

Project Valuation

A firm has more things it could do than money to do them with, and choosing among them is the central task of corporate finance. This chapter sets out how that choice is made: which cash flows belong in the analysis, how they are assembled, how the part beyond the forecast horizon is valued, and which decision criteria are trustworthy.

One idea in this chapter is worth more than the rest combined. Beyond the explicit forecast, most of a project's value usually sits in a **terminal value** — a single number standing in for everything after the last forecast year. It is the least examined figure in most analyses and routinely the largest, and there is one specific way of computing it that is internally inconsistent and produces answers that are wrong by a very large margin. Section 8.4 is about that.

HOW TO USE THIS CHAPTER

This chapter gives the rules for deciding which cash flows are incremental, the general construction of free cash flow, the three terminal-value methods with the algebra that shows when each is consistent, and the failure modes of every decision criterion including the popular ones.

The videos build a complete project model for a new outlet and then compute its terminal value twice — once the way most people do it and once correctly — with the two answers separated by more than half the project's value.

The case asks a management team to decide on that outlet, with the terminal-value question left for them to find.

8.1 The Decision, and the Rule

A project is a set of cash flows: money out now, money in later. The rule is Chapter 6 applied to that set.

$$\text{NPV} = \sum \text{Free cash flow}_t \div (1 + r)^t - \text{Initial investment}$$

Accept a positive net present value. What follows in this chapter is entirely about doing the two hard parts honestly: deciding what goes in the numerator, and choosing r . **Neither is arithmetic, and both decide the answer.**

One framing helps before any of the machinery arrives. A capital-allocation decision is always a comparison, even when it is presented as a yes or no. A firm approving a project is declining whatever else that money could have done, including holding it. The discount rate exists to make that comparison, which is why the rule needs no separate hurdle: **the alternative is already inside the rate.**

8.2 Which Cash Flows Count

Only cash flows that change because the project happens belong in the analysis. This is the incremental principle, and stating it is easy while applying it is where most project models go wrong.

$$\text{Incremental cash flow} = \text{Firm cash flow with the project} - \text{Firm cash flow without it}$$

Six rules follow from that definition, and they are worth learning as a list because each corresponds to a specific and common mistake.

THE SIX RULES

ITEM	TREATMENT	REASON
Sunk costs	Exclude	Already spent; unaffected by the decision
Opportunity costs	Include	The forgone use of a resource is a real cost
Cannibalisation	Include	Lost sales elsewhere are caused by the project
Allocated overhead	Exclude unless it actually changes	Reallocating an existing cost changes no cash
Working capital	Include, and release at the end	Money genuinely tied up, then recovered
Financing costs	Exclude from the cash flows	They are in the discount rate; counting both double-counts

Three of these deserve expansion, because they are the ones argued about in meetings.

Sunk costs. A feasibility study already commissioned is gone whether or not the project proceeds, so it does not belong in the decision. This is correct, deeply unpopular, and consistently resisted, because the person who commissioned the study wants it to have mattered. The discipline is to ask what changes from here, and only that.

Opportunity cost. A project using land the firm already owns is not using a free asset. The land could be sold or let, and that forgone amount is a genuine cost of the project. The most frequent version of this error is treating an owned building as costless, which makes every project inside it look better than it is.

Financing costs. Interest does not appear in project cash flows. This surprises people, and the reason is that the discount rate already represents what the providers of capital require. Deducting interest **and** discounting at a rate that includes the cost of debt charges for the same thing twice. The rule is simple to apply: **build the cash flows as though the project were funded entirely by equity, and put all the financing into the rate.**

IN THE VIDEOS Video 8.2 works through a project where a cannibalisation assumption and an opportunity cost each change the recommendation.

8.3 Building Free Cash Flow

Free cash flow is what the project produces for all providers of capital, after everything the project itself requires. The construction is the same every time.

FREE CASH FLOW, IN GENERAL FORM

LINE	SIGN	NOTE
Revenue	+	Incremental only
Less cash operating costs	–	Excludes depreciation
Equals EBITDA	=	Operating profit before non-cash charges — not cash flow; see below
Less depreciation	–	Non-cash, deducted only to compute tax
Equals EBIT	=	Operating profit
Less tax on EBIT	–	Tax on operating profit, not on profit after interest
Equals NOPAT	=	Operating profit after tax
Add back depreciation	+	It never left; reversed now that tax is computed
Less capital expenditure	–	The actual cash spent on assets
Less increase in working capital	–	Money tied up in receivables and stock
Equals free cash flow	=	Available to lenders and owners

The table also shows, by what it puts below EBITDA, what EBITDA is not. Take a firm with EBITDA of 15, depreciation of 5, tax at 30 percent, capital expenditure of 12, and 4 of additional working capital. EBIT is 10, NOPAT is 7, and free cash flow is $7 + 5 - 12 - 4 = -4$. Now take the same firm growing at nothing: it spends 5 to replace what it wore out and ties up no further working capital, so its free cash flow is $7 + 5 - 5 = +7$.

Identical EBITDA, and cash flow of -4 in one case and $+7$ in the other. **EBITDA is not cash flow and is not close to it, because it is struck before the two items that actually consume cash — the assets a business must buy and the capital its growth traps.** It is a useful measure precisely because it strips those out, which is what allows two firms with different asset bases and different depreciation policies to be compared on operations alone. It becomes a dangerous one the moment anybody treats it as the money available to service a loan.

The depreciation treatment is the part that confuses people, and it is worth a sentence. Depreciation is subtracted and then added back, which looks pointless. It is not: **it is subtracted because it reduces taxable profit, and added back because it never involved a payment.** Its only role in the calculation is to determine the tax bill, and the amount by which it reduces tax — the depreciation charge multiplied by the tax rate — is called the depreciation tax shield. It is real money, and a project whose assets can be written off faster is genuinely worth more.

Two further conventions matter for consistency and both are easy to get wrong.

- **Tax is computed on EBIT, not on profit after interest.** Using the latter smuggles the financing benefit into the cash flows, which then reappears in the discount rate.
- **Working capital is an increase, not a level.** The cash flow item is the change during the period. At the end of a project the accumulated working capital is released and comes back as a positive flow — a real item, frequently omitted, and worth several percent of a project's value.

IN THE VIDEOS Video 8.3 builds five years of free cash flow line by line for a new outlet, including the working-capital movements.

8.4 Terminal Value, and the Capital-Expenditure Rule

Forecasts have horizons; businesses do not. Terminal value is the single figure that stands in for everything after the last forecast year, and in a typical project it carries between half and three-quarters of the total. Three methods are in general use.

THE THREE METHODS

METHOD	FORMULA	ASSUMES
Perpetuity growth	FCF in the year after the horizon / (r - g)	The business continues indefinitely
Exit multiple	A multiple applied to a final-year metric	The business is sold at prevailing prices
Liquidation	The realisable value of the assets	The business stops and is broken up

Each is defensible in the right circumstance, and the choice should follow the intention. A business that will genuinely be run indefinitely warrants a perpetuity; one that is expected to be sold warrants a multiple; a project with a defined end and disposable assets warrants a liquidation figure. **Where two methods are available, compute both.** A perpetuity terminal value can always be converted into the multiple it implies, and if that multiple is far outside what comparable businesses trade at, the growth assumption is doing something it should not.

The rule that makes a perpetuity consistent

Here is the error that this chapter exists to prevent. A perpetuity terminal value is often computed by taking the final year's free cash flow, growing it, and dividing by the rate less growth. Free cash flow in that final year typically shows a capital expenditure figure of zero or near zero, because the project's spending was concentrated at the beginning. The perpetuity therefore assumes a business that **grows forever while never replacing an asset**, which is not a possible business.

The correct statement is algebraic and short. Let a business earn a return on invested capital of ROIC and grow at g. To grow at g, its invested capital must also grow at g, which requires net investment of g multiplied by invested capital. Since invested capital is NOPAT divided by ROIC, net investment is $\text{NOPAT} \times g / \text{ROIC}$, and what remains is:

$$\text{Steady-state free cash flow} = \text{NOPAT} \times (1 - g / \text{ROIC})$$

This single line contains the whole rule. **The term g/ROIC is the fraction of operating profit that must be reinvested to sustain the assumed growth**, and it can never be zero unless growth is zero.

A business growing at eight percent while earning twenty percent on capital must reinvest 40 percent of its profit, permanently.

Read the same relationship the other way and it says something students find more useful still, because it is the form in which the question usually arrives.

$$g = \text{return on invested capital} \times \text{fraction of profit reinvested}$$

A business earning 20 percent on capital and reinvesting half its profit grows at 10 percent; one earning 10 percent and reinvesting the same half grows at 5. **Growth is therefore not an assumption that can be made independently of the rest of the model — it is an output of what the business earns and how much of it is put back.** A forecast that assumes high growth, high margins, and low reinvestment simultaneously is not aggressive; it is arithmetically impossible, and this is the line that exposes it. Chapter 3 has already met the same identity in equity terms, where sustainable growth is return on equity multiplied by the retention rate; this is the whole-capital version of it.

FRACTION OF NOPAT THAT MUST BE REINVESTED

RETURN ON CAPITAL	G = 4%	G = 6%	G = 8%	G = 10%
12%	33%	50%	67%	83%
18%	22%	33%	44%	56%
24%	17%	25%	33%	42%
30%	13%	20%	27%	33%

Read the table as a constraint on optimism. **A high growth assumption is only cheap for a business earning a high return on capital.** Assuming eight percent perpetual growth for a business earning twelve percent means committing two-thirds of its profit to reinvestment forever, and any model showing that combination alongside a large free cash flow is internally inconsistent.

The cost of ignoring this is not marginal. Take a business with NOPAT of 100 growing at eight percent with a twenty percent return on capital. Correctly, free cash flow in the steady state is 60. Assuming no capital expenditure, it appears to be 100. At a discount rate of 18 percent, the terminal value is 600 computed correctly and 1,000 computed carelessly.

THE SIZE OF THE ERROR

TREATMENT	STEADY-STATE FCF	TERMINAL VALUE AT R 18%, G 8%
Capital expenditure assumed to be zero	100	1,000
Capital expenditure sustaining the growth	60	600
Overstatement	67%	67%

Two-thirds too high, on the largest component of the valuation. **This is the most consequential single error in project and company valuation**, and it is not rare — it appears in professionally prepared models regularly, because the arithmetic is correct at every step and only the assumption underneath it is impossible.

The practical check takes ten seconds and should be run on every model you see. **Reconcile terminal total reinvestment to `g/ROIC`, then inspect capex, working capital and economically equivalent investment.** Capex materially below depreciation is an immediate warning, but capex equal to depreciation funds maintenance only; a growing perpetuity also needs growth reinvestment somewhere in the operating asset base.

IN THE VIDEOS *Video 8.4 computes one project's terminal value both ways, and the difference between them exceeds the entire net present value of the project.*

8.5 What the Discount Rate Is Doing Here

The rate applied to a project must compensate for the risks that project carries — not the risks the firm as a whole carries, and not the rate at which the firm happens to be able to borrow this month. Chapter 10 supplies the arithmetic; the question of which risks require compensation can be answered now, and answering it well is most of the work.

WHAT THE RATE MUST COVER

RISK	QUESTION	DIRECTION
Demand	Will the forecast volume appear	The largest source for most projects
Cost	Are inputs contracted or exposed	Imported inputs raise it, see Chapter 5
Execution	Has the firm done this before	A first attempt deserves more
Competitive response	Can others enter easily	Low barriers raise it
Contract quality	Is revenue committed or hoped for	Contracted revenue lowers it materially
Currency	Do costs and revenue share a currency	A mismatch raises it

Two errors follow from ignoring this table, and they are mirror images.

- **Using the firm's rate for an unusual project.** A stable distributor evaluating its first manufacturing venture at its ordinary rate will accept projects it should decline. The rate belongs to the cash flows, as Chapter 6 argued.
- **Using the cost of the debt raised to fund it.** A project funded with a loan at 15 percent is not a 15 percent project. The loan is cheap because it is senior and secured — the equity behind it bears the risk, and the project rate must reflect the whole capital structure.

Because the rate is a judgement, the honest output is a range. Present the answer at several rates and identify where the recommendation reverses. If a project is worth doing from 15 to 25 percent, nobody needs to settle the rate. If it flips at 19, then the rate is the decision, and the meeting should be about that rather than about the revenue forecast.

8.6 Decision Criteria and Their Failure Modes

Four criteria are in common use. One is reliable and three are not, and knowing precisely how each fails is the difference between using them and being misled by them.

THE FOUR CRITERIA

CRITERION	DEFINITION	PRINCIPAL WEAKNESS
Net present value	Value created, in today's money	None, given the right inputs
Internal rate of return	The rate at which NPV is zero	Several; see below
Payback period	Time until the outlay is recovered	Ignores everything after it
Profitability index	PV of inflows / initial outlay	Sensitive to what counts as the outlay

How the internal rate of return fails

The IRR is popular because it is quotable without agreeing a discount rate. It fails in four distinct ways, and three of them are common.

- **Scale.** A project returning 30 percent on 100 creates less value than one returning 25 percent on 500. At a 10 percent required return the first is worth 18 and the second 68. **IRR ranks them the wrong way round**, and mutually exclusive choices are exactly where ranking matters.
- **Timing.** Two projects with similar returns but different cash-flow shapes can be ranked differently by IRR and NPV, because IRR implicitly assumes reinvestment at the IRR itself, which is usually unrealistic.
- **Multiple roots.** A project whose cash flows change sign more than once can have more than one IRR. The flows -100, +230, -132 have a net present value of exactly zero at both 10 percent and 20 percent, and are positive only between them. Asking for "the" IRR of such a project has no answer.
- **No root.** Some flow patterns have no IRR at all.

The rule that follows is straightforward. **Decide with NPV. Quote IRR as a robustness measure** — it tells a reader how far the rate could move before the answer changes, which is genuinely useful information and exactly what a break-even rate is.

Payback and the profitability index

Payback measures how long the outlay takes to come back. It is a liquidity measure rather than a value measure and ignores everything after the payback date, which means it systematically rejects long-lived projects — precisely the ones that create the most value. Its **discounted** variant fixes one flaw by discounting the flows first; a project with four years of 40 against an outlay of 100 pays back in 2.5 years undiscounted and 3.15 years discounted at 12 percent. It still ignores everything afterwards.

Payback nonetheless has a legitimate use. In an environment where capital may be withdrawn, where policy may change, or where the currency may move sharply, **a shorter payback is worth something real that the NPV does not capture**. Treat it as a constraint alongside the decision rather than as the decision.

The profitability index — present value of inflows divided by the initial outlay — is useful in exactly one situation: when capital is genuinely rationed and a set of projects must be ranked for a fixed budget. Then it ranks by value created per unit of scarce capital, which is the right question. In every other situation NPV answers better.

IN THE VIDEOS *Video 8.6 computes all four criteria for the same project, which passes on two of them and fails on the others.*

8.7 Sensitivity, Scenario, and Break-Even

A project model is a machine for turning assumptions into a number, and the assumptions are estimates. Testing them is not a supplementary exercise; it is the analysis.

- **One-way sensitivity.** Move each input by a plausible amount and rank the results. The output is a list of what matters, and the ranking is more useful to a decision-maker than the base case.
- **Break-even.** For each significant input, find the value at which the NPV reaches zero. Then ask whether that value is plausible. "This project fails if volume comes in more than 12 percent below plan" is a sentence an operating manager can respond to; a net present value is not.
- **Scenario.** Move a coherent set of assumptions together in the directions they would actually move. A weak year brings lower volume, more discounting, and slower collection at the same time.

One general finding is worth carrying into every project you look at: **rank the impacts; do not assume the winner.** Volume affects the contribution margin on units gained or lost, while a unit-cost error can affect every unit still sold. Volume tends to dominate where contribution margins are high and fixed costs are large; unit costs can dominate in a thin-margin business with a high variable-cost base. The ranking is an economic description of the project, not a model-validity test.

The output of all this is a description of the project's character rather than a single number. A project that survives a wide range of inputs is robust and can be approved with ordinary care. A project that fails on a modest miss in one input is fragile, and the correct response is not to reject it but to **identify the input, name it, and decide whether it can be controlled or contracted.** That conversation is the reason the analysis was done.

IN THE VIDEOS *Video 8.7 ranks nine assumptions for one project and finds two that move the answer several times further than the rest.*

8.8 Reference

CHAPTER 8 • FORMULA REFERENCE

QUANTITY	FORMULA	NOTE
Incremental cash flow	With the project – without it	The definition of relevance
NOPAT	$\text{EBIT} \times (1 - \text{tax rate})$	Tax on EBIT, never after interest
Free cash flow	$\text{NOPAT} + \text{depreciation} - \text{capex} - \Delta \text{Working capital}$	Excludes all financing
Depreciation tax shield	$\text{Depreciation} \times \text{tax rate}$	Real money
Terminal value, perpetuity	$\text{FCF next year} / (r - g)$	Requires the capex rule below
Steady-state FCF	$\text{NOPAT} \times (1 - g/\text{ROIC})$	The consistent form
Required reinvestment	g / ROIC , as a share of NOPAT	Zero only if growth is zero
Growth identity	$g = \text{ROIC} \times \text{reinvestment rate}$	Growth is an output, not an assumption
Terminal value, exit multiple	$\text{Multiple} \times \text{final-year metric}$	Cross-check the perpetuity against this
Profitability index	$\text{PV of inflows} / \text{initial outlay}$	Only for rationed capital
Break-even rate	The rate at which $\text{NPV} = 0$	The IRR, used honestly

CARRY FORWARD

Financing never appears in the cash flows. Build the project as if funded entirely by equity and put all the financing into the rate, or you will charge for it twice.

A growing business must reinvest g/ROIC of its operating profit, forever. A terminal value that assumes perpetual growth with no capital expenditure is not conservative or aggressive — it is impossible, and it is typically wrong by more than half.

Reconcile terminal total reinvestment to g/ROIC . Comparing capex with depreciation is the first screen, not the complete economic test.

Decide with NPV; quote IRR as robustness. IRR misranks by scale, misranks by timing, and can have several values or none.