

Cost of Capital

Every valuation in the last four chapters used a discount rate that was handed over. This chapter computes one. It comes last deliberately: the rate is the most contested number in finance, and meeting it before you have used it makes it a piece of machinery to memorise rather than a judgement you have already felt the weight of.

By now you have met the rate four times. Chapter 5 isolated inflation. Chapter 6 took a rate apart into time preference, inflation, and a premium for risk, and ranked cash-flow streams by the premium each deserved. Chapter 7 showed a risk premium as a printed market price. Chapters 8 and 9 asked which risks a project's and a firm's rate must cover. What remains is the arithmetic, and one honest observation about how much the arithmetic is worth.

HOW TO USE THIS CHAPTER

This chapter builds the model where its inputs exist, states plainly how well it performs empirically, gives the levering and unlevering relationship that makes a listed beta usable for a private firm, and compares the three country-risk methods that disagree with one another by more than three percentage points on the same company.

The videos work the full calculation in a United States setting, where every input genuinely exists, and then remove them one at a time to see what survives in a market without them.

The case takes one business and asks what it would be worth in two different countries.

10.1 What a Cost of Capital Is

The cost of capital is what the providers of a firm's money require in order to keep providing it. It is not a cost the firm pays out in the way wages are paid; it is a **threshold return**, below which capital should go somewhere else. A business earning less than this threshold is consuming value while still reporting a profit, and section 11.3 shows what that does to what it is worth.

Two properties of that definition do most of the work.

- **It is a required return, so it belongs to the investment rather than to the investor.** Two firms evaluating the same project should use the same rate, because the risk of the cash flows is a property of the project. A firm applying its own average rate to an unusual project will systematically misjudge it.
- **It is forward-looking.** What the firm paid for capital in the past is irrelevant. The question is what would have to be offered today, which is why Chapter 7 insisted that the cost of debt is a current yield rather than a historical coupon.

10.2 The Risk That Is Not Paid For

The central idea of asset pricing is that not all risk earns a return. Some of it can be removed for free by holding several investments instead of one, and risk that can be removed at no cost cannot

command compensation — anyone demanding payment for bearing it would be undercut by someone who had diversified it away.

TWO KINDS OF RISK

TYPE	SOURCE	BEHAVIOUR IN A PORTFOLIO
Idiosyncratic	A fire, a lost contract, a failed product, a lawsuit	Averages away across holdings
Systematic	Recessions, interest rates, commodity prices, currency	Affects everything at once; remains

The mechanism is arithmetic rather than economic. Independent events partly cancel: one holding's bad year offsets another's good one, and the more holdings there are, the more complete the cancellation. What cannot cancel is anything that moves all holdings together.

HOW QUICKLY DIVERSIFICATION WORKS

Individual holdings with 45% volatility; market volatility 19%

NUMBER OF HOLDINGS	PORTFOLIO VOLATILITY	SHARE OF REMOVABLE RISK REMOVED
1	45.0%	0%
2	34.5%	40%
5	26.3%	72%
10	23.0%	85%
30	20.4%	95%
100	19.4%	98%
Infinitely many	19.0%	100%

Two readings matter. **Most of the benefit arrives early** — ten holdings capture about 85 percent of what is available and thirty capture about 95 — which is why the practical advice is that a few dozen positions are enough and why the pursuit of hundreds adds little. And **the curve flattens onto a floor it never crosses**. That floor is systematic risk, and it is the only risk that earns a return.

The implication for a firm is direct and worth stating carefully, because it is counter-intuitive. A company's own specific risks — its key customer, its founder, its single warehouse — do not raise the discount rate that a diversified investor requires, because that investor has diversified them away. This is true, and it is routinely wrong in practice for the markets this course addresses, where the buyer of a private company is frequently not diversified at all. A family investing most of its wealth in one business bears the whole of its risk and rationally demands compensation for all of it. The model assumes an investor who does not exist in that transaction, and section 10.8 returns to what to do about it.

IN THE VIDEOS *Video 10.2 builds the diversification table one holding at a time and identifies the floor the curve never crosses.*

10.3 Beta and the Cost of Equity

If only systematic risk is compensated, it must be measured. **Beta** measures how much a company's returns move when the market moves: 1.0 moves with the market, 1.5 moves half again as much, 0.5 moves half as much. Mechanically it is the slope of a line fitted through a scatter of the company's returns against a broad index, conventionally sixty months of data.

Read it as a sentence rather than a statistic. A beta of 1.2 says that when the market has risen ten percent, this company has historically risen about twelve. **High beta** belongs to businesses whose fortunes amplify the economy — construction, airlines, capital equipment, luxury goods, banks. **Low beta** belongs to businesses people use regardless — utilities, staple food, basic pharmaceuticals. This is Chapter 9's business risk with a number attached.

$$\text{Cost of equity} = \text{Risk-free rate} + \text{Beta} \times \text{Market risk premium}$$

Three inputs, each with its own difficulty.

THE THREE INPUTS

INPUT	WHERE IT COMES FROM	THE DIFFICULTY
Risk-free rate	A government security in the same currency, maturity matched to the cash flows	Which maturity; and in many currencies no long bond exists
Market risk premium	Long historical excess of equities over government paper	Reasonable estimates differ by about two percentage points
Beta	A regression against a broad index	Providers differ by 0.3 or more for the same firm

A worked instance: a risk-free rate of 3.5 percent, a market risk premium of 5.5 percent, and a beta of 0.90 gives a cost of equity of $3.5 + 0.90 \times 5.5$, or **8.45 percent**. Move beta to 1.40 and the answer becomes 11.2 percent. **Nearly three percentage points of difference between a defensive company and a cyclical one**, which is why beta is worth checking against more than one source rather than accepting the first figure found.

How well the model works

Something should be said here that is often left out of an introductory treatment. **The Capital Asset Pricing Model does not perform well empirically**. It predicts that high-beta stocks deliver higher returns than low-beta stocks, and in the data that relationship is far weaker than predicted; over long periods low-beta stocks have often done at least as well, which is the opposite of the prediction. This has been examined for decades and is not a measurement problem that better data resolves.

Why then use it? Because it is a disciplined way of organising a judgement that must be made regardless. It forces you to name a risk-free rate, state a market premium, and take a position on how cyclical the business is — and to do all three in public where they can be challenged. **Its value is that it makes the assumptions explicit, not that it is right**. Treat the output as one estimate with wide error bars rather than as a measurement, and never present it to two decimal places.

IN THE VIDEOS Video 10.3 estimates a cost of equity in a market where all three inputs genuinely exist, which is the only place the model can be seen working as designed.

10.4 Levering and Unlevering Beta

A published beta describes a company as it is currently financed, because a more indebted firm has more volatile equity. Using a listed firm's beta for a private company with a different capital structure therefore imports the wrong financial risk along with the right business risk. The remedy is to strip the financing out and put the correct financing back in.

$$\text{Unlevered beta} = \text{Levered beta} \div [1 + (1 - \text{tax rate}) \times \text{Debt/Equity}]$$

$$\text{Relevered beta} = \text{Unlevered beta} \times [1 + (1 - \text{tax rate}) \times \text{Target Debt/Equity}]$$

The unlevered figure — sometimes called the asset beta — describes the business risk alone and is comparable across firms in the same industry regardless of how each is funded. Take a listed comparable with a beta of 1.30, a debt-to-equity ratio of 0.60, and a 28 percent tax rate.

UNLEVERING AND RELEVERING

STEP	WORKING	RESULT
Observed levered beta	As published	1.3000
Unlever at D/E of 0.60	$1.30 / (1 + 0.72 \times 0.60)$	0.9078
Relever at a target D/E of 0.25	$0.9078 \times (1 + 0.72 \times 0.25)$	1.0712
Cost of equity at rf 3.5%, MRP 5.5%	$3.5 + 1.0712 \times 5.5$	9.39%

This is the standard route to a discount rate for a private company, and it is the practical answer to the objection that a private firm has no share price to regress. **Take the betas of several listed firms in the same business, unlever each, average the unlevered figures, and relever at the subject's own capital structure.** The average is the point of using several: individual betas are noisy, and the industry asset beta is estimated far more reliably than any single firm's.

10.5 The Cost of Debt and the Tax Shield

The cost of debt is the yield a firm would pay to borrow today, and Chapter 7 supplies it directly when the firm has traded bonds. When it does not, two substitutes are used.

- **The rate on recent borrowing**, provided it is genuinely recent and genuinely comparable in seniority and maturity.
- **A synthetic rating.** Estimate the firm's creditworthiness from its interest cover and leverage, map that to a rating category, and take the spread that category currently trades at. Illustratively: cover above eight times might attract 150 basis points, four to eight times 250, two to four times 400, and below one and a half times a thousand or more. These mappings must be refreshed from current market data, and the method is approximate — but it is transparent and can be argued with, which a guess cannot.

Interest is deductible against tax in most jurisdictions, so the firm's actual cost is lower than the rate it pays.

$$\text{After-tax cost of debt} = \text{Yield} \times (1 - \text{tax rate})$$

A firm borrowing 500 at ten percent pays 50 of interest, which reduces its taxable profit by 50 and its tax bill by 14 at a 28 percent rate. The net cost is 36, an effective rate of 7.2 percent. Debt's required return is lower than equity's primarily because lenders hold a senior contractual claim and bear less risk. **The government then funds part of the interest bill**, reducing the firm's after-tax debt cost further. That incremental tax wedge—not debt's basic risk advantage—is the financing benefit created by deductibility.

One condition is worth remembering. The shield is worth nothing to a firm with no taxable profit to shield. A loss-making company, or one already sheltered by allowances, borrows at the full pre-tax cost, and its calculation should say so.

10.6 The Weighted Average Cost of Capital

A firm funded by both debt and equity faces a blended requirement, weighted by how much of each it uses.

$$\text{WACC} = \text{Cost of equity} \times (E/V) + \text{After-tax cost of debt} \times (D/V)$$

With equity of 700, debt of 300, a cost of equity of 14 percent, a pre-tax cost of debt of 9 percent, and tax at 28 percent, the after-tax cost of debt is 6.48 percent and the weighted average is $14 \times 0.70 + 6.48 \times 0.30$, or **11.744 percent**. Notice how little the debt contributes: 1.94 of the 11.74 percentage points. With moderate leverage and a cheap after-tax debt cost, the WACC is dominated by the cost of equity, which is where the estimation effort should go.

The weights must be market values

This is the most common mechanical error in the whole calculation. **The weights are market values, not book values.** The reason is that the cost of capital is a forward-looking required return on what the capital is worth now, and book equity is a historical residual with no bearing on it.

In the worked company, book equity is far below market equity, so book weights overweight the cheaper debt component and understate WACC. A rate understated by 140 basis points then overstates the value of a growing perpetuity by roughly a quarter. **The direction is not a rule.** A distressed firm can have market equity below book equity; book weights can then overweight the expensive equity component and overstate WACC. Use market or target weights because they represent current capital claims, not because the book-value error always points one way.

If debt is cheaper, why not use only debt?

The formula appears to say that shifting weight towards debt lowers the average without limit. It does not, for a reason the formula conceals.

The weights are not independent of the rates. As a firm borrows more, a larger fixed claim sits ahead of the equity, so the equity becomes more volatile, beta rises, and the cost of equity rises

with it — which is exactly what section 10.4's relevering formula describes. The debt also becomes riskier, so lenders demand more. Substituting cheap debt for expensive equity makes **both** more expensive, and the two effects largely offset.

Beyond a point, real costs appear that have nothing to do with rates. A heavily indebted firm faces genuine distress costs: customers hesitate to place long-term orders, suppliers tighten terms, capable staff leave, and management spends its time negotiating with banks instead of running the business. These are not financing costs; they are operating losses caused by the financing.

The resulting shape is worth carrying. **WACC typically falls modestly as a firm takes on its first debt, flattens across a broad range, and rises as distress risk takes over.** The practical reading is that a firm with no debt is probably leaving some value on the table, a firm with a great deal of debt is buying a small tax benefit at the price of real fragility, and across a wide middle range the choice matters much less than either extreme suggests. Where a firm sits within that range is a question of how stable its cash flows are — which returns to Chapter 9's business risk.

A PRACTICAL CAPITAL-STRUCTURE CHECK

MORE DEBT CAN HELP WHEN...	MORE DEBT CAN HURT WHEN...
Cash flows are stable	Cash flows are volatile or cyclical
Interest cover remains comfortable	Refinancing or currency mismatch is material
Tax shields can actually be used	Distress disrupts suppliers, customers, or employees
Debt funds positive-spread investment	Debt funds value-destroying growth

IN THE VIDEOS *Videos 10.4 and 10.5 assemble the full weighted average, and show what using book weights instead of market weights does to a valuation built on top of it.*

10.7 Outside a Deep Market

Everything above assumed a market that supplies its inputs: a liquid government yield curve, a long index history, and published betas. For a private company in a frontier or emerging market, none of the three exists. The methods practitioners actually use are set out below, together with an honest statement of how far apart they land.

WHAT IS MISSING, AND WHAT IS SUBSTITUTED

INPUT	THE PROBLEM	THE SUBSTITUTE
Risk-free rate	Local government paper carries real default risk	A developed-market rate, then adjusted
Market risk premium	Index history is short and thin	A mature-market premium
Beta	The company is not listed	An industry asset beta, relevered
Country risk	Not in the model at all	An explicit premium, three ways

The country premium is where the disagreement lives. Three methods are in common use, and applied to the same company with a sovereign spread of 550 basis points they give:

THREE METHODS, ONE COMPANY

US risk-free 3.5%, mature market premium 5.5%, asset beta relevered to 0.90

METHOD	CONSTRUCTION	COST OF EQUITY
Add the sovereign spread	$3.5 + 0.90 \times 5.5 + 5.5$	13.95%
Scale the spread by 1.5 for equity volatility	$3.5 + 0.90 \times 5.5 + 8.25$	16.70%
Put the spread inside the premium	$3.5 + 0.90 \times (5.5 + 5.5)$	13.40%

Three and a third percentage points separate the highest from the lowest, using identical inputs and three defensible methods. No amount of care narrows this; it is a genuine disagreement about how country risk enters, not an error any of them makes. The professional response is to state which method you used and why, and to show the valuation under the others. Presenting one of these three as the cost of equity, without saying which convention produced it, conceals the largest uncertainty in the calculation.

Currency

A rate built from US inputs is a **dollar** rate, and discounting local-currency cash flows with it is the mismatch Chapter 5 warned about. Converting requires the Fisher relation from that chapter, applied to both currencies.

$$\text{Local rate} = (1 + \text{dollar rate}) \times (1 + \text{local inflation}) \div (1 + \text{US inflation}) - 1$$

Applied to the 13.95 percent dollar figure above, with US inflation at 2.5 percent:

THE SAME DOLLAR RATE, IN FOUR CURRENCIES

LOCAL INFLATION	WORKING	LOCAL-CURRENCY RATE
8%	$1.1395 \times 1.08 / 1.025$	20.06%
12%	$1.1395 \times 1.12 / 1.025$	24.51%
16%	$1.1395 \times 1.16 / 1.025$	28.96%
20%	$1.1395 \times 1.20 / 1.025$	33.40%

These are not high rates by mistake. A rate of 33 percent in a currency inflating at 20 percent is a **real rate of about 11 percent**, which is unremarkable. The figure looks alarming only because it is quoted in a unit that is shrinking, and the alarm is the reason the mistake in the other direction is so common: analysts apply a familiar-looking 18 percent to cedi or naira cash flows because 33 seems absurd, and thereby overvalue the business by a wide margin.

The check that catches this takes one line. **Convert the terminal growth rate into real terms using the same inflation figure.** A model discounting local-currency flows at 18 percent while assuming three percent perpetual growth in a currency inflating at 12 percent is assuming the business shrinks by roughly eight percent a year in real terms forever — which is not what its author intends and is visible only when the conversion is done.

IN THE VIDEOS *Video 10.6 removes the developed-market inputs one at a time and rebuilds the rate for a company in a market that supplies none of them.*

10.8 What the Answer Is Worth

It is worth adding up the uncertainty rather than letting it accumulate silently.

ERROR BARS ON THE INPUTS

INPUT	TYPICAL RANGE OF REASONABLE ESTIMATES	EFFECT ON THE RATE
Market risk premium	About 2 percentage points	Multiplied by beta
Beta	0.3 or more between providers	Multiplied by the premium
Country risk method	3+ percentage points, as above	Direct
Capital structure weights	Depends on the target assumed	Moderate

Assembled honestly, a cost of equity for a private company in a thin market is an estimate with a range of several percentage points, and no methodology closes that gap. Three consequences follow, and they are the practical conclusion of the chapter.

- **Never present a cost of capital to two decimal places.** The precision is not there, and false precision transfers to every number built on it.
- **Show the valuation across a range of rates,** which is what Chapter 9's sensitivity grid is for. The range of rates is the honest input, so the range of values is the honest output.
- **When the buyer is undiversified, say so.** The model prices systematic risk only. A family or a founder putting most of their wealth into one business bears all of it, and a rate that ignores that is answering a question nobody in the room asked. State the model's answer, state the adjustment, and let the decision-maker see both.

10.9 Reference

CHAPTER 10 • FORMULA REFERENCE

QUANTITY	FORMULA	NOTE
Cost of equity	$r_f + \beta \times \text{Market risk premium}$	Wide error bars
Unlevered beta	$\beta_L / [1 + (1-t) \times D/E]$	Business risk alone
Relevered beta	$\beta_U \times [1 + (1-t) \times D/E \text{ target}]$	Applies the subject's structure
After-tax cost of debt	$\text{Yield} \times (1 - t)$	Yield, not coupon; worthless without taxable profit
WACC	$r_E \times E/V + r_D(1-t) \times D/V$	Market-value weights
Country premium, method 1	Add the sovereign spread	The common default
Country premium, method 2	Sovereign spread $\times$ equity/bond volatility	Usually the highest
Country premium, method 3	Add the spread to the market premium, then $\times \beta$	Usually the lowest
Currency conversion	$(1+r\$)(1+\text{local inflation})/(1+\text{US inflation}) - 1$	Never convert by assumption

CARRY FORWARD

Only systematic risk is priced by the model — and only for a diversified investor. When the actual buyer is not diversified, the model is answering a different question and you should say so.

Use market-value weights. Book weights describe accounting history rather than current capital claims; depending on the facts, they can understate or overstate the cost of capital.

Three defensible country-risk methods differ by more than three percentage points. State which you used and show the others.

A rate and its cash flows must share a currency. Convert with the Fisher relation, and check the terminal growth rate in real terms while you are there.