good at estimating value; you may just need to be better than your competitors. Always remember that valuation of risky long-term projects is as much an art as it is a science. And you wouldn't be the first investor or corporate executive who just happened to be saved by Lady Luck, even if the bet was not a particularly good one.

Investment purposes: If you are not a corporate executive looking to determine your project hurdle rate, but a financial investor looking for good investments from the universe of financial instruments, and with an ability to shift your money around every day, then please do not use the CAPM. Although the

CAPM offers the correct intuition that wide diversification needs to be an important part of *any* good investment strategy, there *are* better investment strategies than just investing in the overall market index. These are discussed in advanced investments courses.

Please do not confuse the CAPM with the mean-variance framework discussed in Chapter , although the latter is one of the CAPM's foundations.1]subsubsect:eff-frontierMean-variance optimization in detail Mean-variance optimization is an asset-selection technique for your individual portfolio, and it works, regardless of whether or not the CAPM holds.

Longer-Term Differences: If you are a corporate executive, be especially cautious about discount rates for expected cash flows far in the future. Look at your cost of capital more holistically. And even if you use the CAPM for guidance, remember that the CAPM has two terms.

The first term is the risk-free rate, which applies to all projects, regardless of beta. Fortunately, this one is easy. You should use higher costs of capital for cash flows that will occur in the more distant future. And you have a great estimate of the premium that long-term projects need to offer over short-term projects, based on the Treasury yield curve. You don't even need historical estimates: you can use the prevailing Treasury yield curve. *Use it! It works!*

It is the second term (the beta multiplied by the risk-premium) — i.e., your beta risk-adjustment — that you must be especially suspicious about. If your cash flows will occur in many years, be modest. Do not over-rely on the risk assessment from the CAPM. Cut down extreme estimates. Shrink and shrink again (towards the average rate of return on risky investments). (Of course, do not forget to be similarly humble in your expected cash flow estimates.) Fortunately, you may be ok:

- As a corporate manager, compare the cost of capital on *your equity* vs. the cost of capital on *your debt* for your long-term cash flows. With an equity premium based on the performance of stocks vs. long-term Treasuries of about 1-2% from 1970 to today, it may not matter so much whether your project A has a beta of 0.8 and your project B has a beta of 1.2. The implied cost-of-capital difference between these two projects of under $(1.2 - 0.8) \cdot 2\% \approx 1\%$/year is already small and probably swamped by your expected cash flow estimation error.
- For long-term cash flows — occurring in, say, 10-20 years — your best estimate of your equity market betas should be tilted even further towards 1.0. If you estimate historical OLS market beta to be 0.5 for A and 1.5 for B today, you may well want to use a market beta (a) shrunk to around 0.7 and 1.3 for predicting future market betas; and (b) shrunk to around 0.9 for A and 1.1 for B for use in the CAPM expected-return formula. Think about this: A and B would now have a difference in the implied cost of equity capital of $0.2 \cdot 2\% \approx 0.4\%$. This is way below your noise-and-uncertainty threshold.

  But let's continue. Say your projects are partly debt-financed, too. Now you need to calculate asset market betas rather than equity market betas. Let's say both projects have 50% debt that is almost risk-free. Then your asset beta would be $0.5 \cdot 0.0 + 0.5 \cdot 0.9 = 0.45$ for A and $0.5 \cdot 0.0 + 0.5 \cdot 1.1 = 0.55$ for B. Now you have a project cost-of-capital difference $(0.55 - 0.45) \cdot 2\% \approx 0.2\%$ between A and B.

How does this expected rate-of-return difference between A and B compare to your own uncertainty about your projects' relative expected cash flows? Does the CAPM beta risk-adjustment really matter much in light of your uncertainty?

In sum, cash flows in the more distant future and cash flows that have higher market-betas should likely be discounted more, as already explained in Chapter . But be humble about your capabilities in trying to distinguish between projects that are similar along time and asset-class dimensions.

Taking Advantage of CAPM Violations

Can investors earn free money (through what is called arbitrage) if the CAPM fails? Not really. The universe remains intact even if the CAPM does not hold, and even in a perfect market.

What would happen in the CAPM if one stock offered more than its appropriate expected rate of return? Its price would be too low. It would be too good a deal. Investors would immediately flock to it, and because there would not be enough of this stock in the economy, investors would bid up its price. This would lower its expected rate of return. The price of the stock would settle at the correct CAPM expected rate of return. Conversely, what would happen if one stock offered less than its due expected rate of return? Investors would not be willing to hold enough of this stock: The stock's price would be too high, and its price would fall. Neither situation should happen in a CAPM world.

Is this a free-money arbitrage? No! When stocks do not follow the CAPM formula, buying them remains risky. Yes, some stocks would offer a higher or lower expected rate of return and thus seem to be too good or too bad a deal, attracting too many or too few investors. (Or, the investors may not even flock towards better deals, perhaps because they have other needs, perhaps because they are asleep at the switch.) But these "CAPM-too-cheap" (or "too expensive") stocks would remain risky bets, and smart risk-averse individual investors would want to buy just a little more (or less) of them. They could not earn risk-free profits. There would be no arbitrage here. The market forces working on correcting any CAPM mispricing are just modest.

And also remember that there are good reasons why the CAPM would not hold in the first place. For example, as we have discussed, it relies heavily on many perfect-market assumptions. If investors are taxed or liquidity-constrained (that is, they cannot easily diversify, e.g., because they are invested in a startup or family firm) or do not agree on the inputs, then it is quite plausible that some firms or even sectors (such as "value-type firms" or "growth-type firms") could offer higher or lower expected rates of return than the CAPM suggests. But not all is lost. It may mean that if *you* are an investor with CAPM preferences, you can do a little better than holding the overall market portfolio by tilting your market-like portfolio just a little towards stocks that offer higher expected rates of return than suggested by the CAPM formula and just a little away from stocks that offer lower expected rates of return.

10.6 Is Beta Useless Without the CAPM?

Importantly, market-beta is not useless if the CAPM does not hold. After all, it was never supposed to be a measure of reward. It is a measure of risk contribution for someone holding the stock market. If you are an investor who holds something close to the market portfolio, market-beta gives you the asset's risk-contribution noxiousness. If you use it to tilt your portfolio towards low-beta stocks, you are creating a portfolio that offers similar overall expected rates of return but at a lower overall portfolio standard deviation. It is also useful for "immunizing" your portfolio against market-wide stock returns.

If the CAPM does not hold, what fails is the market-beta's usefulness in a corporate context for measuring the cost of capital. That is, the manager can no longer use it to infer the preferences and behavior of her investors, and therefore can no longer use it to infer what expected rate of return they demand. *C'est la vie.* What to do then? My advice is to fall back on what you learned in the previous chapter. Try to find the closest benchmarks you can instead of fixating on the stock market.

10.7 CAPM Alternatives and Perspectives

You have already learned in the previous chapter about the principal alternative to the CAPM — benchmarking. The CAPM is really based on similar ideas, but it pushed the idea too far. It was too bold and overconfident.

CAPM vs. Benchmarking: Widening and Narrowing Concepts

The CAPM both generalizes and narrows the idea of benchmarking. The generalization is that market beta is a more universal and objective measure of how equity-like any investment asset is than subjective judgment. It should work for any asset — be it bond, stock, one specific stock or fund, equity options, gold, art, etc. The narrowing is that the CAPM is very specific about the fact that it is market beta — and market beta alone — that is the benchmark of the risk that investors care about. No other factors or exposure to other factors matter.

If the CAPM model is correct, then using more benchmark portfolios (*à la* Chapter) than just the stock market would still be just fine. Each benchmark portfolio would be priced according to the CAPM and lie on the SML. It is merely a convenience of the CAPM that you do not *have to* worry about these benchmark portfolios. If you do use these other benchmarks, fine. If you do not, fine, too. You will still obtain the exact same proper expected rate of return.

If the CAPM model is incorrect, then by using it, you would have gone one step too far. You could easily get the wrong answer. For example, say, investors do not care about market risk (and market beta), but only about oil risk, computer technology risk, and biotech risk. It could be the case that because the market portfolio contained some of these risks, it provided a higher expected rate of return. But it would really matter now whether your project and market beta come from oil risk (which may give a higher expected rate of return) or, say, gold risk (which may not). The CAPM would give you the right answer only if your project happened to have the same proportions as the market portfolio in its exposures.

What you really need for value comparisons are the benchmark portfolios that matter most and fit best. Of course, unlike the CAPM, the benchmark portfolio method can also be harder to use: What is the best benchmark asset? But benchmarking would still work in principle — just as long as you give this method all the right benchmark portfolios as input!

In a sense, the CAPM does only one thing. It tells you what the correct benchmark is — market beta — based on the very common-sense notion that investors should like stocks that are less exposed to market risk. If you use the CAPM, you are assuming it is correct; and, if it is, your estimates become easier and possibly even a little more accurate. But if the model is not correct, you have bet on the wrong horse and you may end up working with an incorrect cost of capital.

My Personal Opinion about Costs of Capital

Now I will give you my own educated opinion about good project cost-of-capital estimates. Beware: Different finance professors may come to different conclusions, so do not take my opinion as the gospel.

➤ Solid Inference — Little Disagreement

The following expected-return premia are rock-solid and not primarily a matter of opinion:

- There definitely is a time value of money.
- There definitely is a term structure. Long-term cash flows usually require higher costs of capital than short-term cash flows. Your investors can earn higher expected rates of return elsewhere for longer-term commitments.
- There definitely is a credit component. Assets have to make up for higher probabilities of default with higher promised yields — that is, higher yields when they succeed.

We have not covered the following yet. It will be explained in Chapter 2]ch:imperfectMarket Imperfections.

- Market imperfections play important roles. There are many kinds. Here are a few examples. There seems to be a liquidity premium. Assets that can be quickly liquidated (especially in general market crashes like 1987 or 2008) are more expensive, and different asset classes seem to have different degrees of liquidity. Because of their collateral, mortgage bonds tend to have lower costs of capital than general bonds. Firms with less access to capital markets, such as startups, seem to pay higher costs of capital, although adjusting for default makes this difficult to measure. Investors pay more in personal income tax for interest receipts than they do for capital gains, which makes equities relatively more desirable and reduces their after-tax income. Sentiment and agency considerations also seem to play important roles in equity trading. Many of these market imperfections embody some concepts of risk, but it is not the market beta. Interestingly, courts agree with imperfect-market views. They allow as much as a 20-30% discount for the value of privately held assets relative to publicly traded peers. We may not know what the costs of capital for small, privately held firms are, but we do know that they are usually much higher.

➤ Uncertain Inference — More Disagreement — My Subjective Assessments

I wish I knew the equity risk premium — and for a lot of different reasons. The CAPM is only one of them. Benchmarking is another. Alas, I am not so confident that I have a good assessment. We are dealing with finance (with estimated probabilities), not physics (with known probabilities). And expected rates of return, such as those on the stock market, are notoriously unreliable.

After taking into account the premia just mentioned (which includes premia that are sometimes included in and have to be captured by the risk premium, but which I already have in my number), the remaining risk premium — especially over longer horizons — is probably relatively small (1-2%). However, we do not know for sure. Our uncertainty is much larger than our certainty about its magnitude. And you need to realize that betas for cash flows far into the future are much closer to 1 than historical regressions would suggest. The "CAPM" beta-metric for measuring the project's risk impact and expected rate of return is only of modest (if any) importance.

So what would I do (unless forced to use the CAPM by a superior)? My best alternative cost-of-capital recommendation would start out just like the CAPM: As the first term in a formula, I would recommend that you use the rate of return on bonds of similar maturity as the cash flow that you want to value. Usually, this means that you assign higher costs of capital to cash flows farther in the future. It is only with the second term — the equity risk-adjustment — that I would tinker. Instead of the (shrunk) CAPM market beta multiplied by some historical equity premium (of 1-3%/year geometric above long-term Treasuries), I would recommend a more holistic approach.

- Take into consideration that projects with high volatility and/or with high leverage are more risky. The equity of these projects probably requires a higher expected rate of return to keep your investors happy. Realize that projects with higher idiosyncratic risk are also usually the same projects about which executives tend to be most overly optimistic. (Check again: are you sure your expected cash flows in the NPV numerator are not biased by your overconfidence?)
- As a manager, take into consideration whether you and your owner-shareholders are well-diversified. If the owners have most of their wealth in your firm or project, then you should require higher rates of return for riskier projects—but in this case, it is not "beta risk" that matters, but "total risk (standard deviation)."
- Take into consideration that public investors may "like" certain types of investments — like crypto or environmental, social, governance[environmental, social, governances] (ESG) related investments — and are often willing to pay higher prices and thus accept lower average rates of return for some such projects. If they are willing to give you money at lower expected rates of return, take this into consideration.

There is probably little harm if you calculate a (repeatedly shrunk) CAPM market beta and apply it to a relatively low equity premium (say, 2%/year) for some heuristic orientation. Assess whether any other non-CAPM cost-of-capital assessments seem reasonably similar to such a CAPM assessment. In this sense, the CAPM can still be informative.

➤ **If Forced**

- If I ran a large publicly-traded firm with good access to capital markets and I needed to assess the cost of capital for a typical medium-term project, I would assume an equity premium of 1-3% per annum and apply the resulting expected rate of return on the stock market (i.e., adding back the fixed-income borrowing rate) to the equity components of all my long-term cash flows. The exception would be projects for which I would have strong prior knowledge that they are very unlike stocks.

  I would consider long-term corporate debt to have a higher cost of capital than equivalent Treasuries, but a lower cost of capital than my own equity — the latter primarily because debt provides a corporate income tax shield (as you will learn in Chapter) and not because the equity premium over long-term corporate bonds is high.

  And I would distinguish between the cost of capital and the hurdle rate (which one should not do a perfect capital market). We will discuss this further in the next three chapters.

- If I ran a startup firm funded by myself or other-not-so-rich investors (so that the startup would risk a large fraction of personal wealths), I would never even think about the CAPM. Risk would definitely play a central role, but not in the market-beta and comovement sense. I would be more concerned about total risk and funding uncertainty.

  I would perhaps assume a cost of capital of 2% to 6% above the *expected* rate of return on my uncollateralized debt, depending on the project risk and my risk aversion. I would probably require an even higher hurdle rate — it could easily be in the double digits. Even more radically, I might even abandon NPV-based models altogether and worry more about the downside of losing my shirt.

  My perspective would change again when I was about to sell the firm to better-diversified investors (such as a venture capitalists). I would try to figure out not just my own but also their perspective. They often base their valuations on alternative projects in the same industry and do a lot of benchmarking and comparables, the subjects of Chapters and .1]ch:compsComparables However, it is here that I would begin relying more on traditional NPV modeling, too.

And I would never use any of my schemes here (or the CAPM) for the pricing of bonds, derivatives, real-estate, human capital, or other non-stock-like kinds of projects.

Am I the only professor who recommends against using the CAPM? No. Many do so in conversations among themselves and with other experts, and even more do when their own money is on the line. However, many have been fainthearted to admit to our *collective* ignorance in front of students. It is easier to teach the intuition and beauty of the CAPM and stop than it is to dismiss it right after teaching it. (And teach it, we must!) Let me appeal to a higher authority for backup: Eugene Fama — the most famous finance professor alive, winner of a Nobel prize, and partly responsible for the original spread of the CAPM — nowadays strongly recommends against the CAPM, too. His view is that using the CAPM expected rate of return as your cost of capital in an NPV calculation effectively divides one bad uncertain number by another bad uncertain number. This practice convolutes errors and uncertainty about expected cash flows in the numerator with errors and uncertainty about expected returns in the denominator. If you get lucky, your errors cancel. If not, they do not. Yikes 😬! "Gene" prefers comparables.

Begin Important

- The CAPM is *the* benchmark model in the real world. Most corporations use it.
- Interviewers will likely expect you (a student applying for a job) to understand the CAPM. Regardless of whether the model holds or not, you have to know it.
- The empirical evidence suggests that the CAPM is not a good model for predicting expected rates of return.
- The first CAPM term (the time adjustment) seems to hold better than the second CAPM term (the risk adjustment).
- Market betas tend to revert back towards 1 over time. This requires you to shrink ordinary OLS beta estimates very aggressively towards 1.
- The geometric equity premium above long-term Treasuries (for evaluating long-term cash flows) has been — and is unlikely to be more in the future than — 2-3% per annum.
- The CAPM never offers great accuracy. Do not lean on or trust the CAPM.
- Mean-variance optimization (Section) works even if the CAPM does not. If you are an investor who dislikes risk, the immediate implication is that instead of the typical market portfolio, you should tilt towards stocks with lower betas and away from stocks with higher betas. This will give you an overall portfolio that has lower risk for the same amount of expected rate of return.
- Peer portfolio benchmarking (Chapter) works regardless of whether the CAPM does or does not work.
- You may or may not want to immunize your project against equity-premium risk and estimation uncertainty, using its beta estimate. Immunized projects have much clearer cost-of-capital benchmarks than unimmunized projects.

End Important

(Omitted solvenow)

Summary

This chapter covered the following major points:

- The CAPM provides an "opportunity cost of capital" for investors, which corporations can use as the cost of capital in the NPV formula. The CAPM formula is

  (Omitted eq)

  Thus, there are three inputs: the risk-free rate of return (r_F), the expected rate of return on the market ($E(r_M)$), and the project's or firm's market beta (β_i). Only the latter is project-specific.
- The line plotting expected rates of return against market beta is called the security market line (SML).
- The CAPM provides an expected rate of return, consisting of the term premium and the risk premium. It ignores the default premium. In the NPV formula, the default risk and default premium work through the expected cash flow in the numerator, not through the expected rate of return (cost of capital) in the denominator.
- For r_F, you should use bonds that match the timing of your project's cash flows. Thus, cash flows farther in the future often require higher opportunity costs of capital. Even if you do not believe in the CAPM, term adjustment is important.
- The expected rate of return on the stock market is a critical CAPM input if the project's market beta is high — but the best estimate of the equity premium itself remains elusive. There are many guesstimation methods, but no one really knows which one is best. Reasonable estimates for the equity premium ($E(r_M) - r_F$) can range from about 1% to 8% per annum, although 2-3% seems most common for cash flows more than a few years into the future.
- There are a number of methods to estimate market beta. Don't be too confident about betas far from 1, especially for long-term project cash flows.
- If you combine a short position in the stock market with a positive-beta project, the combined project is a lot easier to price than a project with a positive beta. By effectively creating a combined project with a zero market-beta, you can neutralize the effects of "CAPM is wrong model" and equity-premium errors.
- Never believe the CAPM blindly. Its estimates are poor. Use them more for "general direction" than as "accurate guides." Think compass, not GPS.
- Even though its estimates are poor, understand the CAPM well. Everyone will expect you to.

This negative perspective on the CAPM is so uncommon in a textbook (but not among the experts actually studying the models) that it is important that you don't misunderstand what this chapter says. So let's end this chapter with a FAQ:

- **Q**: Should riskier cash flows not require higher promised rates of return?

A: Riskier projects have to promise higher rates of return, i.e., offer higher default premiums. This is not the same as higher risk premiums in the CAPM sense. In NPV applications, make sure to reflect the default risk in the expected cash flow numerator. Riskier projects need to pay off a lot more when they succeed, just to make up for the fact that they fail more often.

- **Q**: Should long-term and therefore riskier cash flows not require higher expected rates of return?

  A: Long-term projects command term premiums. Thus, in NPV applications, you should usually use higher required costs of capital for more distant cash flows. You should not use the CAPM for this. The U.S. Treasury Yield Curve gives you a working first estimate about how much extra premium long-term cash flows should require above short-term cash flows.

- **Q**: Besides a term premium (and perhaps a leverage-structure adjustment to account for debt capacity), should riskier stocks and corporate cash flows have higher expected discount rates?

  A: Maybe, but be careful. First, remain humble and make your estimate err on the side of modesty. Don't be overconfident in your ability to judge equity risks. If you can judge the risks well, make sure your estimates first flow into your expected cash flows in the NPV numerator. Second, don't be too wedded to the CAPM for the extra "risk-premium kicker." Instead, combine your cost-of-capital estimate with judgment, perhaps also looking at factors such as volatility (especially if your owners are not fully diversified). If I am vague here, it is deliberate. I cannot give you clear universal guidance.

- **Q**: If CAPM cost-of-capital estimates are so bad, why has this not hurt companies that rely on the CAPM more badly?

  A: Most companies use a hurdle rate considerably higher than their cost of capital estimates. That is, even though they pay lip service to the CAPM, they don't really use it as their end-all.

Preview of the Chapter Appendix in the Companion

The appendix to this chapter explains:

- How the "certainty equivalence value" (CEV) allows you to use the CAPM for projects that you are not buying at the appropriate equilibrium price. For example, you would need the CEV to work out how to value an inheritance that will be higher if the economy does well. (Just because the inheritance is "free" to *you* does not mean that there is a zero value to it.)
- How to use the CEV formula to estimate the value of a project for which you have historical cash flows, but no market value information.
- How the CAPM is derived from the fact that the optimal portfolio is always the combination of two portfolios, one of which may be the risk-free asset.
- What a few more CAPM alternatives are and how to use them. The first alternative is the APT (arbitrage pricing theory) and its relative, the Intertemporal CAPM. The second alternative is a "Fama-French"-style model, which uses factors such as value, growth, momentum, investment, and robustness. This

Fama-French model seems to predict better than any alternatives, but it is less grounded in theory (or, you may say, reason) than the former. It also often gives counterintuitive results — e.g., that small growth stocks are safer than large value stocks and therefore that managers should use *lower* discount rates on, say, risky tech ventures.

Compiled: Thursday 3rd April, 2025