Course Companion

University Business

Every topic handout in one volume

exercises.lol

Contents

1. Introduction to Business and Management 2

2. Topic 2 13

3. Topic 3 24

4. Topic 4 35

5. Topic 5 45

6. Topic 6 57

7. Topic 7 66

8. Topic 8 77

9. Topic 9 88

10. Topic 10 98

11. Topic 11 110

12. Topic 12 124

Introduction to Business and Management

University Business

What this course is, and how to use it

This is the first course in University Business: an **evidence-based introduction to business and management** . It builds the vocabulary every later course assumes – accounting, economics, marketing, operations – while training one habit above all: separating **facts** from **assumptions** , and both from **conclusions** .

- Units 1.1 to 1.3 map the territory: organisations, stakeholders, ownership and the external environment.
- Units 1.4 to 1.6 look inside: decisions, business models and entrepreneurship.
- Unit 1.7 assembles the course's central skill – the evidence-based case analysis.

The running case in this handout is **Nimbus Bakery** , a fictional two-shop bakery: small enough to reason about completely, real enough that every framework earns its place. Numbers in worked examples are illustrative and say so. Each unit opens with a story, shows a picture before the formal idea, defines terms precisely, works two escalating examples whose arithmetic is machine-verified, names the mistakes people actually make, and closes with a three-line summary.

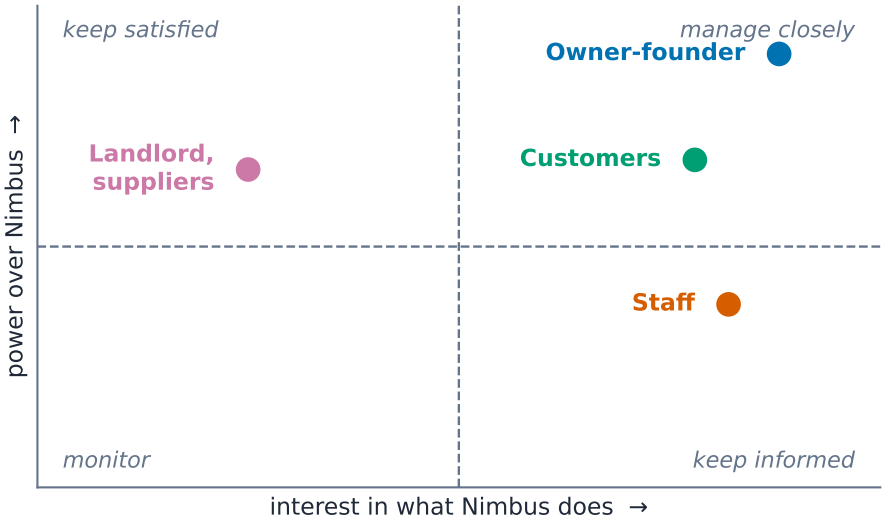
Organisations and stakeholders

Two friends bake brilliant bread and must decide: sell at a weekend market, contract independently with cafés, or open one shop together. The last option creates an **organisation** – and this unit explains when that is the cheap choice, and who then holds power over it.

An organisation is a group of people pursuing a purpose that outlasts any one member – firms, charities, public agencies. Why they exist at all: when transactions inside one firm are cheaper than contracting every task separately in the market (hiring a baker once beats negotiating every loaf), coordination through an organisation wins. The boundaries move – the same bakery might outsource delivery – wherever the comparison flips.

Stakeholders are the people and groups who affect, or are affected by, the organisation. Map them honestly – influence and interest differ:

Stakeholder	Main interest	Power over Nimbus
Owner-founder	profit, reputation, growth	decides
Customers	quality, price, convenience	choose competitors
Staff	fair pay, stable hours	can leave
Landlord, suppliers	paid on time	can withhold



Where each stakeholder sits decides how you manage them

Worked example 1 (read the map). Nimbus plans Sunday opening. Reading the grid: customers (high interest, moderate power – they can simply not come) need a reason – a promoted Sunday item; staff (high interest, low power) need notice and pay, and their interest can become power if two of four resign; the landlord (high power, low interest) needs nothing – Sunday opening is inside the lease. The map ranks the conversations: customers first, staff second, landlord not at all.

Worked example 2 (map the resistance, escalated). The same bakery considers a supplier switch to a cheaper flour mill. New stakeholder positions: the incumbent supplier (high interest – this is their revenue; low-to-moderate power – a contract with 60 days' notice); the new mill (low current interest, rising if it wins the tender); the owner's bank (low interest unless margins move the loan covenant). The map's decision content: the veto-free path exists, and the real risk is quality regression – a *customer* stakeholder effect, not a supplier one. A stakeholder map that only lists friends and enemies has missed its job: predicting where the surprise comes from.

Common mistakes. Treating every stakeholder as equally important – the whole point is ranking. Forgetting regulators and lenders until they object. Listing internal departments as external stakeholders.

Summary: organisations exist when internal coordination beats market contracting; stakeholders map on interest and power; the map ranks whose expectations to manage first and predicts where resistance will actually come from.

Ownership and governance

The friends open the shop together – and immediately face the question that shapes every firm's future: whose money is at risk, who decides, and who can sell out? The three main **ownership forms** differ on exactly those three axes – liability, finance and control:

Form	Liability	Finance	Control
Sole trader	unlimited, personal assets at risk	owner's savings, loans	one person
Partnership	joint and several (usually)	partners' capital	partners agree
Company	limited to what is invested	shares can be sold	shareholders elect directors

	Sole trader	Partnership	Company
Who decides?	the owner	the partners	directors, elected by shareholders
What can owners lose?	everything, home included	everything, even for a partner's debt	only what they put in
Money from	own savings, loans	partners' capital, loans	selling shares, loans
Paperwork	little	a partnership agreement	public accounts, a board

capped risk and outside money arrive together; the price is sharing control

Limited liability – shareholders can lose only what they invested – is the invention that made large-scale risk capital possible: strangers will fund a ship voyage if the worst case is their ticket, not their house. Its cost: **governance**, the system of boards, voting and disclosure that runs the company on the owners' behalf. **Agency problems** arise because managers (the agents) have their own interests – salaries, safety, status – and governance exists to keep them answerable to the shareholders whose money they spend.

Worked example 1 (the form follows the risk). Three bakery ownership options, priced: sole trader – the founder's flat is exposed if the business fails; partnership – each partner can bind the others (joint and several liability), fine while everyone bakes, lethal when one signs a flour contract alone; company – the founder pays incorporation and accounting costs but caps the loss at the investment. For a two-shop bakery with one lease and four staff, the company's protection usually earns its cost – but that is a judgement about *their* risk tolerance, and the analysis says so.

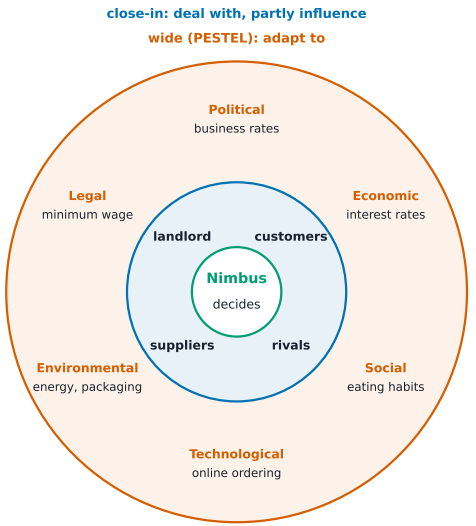
Worked example 2 (the arithmetic of each form, escalated). Year-end profit is £86,500. Partnership sharing 3:2 (agreed in writing): £51,900 and £34,600 – tax arrives on each partner's share whether or not it was drawn. As a company: the same £86,500, of which 25% (£21,625) is retained; the board declares the remaining £64,875 as dividend on 60,000 £1 shares – £1.08 per share. Same profit, different cash timing, different tax treatment, different control (the 3:2 partners each have a vote; the 60/40 shareholders elect the board). Ownership form is not paperwork – it reallocates money and power, and the numbers show how.

Common mistakes. Confusing limited liability with no responsibility – directors can still be personally liable for wrongful trading. Assuming the 50/50 partnership is safe: equal votes plus no tie-break clause equals deadlock. Reading "the shareholders own the company" as "the shareholders run the company".

Summary: liability, finance and control are the three axes; limited liability bought scale at the price of governance; profit splits and dividends are where the form shows up in arithmetic; agency problems are why governance exists.

The external environment

No firm chooses its weather. Bread demand can fall in a recession, a new low-carb trend can cut breakfast sales, a sugar tax can change the price list – none of it decided by the bakery. The **external environment** splits into the close-in (competitors, suppliers, customers, regulation) and the wide (economic cycles, technology, social change, law). Frameworks like PESTEL are checklists, not theories: their value is preventing you from forgetting a factor, not predicting it.



The discipline is evidence: claims about the environment need dates, sources and definitions. "Bread demand is falling" is a hypothesis until it carries a number, a period and a geography – and even then, ask what would have to be true for the trend to reverse. And ask of every driver: can the firm influence it, or only adapt to it? A bakery cannot move the business cycle; it can move its supplier list.

Worked example 1 (grade the claims). Three claims from a trade article: "Bakery closures up 12% (source: trade association, 2025, UK)" – a fact, but check the denominator and the definition of closure; "Consumers want more artisan products" – an assumption dressed as fact: which consumers, measured how?; "Rent inflation will continue" – a forecast, i.e. an assumption with a confidence interval nobody wrote down. Sorting those three is the unit's entire skill: facts carry sources, assumptions carry owners.

Worked example 2 (driver → decision, escalated). Take the two clean facts – closures up 12%, and (invented for the case) supermarket in-store bakery prices down 4% – and walk them to a decision. If closures release skilled bakers, Nimbus' staffing risk falls (a hiring opportunity); if supermarket prices fall, Nimbus cannot win a price war and must win on freshness and pre-order convenience – a *positioning* conclusion. Each link in that chain is marked: closures → labour supply (fact → assumption about mobility); prices → positioning (fact → strategic assumption). The environment unit ends not with a list of forces but with options the forces have priced.

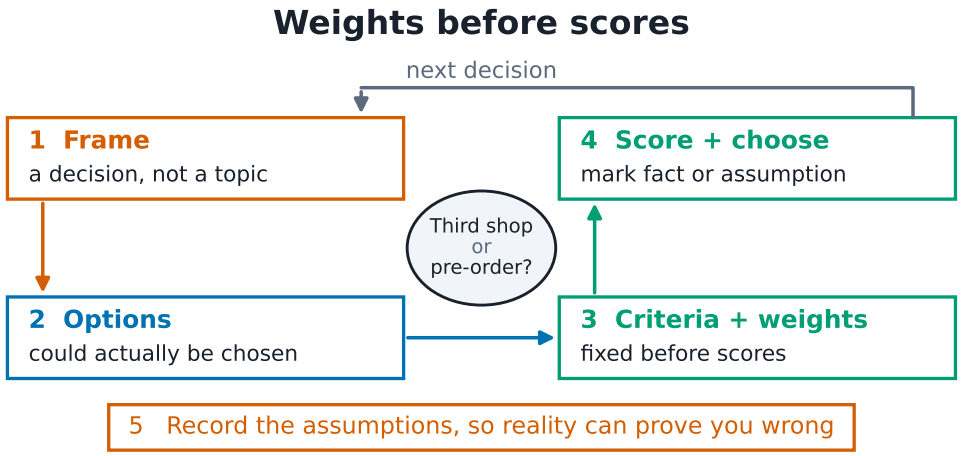
Common mistakes. Treating PESTEL as an analysis rather than a checklist. Quoting a percentage without its period, geography and source. Confusing "influence" with "predict" – the point is to prepare responses, not to foretell.

Summary: close-in versus wide environment; frameworks prevent forgetting, they do not explain; every environmental claim needs a number, a period, a geography and a source; walk each driver to a decision and mark the factual and assumptional links.

Managers, management and decisions

The bakery's founder faces the course's favourite dilemma: open a third shop, or expand pre-order? Both sound strategic; both have advocates; the choice will be made on Tuesday. **Management** is exactly this – deciding what to do and making it happen: planning, organising, leading and controlling – a loop, not a hierarchy. The course trains one piece deliberately: **structured decisions** .

1. Frame the question precisely (a decision to make, not a topic to discuss).
2. List options that could actually be chosen.
3. Name the criteria and their weights before seeing the options' scores.
4. Score the evidence, mark each score as fact or assumption.
5. Choose, then record what the decision assumes – so reality can prove you wrong.

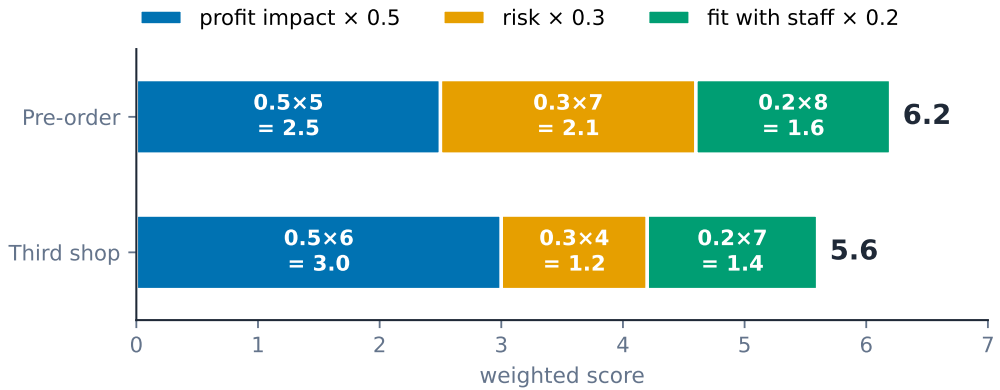


Weights fixed before scores, every score marked fact or assumption

Bias is the standing enemy: anchoring on the first number, confirming what you already believe, weighting the most recent event too highly. The structure above exists to slow those down, not because managers enjoy paperwork.

> **Worked example 1 (the decision table, verified).** Nimbus weighs a third shop against a pre-order expansion. Criteria and weights fixed first: profit impact (0.5), risk (0.3), fit with staff skills (0.2). Third shop scores 6, 4, 7 → $0.5 \times 6 + 0.3 \times 4 + 0.2 \times 7 = 5.6$. Pre-order scores 5, 7, 8 → $0.5 \times 5 + 0.3 \times 7 + 0.2 \times 8 = 6.2$. Pre-order wins – not because it is exciting, but because the weights said risk and fit matter.

> **Worked example 2 (the sensitivity that makes the table honest, escalated).** The founder's sister disputes the profit weight: "It's a business – weight profit 0.6." Recompute with weights 0.6 / 0.25 / 0.15: third shop $0.6 \times 6 + 0.25 \times 4 + 0.15 \times 7 = 5.65$; pre-order $0.6 \times 5 + 0.25 \times 7 + 0.15 \times 8 = 5.95$ – pre-order still wins. Solve for the flip point: with fit fixed at 0.2, the weighted totals cross at profit weight **0.65** (both 5.9). So the recommendation is robust to any profit weight below 0.65 and fragile above it – one number that tells the family exactly how far their disagreement has to go before the answer changes. That is what a decision table is for: making the disagreement discussable.



The weights, not the excitement, decide the winner

Common mistakes. Choosing weights after seeing scores – the table then only decorates the verdict. Scoring on adjectives ("strong") instead of anchored scales ("6 = payback under 3 years"). Forgetting step 5: a decision without recorded assumptions cannot be audited when it fails.

Summary: management is a loop of planning, organising, leading and controlling; structure the big decisions – frame, options, weights first, scored evidence, recorded assumptions; sensitivity analysis converts disagreement into one number.

Business models and value creation

A second shop in a richer district, a market stall on Saturdays, a pre-order subscription – each is a different **business model** : a description of how the organisation creates, delivers and captures **value** . The classic lens is the **value chain** – the sequence of activities from inbound ingredients to a served customer:

inbound logistics → operations → outbound logistics → marketing and sales → service



Nimbus' value chain, link by link

Nimbus' model: fresh-daily baking (operations) sold at premium prices through two shops and pre-order (marketing and channels), with staff cross-trained (service). Every link must fit: pre-order only pays if operations can batch it cheaply.

The test of a model is arithmetic, not adjectives: revenue per customer × customers must cover variable costs, fixed costs and the owner's minimum return – otherwise the model is a hobby.

> **Worked example 1 (the model as arithmetic, verified).** Nimbus' numbers: average basket £9.40; about 2,100 baskets per shop per week – 4,200 baskets weekly for two shops open six days. Weekly revenue $4{,}200 \times 9.40 = £39{,}480$. Ingredients and packaging run 42% of sales (variable), so contribution is $39{,}480 \times 0.58 = £22{,}898$. Fixed costs weekly – rent £6,200 (two leases), staff £2,600, energy and other £800 – total £9,600. Weekly profit $22{,}898 - 9{,}600 = £13{,}298$. Break-even: daily fixed $9{,}600/6 = £1{,}600$; contribution per basket $9.40 \times 0.58 = £5.45$; the two shops together need $1{,}600/5.45 \approx 294$ baskets a day, about 147 each, against the 350 each they sell. The model is a model only because that arithmetic closes – and every one of those figures is marked assumption until the accounts say otherwise.

> **Worked example 2 (profit is not cash, escalated).** A quieter month: sales made £34,000, of which £4,000 is a corporate customer still on invoice; expenses incurred £28,500, of which £600 is a supplier not yet paid. Profit (matching incurred against earned) is $34{,}000 - 28{,}500 = £5{,}500$. Cash movement is $(34{,}000 - 4{,}000) - (28{,}500 - 600) = £2{,}100$. Both numbers are true, both are needed, and the £3,400 gap is exactly the £4,000 receivable the corporate customer is holding, less the £600 Nimbus itself still owes the supplier. A profitable firm without cash fails just as dead: the value chain feeds the profit line, but the calendar of payments decides survival.

Common mistakes. Describing the model in adjectives with no unit economics. Confusing revenue (sales made) with profit (revenue minus expenses incurred) with cash (money in the bank). Testing a model's links separately – the chain only pays when the links fit together.

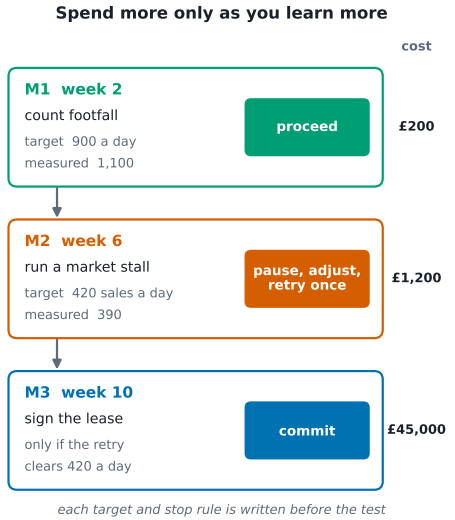
Summary: a business model is create-deliver-capture, mapped link by link on the value chain; the test is arithmetic – baskets × contribution vs fixed costs; revenue, profit and cash are three different things and all three must survive.

Entrepreneurship

Nimbus' founder is offered a market stall in a neighbouring town: six weekends, a few hundred pounds of fees, and no idea whether anyone there wants £9 hampers of sourdough. **Entrepreneurship** is pursuing an opportunity under uncertainty with resources you do not fully control – and this unit is the discipline that keeps the pursuit honest.

- Evaluate the **opportunity** before the passion: who exactly has the problem, how do they solve it today, and why would they switch?
- Size the resources honestly: what must be true for this to work, and which of those things is an assumption rather than a fact?
- Plan to learn: small tests first (a market stall before a second shop), milestones with dates, and an honest statement of what failure would look like.

Risk is measurable (probabilities known); **uncertainty** is not. Entrepreneurship lives in the second, which is why milestone planning beats five-year forecasting.



Worked example 1 (the stall test, priced honestly). Six weekends at the neighbouring market: stall fees £350 per weekend; flour, packaging and transport run 45% of sales. If takings are £3,800 across the six weekends, the result is $3{,}800 \times 0.55 - 350 \times 6 = 2{,}090 - 2{,}100 = -£10$ – break-even to the pound. The test was worth running precisely because the answer was marginal: it bought real demand data (which items sold, at what hours, to whom) for £10, a price no consultant would quote. Marked honestly: the £3,800 is a fact after the sixth weekend, an assumption before the first.

Worked example 2 (milestones and runway, escalated). The founder has £26,000 of savings allocated to the second shop and estimates burn (rent deposit fit-out, pre-opening wages) of £2,200 per week: runway $26{,}000/2{,}200 \approx 11.8$ weeks. The milestone plan that follows: signed lease by week 3 (or walk away with £19,400 intact); fit-out passed by week 7; opening stock and staff trained by week 10; open by week 11 with £1,800 contingency. Each milestone is a real-option: miss week 3 and the honest decision is to stop, not to compress every later stage into the remaining money. Entrepreneurship's arithmetic is the arithmetic of stopping rules.

Common mistakes. Falling in love with the solution before pricing the problem. Treating the business plan as a prediction rather than a list of testable assumptions. Spending to the last pound of runway because stopping feels like failure – running out mid-build is the actual failure.

Summary: opportunities are evaluated, not felt; small tests buy information at known prices; milestones with dates turn a runway into a sequence of stopping rules; risk is measurable, uncertainty is not, and entrepreneurship lives in the second.

Evidence-based case analysis

The course's final skill assembles all of the above into a **case analysis** :

1. **Situation**: what are the facts, with sources? What is merely assumed?
2. **Problem**: the decision to be made, stated as a question with options.
3. **Analysis**: apply the frameworks that earn their place – stakeholder map, value chain, decision table – each carrying evidence, not adjectives.
4. **Recommendation**: one option, chosen against stated criteria, with its risks named and its assumptions listed for monitoring.
5. **Review**: the metrics that would tell you the recommendation is failing.

Worked example 1 (the case in five lines). Nimbus' Sunday-opening case in the full frame: Situation – two shops, four staff, weekend competitors opening (sourced trade listing); Problem – "Should Nimbus open Sundays?" with options {open both shops, open one, not open}; Analysis – the decision table (worked example in 1.4) plus staffing interest from 1.1; Recommendation – open one shop, six-week trial; Review – Sunday baskets/day vs the 147-per-shop break-even, staff turnover, cannibalisation of Saturday. Every framework that appeared earned its place; none arrived as decoration.

Worked example 2 (assumption audit, escalated). Take the recommendation above and list what it assumes, each with a monitoring metric: customers will travel on Sundays (metric: first-four-Sunday baskets vs forecast); the trial will not shift Saturday demand (metric: Saturday baskets year-on-year); staff will accept the rota (metric: rota refusals). Three assumptions, three numbers, three dates. When a classmate says "the analysis was wrong", the audit answers *which part* – and that answer, not the original verdict, is what improves the next decision.

Common mistakes. Summarising the case instead of deciding in it. Recommending "more research" – name the assumption the research would test. Writing a review section with no numbers or dates; unfalsifiable metrics protect nobody.

Summary: situation-problem-analysis-recommendation-review; frameworks carry evidence or stay home; recommendations travel with named assumptions and dated, falsifiable review metrics – and every claim you meet, in a case, a report or a meeting, gets sorted into fact, assumption or conclusion.

Topic 02

University Business

What this course is, and how to use it

This is the second course in University Business: **financial accounting** – the discipline that turns a year of transactions into three statements a stranger can trust. Accounting is not arithmetic; it is a set of definitions applied consistently, and its central machine is the check: every statement in this handout **ties by construction** – the balance sheet balances because the arithmetic has no other choice, and the machine verification of every worked example confirms it.

The running case continues course 1: **Nimbus Bakery Ltd**, the two-shop bakery from course 1, now incorporated. On 1 January the founders buy the company back from themselves: they issue 20,000 £1 shares for cash and sign a £15,000 bank loan at 6% a year. Every unit works the same company forward – transactions, ledger, adjustments, statements, cash flow, ratios – so the numbers you meet in unit 2.4 are the ones you built in unit 2.1.

- Units 2.1–2.2 build the machine: the accounting equation and double entry.
- Units 2.3–2.4 apply time: accruals, adjustments, and the linked statements.
- Units 2.5–2.6 read the result: cash flow and ratios – and their limits.

How to read each unit: a story first, then a picture, then precise definitions, then two escalating worked examples whose numbers are machine-verified, the mistakes people actually make, and a three-line summary. IFRS-based framing; jurisdiction notes appear where they matter.

The accounting equation and elements

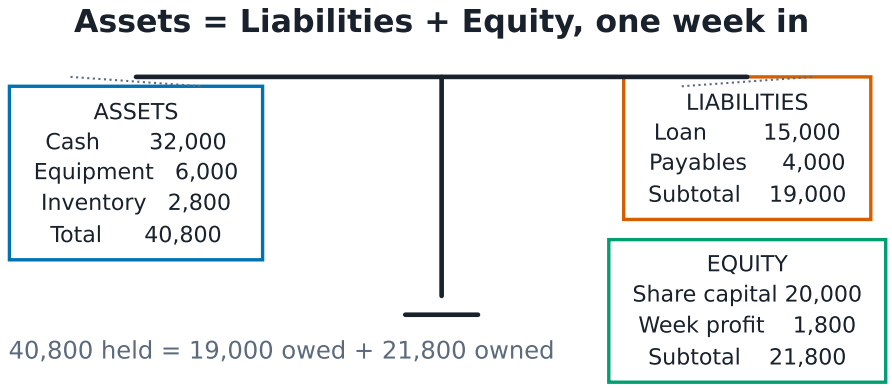
Nimbus's founders have £35,000 of starting money between them – savings and a loan approval. The moment it enters the company, a puzzle appears: the company *has* £35,000 and *owes* £35,000. To whom? To the people who funded it. That split – what the business holds against who claims it – is the whole of balance-sheet accounting, drawn as a scale that must tip level: the **accounting equation**

.

The five **elements** are defined precisely, not intuitively:

- An **asset** is a present economic resource controlled by the entity as a result of past events (the oven it bought, the cash it holds, the flour it will turn into bread).
- A **liability** is a present obligation to transfer an economic resource (the bank loan, the supplier not yet paid).
- **Equity** is the residual: assets minus liabilities – the funders' claim.

- **Income** increases economic benefit and grows equity; an **expense** decreases it. Profit = income − expenses, and profit flows into equity.



The scale always tips level: what is held equals who claims it

The derivation. Every transaction is recorded by its **double effect** on the equation – that is not a rule someone invented but the only way the equation can survive. Issuing shares: an asset rises (cash) and equity rises (share capital) – both sides grow. Buying an oven for cash: one asset swaps for another – the total is unchanged. Buying flour on credit: an asset rises and a liability rises. Selling bread: cash rises (asset) and profit rises (equity), while the flour used leaves the business (inventory falls, equity falls through cost of sales). Five moves, one invariant: Assets = Liabilities + Equity, always.

> **Worked example 1 (the first week, one line at a time).** January, day by day:

1. Issue 20,000 shares at £1: cash +20,000, equity +20,000. Check: 20,000 = 0 + 20,000 ✓.
2. Take the £15,000 loan: cash +15,000, liability +15,000. Check: 35,000 = 15,000 + 20,000 ✓.
3. Buy the oven, £6,000 cash: equipment +6,000, cash −6,000. Check: 35,000 = 15,000 + 20,000 ✓ (an asset swap changes no total).
4. Buy flour and ingredients, £4,000 on credit: inventory +4,000, payables +4,000. Check: 39,000 = 19,000 + 20,000 ✓.
5. The week's bread sells for £3,000 cash, using £1,200 of ingredients: cash +3,000, inventory −1,200, equity +1,800 (revenue 3,000 − cost of sales 1,200). Check: 40,800 = 19,000 + 21,800 ✓.

Worked example 2 (escalated: the end-of-week balance sheet). Post the five transactions to the closing position: **assets** = cash 32,000 + equipment 6,000 + inventory 2,800 = 40,800; **liabilities** = loan 15,000 + payables 4,000 = 19,000; **equity** = share capital 20,000 + week's profit 1,800 = 21,800. The tie-out: 19,000 + 21,800 = 40,800 ✓ (machine-verified). Two readings of the same number: the £40,800 is both what Nimbus *holds* and who *claims* it – the bank and supplier claim £19,000; the founders' stake has grown from 20,000 to 21,800 in one week, and the growth is exactly the profit. The equation is not a formula to memorise; it is two lists of the same company.

Common mistakes. Calling the owner's investment an income (it is a capital contribution – equity, not revenue). Reading "asset" as "valuable" (the definition requires control and past events – a brilliant baker's reputation is not yet an asset). And thinking the equation can go out of balance through error – it cannot; a transaction misposted breaks one side, which is exactly how errors are found (unit 2.2).

Summary: five elements, one invariant – assets equal liabilities plus equity; every transaction is a double effect; profit is income minus expenses and lands in equity; the balance sheet is two lists of one company.

Double entry and ledgers

The equation says a transaction has two effects; the **ledger** makes the two effects *postable*. Each element gets accounts shaped like a T: the left side is the **debit**, the right the **credit**. The rule that falls out of the equation: **assets and expenses grow on the debit side; liabilities, equity and income grow on the credit side** – and every transaction posts equal debits and credits. That symmetry, not habit, is why the method is called **double entry**.

The week posted: debits left, credits right

Cash	
20,000	6,000
15,000	
3,000	
38,000	**6,000**

Equipment	
6,000	
6,000	**0**

Inventory	
4,000	1,200
4,000	**1,200**

Cost of sales	
1,200	
1,200	**0**

Loan	
	15,000
0	**15,000**

Payables	
	4,000
0	**4,000**

Share capital	
	20,000
0	**20,000**

Revenue	
	3,000
0	**3,000**

Trial balance: Dr 32,000 + 6,000 + 2,800 + 1,200 = 42,000
= Cr 15,000 + 4,000 + 20,000 + 3,000 = 42,000

Debits left, credits right; the trial balance adds each column and demands equality

The derivation. Write the accounting equation as a grid with assets and expenses naturally on the left, liabilities, equity and income on the right. A transaction that increases an asset posts a left-side (debit) entry; one that increases a liability posts a right-side (credit) entry; decreases post on the opposite side. Because every transaction preserves the equation, the sum of all debit entries must equal the sum of all credit entries – the **trial balance** is that sum, and its equality is a theorem, not a hope.

Worked example 1 (post the week, extract the trial balance). The five transactions of unit 2.1, posted:

- Dr Cash 20,000 / Cr Share capital 20,000; then Dr Cash 15,000 / Cr Loan 15,000; then Dr Equipment 6,000 / Cr Cash 6,000; then Dr Inventory 4,000 / Cr Payables 4,000; then Dr Cash 3,000 / Cr Revenue 3,000 and Dr Cost of sales 1,200 / Cr Inventory 1,200.

Closing balances: Cash 32,000 Dr; Equipment 6,000 Dr; Inventory 2,800 Dr; Cost of sales 1,200 Dr; Loan 15,000 Cr; Payables 4,000 Cr; Share capital 20,000 Cr; Revenue 3,000 Cr. **Trial balance: debits 42,000 = credits 42,000** ✓ (machine-verified – 32,000 + 6,000 + 2,800 + 1,200 against 15,000 + 4,000 + 20,000 + 3,000).

Worked example 2 (escalated: the errors a trial balance catches – and the ones it cannot). Three deliberate errors, each diagnosed:

- *Transposition.* A café sale of £3,100 posted to revenue as £1,300 (cash took the right 3,100): the trial balance differs by 1,800. The classic clue: the difference is divisible by 9 (1,800 = 9 × 200) – swapping two digits changes a number by a multiple of 9, so a ÷9 difference sends you hunting for transposed digits.
- *One-sided omission.* The £4,000 flour purchase never posted to payables: the difference equals the missing entry exactly (4,000), pointing at one account's missing line.
- *The invisible one.* Wages of £52,000 posted twice: debits and credits both grow by 52,000 – the trial balance still ties, and the profit is understated by 52,000 with no arithmetic alarm. The lesson: the trial balance catches *posting asymmetry*, not misclassification or duplication – the review layer (unit 2.3's reconciliations) must catch the rest.

Common mistakes. Reading "debit" as "minus" – the words mean left and right, nothing more; whether a debit grows or shrinks an account depends on which side of the equation the account lives. Chasing a difference without first testing it: a difference equal to one entry suggests that entry's other side was never posted; a difference whose half equals an entry suggests that entry was posted on the wrong side (two debits, say); a difference that divides by 9 suggests transposed digits. And trusting a balanced trial balance as proof the books are right – example 2's third error is the counter-example.

Summary: T-accounts, debits left and credits right; assets and expenses grow on the debit, the other three on the credit; the trial balance is a theorem about symmetry; it

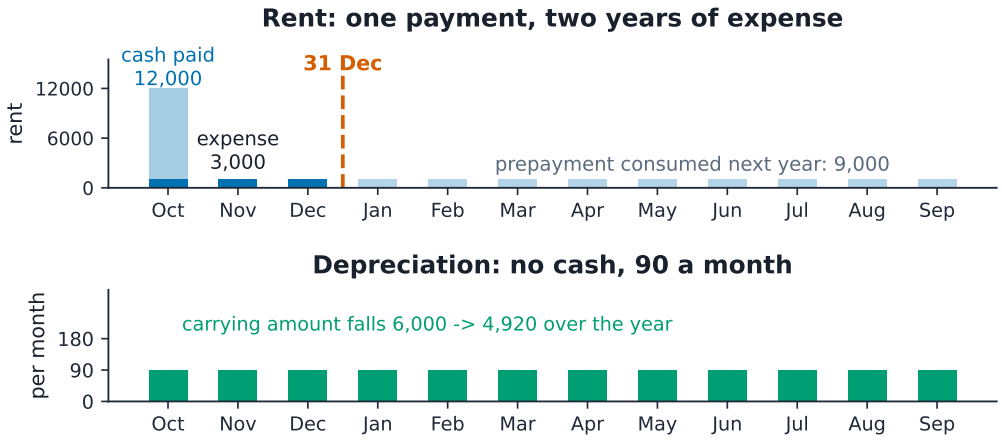
catches asymmetry, never misclassification.

Accruals and adjustments

The ledger records what happened; the year-end question is *when* it happened. Nimbus pays a year's shop rent on 1 October; the **accrual basis** says the expense belongs to the months the shops were occupied, not the month the cash left. Adjusting is the act of moving entries from the cash's timeline to the work's timeline.

The four recurring adjustments, defined:

- **Prepayment**: cash paid before the expense – the unused part is an asset (Nimbus has bought a right to occupy).
- **Accrual**: expense incurred before the cash is paid – a liability (wages earned by staff in December, paid in January).
- **Depreciation**: the cost of a long-lived asset spread over its useful life – Nimbus's oven, straight-line: (cost − residual) ÷ years.
- **Revenue receivable**: work done but not yet billed or collected – an asset, treated here only when Nimbus starts selling to cafés on credit.



Cash has a timetable; expense has another – adjustments reconcile them

The derivation. Straight-line depreciation is arithmetic honesty: the oven cost £6,000 and will be worth £600 after five years, so five years consume £5,400 of service – £1,080 a year, £90 a month. The accrual logic is the matching principle : expenses are recognised when the associated revenue is earned or the benefit is consumed – which is why the £12,000 rent payment on 1 October leaves only £3,000 as this year's expense, with £9,000 standing on the balance sheet as a prepayment asset.

Worked example 1 (three adjustments in isolation).

- Rent: annual £12,000 paid 1 October for the year to next 30 September; the year end is 31 December. Expense for the three months = 12,000 × 3/12 = **3,000**; prepayment asset = 12,000 × 9/12 = **9,000**. Check: 3,000 + 9,000 = 12,000, the cash paid ✓.
- Depreciation: (6,000 − 600)/5 = **1,080** a year; the oven carrying amount at 31 December = 6,000 − 1,080 = **4,920**.
- Loan interest: 15,000 × 6% = **900** for the year, payable each 1 January – an accrual liability at the year end because the bank's service was consumed before it is paid.

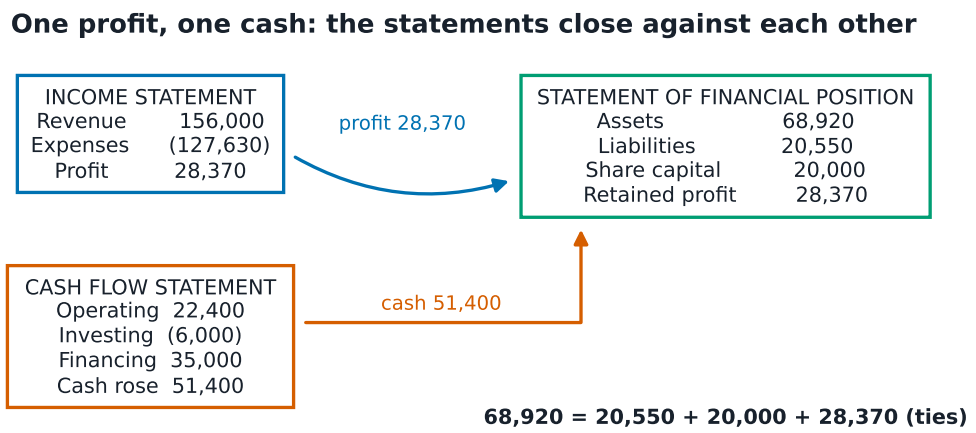
Worked example 2 (escalated: the full year-end adjustment set). Year 1 in full: cash sales 156,000; inventory purchased 66,000, of which 3,600 remains (cost of sales = 66,000 − 3,600 = 62,400); wages 52,000 of which 400 is accrued (staff December rota paid in January); rent 12,000 paid as above; other running costs 8,000 paid; the electricity bill for the quarter, 250, received but unpaid; the oven's first-year depreciation 1,080; loan interest 900 accrued. **Profit for year 1** = 156,000 − 62,400 − 52,000 − 3,000 − 8,000 − 1,080 − 250 − 900 = **28,370** (machine-verified). Every adjustment is visible in it: the rent is 3,000, not 12,000; the oven costs 1,080, not 6,000; the electricity and interest appear although no cash has moved. One trap: the wages figure of 52,000 *already includes* the unpaid 400, so the accrual changes the cash paid (51,600), not the expense. Subtracting another 400 is the classic year-end slip, and it understates profit by exactly 400.

Common mistakes. Double-counting an accrual already inside a stated expense total (the difference between "wages paid 51,600" and "wages expense 52,000" is the accrual – state which you were given). Depreciating to zero rather than to residual value. And calling the prepayment an expense because "the cash has gone" – cash timing is unit 2.5's business; this unit matches work.

Summary: adjustments move entries from the cash timeline to the work timeline; prepayment is an asset, accrual a liability, depreciation spreads cost over life; profit is what the matching principle leaves behind.

Financial statements

Three statements, one company, zero independent numbers: the **income statement** answers "what did the year earn", the **statement of financial position** answers "what stands at the year end", and the link between them is a single number – profit – flowing into equity as retained earnings. Build them for Nimbus's year 1 and watch the tie-out.



One profit, one cash: the statements share numbers by construction

The derivation. The income statement is the equity side of the year in motion: opening equity 20,000, plus income 156,000, minus expenses 127,630, equals closing equity before dividends 48,370 – of which 28,370 is the year's retained profit. The statement of financial position is the same equation at a date, with the profit already inside: assets 68,920 = liabilities 20,550 + equity 48,370. The cash flow statement (unit 2.5) reconciles the other balance-sheet line that moved – cash.

Worked example 1 (the income statement, in a form a non-accountant reads).

Nimbus Bakery Ltd – income statement, year 1	£
Revenue	156,000
Cost of sales	(62,400)
Gross profit	**93,600**
Wages	(52,000)
Rent (3 of 12 months)	(3,000)
Other running costs	(8,000)
Depreciation	(1,080)
Electricity accrued	(250)
Loan interest accrued	(900)
Profit for the year	**28,370**

Check: gross profit 60.0% of revenue; the eight expense lines re-derive 28,370 exactly ✓.

Worked example 2 (escalated: the statement of financial position that closes to the penny). Track the cash: inflows 20,000 shares + 15,000 loan

- 156,000 sales = 191,000; outflows 6,000 oven + 62,000 paid of purchases (66,000 bought, 4,000 owing) + 51,600 paid of wages + 12,000 rent + 8,000 other = 139,600; **closing cash 51,400**. The statement of financial position at 31 December:

Assets	£	Liabilities and equity	£
Cash	51,400	Loan	15,000
Inventory	3,600	Payables	4,000
Prepaid rent	9,000	Accrued wages	400
Equipment (6,000 − 1,080)	4,920	Electricity owing	250
		Accrued interest	900
		Liabilities	**20,550**
		Share capital	20,000
		Retained profit	28,370
		Equity	**48,370**
Total	**68,920**	**Total**	**68,920**

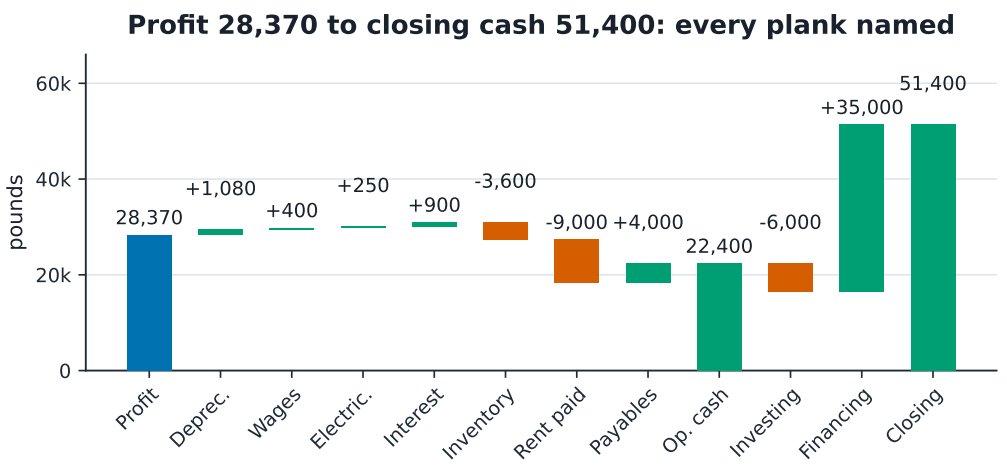
The tie: 20,550 + 48,370 = 68,920 ✓ (machine-verified). Every number is one you built: the 9,000 is unit 2.3's prepayment, the 4,920 its carrying amount, the 28,370 its adjusted profit. Nothing new is invented at statement stage – the statements only rearrange the ledger.

Common mistakes. Dropping the depreciation from the balance sheet's equipment line (it belongs to both statements: expense and carrying-amount reduction). Putting the prepayment under liabilities because it feels like "money already spent". And hunting for the "missing" profit number when the balance sheet will not tie – it is always one of: an unposted adjustment, a double-counted accrual, or dividends (which reduce equity *outside* the income statement, as unit 2.6 will show).

Summary: income statement = the year's revenue and matched expenses; balance sheet = the year-end equation with profit inside equity; the two share every number by construction; dividends reduce equity directly.

Cash flow statements

Nimbus made 28,370 of profit and its cash rose by 51,400 – and neither number explains the other until you build the bridge. The **cash flow statement** classifies every movement into **operating** (the bread business), **investing** (buying long-lived assets) and **financing** (shares and loans), and reconciles profit to operating cash – the single most revealing arithmetic in the statements.



Profit to cash is a bridge with named planks: depreciation, accruals, working capital

The derivation. Profit is measured on the accrual basis; cash is not. The bridge undoes the accruals one plank at a time: add back depreciation (an expense that moves no cash), add accruals (expenses recognised before paid), subtract the growth in inventory and prepayments (cash spent on things not yet consumed), add the growth in payables (goods consumed but not yet paid). What remains is the cash the operations actually generated.

> **Worked example 1 (the bridge, plank by plank).** Year 1, indirect method: profit 28,370 + depreciation 1,080 + accrued wages 400 + electricity 250 + accrued interest 900 − growth in inventory 3,600 − growth in prepaid rent 9,000 + growth in payables 4,000 = **operating cash 22,400** (check the signs: expenses-without-cash add back; cash-without-expense subtracts; payables growing is cash kept). Then investing: oven −6,000. Financing: shares +20,000, loan +15,000. Net change 22,400 − 6,000 + 35,000 = **51,400** – and the balance sheet's cash line says 51,400 ✓ the statements close against each other, machine-verified.

> **Worked example 2 (escalated: why profit and cash diverge).** Three divergence engines, each visible in Nimbus's numbers. *Timing*: the rent prepayment moved 9,000 of cash a year early. Profit 28,370 against operating cash 22,400 – a 5,970 gap, and the bridge accounts for it exactly: the 9,000 prepaid rent and 3,600 inventory build (12,600 of cash absorbed) against 1,080 of depreciation, 1,550 of accruals and 4,000 of payables growth (6,630 of non-cash offset), and 12,600 − 6,630 = 5,970 ✓ (machine-verified). *Growth*: a business that doubles its inventory doubles the cash trapped in it while profit looks untouched. *Capital structure*: the 35,000 of financing cash has nothing to do with profit at all. Year 2 will exercise the timing engine hard when Nimbus starts selling to cafés on 24-day credit.

Common mistakes. Adding back the whole rent payment as "an adjustment" (only the depreciation and accruals are added back; the prepayment lives in the working-capital lines). Classifying loan repayment as operating (it is financing; only the *interest* is operating). And reading a strong operating cash number as proof of health without asking what financing is doing – a company can look cash-rich while the bank is calling the loan.

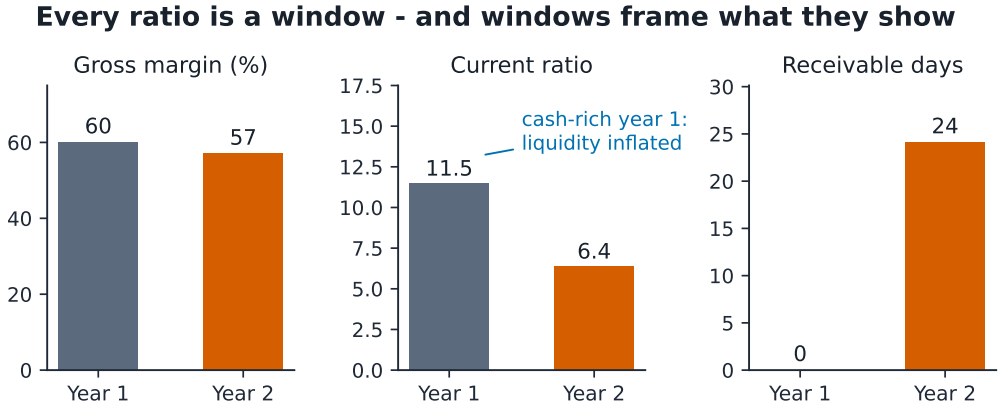
Summary: three classes – operating, investing, financing; the bridge from profit adds non-cash expenses and accruals, then moves working capital; divergence has named engines: timing, growth, capital structure.

Ratios and their limits

The statements are built; ratios are the reading. A ratio is one statement number divided by another – margin, liquidity, gearing, return – and each is a *question*, not a verdict. Nimbus's year-1 numbers make the point forcefully, because a young company's ratios lie in both directions.

The core set, defined:

- **Gross margin** = gross profit ÷ revenue – pricing power and production cost discipline.
- **Operating margin** = profit ÷ revenue – what is left of each pound after the whole running of the business.
- **Current ratio** = current assets ÷ current liabilities – short-term cover.
- **Gearing** = debt ÷ (debt + equity) – how much of the company is loan-financed.
- **Return on equity** (ROE) = profit ÷ equity – the funders' rate of return.



Every ratio is a window – and windows frame what they show

Worked example 1 (the year-1 read, computed). From the statements: gross margin = 93,600/156,000 = **60.0%**; net margin = 28,370/156,000 = **18.2%**; current assets = 51,400 + 3,600 + 9,000 = 64,000 against current liabilities of 4,000 + 400 + 250 + 900 = 5,550, so current ratio = **11.5**; gearing on capital employed = 15,000/(15,000 + 48,370) = **23.7%**. The immediate smell test: a current ratio of 11.5 is not "excellent", it is *absurd* – £64,000 of short-term resources against £5,550 of bills means most of that cash is idle, and idle cash earns nothing. The ratio did not lie; the reader who stopped at "above 2 is good" did.

Worked example 2 (escalated: year 2 changes the story, and the limits bite). Year 2: Nimbus takes café contracts – revenue 182,000, cost of sales 78,300 (gross margin **57.0%** – the wholesale price of café bread is below retail), profit **30,100**; at the year end receivables are 12,000, inventory 4,400, cash 22,000 (a £20,000 dividend was paid and £5,000 of the loan repaid – so equity = 48,370 + 30,100 − 20,000 = **58,470** and the loan stands at 10,000, gearing **14.6%**). Working-capital ratios wake up: receivable days = 12,000/182,000 × 365 = **24.1 days** (the credit term is 24 – collection is on target); inventory days = 4,400/78,300 × 365 = **20.5 days**. Now the limits, honestly: the year-1 ROE of 28,370/48,370 = 58.7% says more about the founders' small £20,000 stake than about the bakery: a first year's profit set against a capital base that existed only from day one inflates the return (on average equity it is 83%, no better); the year-2 56.3% on average equity is still distorted by the founders' small stake. Ratios need *time* (trends, not levels), *peers* (a 60% gross margin means one thing in bread, another in software), and *definitions* (check which margin, which equity, which debt before comparing a single number).

Common mistakes. Averaging ratios instead of computing on averages. Comparing across sectors or definitions without checking them. And the year-one trap this unit's numbers display: reading liquidity and return off a balance sheet whose cash was raised, not earned. Every ratio in this unit was computed from statements you built – the discipline is the same at any size.

Summary: margins ask about pricing, liquidity about cover, gearing about funding, return about the funders; a ratio is a question asked of two numbers – time, peers and definitions decide whether it can be answered.

Topic 03

University Business

What this course is, and how to use it

This is the third course in University Business: **microeconomics** – the study of how people and firms choose under scarcity, and how those choices meet in markets. The subject is a pair of habits, not a pile of curves: state the **assumption** before using a model, and check the **incentive** before predicting a response. Every diagram in this handout is paired with the algebra behind it, and every worked example is **machine-verified** – the equilibrium, the elasticity, the cost table and the welfare triangle all recompute exactly as printed.

The running case stays with **Nimbus Bakery Ltd** and its city. The market is artisan bread: daily demand and supply curves for a loaf, then Nimbus's own ovens and bakers as the production story, then the city's market structure, and finally the flour mill upstream whose smoke makes the textbook's market failure concrete. One set of numbers runs through all six units – the £4 loaf you derive in unit 3.2 is the price Nimbus takes as given in unit 3.5.

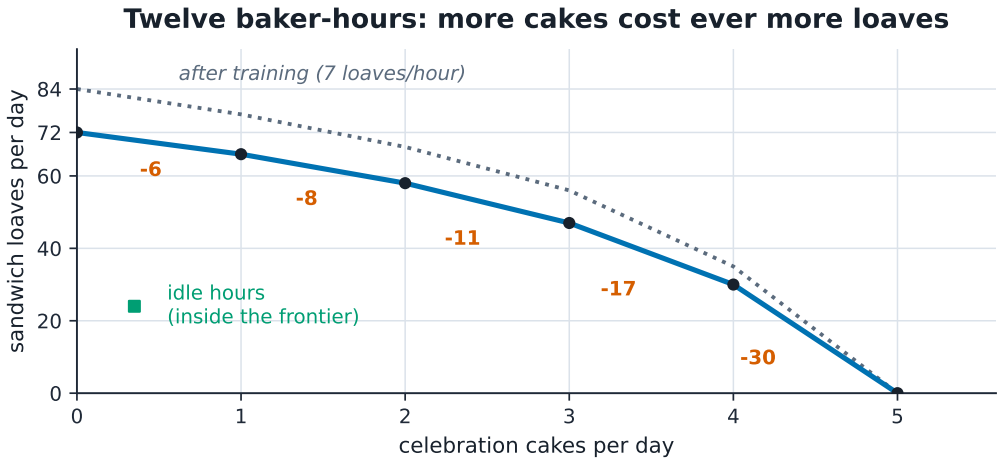
- Units 3.1–3.3 build the market: choice, demand and supply, elasticity.
- Units 3.4–3.5 go inside the firm: costs, then pricing by market structure.
- Unit 3.6 asks when the market answer is wrong – and what intervention can and cannot repair.

How to read each unit: a story first, then a picture, then precise definitions, then two escalating worked examples whose numbers are machine-verified, the mistakes people actually make, and a three-line summary.

Scarcity and opportunity cost

Why it matters. Nimbus has 12 baker-hours a day and two products that compete for them. Everything microeconomics will say later rests on this one situation: to choose anything is to give something up, and the cost of a choice is the best alternative it displaces – not the money, the alternative.

The idea, with a picture. The **production possibilities frontier** draws every combination of celebration cakes and sandwich loaves the 12 hours can produce, given the bakery's recipes and skill. Points on the curve use everything; points inside waste time; points beyond need more hours or better skill. The curve bows outward because bakers are not identical: the first cake is made by the baker best at cakes, the last cake would have to be made by someone brilliant at loaves.



Each extra cake costs more loaves than the one before – the slope is the opportunity cost

The definitions. **Scarcity** means wants exceed resources at any moment. **Opportunity cost** is the value of the next-best alternative given up – the loaf hours, not the flour bill. A model's **ceteris paribus** clause ("all else equal") is a promise about what is being held fixed, and every prediction in this course inherits it.

The derivation. With 6 loaves per hour, loaves = 6 × hours left after cakes. Cake hours rise with skill mismatch, so the frontier through the five verified points (cakes, loaves) = (0, 72), (1, 66), (2, 58), (3, 47), (4, 30), (5, 0): each extra cake gives up 6, then 8, then 11, then 17, then 30 loaves – **increasing opportunity cost**, which is the same statement as the outward bow.

> **Worked example 1 – the second cake.** Nimbus currently bakes one cake a day. The festival asks for a second. The cost is not flour and butter; it is 8 loaves – the hours the second-best cake baker would have spent at 6 loaves an hour. Priced at £4 a loaf, the second cake's opportunity cost is £32, and it is worth baking only if the cake sells for more than £32 plus its ingredient bill. The decision rule generalises: expand until marginal benefit equals opportunity cost.

> **Worked example 2 – escalate to growth.** The council offers a training grant that lifts every baker to 7 loaves an hour. The frontier shifts out to (0, 84), (1, 77), (2, 68), (3, 55), (4, 35), (5, 0) – check the first segment: giving up 7 loaves, not 6. Growth does not remove scarcity: the curve still binds, choices still trade off, and the opportunity cost of cake number 3 has risen from 11 loaves to 13. What changed is the menu, not the necessity of choosing.

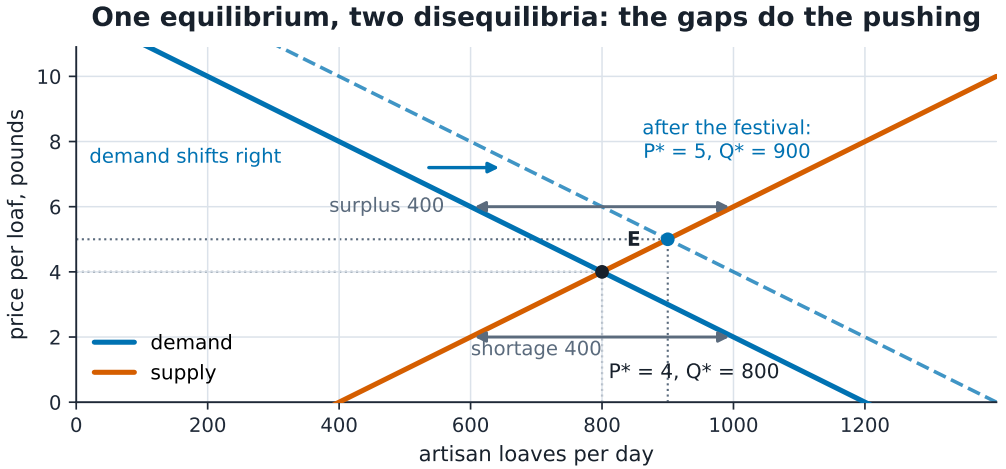
Common mistakes. Counting the money spent instead of the alternative forgone (the flour was bought either way – sunk). Drawing a straight-line frontier and then talking about increasing cost – a straight frontier has constant cost. And reading "efficiency" as one point: every point ON the curve is efficient; which one is best depends on the customer, not the curve.

Summary. Cost = next-best alternative, not cash out. The frontier's slope is opportunity cost; the bow is increasing cost. Growth shifts the menu; choice remains.

Demand and supply

Why it matters. The city's artisan-bread market is where Nimbus's £4 loaf comes from. Demand and supply is a machine with three parts – two curves and their interaction – and almost every microeconomic question is a comparative statics exercise on it: shift one curve, recompute, compare.

The idea, with a picture. Daily demand for artisan loaves is $Q_d = 1200 - 100P$; daily supply is $Q_s = 400 + 100P$ (prices in £). The curves cross once: at £4, buyers take 800 loaves and bakers offer exactly 800. Below £4 the plans conflict – buyers want more than bakers will bake – and the shortage bids the price up; above £4 a surplus pushes it down. The crossing is the only price where nobody has an incentive to change behaviour.



One equilibrium, two disequilibria: the gaps do the pushing

The definitions. **Demand** is the whole relationship between price and quantity buyers plan to buy; **quantity demanded** is one point on it. A price change MOVES ALONG the curve; anything else that shifts buyer plans (income, tastes, other prices, expectations) SHIFTS the curve – the two vocabulary errors this unit exists to kill. **Supply** and **quantity supplied** split the same way for sellers. **Equilibrium** is the price where plans agree; **comparative statics** is the discipline of changing one thing, holding the rest, and comparing equilibria.

The derivation. Set $Q_d = Q_s$: $1200 - 100P = 400 + 100P$, so $800 = 200P$ and

$P^* = £4$, $Q^* = 800$. The same algebra finds any new equilibrium after a shift – and the shortage at £2 ($Q_d - Q_s = 1200 - 200 - (400 + 200) = 400$ loaves) and the surplus at £6 (also 400 the other way) come from the same two lines.

> **Worked example 1 – the festival raises incomes.** A week-long festival lifts city incomes; at every price buyers now want 200 more loaves, so demand becomes $Q_d = 1400 - 100P$. New equilibrium: $1400 - 100P = 400 + 100P$ gives $P^* = £5$, $Q^* = 900$ – both price and quantity rise, which is the signature of a demand shift outward along an upward supply. Nimbus, watching its till, sees the £1 without needing to know why – the market aggregated the festival for it.

> **Worked example 2 – escalate to a supply shock.** Flour prices fall; at every price bakers offer 200 more loaves: $Q_s = 600 + 100P$. New equilibrium: $1200 - 100P = 600 + 100P$ gives $P^* = £3$, $Q^* = 900$ – quantity rises again but the PRICE now falls, the signature of a supply shift. The two shocks together (both curves shift right) give $P^* = £4$ again with $Q^* = 1000$: quantities add, price effects cancel – a check students can run in one line before drawing anything.

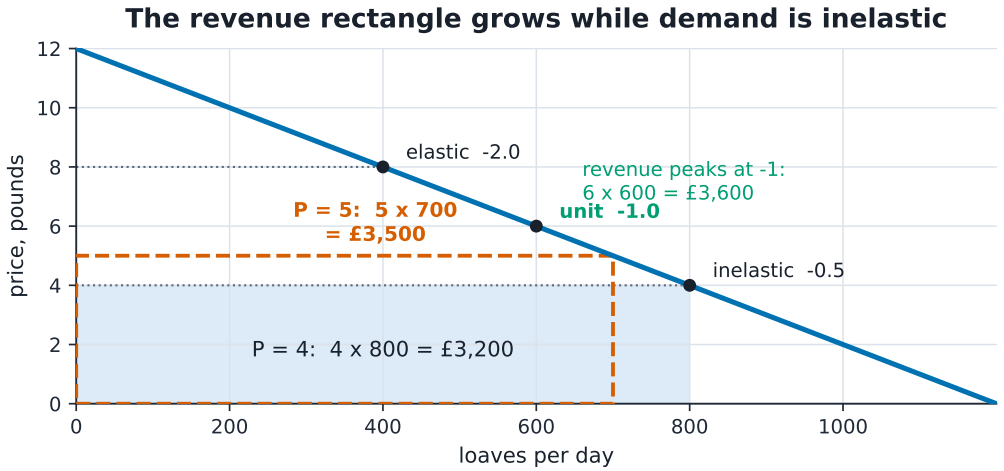
Common mistakes. Saying "demand rose so the demand curve shifted up" and then ALSO moving along it – pick one mechanism per event. Confusing a shift of supply with a movement along supply caused by the new price. And predicting quantity without the second curve: neither curve alone pins down an equilibrium.

Summary. Price changes move along; other changes shift. Equilibrium solves two equations. Each single-curve shock moves price and quantity in its own signature direction.

Elasticity

Why it matters. Direction is not enough. When the festival pushed the price to £5, Nimbus's revenue ROSE – but a taxi firm raising fares the same 25% might lose revenue. Elasticity is the number that tells you which world you are in, and it is the firm's first pricing question.

The idea, with a picture. The **price elasticity of demand** is the percentage response of quantity to a percentage change in price, $\varepsilon = \frac{\Delta Q/Q}{\Delta P/P}$. On the bread demand at £4: $\varepsilon = (-100 \times 4)/800 = -0.5$ – quantity responds half as fast as price in percentage terms, so a higher price RAISES revenue: the price gain beats the quantity loss. The revenue rectangles make it visible: at £4 the rectangle is £4 × 800 = £3,200; at £5 it is £5 × 700 = £3,500 – taller and barely narrower.



Inelastic demand: the revenue rectangle grows as the price rises – until the midpoint

The definitions. Demand is **inelastic** when $|\varepsilon| < 1$ (quantities respond less than proportionally), **elastic** when $|\varepsilon| > 1$, **unit elastic** at exactly 1. On a straight-line demand, elasticity falls from infinity at the top to zero at the bottom – "elastic" and "inelastic" label POINTS, not whole curves. The **income elasticity** and **cross elasticity** apply the same machinery to income changes and other goods' prices – they classify goods: normal or inferior, substitutes or complements.

The derivation. For a linear demand $Q = a - bP$, the point elasticity is $\varepsilon = -b \cdot P/Q$. Since Q shrinks as P rises, $|\varepsilon|$ grows along the curve; total revenue $P \cdot Q$ is maximised exactly where $\varepsilon = -1$ (the rectangle can no longer grow) – which is the midpoint of a linear demand. Elasticity is a ratio of percentage changes, so it has no units and can compare bread with taxi rides.

> **Worked example 1 – the bakery's pricing floor test.** At £4, $\varepsilon = -0.5$; at £5, $\varepsilon = -100 \times 5/700 = -0.71$ – still inelastic, and revenue indeed rose £3,200 → £3,500. Push to £8: $Q = 400$, $\varepsilon = -2.0$ – elastic, and revenue falls to £3,200 again. The revenue-maximising price sits where $\varepsilon = -1$: $-100P/(1200 - 100P) = -1$ gives $P = £6$, $Q = 600$, revenue £3,600 – the largest rectangle the demand allows. Elasticity is not a curiosity; it is the pricing dial.

> **Worked example 2 – escalate to classification.** Festival data lets Nimbus classify its neighbours. Incomes rise 10% (from £4,000 to £4,400 a month) and bread quantity rises 15% (800 → 920): income elasticity = 0.15/0.10 = 1.5 – a normal good, expenditure-share LUXURY. Then the bread price moves £4 → £5 and the croissant stall's quantity falls 200 → 160: midpoint cross elasticity = $(-40/180)/(1/4.5) = -1.0$ – complements, and strongly so. Three elasticities, three business decisions: price, target customers, and what to bundle.

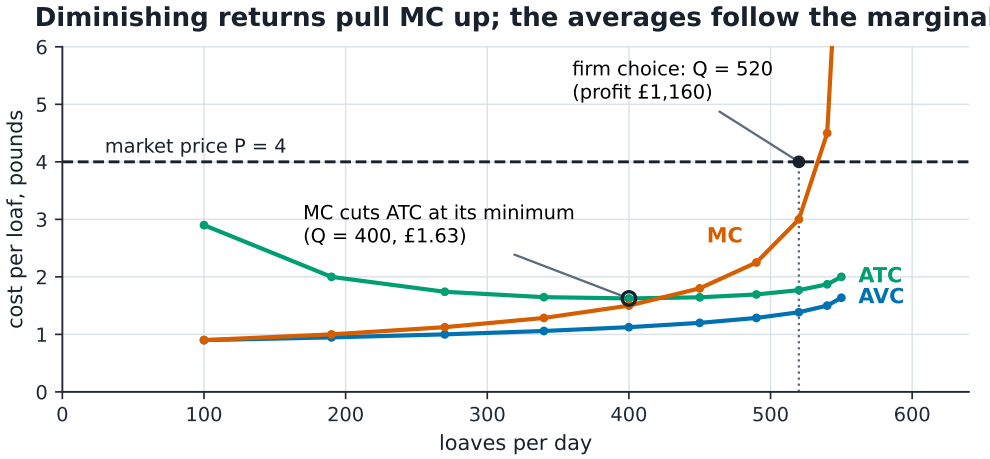
Common mistakes. Using slope instead of elasticity (slope depends on units; elasticity does not). Quoting "the elasticity of demand" as a single number for a whole linear curve. And forgetting the sign convention: demand elasticity is negative; saying "elasticity is 0.5" without the minus is how price predictions flip.

Summary. $\varepsilon = -bP/Q$ on a linear demand; revenue rises with price while inelastic, peaks at -1, falls when elastic. Income and cross elasticities classify goods into normal/luxury and substitute/complement.

Production and costs

Why it matters. To explain Nimbus's supply behaviour in the next unit you need its cost structure – and costs come from production, not from an accountant's spreadsheet. One production table generates every cost curve in the textbook, and the connections between them are the unit's real content.

The idea, with a picture. Nimbus's day: ovens cost £200 whatever happens (the **fixed cost**), each baker is paid £90, and output per baker falls as the small shop fills: with 1–10 bakers the daily loaves are 100, 190, 270, 340, 400, 450, 490, 520, 540, 550. The **marginal product** of successive bakers is 100, 90, 80, … – **diminishing returns** , the physical law behind every rising marginal cost curve.



Diminishing returns pull MC up; the averages follow the marginal

The definitions. Variable cost changes with output (here wages: $90L$); total cost adds the fixed £200. **Marginal cost** is the cost of one more loaf: $MC = w/MPL$ – the wage divided by what the last baker adds. **Average total cost** is cost per loaf, $ATC = (90L + 200)/Q$. **Economies of scale** describe the LONG run, when the shop itself can grow – do not confuse falling ATC from spreading a fixed cost with scale economies.

The derivation. The table is exact by construction: at $L = 5$, $Q = 400$, $MC = 90/60 = £1.50$, $AVC = 450/400 = £1.125$, $ATC = 650/400 = £1.625$; at $L = 8$, $Q = 520$, $MC = 90/30 = £3.00$, $ATC = £1.77$. Two structural facts follow from the arithmetic alone: MC rises because MPL falls, and MC cuts ATC at ATC's minimum (about £1.63 at $Q = 400$) – when the marginal loaf costs less than the average, the average falls, and vice versa: the grade-point-average logic.

> **Worked example 1 – the ninth baker.** The eighth baker adds 30 loaves (£3.00 each against the £4 market price – worth it). The ninth adds 20 loaves at £4 = £80 of revenue for £90 of wage: marginal cost $90/20 = £4.50 > £4$. Nimbus stops at eight bakers, 520 loaves, profit $4 \times 520 - (90 \times 8 + 200) = 2080 - 920 = £1{,}160$. The rule "produce while $MC \leq P$" and the rule "hire while the baker's value exceeds the wage" are the same sentence.

> **Worked example 2 – escalate to the shutdown question.** A supermarket price war drops the loaf to £0.85. That is below the smallest average variable cost on the table – £0.90, the first baker's 100 loaves at £90 of wages – so no scale of baking even pays the bakers who do it: the best candidate, one baker, brings £85 of revenue against £90 of wages. Shut the ovens for the price war's duration and lose only the £200 fixed cost. The boundary is worth one more look: had the price settled at £1.00, the answer flips – the first 190 loaves cover their £180 of wages and £10 of the rent, so bake with two bakers and lose £190 rather than £200. Produce when $P \geq \min AVC$; when $P < \min AVC$, every loaf deepens the loss.

Common mistakes. Treating fixed cost as irrelevant to decisions (it is irrelevant to the MARGINAL decision – and decisive for the shutdown one). Averaging first and differencing after (MC comes from the marginal product, never from a change in ATC). And reading scale economies from a falling short-run ATC – the shop never changed size.

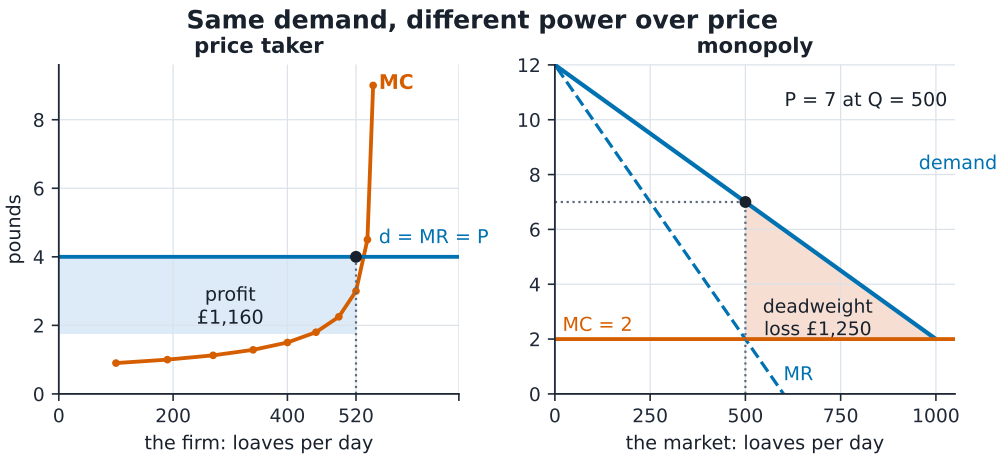
Summary. $MC = w/MPL$; diminishing returns make it rise. MC cuts ATC at its minimum. Produce while $P \geq MC$, shut down below min AVC.

Market structures

Why it matters. The same bakery, facing a different market, prices differently – and the city's welfare differs with it. Structure is a spectrum defined by two questions: does the firm face a given price, and can it keep its profits from being competed away?

The idea, with a picture. In **perfect competition** each of the city's many bakeries is a price taker at £4: its demand curve is flat, it supplies 520 loaves (unit 3.4's answer), and free entry grinds profit toward zero in the long run. A **monopoly** – one bakery with the whole market and a blocked entry – faces the market demand curve and chooses the point on it that maximises profit: $Q = 500$, £7 a loaf. Between them sit **monopolistic competition** (many sellers, differentiated loaves, free entry – Nimbus's actual world) and **oligopoly**

(few sellers whose prices depend on each other's choices).



Same demand, different power over price: the markup and the triangle are the price of monopoly

The definitions. A firm's **market power** is measured by the Lerner index $(P - MC)/P$ – zero for a price taker. The monopolist sets **marginal revenue equal to marginal cost** and charges the demand price of that quantity; because it faces the whole downward demand, MR lies below the price. The **deadweight loss** is the trades between Q_m and the efficient quantity that mutually beneficial exchange would have made but the markup priced out.

The derivation. Market demand inverted is $P = 12 - Q/100$, so revenue is $TR = 12Q - Q^2/100$ and $MR = 12 - Q/50$. With the industry's marginal cost flat at £2 (a simplification – the point here is the markup rule, not the cost shape), $MR = MC$ gives $12 - Q/50 = 2$, so $Q_m = 500$ and $P_m = 12 - 5 = £7$. Lerner index $(7-2)/7 = 0.71$. At the competitive benchmark of this diagram, $P = MC = £2$ would sell $Q = 1000$, so the markup kills the trades from loaf 500 to loaf 1000 – a deadweight loss of $\frac{1}{2}(7-2)(1000-500) = £1{,}250$ a day against this cost curve.

> **Worked example 1 – the canal concession.** Suppose Nimbus buys every rival and the council grants the towpath concession exclusively. Setting $MR = MC$: 500 loaves at £7, revenue £3,500, cost at £2 a loaf £1,000, profit £2,500 – the largest the demand and this cost curve allow. Every buyer between the 500th and the 1,000th loaf values a loaf above its £2 marginal cost but below £7 – those trades die, and their surplus (the £1,250 triangle) is nobody's. The monopoly profit is a transfer; the triangle is a loss.

> **Worked example 2 – escalate to the real middle.** Nimbus's actual market is monopolistic competition: its sourdough is distinct, so it has a small markup – but the festival's profits attract a rival bakery, whose entry shifts Nimbus's own demand curve left until economic profit is zero. The equilibrium condition is not $P = MC$ (differentiation costs that) but $P = ATC$ with $P > MC$ – markup exactly eaten by the cost of variety. Check the logic on the numbers: entry ends when Nimbus's profit £1,160 has been competed to zero, and the customer's choice set has grown by one bakery – the welfare verdict on variety versus markup is a genuine argument, not a formula.

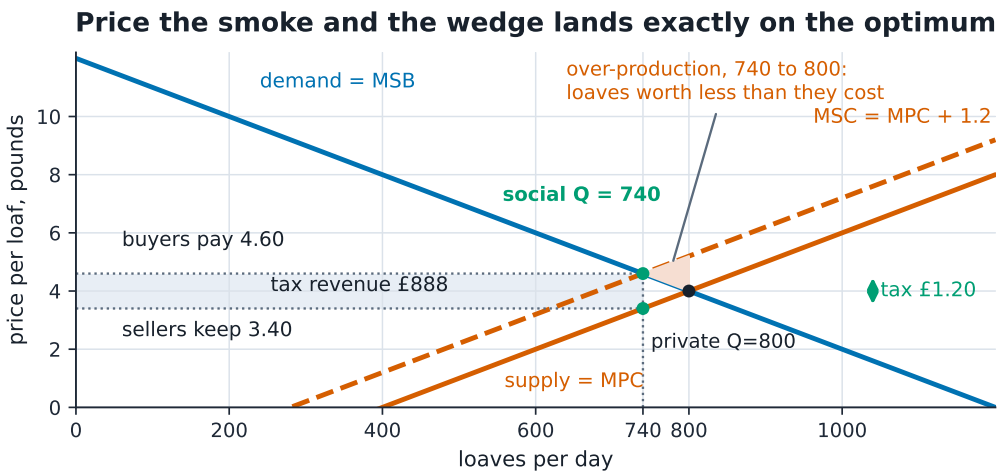
Common mistakes. Saying the monopolist "charges the highest price it can" – it charges the price of its MR = MC quantity, which is well below the top of the demand curve. Reading deadweight loss off the wrong pair of quantities. And treating the four structures as a moral ladder – they are predictions about price and output, evaluated case by case.

Summary. Price taker: $P = MC$, zero long-run profit. Monopolist: $MR = MC$, markup by the Lerner rule, deadweight loss between Q_m and Q_c. Monopolistic competition: markup yes, economic profit no.

Market failure and intervention

Why it matters. The flour mill upstream burns cheap coal; the smoke drifts over the city. Every loaf's flour carries a cost nobody's bill shows – and the £4 equilibrium from unit 3.2 quietly over-produces bread relative to what the city actually wants. This unit is about the one situation every market failure shares, and what a tax can and cannot do about it.

The idea, with a picture. An **externality** is a cost or benefit that falls on people outside the transaction. The mill's smoke makes each loaf's true **marginal social cost** exceed the bakers' private supply curve by £1.20 a loaf. The private market clears at £4/800; the social optimum is where demand meets MSC – 740 loaves. The gap between the two quantities is the over-production, and the triangle between demand, supply and MSC over that gap is the social loss.



The wedge moves the quantity to the social optimum – the rectangle is a transfer, the triangles are the stakes

The definitions. **Marginal social cost** adds the private marginal cost and the **marginal external cost**; efficiency requires MSB = MSC. A **Pigouvian tax** set equal to the external cost re-creates the efficient quantity – the **tax wedge** splits the price buyers pay from the price sellers keep. **Incidence** says who really bears a tax: it depends on elasticities, not on who hands over the money. **Public goods** (non-rival, non-excludable – street lighting) and **information failures** complete the family of failures a price system cannot price.

The derivation. With the £1.20 external cost, the efficient quantity solves demand against cost-plus-wedge: buyers face P_c, sellers keep $P_p = P_c - 1.20$, and $1200 - 100P_c = 400 + 100P_p$ gives $P_c =$ £4.60, $P_p =$ £3.40, $Q = 740$ – exactly the MSB = MSC quantity. A tax of £1.20 achieves it: revenue £1.20 × 740 = £888, and the burden splits £0.60 to buyers (4 → 4.60) and £0.60 to sellers (4 → 3.40) – the equal split is no accident: the curves' elasticities at the equilibrium match.

Worked example 1 – the smoke, priced. The council prices the mill's damage at £1.20 a loaf's worth of flour. Before: 800 loaves, and loaves 741–800 each cost the city more in smoke damage than they are worth to buyers – check loaf 780: a buyer values it at £4.20 and the baker's private cost is £3.80, which is why the trade happened, but the social cost is £5.00 – the city pays £1.20 that no bill showed. After the £1.20 tax: 740 loaves, buyers pay £4.60, bakers keep £3.40, the mill pays £888 into the clean-air fund, and the damage on the 60 loaves no longer baked never happens. The over-production triangle – roughly $\frac{1}{2} \times 1.2 \times 60 =$ £36 a day – is the pure gain.

Worked example 2 – escalate to the limits. Now the council tries the same £1.20 as a plain revenue tax with the mill unpriced – same wedge, wrong target: quantity still falls to 740, but the externality is untouched and the £888 buys roads. Intervention matches the failure or misfires. Then the honest second list: the £1.20 damage estimate is a measurement (the true MSC curve is not printed anywhere); enforcement costs eat the £36; and if bakers' supply were perfectly inelastic the entire wedge would land on sellers regardless of the paperwork. State the model, state its limits – the unit's last lesson is about intellectual honesty, not triangles.

Common mistakes. Confusing the tax REVENUE rectangle with a welfare gain – it is a transfer, and the deadweight triangles are the stakes. Assuming the statutory payer bears the incidence – elasticities decide. And concluding "externality, therefore tax": the number has to be measured, and measurement is where most real policies fail.

Summary. MSC = MPC + external cost; efficient quantity sets MSB = MSC. A tax equal to the external cost restores it; incidence follows elasticities. Every intervention claim carries a measurement and its limits.

Topic 04

University Business

What this course is, and how to use it

This is the fourth course in University Business: **macroeconomics** – the study of whole economies: what a nation produces, why it grows, what money does, and what policy can and cannot fix. The microeconomics of course 3 watched one market; this course watches the sum of them all, where one household's spending is another's income and the machine verification of every worked example keeps the accounting honest.

The running case is **Aurora**, a small open economy of ten million people, followed through one decade of data. Aurora's 2024 accounts tie by construction on both sides (expenditure and income both close at £1,320bn); its growth decomposes exactly; its price index adds up item by item; its banks multiply deposits to the last pound; and its policy episodes carry their trade-offs stated openly.

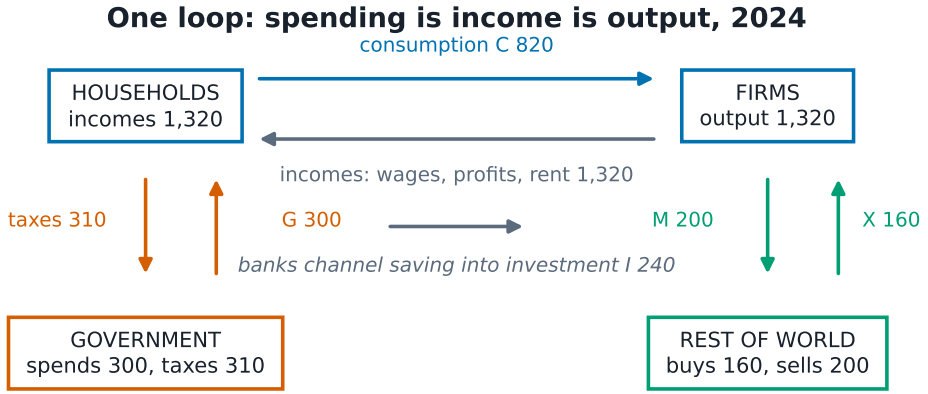
- Units 4.1–4.3 measure the economy: output, growth, money and jobs.
- Units 4.4–4.5 explain the levers: banks, the central bank, fiscal policy.
- Unit 4.6 opens the economy: trade, exchange rates and external balance.

How to read each unit: the question first, then the machinery, then two escalating worked examples whose numbers are machine-verified, the mistakes people actually make, and a three-line summary.

National accounts

Why it matters. "The economy grew by 2.2%" is a sentence built from definitions, and every debate about growth, austerity or productivity is a debate about those definitions. This unit builds GDP from both sides and shows why the two sides cannot disagree.

The idea, with a picture. Every pound of output IS somebody's income, which becomes somebody's spending, which buys output: the **circular flow** . GDP can be counted at any point on the loop – the spending that buys final goods, or the incomes generated producing them – and the two counts must meet.



One loop, five leakages and injections: the loop closes by construction

The definitions. Gross domestic product is the market value of final goods and services produced in a period. By expenditure: $Y = C + I + G + (X - M)$. **Nominal** values use current prices; **real** values hold prices fixed at a base year; the **GDP deflator** GDP = nominal/real × 100 separates the two. Intermediate goods are counted once – through the buyer of the final good, or as **value added** at each stage.

The derivation. Aurora 2024, in £bn: $C = 820$, $I = 240$, $G = 300$, $X = 160$, $M = 200$, so $Y = 820 + 240 + 300 + 160 - 200 = 1{,}320$. The income side: compensation 700 + mixed income 100 + profits 330 + rents 78 + net interest 42 + taxes on production 120 − subsidies 50 = **1,320** – the same number, by the logic of the loop, not by luck. With 2020 as the base year, real GDP in 2024 is 1,100, so the deflator is 1,320/1,100 × 100 = 120.0: cumulative nominal growth of 32.0% over the four years decomposes into 10.0% real growth and 20.0% prices.

> **Worked example 1 – the two sides meet.** A statistician sums Aurora's incomes and gets 1,320; the spending count also gives 1,320 ✓. She then notices the bakery sold £50 of flour to a bakery that sold £200 of bread. Counting both would double-count £50: flour is intermediate – it sits inside the £200 as value added. Check with value added: the mill added £50 (income to the mill), the bakery added £150 (wages and profit): £200 total, the same as the final sale ✓.

Worked example 2 – escalate to deflators. Aurora's nominal GDP grew from 1,000 in 2020 to 1,320 in 2024 (+32.0%), yet well-being did not rise by nearly a third: the same physical basket costs 20% more (deflator 120.0), so REAL GDP grew 1,320/1.2 = 1,100, exactly +10.0%. Per person, with population up 4.0% (10.00 → 10.40 million), real GDP per capita grew ≈ 5.8% over the four years – the number that actually answers "are we better off?". Every growth claim in the news decomposes into these three arithmetic steps.

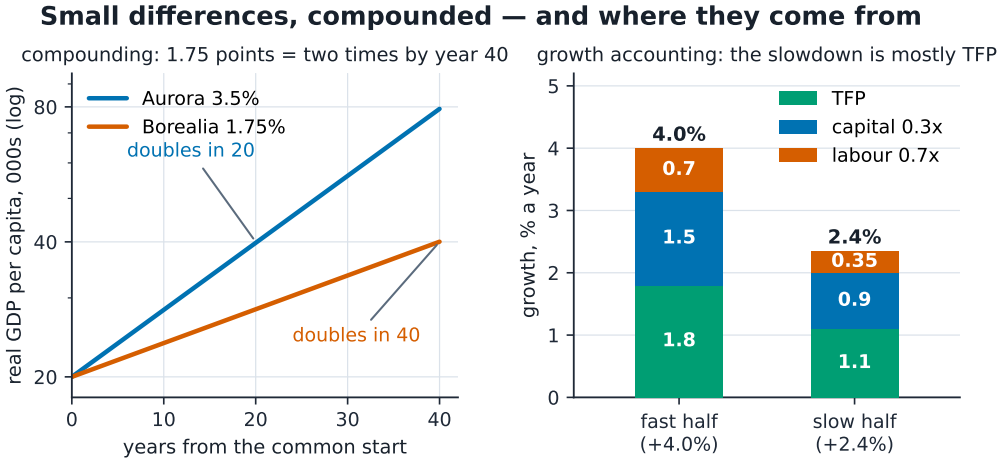
Common mistakes. Adding intermediate sales to GDP (double counting). Comparing nominal figures across years and calling it growth. And forgetting that M enters with a MINUS sign because imports are already inside C, I and G.

Summary. $Y = C + I + G + X - M$; the income side must equal it. Real = nominal/deflator; per-capita real is the living-standard measure.

Growth and productivity

Why it matters. Growth is the whole difference between the economy your grandparents lived in and yours. Small percentage differences, compounded, do all the work – and growth accounting tells you whether they came from more machines, more people, or doing things better.

The idea, with a picture. Output comes from capital, labour and ideas: $Y = A \cdot K^{0.3} L^{0.7}$. In growth rates: $\Delta Y/Y = \Delta A/A + 0.3\,\Delta K/K + 0.7\,\Delta L/L$. The residual $\Delta A/A$ – **total factor productivity** – is the part growth accounting cannot see directly, which is exactly why it is interesting.



On a log scale a straight line is a constant growth rate – the slope tells the story

The definitions. Capital is produced means of production; the exponent 0.3 is capital's income share. **Labour** force growth matters through participation and hours, not just population. The **rule of 70** 70 : at g% a year, something doubles in about 70/g years. Proximate causes (capital, schooling) are the arithmetic; fundamental causes (institutions, geography, culture) explain why the proximate ones differ.

The derivation. Log-differentiate $Y = AK^{0.3}L^{0.7}$ and the weights are the income shares: $\Delta Y/Y = \Delta A/A + 0.3\Delta K/K + 0.7\Delta L/L$. Aurora's fast half-decade: output +4.0%, capital +5.0%, labour +1.0% → TFP = 4.0 − 1.5 − 0.7 = 1.8%. Slow half: +2.4%, capital +3.0%, labour +0.5% → TFP = 2.4 − 0.9 − 0.35 = 1.15 ≈ 1.1%. The slowdown is one third fewer machines, two thirds less improvement.

Worked example 1 – the decomposition. A minister claims the fast years came from investment. The ledger says: of the 4.0%, capital deepening delivered 1.5, labour 0.7, TFP 1.8 – nearly half the growth (1.8 of 4.0) was doing things better, not owning more things. And check the logic both ways: TFP is a RESIDUAL – whatever growth is left after inputs – so measurement error in K or L lands in A and in no one's confession.

Worked example 2 – escalate to compounding. Two neighbours both start at real GDP per capita of 20,000. Aurora compounds at 3.5%; Borealia at 1.75%. Rule of 70: Aurora doubles in 20 years, Borealia in 40. After 40 years: $20{,}000 \times 2^2 = 80{,}000$ against $20{,}000 \times 2 = 40{,}000$. A 1.75-point gap, held for one working life, becomes a two-times living standard. This is why growth policy dwarfs every business-cycle debate.

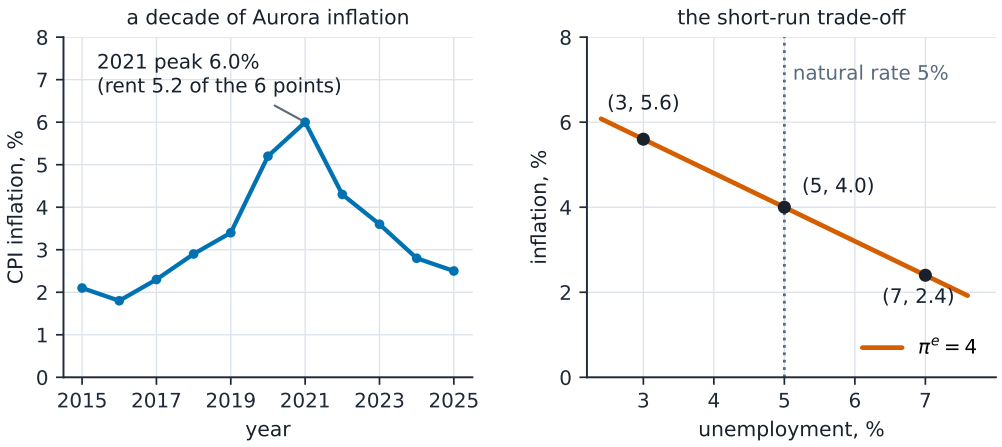
Common mistakes. Reading TFP as "technology only" – it is everything inputs fail to explain, including mismanagement and measurement. Comparing growth rates over different windows without saying so. And extrapolating the fast phase forever – convergence is conditional, not automatic.

Summary. Growth = inputs weighted by income shares + the residual. The rule of 70 turns rates into doublings. Fundamental causes decide the long run; the residual keeps the discipline honest.

Inflation and unemployment

Why it matters. Inflation is the tax nobody voted for; unemployment is the waste nobody booked. Reading a price index and a jobs report correctly is a civic skill, and the trade-off between the two is the oldest argument in macro.

The idea, with a picture. A **consumer price index** prices a FIXED basket: 10 loaves, a month of rent, 40 bus rides. If the basket cost £920 last year and £975 this year, the index reads 106.0 and inflation is 6.0% – with rent contributing 5.2 of the 6 points. An index is only as honest as its basket weights.



The short-run trade-off: points scattered around $\pi = \pi e - 0.8(u - 5)$

The definitions. Inflation is the rate of change of a price index; the **deflator** covers all GDP, the CPI a household basket. The **Fisher equation**: real rate $= (1+i)/(1+\pi)-1$. The **unemployment rate** = unemployed/labour force; a person is unemployed only if seeking work – discouragement shrinks the rate without creating a job. **Phillips curve**: $\pi = \pi^e - 0.8(u - u^n)$ with natural rate $u^n = 5\%$.

The derivation. The basket: bread $10\times(2.2-2.0) = +2$; rent $1\times(848-800) = +48$; bus $40 \times (2.625 - 2.5) = +5$; total +55 on a base of 920 → CPI 106.0, inflation 6.0%, and rent's share $48/920 = 5.22$ points of it ✓. Fisher: at 5% nominal and 3% inflation the real rate is $1.05/1.03 - 1 = 1.94\%$ – not the naive 2%. Phillips with $\pi^e = 4$: u = 7 → $\pi = 2.4$; u = 5 → 4.0; u = 3 → 5.6 – each point of unemployment below 5 buys 0.8 of inflation, until expectations catch up.

Worked example 1 – the index under the headline. Aurora's 2022 headline: 6.0%. A tenant reads "rent rose 6%" – false: the BASKET rose 6%, of which rent contributed 5.22 points; rent itself rose $48/800 = 6.0\%$, bread 10.0%, bus fares 5.0%. Weights turn item rates into basket rates: $\pi = \sum w_i\pi_i$ with w_i the basket shares: bread 2.17% weight × 10.0%, rent 86.96% × 6.0%, bus 10.87% × 5.0% → 0.22 + 5.22 + 0.54 = 5.98 ≈ 6.0 ✓ (rounding the last point).

Worked example 2 – escalate to the Fisher trap. A saver takes a 5% deposit rate expecting 3% inflation: the ex-ante real return is 1.94%. If inflation turns out to be 6%, the ex-post real rate is $1.05/1.06 - 1 = -0.94\%$ – the saver paid for the surprise, the borrower gained, and the contract was silent about it. That redistribution is why central banks target stability rather than zero: an EXPECTED inflation is a tax known in advance; an unexpected one is a lottery.

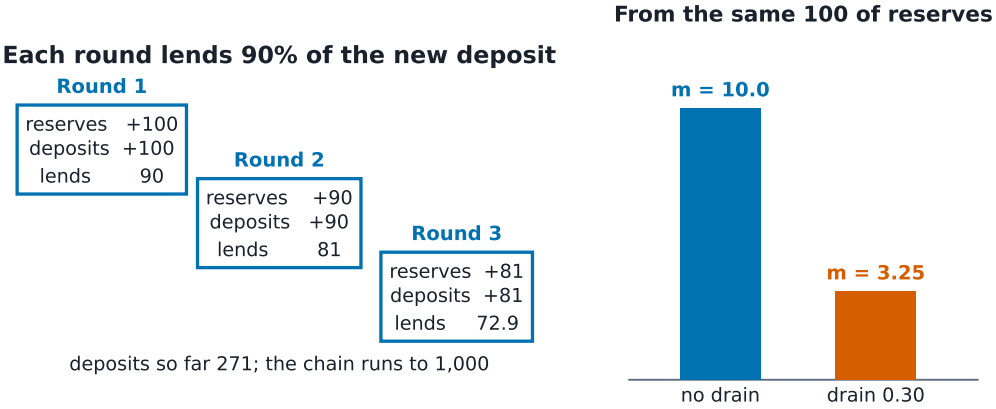
Common mistakes. "The CPI went up 106%" – the index LEVEL rose from 100 to 106; inflation was 6%. Subtracting inflation from the nominal rate (the 2% error at moderate rates, worse when rates rise). And reading a falling unemployment rate as pure good news without checking participation.

Summary. Index = fixed basket priced twice; contributions are weighted item rates. Real $= (1+i)/(1+\pi) - 1$. Phillips trades inflation against unemployment only until expectations move.

Money and banking

Why it matters. Banks do not store money; they create it by lending, and the central bank steers the whole system with a handful of tools. Every mortgage rate and every bank run lives inside this machinery.

The idea, with a picture. A deposit is both a liability (the bank owes you) and the raw material of lending. With a 10% **reserve ratio** and no cash drain, £100 of reserves supports £1,000 of deposits: the **money multiplier** at work, balance sheet by balance sheet.



Each round lends 90% of the new deposit; the currency leak shortens the chain

The definitions. M1/M2 money aggregates: currency plus deposits of increasing breadth. The **multiplier** with cash drain c and reserve ratio r_d: $m = (1+c)/(r_d+c)$. Central-bank tools: the **policy rate**, open-market operations, reserve requirements, and the lending facility – the floor system in one sentence. **Nominal vs real rates** is the Fisher equation of unit 4.3; the central bank moves the nominal short rate, and the economy responds to the real one.

The derivation. The chain: bank receives 100, keeps 10, lends 90; the 90 returns as a deposit, 9 is kept, 81 lent… total deposits $= 100 \times 1/(1 - 0.9) = 1{,}000$ ✓. With households holding cash of 30 cents per deposit dollar ($c = 0.3$): $m = 1.3/0.4 = 3.25$, so the same 100 of respires supports 325 of money – the drain is why multipliers in

real economies are 3-ish, not 10.

> **Worked example 1 – three balance sheets.** Round 1: reserves +100, deposits +100; lend 90 → borrower spends; the seller deposits 90. Round 2: reserves +90, deposits +90; lend 81. Round 3: lend 72.9. After three rounds deposits are 100 + 90 + 81 = 271 and the chain continues to 1,000 – check the remaining 729 is $72.9/(1-0.9)$, the tail of the geometric series ✓. Nothing here is a vault being emptied; every loan CREATES its own deposit.

> **Worked example 2 – escalate to transmission.** The central bank raises the policy rate from 5% to 6%. Transmission: banks' funding costs rise → lending rates rise → the real rate (expected inflation 4%) rises from 0.96% to 1.92% → investment and durable spending fall → output gap negative → inflation drifts down – each arrow a mechanism with a lag, and each lag a reason policy acts on forecasts, not readings. With the multiplier at 3.25, a 10 open-market SALE drains 32.5 of money – tools are interlocking, not independent.

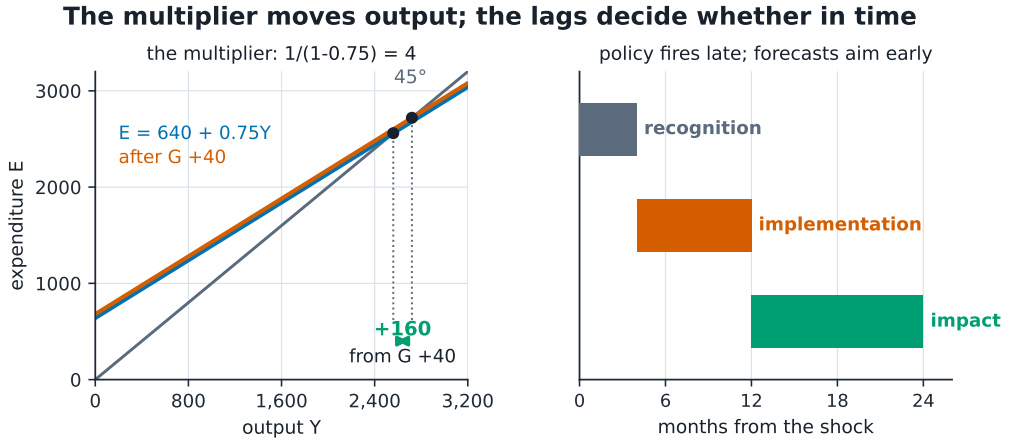
Common mistakes. "Banks lend out reserves" – lending creates deposits; reserves settle payments between banks. Confusing the multiplier's arithmetic limit with the system's behaviour (banks lend when profitable loans exist, not automatically). And quoting nominal rates as if they were the cost of borrowing.

Summary. $m = (1+c)/(r_d+c)$; deposits are created by lending. Policy rate → real rate → spending → inflation, all with lags.

Macroeconomic policy

Why it matters. When a recession hits, somebody decides. This unit is about what the decision-makers can see, what they can move, and when their answer arrives – the multiplier arithmetic and the lags, with the model's limits stated as part of the answer.

The idea, with a picture. In the simple Keynesian cross, spending multiplies: with **marginal propensity to consume** $MPC = 0.75$, the **multiplier** is $1/(1-0.75) = 4$. A £40bn spending rise lifts income by £160bn – but the picture's static arrow hides the clock.



The multiplier moves output; the lags decide whether it moves in time

The definitions. The **spending multiplier** is $1/(1-MPC)$; the **tax multiplier** is $-MPC/(1-MPC)$; the **balanced-budget multiplier** is exactly 1. **Crowding out**: government borrowing raises rates and displaces private investment – the multiplier's fine print. **Lags**: recognition (seeing the shock), implementation (passing and executing), impact (the multiplier working through) – monetary policy is quicker to fire and slower to land than fiscal.

The derivation. $MPC = 0.75$: spending multiplier $1/0.25 = 4$ ✓; $G+40 \to Y+160$. Tax cut $T = 40$: first-round spending is only the spent part, $40 \times 0.75 = 30$, so $\Delta Y = 30 \times 4 = 120$ ✓ – the tax multiplier is -3. Equal rises in G and T: $+50 \times 4$ and -50×3 net to exactly +50 – the balanced-budget multiplier of 1, a rare theorem you can prove in one line.

> **Worked example 1 – the recession response.** Aurora's output gap is −60bn. Options: G +15 (impact +60, adds to debt), T −20 (impact +60, costs 20 of revenue against the G route's 15), or a mix. Each closes the gap arithmetically; the choice is about debt, composition and speed. Now the fine print: if the central bank simultaneously raises rates to defend inflation, part of the 60 is crowded out; if imports leak, the multiplier is smaller than 4 – the model's limits belong in the answer, not a footnote.

Worked example 2 – escalate to the dated episode. The 2021–23 global inflation: supply shocks (energy, shipping) met stimulus-strengthened demand; UK CPI peaked at 11.1% in October 2022 (ONS series), unemployment stayed near 3.7–4.0%, and central banks raised policy rates from near zero to 4–5% through 2023. Read it with this course's tools: a negative supply shock moves inflation AND output the wrong way at once (the Phillips trade-off temporarily fails), lags meant the tightening landed in 2023, and fiscal support for energy bills worked THROUGH prices, not around them. One episode, every unit of this course.

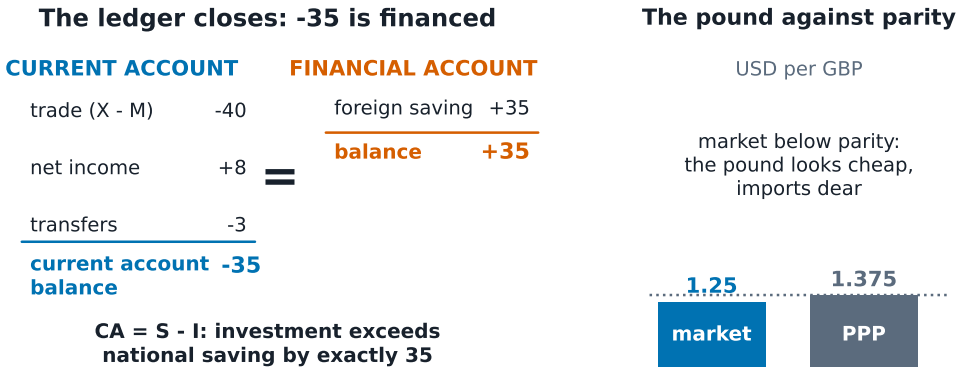
Common mistakes. Quoting the multiplier as a fact rather than a model result (leakages make real ones smaller). Comparing a tax cut and a spending rise pound-for-pound (the first round differs). And treating "policy responded" as "policy worked" without asking when the effect landed.

Summary. Multipliers: spending $1/(1-MPC)$, tax $-MPC/(1-MPC)$, balanced budget exactly 1. Lags are why policy targets forecasts; limits are part of the answer.

International linkages

Why it matters. Aurora imports a fifth of its GDP, its currency floats and its firms borrow abroad. The balance of payments is the place where domestic policy and the rest of the world settle their accounts – and where a current-account deficit is a DECISION, not a sin.

The idea, with a picture. The **balance of payments** always balances: the **current account** (trade in goods and services, net income, net transfers) is mirrored by the **capital and financial account** – a deficit on one is a surplus on the other, which is the same as saying the country is borrowing from abroad.



A current-account deficit is foreign saving financing domestic investment – the ledger makes the link explicit

The definitions. The **nominal exchange rate** is the price of one currency in another; the **real exchange rate** adjusts for price levels. **Purchasing power parity** says the rate should equal the ratio of basket prices. **Uncovered interest parity** : the currency with the higher interest rate is expected to depreciate, *ceterisparibus*.

The derivation. Aurora's external ledger (£bn): exports 160, imports 200 → net exports −40; net investment income +8; net transfers −3 → current account **−35**, matched by a financial account **+35** – national accounts arithmetic again: $CA = S - I$, so a −35 current account means domestic investment exceeds national saving by exactly 35. PPP: the same basket costs 800 in pounds at home and 1,100 in dollars abroad → PPP rate 1,100/800 = 1.375 USD/GBP; the market rate at 1.25 makes the pound CHEAP relative to parity – foreign goods look expensive to Aurora's consumers, exactly the force that narrows the trade gap.

Worked example 1 – the ledger closes. An Aurora journalist writes "our trade deficit is £40bn and the balance of payments must be in deficit". Both halves are wrong: the trade balance is −40, but the current account is −35 (income and transfers cut in), and the BALANCE of payments is zero by construction – the −35 is financed by a +35 financial surplus, i.e. foreign saving. The decision to worry or not is about WHAT the borrowing buys (investment versus consumption), not the sign.

Worked example 2 – escalate to parity. Aurora's rate is 1.25 while PPP says 1.375: at the margin, shopping abroad is dear, exports are cheap, and the gap should close – but slowly, because baskets are not identical and transport is real. Interest rates: Aurora 5%, the world 4%. UIP says the pound must be expected to depreciate about 1%: $1.25 \times 1.04/1.05 = 1.238$ – a saver choosing between the two currencies is, by the parity logic, indifferent once the expected move is priced. Check it like everything else in this course: the equation is a model with famous failures; carry-trade profits exist because UIP fails in the short run.

Common mistakes. Adding current and capital accounts and expecting a deficit (the sum is identically zero). Calling a depreciation bad for everyone (exporters gain; importers pay). And treating PPP as a forecast – it is a long-run anchor with decades of patience.

Summary. The BOP balances by definition; CA = S − I is the content. PPP is the long-run anchor, UIP the short-run model – both with stated limits.

Topic 05

University Business

What this course is, and how to use it

This is the fifth course in University Business: **business statistics** – the discipline of turning data into decisions without lying to yourself. One reproducible dataset runs the whole course: **40 takeaway delivery times** from Nimbus's weekend service, in minutes, with every statistic in this handout recomputable from the data page at the end.

The course's spine is a single habit: **state the claim, state the uncertainty, and say what would change your mind.** Means carry standard errors; tests carry their assumptions; correlations carry disclaimers; and every worked example's numbers are machine-verified against the dataset.

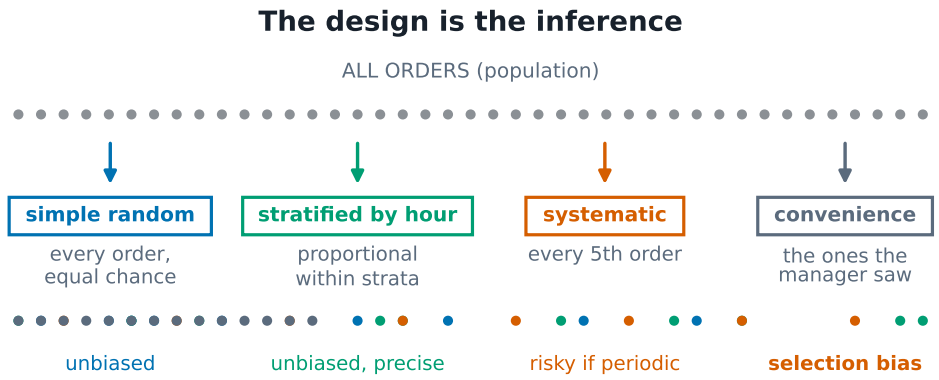
- Units 5.1–5.2 collect and describe: sampling, bias, honest pictures.
- Units 5.3–5.5 reason under chance: probability, estimation, testing.
- Units 5.6–5.7 relate and decide: regression, correlation, and the limitations that separate analysis from persuasion.

How to read each unit: the question first, then the machinery, then two escalating worked examples, the mistakes people actually make, and a three-line summary.

Data and sampling

Why it matters. Every number in a business report came from SOME selection process. Before you interpret the number, you must know what was selected, how, and who was missed – or the analysis downstream is decorating a guess.

The idea, with a picture. The 40 delivery times are a **sample** from the weekday service – all orders in a period would be the **population**. The sampling design decides what the sample can honestly say: a **simple random sample** gives every order the same chance; convenience samples (the orders the manager happened to see) do not.



size cannot cure bias: 10,000 convenience orders still say nothing about lunch

The design IS the inference: what you sampled is what you can claim about

The definitions. Data is **quantitative** (delivery minutes) or **categorical** (which rider); scales are **nominal, ordinal, interval, ratio** – minutes are ratio (zero means zero). **Bias** is systematic error from the design: **selection bias** (who is reachable), **non-response bias** (who answers), **survivorship** (who remains). A **data dictionary** records each field's meaning, unit and coding – reading one critically is step zero of any analysis – the **data dictionary** is the dataset's passport.

The derivation. Expectations: if the sample is random, the sample mean is an unbiased estimator of the population mean – its long-run average hits the truth. If the design is biased, no sample size rescues it: a convenience sample of 10,000 evening orders still says nothing about lunch. Bias is a property of the DESIGN, not the size.

> **Worked example 1 – read the dictionary.** The dataset's dictionary: `delivery_min` (ratio, minutes, one decimal, missing when the rider forgot to stop the timer – 3 cases); `hour` (ordinal, 11–21); `rider_id` (nominal). Before any average: the 3 missing values are NOT zero minutes and not random – riders forget more when busy. Decide: report the mean on the 37 complete cases AND the missingness, never a silent average.

> **Worked example 2 – escalate to design.** Nimbus wants to know if weekend service is slower. Design A: analyse the 40 times (all from Saturday). Design B: sample 20 Saturday + 20 Wednesday orders – a **stratified** sample – because the question compares days. Design A answers "what is Saturday like"; only B answers the question asked. Check the trap in reverse: sampling 20+20 and then pooling into one average weights the days equally whether or not the business does – the design must match the claim, in both directions.

Common mistakes. Treating a convenience sample as random because it is large. Confusing non-response with absence (the customers who did not answer the survey

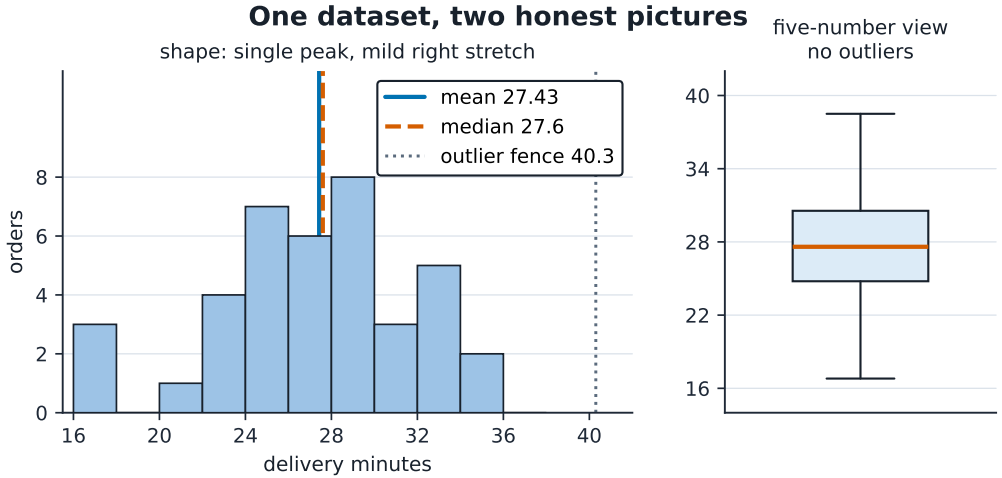
are data). And reading a dataset without its dictionary – units and codings are where analyses die.

Summary. The design is the inference. Bias is systematic and sample size cannot cure it. The dictionary is part of the data.

Descriptive measures and visualisation

Why it matters. "Average delivery: 27.4 minutes" is one sentence that hides a distribution. Description is not a formality before the "real" statistics – it is where most business insight actually lives, and where most business charts lie.

The idea, with a picture. Centre, spread, shape, outliers: the four questions every dataset answers. The 40 times: **mean** 27.43, **median** 27.6, **standard deviation** 4.91, quartiles 24.5 / 27.6 / 30.9, **IQR** 6.3 – and a histogram shows a roughly symmetric single peak with mild right stretch (38.5 is high but within the fence: Q3 + 1.5×IQR = 40.3).



One dataset, two honest pictures: shape from the histogram, five numbers from the box

The definitions. Mean = sum/n; median = middle order; sd = $\sqrt{\sum(x_i-\bar{x})^2/(n-1)}$; IQR = Q3 − Q1. **Outliers**: beyond Q1 − 1.5×IQR or Q3 + 1.5×IQR. Shape: symmetric, skewed, multimodal. The mean follows the tail; the median resists it – with a right-skewed income or delay distribution, report the median and say why.

The derivation. For the data: $\bar{x}$ = 1,097.1/40 = 27.43 ✓; the sd's numerator is $\sum(x_i-\bar{x})^2 = 940.9$, so $s = \sqrt{940.9/39} = 4.91$ ✓. Why divide by $n-1$: the deviations of a sample are constrained ($\sum(x_i-\bar{x}) = 0$), so n understates spread – one lost degree of freedom, paid to estimate the mean first.

> **Worked example 1 – the honest picture.** A marketing slide shows the delivery chart with a y-axis starting at 24 minutes: the 17–23 minute orders vanish and differences look tripled. Redraw from zero (or a stated broken axis): the 27.4 average sits in an honest landscape. Axis truncation is the oldest lie in the book – the numbers did not change, the picture did.

> **Worked example 2 – escalate to outliers.** Add one data-entry error: a 168-minute delivery (the timer never stopped). New mean (41 values): 30.86 (up 3.43 minutes!); new median: 27.7 (up 0.1); new sd: 22.5 (up 17.6). One keystroke moved the mean by more than 3 minutes and multiplied the sd by four and a half, but barely touched the median. Check it against the fence: with the error, the quartiles barely move, Q3 + 1.5×IQR is still about 40, and 168 sits far outside – the fence CATCHES it, and the dictionary explains it. Outlier handling: investigate, document, then decide – never silently delete.

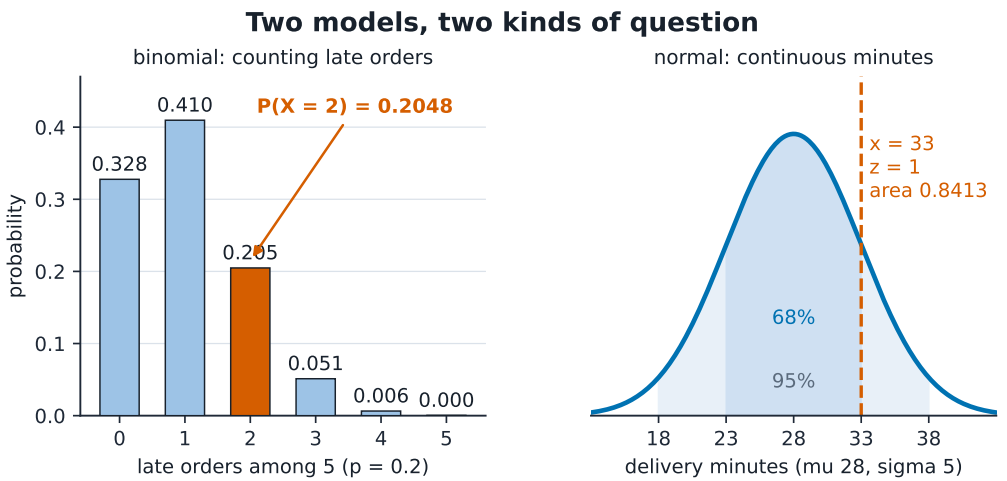
Common mistakes. Quoting a mean for skewed data without saying so. Deleting outliers before investigating them. And choosing chart types by prettiness: a bar chart of a single continuous variable hides everything a histogram shows.

Summary. Centre (mean vs median), spread (sd, IQR), shape, outliers – report all four. The picture can lie even when the numbers are true.

Probability and distributions

Why it matters. "What is the chance this order is late?" is a probability question, and the two workhorse models – binomial and normal – answer most business versions of it. The skill is matching the question to the model and stating why.

The idea, with a picture. If 20% of orders run late, the number of late orders in the next 5 is **binomial** : $P(X{=}2) = \binom{5}{2}(0.2)^2(0.8)^3 = 0.2048$. Delivery times themselves cluster around a centre with symmetric tails – modelled by the **normal** curve, where 68/95/99.7% of values sit within 1/2/3 sd.



Two models, two questions: counts of late orders, and continuous minutes

The definitions. Probability rules: $P(A \cup B) = P(A) + P(B) - P(A \cap B)$; independence means $P(A|B) = P(A)$. **Binomial**: n independent trials, success probability p. **Normal**: symmetric, bell-shaped, fully described by μ and σ. **Standardising**: $z = (x - \mu)/\sigma$ converts any normal to the standard curve.

The derivation. Binomial probabilities come from counting orders: $\binom{5}{2} = 10$ arrangements $\times\, 0.2^2 \times 0.8^3 = 0.02048$ each $\to 0.2048$ ✓; and $P(X \le 1) = 0.8^5 + 5(0.2)(0.8^4) = 0.7373$ ✓. Normal: with $\mu = 28$, $\sigma = 5$, an order at 33 minutes is $z = 1$ above the mean $\to \Phi(1) = 0.8413$ below it ✓ – the table's most-used number.

Worked example 1 – the service promise. Nimbus promises "30 minutes or your dessert is free". Under the model $\mu = 28$, $\sigma = 5$: $P(X > 30) = P(z > 0.4) = 1 - 0.6554 = 0.3446$ – a THIRD of orders would claim. At $\mu = 26$: $P(X > 30) = P(z > 0.8) = 0.2119$ still one in five. The promise's cost is a probability statement, and only a faster kitchen (a smaller μ), not a longer promise, changes the tail's shape.

Worked example 2 – escalate to the matching. Three questions, three models – match them. (a) "Of the next 20 orders, chance at least 4 are late (p = 0.2)": binomial, $P(X \ge 4) = 1 - P(X \le 3) = 1 - 0.4114 = 0.5886$. (b) "Chance one order takes over 35 minutes" with $\mu = 28$: normal, $z = 1.4$, tail 0.0808. (c) "Kitchen staff checked 2 of today's 3 Desserts – chance both were the slow ones": counting with combinations, $1/\binom{3}{2} = 1/3$. Wrong model, confidently computed, is still wrong – the justification is part of the arithmetic.

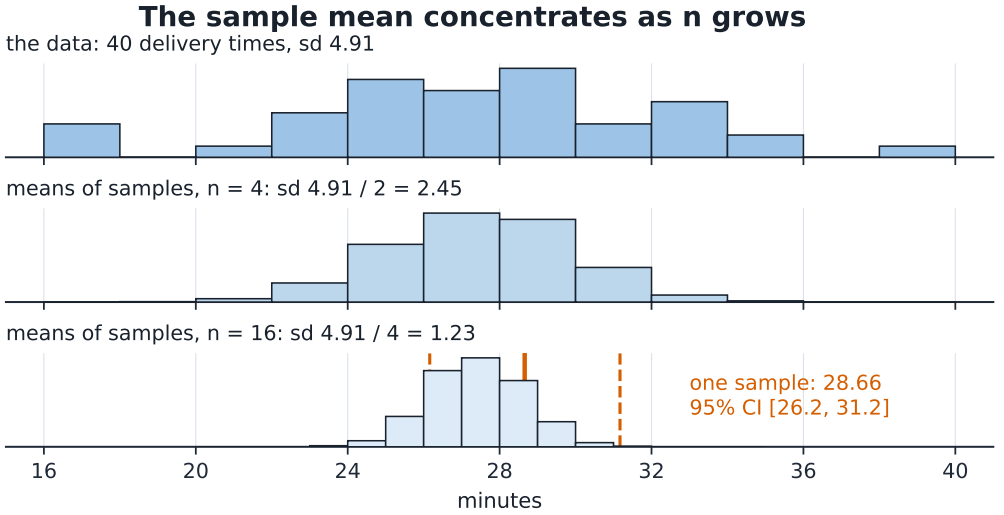
Common mistakes. Adding non-exclusive probabilities without subtracting the overlap. Using the normal for counts (binomial) or skewed delay data without saying so. And reading the table's tail as the body – always draw the curve and shade before you look up.

Summary. Counts → binomial; symmetric continuous → normal; $z = (x - \mu)/\sigma$ is the translation. State the model before computing the number.

Sampling distributions and estimation

Why it matters. The 40 times are one sample; Nimbus could have measured a different Saturday and got a different mean. Estimation is the discipline of saying what the ONE number you have says about the truth you cannot see – with an honest band around it.

The idea, with a picture. Draw a random sample of 16 orders, take its mean; draw again; again – the means themselves form a distribution: narrower than the data, centred on the truth. That is the **central limit theorem** , and its width shrinks with $\sqrt{n}$.



The sample mean is itself a random variable - and it concentrates as n grows

The definitions. The **standard error** of the mean is $SE = s/\sqrt{n}$ – the sd of the sample mean, not of the data. A 95% **confidence interval** : $\bar{x} \pm t \times SE$. The **confidence level** is a property of the METHOD: 95% of intervals so constructed capture the truth – not "95% probability the truth is in this one".

The derivation. The course's seed-11 random sample of 16 orders: mean 28.66, sd 4.70, so $SE = 4.70/\sqrt{16} = 1.175$. With $t_{15} = 2.131$: the 95% interval is 28.66 ± 2.50, i.e. **[26.2, 31.2]** ✓. Quadruple the sample to 64 and the SE halves – precision is bought at $\sqrt{n}$, the most expensive exchange rate in statistics.

Worked example 1 – read the interval properly. "Average delivery is between 26.2 and 31.2 minutes, 95% confident." Three readings: (a) wrong – "the true mean has a 95% chance of being in [26.2, 31.2]"; (b) wrong – "95% of orders take 26–31 minutes" (that is a prediction interval's job, and it would be far wider); (c) right – "the METHOD that produced this band captures the true mean in 95% of its uses". The interval is about the MEAN, and the confidence is about the PROCEDURE.

Worked example 2 – escalate to the design's limit. The manager wants the mean known to ±0.5 minutes. Required n: $t \cdot s/\sqrt{n} \leq 0.5 \rightarrow$ at that size $t \approx 1.97$ (not the 2.131 of $n = 16$), so $\sqrt{n} \geq 1.97 \times 4.7/0.5 = 18.5 \rightarrow n \geq 342$ orders. Five times the precision (±2.5 down to ±0.5) costs over twenty times the data (16 → 342) – and if the sample were convenience-sampled, 342 biased observations are precisely 342 times wrong. Precision and unbiasedness are different purchases.

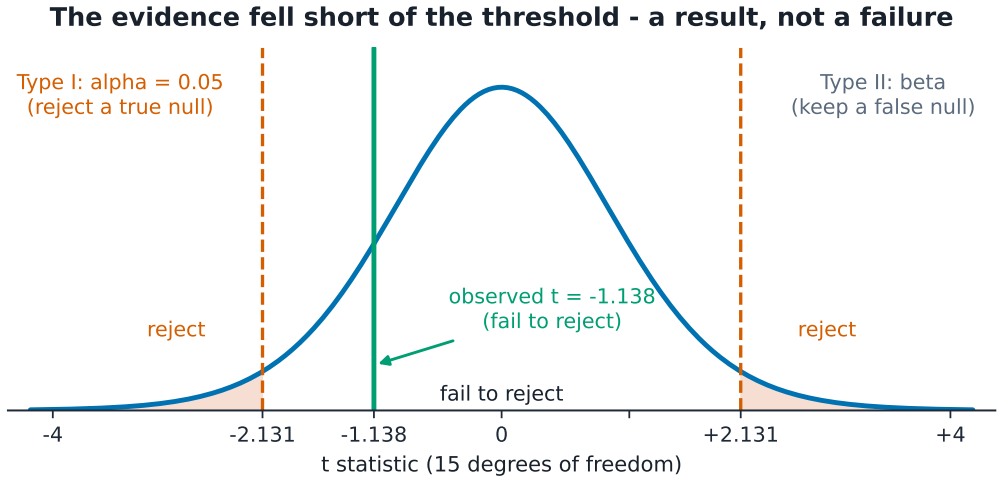
Common mistakes. Confusing sd (data spread) with SE (mean's uncertainty). Reading the confidence level as a probability about this interval. And forgetting the interval assumes the sampling design was random – the arithmetic cannot repair the design.

Summary. $SE = s/\sqrt{n}$; the CLT makes means roughly normal; $\bar{x} \pm t \cdot SE$ with the confidence read as a property of the method.

Hypothesis testing

Why it matters. "Is the weekday service actually slower than the weekend?" A test is a formal way to ask whether the data would be surprising if nothing were going on. Used honestly, it disciplines claims; used ritually, it manufactures them.

The idea, with a picture. Set the **null hypothesis** (H_0: mean = 30) and its alternative; compute how extreme the sample would be under H_0; reject only if the extremeness is past the threshold. The **p-value** p : the chance of data at least this extreme IF H_0 were true.



The evidence fell short of the threshold: fail to reject - a result, not a failure

A hypothesis test compares the data's extremeness to a chosen threshold. Test statistic $t = (\bar{x} - \mu_0)/(s/\sqrt{n})$. Significance level α – the Type I error rate WE CHOOSE (usually 0.05). Type II error β – missing a real effect; **power** $= 1 - \beta$. "Statistically significant" means p $< \alpha$, no more.

The derivation. Is the mean 30 minutes? With the seed-11 sample: $t = (28.66 - 30)/1.175 = -1.14$. The critical value at $\alpha = 0.05$, two-sided, 15 df: $|t| > 2.131$. Since $1.14 < 2.131$, we **fail to reject** – the data are consistent with a 30-minute mean AND with plenty else. Note what this does NOT say: not "the mean is 30", not "there is no difference" – only "this sample would not be surprising if it were".

Worked example 1 – the honest p-value. A rider claims their Saturday route is faster: their 16 runs average 26.9 against the shop's 28.7, $s = 4.0$, so $t = (26.9 - 28.7)/(4.0/4) = -1.8$, p ≈ 0.09. "Almost significant" is not a category – report p = 0.09 and let the decision weigh the cost of being wrong. If free desserts ride on the call, $\alpha = 0.10$ might be right; if re-training rides on it, 0.05 or 0.01. α is a business choice wearing a Greek letter.

Worked example 2 – escalate to significance vs importance. With n = 400 orders per day, a 0.4-minute difference is p < 0.01 – statistically significant and operationally invisible (no customer notices 24 seconds). With n = 8, a 3-minute difference can be p = 0.3 – invisible to the test yet real to every customer. The fix is not a bigger p-value but a reported EFFECT SIZE with its interval: "0.4 minutes, 95% CI [0.1, 0.7] – significant and trivial". Significance asks "is it there?"; the effect size asks "does it matter?" – business needs the second answer.

Common mistakes. "Accepting" the null (you fail to reject; absence of evidence is not evidence of absence). Confusing p with the probability the hypothesis is true.

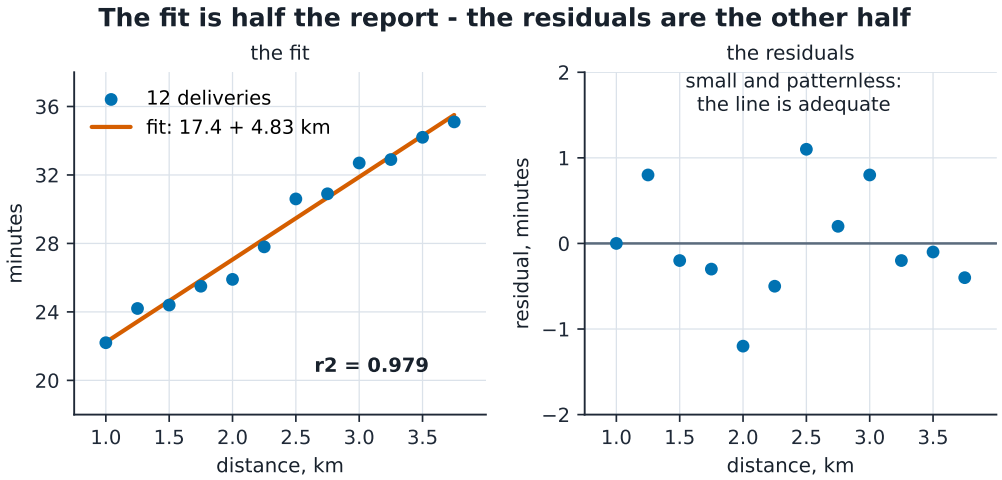
And p-hacking – testing until something clears 0.05 – which unit 5.7 unpacks.

Summary. $t = (\bar{x} - \mu_0)/SE$; reject past the threshold chosen as your error budget. p is about the data, not the hypothesis. Significance ≠ importance: always report the effect size.

Regression and correlation

Why it matters. "Do longer deliveries come from farther addresses?" Regression turns pairs of numbers into a line, quantifies how tightly they move together, and – if you are careless – implies a cause. This unit fits the line, reads its residuals, and polices the causal claim.

The idea, with a picture. Twelve tracked deliveries: distance (1.0–3.75 km) and time. The scatter is tight and upward; the **least squares** line: $\hat{y} = 17.4 + 4.83x$ – each extra kilometre costs 4.83 minutes; **correlation** $r = 0.989$, $r^2 = 0.979$: 97.9% of the variance in time travels with distance.



The fit is only half the report - the residuals are the other half

The definitions. Slope: change in $\hat{y}$ per unit x; intercept: $\hat{y}$ at x = 0 (often meaningless – 0 km is not a delivery). **Residual** : $y - \hat{y}$, the part the line misses. r^2: the share of y-variance the line explains. **Extrapolation** beyond the data's range is modelling fiction.

The derivation. Least squares minimises $\sum(y_i - a - bx)^2$; the normal equations give $b = r\, s_y/s_x$ and $a = \bar{y} - b\bar{x}$. For the 12 deliveries: $b = 4.83$, $a = 17.4$ ✓; residuals run from −1.2 to +1.1 minutes with no pattern ✓ – a straight line is adequate here, and the residual plot is the evidence.

> **Worked example 1 – read the fit.** $\hat{y} = 17.4 + 4.83x$: a 2 km delivery is predicted $17.4 + 9.7 = 27.1$ minutes; a 3 km one 31.9 ✓. The intercept's 17.4 minutes at 0 km is kitchen+handover time – meaningful here ONLY because the physics suggests a fixed prep component; state such readings, do not inherit them. And a 7 km prediction (51 minutes) is extrapolation: no data beyond 3.75 km has vouched for the line.

> **Worked example 2 – escalate to the causal police.** Correlation of 0.989 is as strong as business data gets – and still proves no direction. Third variables: both distance and time grow with traffic density; the ROUTE assignment (riders take far orders at off-peak) could make distance correlate with time without causing it at the fitted rate. The honest sentence: "distance predicts time in this sample ($r^2 = 0.98$); a causal rate would need route assignment held fixed". Correlation is a statement about THIS dataset's joint variation; causation is a claim about what would change if you intervened – different verbs, different evidence.

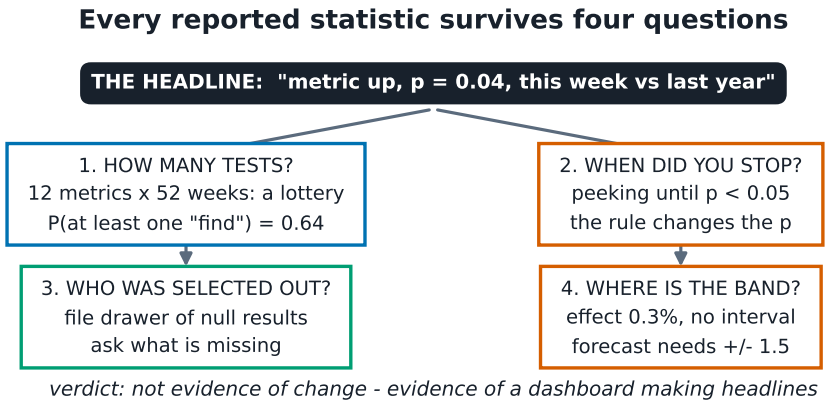
Common mistakes. Reading r^2 as "99% accurate predictions". Interpreting the intercept when x = 0 is outside meaning. Extrapolating past the data. And the classic: correlation → causation, with no mechanism, no experiment, no caution.

Summary. $\hat{y} = a + bx$ by least squares; the residuals must be examined, not assumed. r^2 measures association; only design earns causation.

Decision limitations

Why it matters. Everything before this unit can be done correctly and still produce a lie – by asking many questions, stopping when convenient, or reporting certainty that never existed. This final unit is the checklist that keeps statistics a discipline instead of a rhetoric.

The idea, with a picture. The traps in a map: **multiple testing** (20 questions at $\alpha = 0.05$ expect one "significant" fluke by design: $1 - 0.95^{20} = 0.64$ chance of at least one), **stopping rules** (peeking at data and continuing until $p < 0.05$), **selection of outcomes** (reporting the one metric that moved), and **forecast certainty** (point predictions without intervals).



Every reported statistic survives four questions - or it is a story

The definitions. p-hacking: trying many specifications until one clears the threshold. **HARKing**: hypothesising after results are known. **Publication/selection bias**: the file drawer of null results nobody shows you. A forecast without an interval is a mood.

The derivation. Multiple testing arithmetic: 20 independent tests at $\alpha = 0.05 \rightarrow P(\text{at least one false positive}) = 1 - 0.95^{20} = 0.64$ – a likely fluke, not a discovery. The Bonferroni discipline: for 20 tests at family $\alpha = 0.05$, each needs $p < 0.0025$. The 5.5 example now reads both ways: $p = 0.09$ is "not significant" AND one of the few tests actually reported – the number of unreported tests is part of the number's meaning.

> **Worked example 1 – the dashboard audit.** A dashboard tracks 12 metrics weekly; this week one shows $p = 0.04$ against last year. Question 1: how many comparisons? 12 metrics × 52 weeks – this "finding" was drawn from a lottery with hundreds of tickets. Question 2: was it declared in advance? No. Question 3: does the effect size matter? 0.3%. Verdict: not evidence of change – evidence of a dashboard's ability to generate headlines. The correct response is a pre-registered confirmation on fresh data.

> **Worked example 2 – escalate to the honest forecast.** The manager forecasts "next Saturday's average delivery: 27.4 minutes". The dataset's answer: the mean of a fresh 40-order Saturday has SE $\approx 4.91/\sqrt{40} = 0.78$, and the 27.4 it is forecast from carries the same error again, so the forecast error is $4.91\sqrt{1/40 + 1/40} = 1.10$ and the forecast should read **"27.4 ± 2.2, 95%"** – and a single Saturday's average will fall outside that band one time in twenty. Decision consequence: staff to the band's upper end (29.6), not the point. A number without its uncertainty cannot be managed, only hoped at.

Common mistakes. Treating one significant result among many as a finding. Confusing "we failed to reject" with "we proved the null" – the mirror trap of unit 5.5. And demanding certainty from forecasts instead of demanding the interval.

Summary. Count the tests, declare the rules, show the file drawer, band every forecast. Statistics earns decisions by admitting what it does not know.

The dataset (reproducible)

40 takeaway delivery times, minutes, one decimal, sorted:

16.8, 17.2, 17.3, 21.1, 22.8, 22.9, 23.0, 23.0, 24.3, 24.4, 24.9, 25.2, 25.2, 25.2, 25.5, 26.0, 26.1, 26.9, 27.4, 27.5, 27.7, 28.2, 28.4, 28.7, 28.7, 28.9, 28.9, 29.2, 29.3, 30.4, 31.0, 31.4, 32.2, 32.5, 33.2, 33.3, 33.7, 34.8, 35.5, 38.5

Verified statistics: n 40; mean 27.43; median 27.6; sd 4.91; quartiles 24.5 / 27.6 / 30.9; IQR 6.3; no points beyond the 1.5×IQR fences (15.0, 40.3). The estimation sample (seed 11, n = 16): mean 28.66, sd 4.70, SE 1.175, 95% CI [26.2, 31.2], t vs 30 = −1.14. The regression pairs (12 deliveries, km vs minutes): slope 4.83, intercept 17.4, r 0.989, r^2 0.979. All figures in this handout recompute from these values.

Topic 06

University Business

What this course is, and how to use it

This is the sixth course in University Business: **managerial accounting** – the numbers managers use to decide: what a product costs, what volume breaks even, what a budget commits, which variances to chase, and which costs a decision actually changes. Course 2's financial accounting reported to strangers; this course arms insiders.

The running case is **Nimbus Kitchen**, the founders' frozen-meals venture: £6.50 meals, £3.10 variable cost, £4,200 monthly fixed cost. Every unit works the same sheet forward – behaviour, volume, costing, budgets, variances, decisions – and every worked example's numbers are machine-verified.

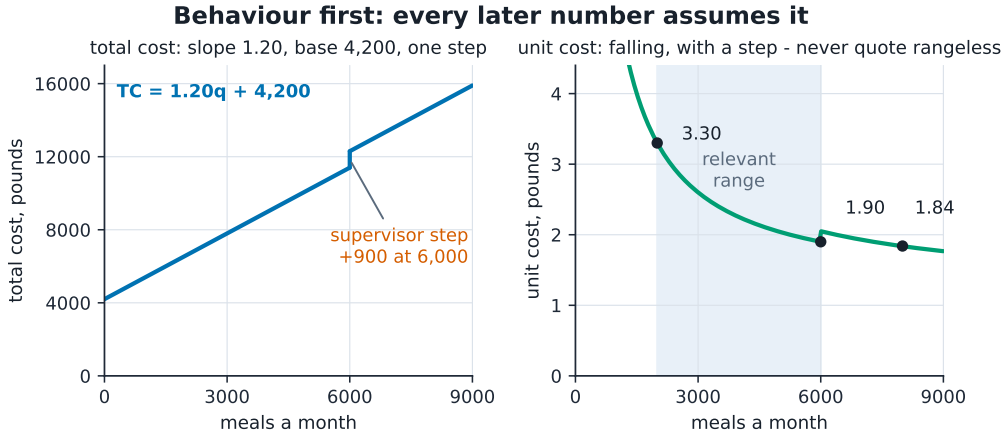
- Units 6.1–6.2 build the model: how costs behave, and how volume turns contribution into profit.
- Units 6.3–6.5 run the firm: costing systems, the budget set, variance analysis.
- Unit 6.6 closes with decisions: which numbers a choice actually changes.

How to read each unit: the question first, then the machinery, then two escalating worked examples, the mistakes people actually make, and a three-line summary.

Cost behaviour

Why it matters. Every managerial number downstream – break-even, product cost, variances, outsourcing – assumes a cost model. Get the behaviour wrong and the arithmetic is right about the wrong world.

The idea, with a picture. Nimbus Kitchen's monthly cost: £1.20 variable per meal (ingredients and packaging), £2,400 kitchen rent (**fixed** at any volume in the lease's range), and a supervisor at £1,800 **stepping** to £2,700 past 6,000 meals. The **relevant range** is the band where each behaviour claim holds.



Total cost bends at the step; unit cost keeps falling - never quote either without its range

The definitions. Variable cost changes with volume (£1.20/meal); **fixed cost** does not, within the range (£2,400); step costs jump at thresholds. **Mixed costs** have both parts. **Traceability**: direct costs attach to one product; **indirect** (overhead) costs need allocation – unit 6.3's business.

The derivation. Total cost $= 1.20q + 4{,}200$ up to 6,000 meals, then $1.20q + 5{,}100$: at 2,000 meals £6,600 ✓, at 6,000 £11,400 ✓, at 8,000 £14,700 ✓. Unit cost: £3.30 at 2,000 falling to £1.84 at 8,000 ✓ – the fixed £4,200 spreads over more meals. The **high-low** estimate on 2,000–6,000 (within one step) recovers exactly £1.20/meal ✓; blind across the step (2,000–8,000) it returns £1.35 – a step cost laundered into a variable one.

> **Worked example 1 – estimate the function honestly.** Four months of data: (2,400, £7,080), (3,300, £8,160), (4,500, £9,600), (5,200, £10,440). High-low: $(10{,}440 - 7{,}080)/(5{,}200 - 2{,}400) =$ £1.20 per meal ✓, fixed $7{,}080 - 1.20 \times 2{,}400 =$ £4,200 ✓. A scatter inspection confirms linearity within the range – the estimate is not just arithmetic, it is an assumption you checked.

> **Worked example 2 – escalate to the step.** July plans 6,400 meals. June's £4,200 fixed no longer holds: the supervisor steps (+£900), so fixed becomes £5,100 and total $= 1.20 \times 6{,}400 + 5{,}100 =$ £12,780 ✓. The unit cost quoted to retail partners must say WHICH range: £2.00 at 6,400 meals (12,780/6,400), not the £1.86 that June's £4,200 fixed cost, blindly extrapolated, would give. Behaviour statements expire at thresholds – date them like milk.

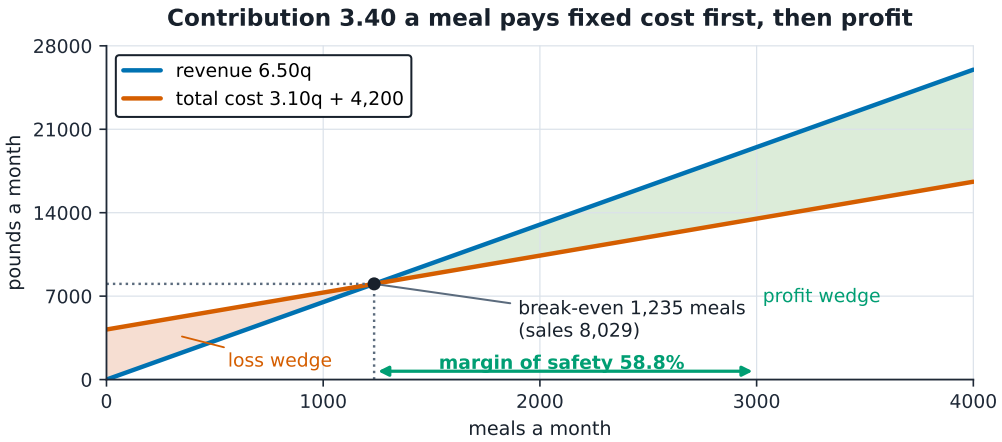
Common mistakes. Quoting unit fixed costs as if they were costs of one more unit (marginal cost is £1.20, not £1.90). Extrapolating a cost function past its relevant range. And treating step costs as either purely fixed or purely variable because the spreadsheet prefers it.

Summary. $TC = vq + F$ within the range; steps break the range. Marginal cost is v; unit cost is $v + F/q$ – different animals.

Cost-volume-profit analysis

Why it matters. "How many meals to break even?" "What volume for £5,000 profit?" "What if the price falls 5%?" CVP is the manager's flight simulator, and its one big simplification – linear behaviour in the relevant range – is also its licence.

The idea, with a picture. At £6.50 a meal with £3.10 variable, the **contribution margin** is £3.40 per meal (52.3% of price). Fixed costs of £4,200 need 4,200/3.40 = 1,235 meals – the **break-even point**, £8,029 of sales.



Loss wedge left of break-even, profit wedge right - the margin of safety measures the runway

The definitions. $CM = P - v$; CM ratio = CM/P. Break-even: $q^* = F/CM$. **Margin of safety**: $(q_{plan} - q^*)/q_{plan}$. **Operating leverage**: CM/profit – how hard profit swings with volume.

The derivation. Profit = $CM \cdot q - F$: at 3,000 meals, $3.40 \times 3{,}000 - 4{,}200$ = £6,000 ✓. Break-even 1,235 meals ✓ (check: $1{,}235 \times 6.50$ = £8,028 sales, cost = $1{,}235 \times 3.10 + 4{,}200$ = £8,029 – within rounding ✓). Target profit £5,000: $q = (4{,}200 + 5{,}000)/3.40$ = 2,706 meals ✓. Margin of safety at the 3,000-meal plan: 1,765/3,000 = 58.8% ✓.

> **Worked example 1 – target profit with tax.** The founders want £6,000 AFTER a 25% tax: pre-tax profit must be 6,000/0.75 = £8,000, so $q = (4{,}200{+}8{,}000)/3.40$ = 3,588.2, rounded up to 3,589 meals (3,588 falls £1 short) ✓. The tax rate enters only through the profit target – divide, then proceed. Check the sanity of the answer against capacity: at 3,589 meals the kitchen is within the 6,000-meal step ✓ – the CVP world still holds.

> **Worked example 2 – escalate to the mix.** Two meals: Alpha at £5.00 (CM £2.60) and Beta at £4.00 (CM £1.80), sold 60:40. The bundle: price $0.6 \times 5.00 + 0.4 \times 4.00$ = £4.60, CM = $0.6 \times 2.60 + 0.4 \times 1.80$ = £2.28 ✓. Bundle break-even: 4,200/2.28 = 1,842 bundles = 1,105 Alpha + 737 Beta ✓. Now the trap: if the mix drifts to 50:50, the bundle CM FALLS to $0.5 \times 2.60 + 0.5 \times 1.80$ = £2.20 – because Beta carries the lower margin, a mix shift toward Beta lowers the bundle CM and raises break-even. Mix is a margin in disguise: state it, then watch it.

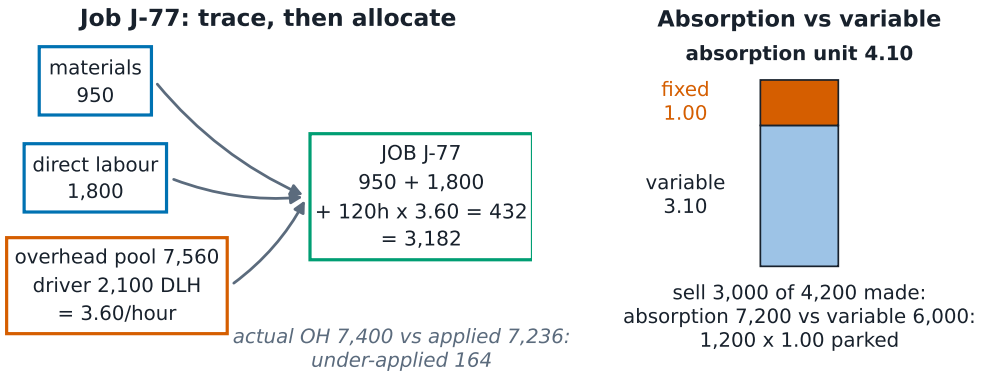
Common mistakes. Using unit cost (with fixed spread) in CVP instead of the variable cost. Forgetting the step when the plan crosses a threshold. And quoting break-even without the mix that produced it.

Summary. $q^* = F/CM$; target profit divides by CM after tax adjustment; mix shifts move the bundle CM – the silent variable.

Costing systems

Why it matters. "What does one meal cost?" has three answers – variable, absorption, and the job-order price – and each answers a different question. Choosing the costing system is choosing which question the number answers.

The idea, with a picture. Direct costs trace (ingredients £1.20, direct labour £1.10 per meal); **overhead** – variable £0.80 per meal plus the £4,200 fixed block – needs a driver. At normal capacity 4,200 meals (2,100 direct labour hours at 0.5 h/meal), the pool is $0.80 \times 4{,}200 + 4{,}200$ = £7,560 → **£3.60 per labour hour**.



Trace what you can, allocate what you must - and state the driver

The definitions. Job costing: costs attach to a unit or batch (the catered order J-77). **Process costing**: costs average over continuous production (the meal lines). **Absorption vs variable costing**: absorption includes fixed overhead in unit cost (here £4,200/4,200 = £1.00 per meal); variable expenses it. **Over/under-applied**: the gap between applied and actual overhead.

The derivation. Job J-77 (a 120-hour catering batch): materials £950 + direct labour £1,800 + overhead 120 × £3.60 = £432 → **£3,182** ✓. If the month's actual overhead is £7,400 at 2,010 actual hours (4,020 meals × 0.5), applied is 2,010 × 3.60 = £7,236 → **under-applied by £164** ✓ – the rate's estimate was low, and the costing income statement must reconcile it.

> **Worked example 1 – absorption vs variable income.** Produce 4,200 (normal capacity – so no volume variance), sell 3,000, close 1,200. Absorption unit cost = 3.10 + 1.00 = £4.10; absorption profit = 3,000 × (6.50 − 4.10) = £7,200; variable profit = 3,000 × 3.40 − 4,200 = £6,000 ✓ – the £1,200 gap is exactly the 1,200 unsold meals × £1.00 of fixed cost parked in inventory. Production inflates absorption profit: know which arrow you are reading.

> **Worked example 2 – escalate to the driver choice.** Overhead is really 60% machine-driven. A single labour-hour rate loads the hand-finished Alpha line with the automation's cost. Two pools: machine-driven £4,536 (60% of £7,560) over 1,800 machine hours → £2.52/hr; the remaining £3,024 over 2,100 labour hours → £1.44/hr ✓. Job J-77 (120 labour hours, 0 machine hours) falls from £432 to £173 of overhead – the costing system just changed the product's apparent profitability. Drivers are not plumbing; they are pricing decisions about who pays for capacity.

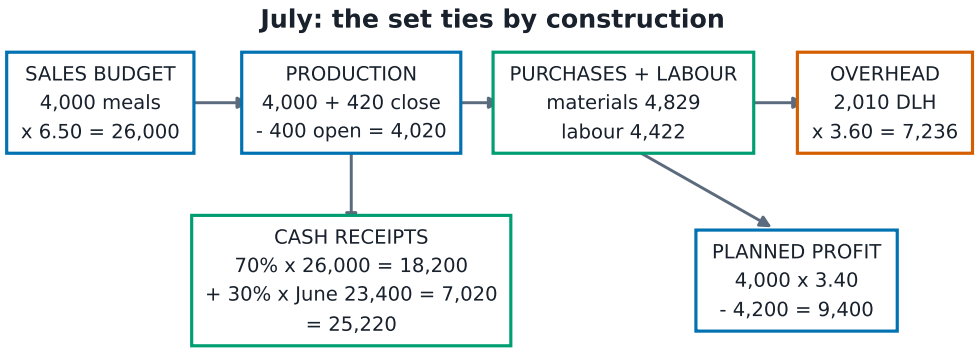
Common mistakes. Reconciling over/under-applied overhead by ignoring it. Reading absorption profit as performance when inventory moved. And choosing drivers by data availability alone – convenience costing is still an allocation with an attitude.

Summary. Trace direct, allocate indirect through a stated driver. Absorption parks fixed cost in inventory; variable exposes it. The driver decides who pays.

Budgeting

Why it matters. A budget is a plan with numbers that commits spending, coordinates departments, and – like any target – changes the behaviour of the people measured by it.

The idea, with a picture. The operating **budget set** builds left to right: sales budget → production budget → direct materials and overhead → cash budget. Nimbus Kitchen plans May–July sales of 3,200 / 3,600 / 4,000 meals at £6.50.



Each box consumes the previous one - the set ties by construction

The definitions. The production budget: sales + closing inventory − opening inventory. The **cash budget** converts sales to receipts by collection pattern (70% in month, 30% following). **Flexible budgets** restate plans at actual volume for variance work (unit 6.5). **Budget slack** and gaming are behavioural facts, not errors.

The derivation. July production: sales 4,000 + closing inventory (10% of August's 4,200 = 420) − opening (10% of July = 400) = **4,020 meals** ✓. July cash receipts: 70% of July sales (0.7 × 4,000 × 6.50 = £18,200) + 30% of June (0.3 × 3,600 × 6.50 = £7,020) = **£25,220** ✓. The profit budget ties back to CVP: July's planned profit = 4,000 × 3.40 − 4,200 = £9,400 ✓.

> **Worked example 1 – the set ties.** Build July's purchases: production 4,020 meals × £1.20 materials = £4,824; materials inventory at 5% of next month (August 4,100 meals × £1.20 × 5% = £246 closing; opening 5% of July = £241) → purchases **£4,829** ✓. Direct labour £1.10 × 4,020 = £4,422; overhead applied at 2,010 hours × £3.60 = £7,236 – inside which sits the £4,200 fixed block. Every figure consumes another's output; change the sales plan and the whole sheet breathes. That coupling is the point.

> **Worked example 2 – escalate to behaviour.** The production manager is bonused on hitting the 4,020-meal budget. August demand collapses to 3,700. The honest flexible budget says 3,730 production (3,700 + 370 − 340); the manager's incentive says overproduce toward 4,020 to "hit target", parking £1.00/meal of fixed cost in inventory (unit 6.3's absorption trap) and reporting profit the cash cannot find. The budget's arithmetic was never wrong – the target was. Participative budgets and rolling reforecasts exist precisely because plans and incentives drift.

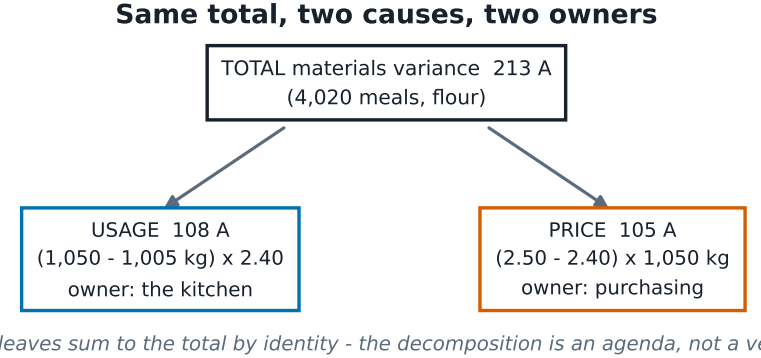
Common mistakes. Budgeting production as sales (inventory policy disappears). Treating the cash budget as the profit budget with a lag – timing and measure differ. And evaluating volume-variance performance against a budget nobody re-based after the world moved.

Summary. The set ties: sales → production → purchases/labout/overhead → cash. Inventory and collection policies are the hinges. Budgets shape behaviour as surely as they record it.

Variance analysis

Why it matters. The budget said £2.40/kg of flour at 0.25 kg per meal; July said otherwise. Variances decompose the gap into price, usage and volume pieces, each with an owner – the antidote to "costs were up".

The idea, with a picture. July: 4,020 meals produced, 1,050 kg of flour used (standard: $4{,}020 \times 0.25 = 1{,}005$ kg), paid £2.50/kg (standard £2.40). The **flexible budget** restates the plan at 4,020 meals; the variances measure the rest.



Same total, two causes, two owners - the kitchen and the buyer

The definitions. Each **variance** compares actual with standard: the usage variance is $(AQ-SQ)\times SP$, the price variance is $(AP-SP)\times AQ$; and a **favourable variance** means below-standard cost or above-standard revenue – not "good": a favourable price variance from lower-grade flour can buy an adverse usage variance plus a quality problem. **Volume variances** compare actual and budgeted activity.

The derivation. Usage: $(1{,}050-1{,}005)\times 2.40 = £108$ adverse ✓ – the kitchen used 45 kg more than the standard allows. Price: $(2.50-2.40)\times 1{,}050 = £105$ adverse ✓ – the buyer paid 10p over. Total materials variance: £213 adverse ✓, and the two leaves sum exactly to it – the decomposition is an identity, not an opinion.

> **Worked example 1 – who owns what.** The £108 usage variance belongs to the kitchen (yield, training, recipe drift); the £105 price variance to purchasing (supplier, timing, grade). A manager folding them into one number owns nothing and learns nothing. And the two can trade: the buyer's cheaper flour (favourable price) thickens the sauce (adverse usage). Variances are a conversation between centres, and the tree is the agenda.

> **Worked example 2 – escalate to honest interpretation.** Sales volume variance for July: actual 3,700 meals against the 4,000 budget at CM £3.40 → **£1,020 adverse** ✓. But a supplier strike took the kitchen down for three days; the flexible-budget revenue cost of those days is roughly £2,550 of contribution – the "variance" is mostly a diagnosis, not a performance. Three questions before judging any variance: is it material (a £12 rounding variance is not a finding), is it controllable by its owner, and is it offset elsewhere (a favourable price buying an adverse usage)? Variances start inquiries; they do not end them.

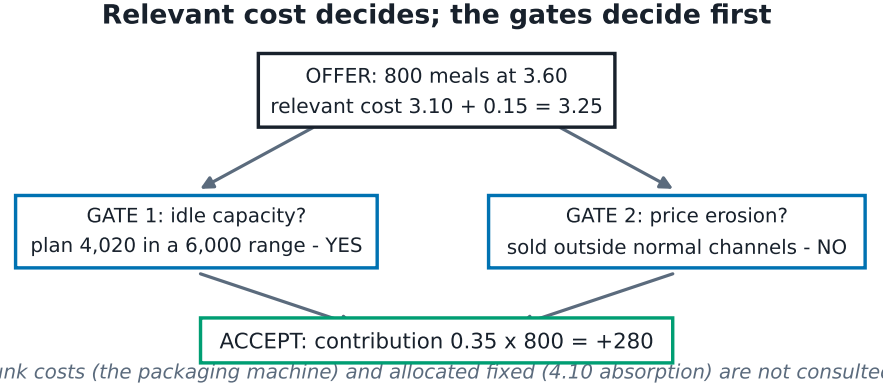
Common mistakes. Computing variances against the STATIC budget when volume moved (the flexible budget is the baseline). Summing variances into one number and delegating it to nobody. And reading favourable as virtue – cheap can be expensive.

Summary. Usage × standard price; price gap × actual quantity; the leaves sum to the total. Every leaf has an owner; every owner gets the conversation, not the verdict.

Relevant-cost decisions

Why it matters. The special order arrives at £3.60 against a £6.50 list price, the outsourcer bids £4.40 against a £4.10 absorption cost, and the sunk machine sits there whispering. Decisions need only the numbers that CHANGE – and the discipline to ignore the rest.

The idea, with a picture. A **relevant cost** is one that differs between the alternatives: always future, always differential. Sunk costs never qualify; allocated fixed costs rarely do.



Two gates before the verdict: does capacity exist, does the price survive elsewhere?

The definitions. Sunk cost: already incurred, ignores you. **Avoidable cost**: disappears with the decision. **Opportunity cost**: the best alternative forgone. The decision rule: compare incremental revenue against incremental cost, including opportunity.

The derivation. The special order: 800 meals of idle capacity, offer £3.60. Relevant cost per meal: variable £3.10 + £0.15 special packaging = £3.25 → contribution £0.35 × 800 = **£280, accept** ✓. The £4.10 absorption unit cost is irrelevant – most of it is fixed cost that exists either way, plus an allocation with an attitude.

> **Worked example 1 – outsourcing with a freed cost.** A co-packer bids £4.40 per meal against Nimbus's £4.10 absorption cost. Looks like "buy" – but only £3.10 of the internal cost is avoidable, plus £0.30 of rent actually freed by shutting one kitchen line: avoidable £3.40 < £4.40 → **keep making** ✓. The remaining £0.70 per meal is fixed or allocated – the decision does not change it, so the decision does not consult it. Every outsourcing bid is really two questions: what truly disappears, and what could the freed space earn instead?

> **Worked example 2 – escalate to product mix.** Flour is short: 1,000 kg available against next month's 4,040-meal plan (1,010 kg) – a 10 kg squeeze, or with Alpha needing 0.30 kg and Beta 0.20 kg, the real question is which meal earns more per kg of flour. Alpha: CM £2.60 / 0.30 kg = £8.67/kg; Beta: £1.80 / 0.20 = £9.00/kg – **Beta earns more per scarce kg**, so in a true constraint, bias the mix toward Beta and let Alpha wait ✓. General rule: rank products by CM per unit of the BINDING constraint, not per unit of product. Constraints move; the ranking discipline travels.

Common mistakes. Letting sunk costs argue ("we've spent too much to stop now"). Absorption unit costs masquerading as decision costs. Forgetting opportunity costs of freed capacity – zero is an assumption, not a law.

Summary. Relevant = future + differential. Sunk never, allocation rarely, opportunity always. Rank by CM per binding constraint.

Topic 07

University Business

What this course is, and how to use it

This is the seventh course in University Business: **marketing** – the discipline of creating, communicating and delivering value, measured honestly. The tool of the trade is a claim with a clean denominator: every metric in this handout says what it counts over what base, because a marketing number without its denominator is a mood.

The running case is **Nimbus Kitchen's subscription launch**: the founders take the £6.50 frozen meal into a direct-to-consumer weekly bundle. One campaign runs through all seven units – research, segmentation, product, price, channels, promotion and ethics – with every metric recomputable from the campaign sheet at the end.

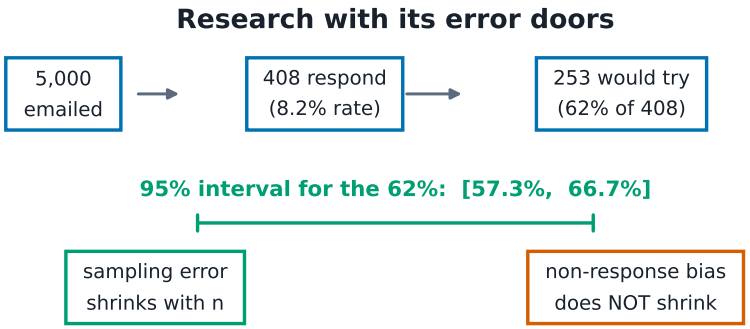
- Units 7.1–7.2 aim the campaign: research that can be wrong, segments that can be acted on.
- Units 7.3–7.5 build the offer: the product levels, the price, the channel economics.
- Units 7.6–7.7 run and judge it: the funnel with honest denominators, and the ethics that keep the numbers worth having.

How to read each unit: the question first, then the machinery, then two escalating worked examples, the mistakes people actually make, and a three-line summary.

Marketing research

Why it matters. The campaign is about to spend £3,600 on the belief that "people want weekly meal bundles". Research exists to make that belief testable – and to kill it cheaply if it is wrong.

The idea, with a picture. A **research question** must be answerable by some data you could actually collect; a hypothesis makes it falsifiable. Nimbus emails 5,000 past customers; 408 respond (an 8.2% **response rate** – already a warning: non-respondents differ), and 62% say they would try a weekly bundle.



Every number carries its error doors - sampling on the interval, design in the response rate

The definitions. **Sampling error** shrinks with $\sqrt{n}$; the 95% interval for a proportion is $p \pm 1.96\sqrt{p(1-p)/n}$. **Non-response and selection bias** do NOT shrink – they are design facts. **Exploratory** research generates hypotheses; **descriptive** measures; **causal** needs experiments (A/B tests), not surveys.

The derivation. For $p = 0.62$, $n = 408$: $1.96\sqrt{0.62 \times 0.38/408} = 0.047 \rightarrow$ the interval is **[57.3%, 66.7%]** ✓. Stated honestly: "somewhere between 57 and 67 percent of people like the ones who answered would try it" – two hedges, both earned.

> **Worked example 1 – the question that could not fail.** First draft: "Do customers love our meals?" – no data could refute it. Second draft: "Of past customers, what share would order a 5-meal bundle at £29 within 30 days?" – measurable, falsifiable, and the survey's 62% intent becomes a hypothesis the launch will TEST (intent overstates purchase; unit 7.6's 15% trial rate is the real number). A question you cannot be wrong about is not research.

> **Worked example 2 – escalate to the design's limits.** The 8.2% response rate: who answers a food survey? Regulars – the 62% intent likely overstates the customer base. Two repairs: (a) weight by order frequency if the frame allows; (b) better, a small **A/B test** – offer the bundle to 200 random past customers and watch behaviour. 27 accept (13.5%): the behavioural number lands far below the attitudinal 62%, exactly as intent literature predicts. Survey for direction; experiment for truth.

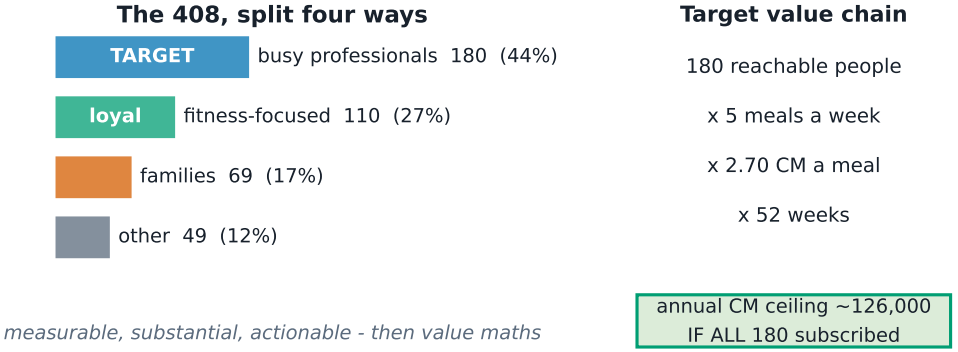
Common mistakes. Quoting 62% without the interval or the response rate. Treating a convenience sample of responders as the customer base. And asking "would you buy?" – intention is not behaviour, and the gap is the most replicated finding in market research.

Summary. Questions must be falsifiable; intervals carry sampling error; design carries bias. Surveys aim, experiments decide.

Segmentation and positioning

Why it matters. "Everyone" is not a target; it is the absence of one. Segmentation splits the market into groups the campaign can actually reach, and positioning chooses the race the product can actually win.

The idea, with a picture. The 408 responses split four ways: busy professionals 180 (44.1%), fitness-focused 110 (27%), families 69 (17%), other 49 (12.0%). A segment is useful only if it is **measurable, substantial, actionable** – the criteria, not the demographics.



Measurable, substantial, actionable - then the value maths decides the target

The definitions. **Segmentation** by behaviour, needs or demographics; **targeting** picks by attractiveness (size × growth × fit); **positioning** is the claim held in the customer's mind against competitors – defined by differences the firm can deliver. A **value proposition** states for whom, what benefit, why believed.

The derivation. Busy professionals: 180 respondents, 27% say they would pay for convenience NOW; the launch target is the reachable slice – call it 180 in the pilot city. Weekly value if every one of the 180 subscribed: $180 \times 52 \times £2.70 \times 5 = £126{,}360$ a year – the full-subscription ceiling, not the forecast (the 27% trial rate and churn do the discounting) ✓. The fitness segment (27%) is smaller but more loyal (64% weekly vs 44%); at scale it may beat the professionals on lifetime value. Segment value = size × conversion × retention × margin – never size alone.

> **Worked example 1 – the actionable test.** "Left-handed bakers" is measurable (rarely substantial). "People who hate cooking" is substantial (rarely actionable – no list, no channel). Busy professionals pass all three: countable (LinkedIn targeting), substantial (44% of responders), actionable (a channel and a message: "evenings back"). The criteria are a sieve, and each failure mode has a name.

Worked example 2 – escalate to positioning. The competitor's frozen meal sits at £6.90 on shelf, positioned on "healthy". Nimbus cannot out-health the health brand; it CAN out-convenience everyone: the weekly bundle, the subscription, the 6 pm delivery slot. Position: "Tuesday evenings, solved" – a specific moment, a credible difference, a reason to believe (5 dinners, one decision). Check with a **perceptual map**: convenience (x) against health (y) puts the competitor high-left; Nimbus claims high-right. A position nobody can attack is a difference nobody can copy quickly.

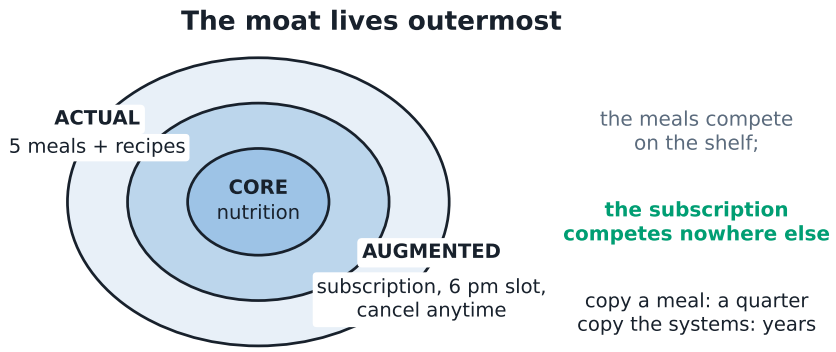
Common mistakes. Segmenting by demographics when the BEHAVIOUR differs. Choosing a segment the firm cannot reach (fails actionable). And positioning on "quality" – everyone says quality; a difference everyone claims is not a position.

Summary. Measurable, substantial, actionable – then value = size × conversion × retention × margin. Position on a deliverable difference, owned against the map.

Product and value

Why it matters. "The product" is more than the meal in the box: the bundle, the subscription, the delivery slot and the cancel-anytime promise are all product. Value is what the customer gets; the price (and unit 7.4) merely captures part of it.

The idea, with a picture. The **three product levels** : core (nutrition), actual (5 frozen meals, recipes, packaging), augmented (subscription flexibility, the 6 pm slot, the app). Nimbus's differentiation lives almost entirely in the augmented layer – the meals compete, the subscription does not.



The meals compete on shelf; the subscription competes nowhere else

The definitions. Core benefit; actual product; augmented product. **Value** = perceived benefits − perceived costs (money, effort, risk). The **value proposition** names the segment, the benefit and the proof. A **brand** is the promise plus the evidence kept.

The derivation. Customer's ledger for the bundle: benefits – 5 dinners solved (£35 of takeaway avoided), one decision a week, zero spoilage; costs – £29, remembering to freeze, novelty risk. Economic value ≈ £35 − £29 = £6 visible, plus time and mental-load benefits the price does not charge for – the surplus a price rise could later capture. Value math is the ceiling of price; the next unit stands on it.

Worked example 1 – the augmentation that carries the margin. The competitor can copy a frozen meal in a quarter. Copying the 6 pm slot, the subscription logistics and the cancel-anytime trust takes years of systems. So the campaign LEADS with augmented features – "skip a week in one tap" is a meal the rival cannot serve. Product policy follows the defensibility, not the recipe.

Worked example 2 – escalate to the value test. "Meal-kit fatigue" is a known churn driver. Nimbus's augmented answer: week-6 menu rotation and a two-minute "change my box" flow – designed against the churn data, not the chef's pride. The test: does the augmentation move retention? The pilot's month-2 retention is 66.4% (142 of 214 – unit 7.6); each +5 points of retention is worth roughly 214 × 0.05 × 5 × £2.70 × remaining-months – product decisions priced in cohort math, not adjectives.

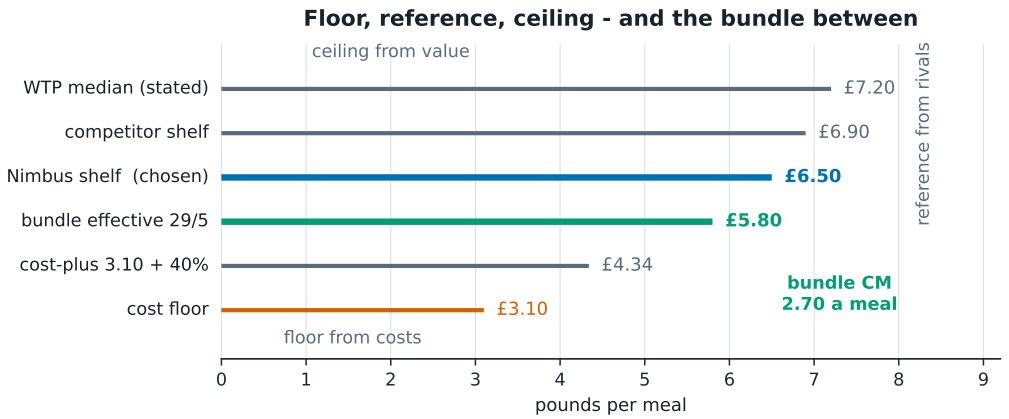
Common mistakes. Competing on the actual product when the augmentation is the moat. Confusing features with value (a benefit the customer does not perceive is a cost). And pricing before counting value – the ceiling comes first.

Summary. Core, actual, augmented – the moat lives outermost. Value = benefits − costs, and it caps the price. Augment against churn, priced in cohort math.

Pricing

Why it matters. Price is the only marketing lever that is pure revenue: a 1% price improvement flows almost untouched to profit. It is also where the cost, the competition and the value all meet – and where the honest marketer states which one they are following.

The idea, with a picture. Nimbus's ladders: cost-plus says £3.10 + 40% = £4.34 (underbids the value); the competitor sits at £6.90; the survey's median willingness-to-pay is £7.20. The shelf price stays £6.50 – but the SUBSCRIPTION bundle prices at £29 for 5 meals (£5.80 a meal, CM £2.70) with a two-meal minimum.



The floor from costs, the ceiling from value, the reference from rivals - and the bundle lives between

The definitions. **Cost-plus** covers cost plus a margin; **competition-based** follows rivals; **value-based** captures a share of the value created. **Penetration pricing** buys volume; **price skimming** skims the eager. Bundles trade price for commitment.

The derivation. Bundle economics: £29/5 = £5.80 effective per meal; CM = 5.80 − 3.10 = £2.70 ✓. The two-meal minimum protects the picking-apart of the bundle (buying 2 of 5 favourite meals at £11.60 undercuts nothing). Shelf £6.50 keeps the anchor high so the bundle reads as the deal: £32.50 à la carte → £29 saves £3.50 and one decision. Price architecture is arithmetic in the customer's head – use theirs, not yours.

> **Worked example 1 – which basis, when.** Cost-plus £4.34 leaves £2.86 of created value uncaptured and signals "cheap" against the £6.90 rival. Matching the rival at £6.90 forfeits the subscription story. Value-based: the bundle's visible surplus is £6 (unit 7.3) – capture part of it (£3.50) in the bundle discount, bank the rest as the customer's reason to subscribe. The honest sentence: "the floor is cost, the ceiling is value, the rival sets the reference – we set the architecture."

> **Worked example 2 – escalate to penetration with a deadline.** Launch month: first 100 subscribers get the 5-meal bundle at £24.80 (£4.96 a meal, CM £1.86). The maths: 100 × 5 × (2.70 − 1.86) = £420 of contribution deliberately spent to seed reviews, referrals and the cohort the funnel (unit 7.6) will judge. Penetration without a reversion date is just a price cut with a story; the sheet says "reverts to £29 in month 2" – and the churn at reversion is itself a measurement.

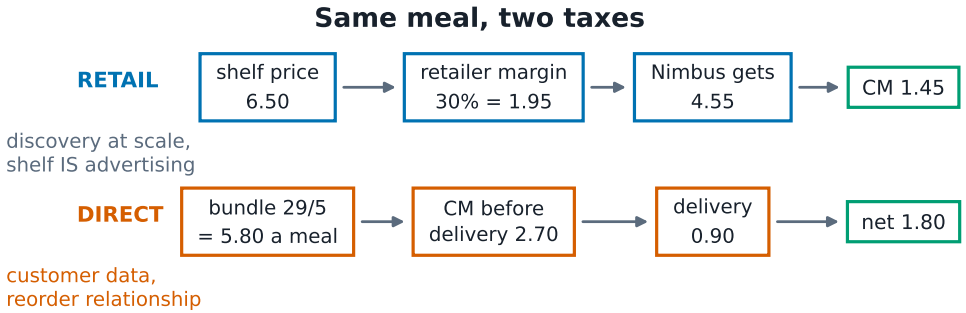
Common mistakes. Pricing the bundle by averaging shelf prices without the minimum-order protection. Confusing WTP STATED (survey £7.20) with WTP revealed (behaviour – always lower). And discounting forever: promotions are loans against the brand.

Summary. Floor = cost, ceiling = value, reference = competition. The bundle trades price for commitment at £2.70 CM. Penetration has a reversion date.

Channels

Why it matters. The customer meets the product through a channel, and the channel's margin is a tax on the story: retail takes 30%, direct takes logistics. Channel choice decides which promise survives contact with economics.

The idea, with a picture. Nimbus's two routes: retail (shelf £6.50, retailer pays £4.55 after 30% margin → CM £1.45 a meal) and direct subscription (£5.80 effective, CM £2.70 but with delivery costs). Same meal, different economics.



Same meal, two taxes - the retailer takes margin, the subscription takes logistics

The definitions. A **channel** moves product, information, promotion and money. **Channel margin** is each layer's cut. **Disintermediation** goes direct; **multichannel** runs both. **Conflict** arises when channels compete on price – the classic shelf-vs-subscription war.

The derivation. Retail: £6.50 × 0.70 = £4.55 to Nimbus; CM = 4.55 − 3.10 = **£1.45** ✓. Direct: £2.70 CM minus £0.90 per-box delivery = £1.80 net – richer than retail AND owns the customer data, the reorder relationship and the feedback loop. Retail still earns its place: discovery at scale, zero logistics, and the shelf IS advertising. The channel P&L prices each job the channel does.

> **Worked example 1 – the margin map.** One city, one retailer, 40 shelf slots at £6.50: 40 × £1.45 = £58 a restock of contribution, plus 600 passer-by impressions a week. The subscription: 214 subscribers × 5 × £1.80 = £1,926 a month. Different jobs, different scales – the map decides where each campaign dollar works, and the honest marketer prices the SHELF's advertising value before cutting retail for "low margin".

> **Worked example 2 – escalate to conflict.** The retailer sees the £29 bundle (£5.80/meal) undercutting its £6.50 shelf and threatens delisting. Three answers, honestly priced: (a) differentiate the bundle (5-meal commitment, delivery – a different product, defensible); (b) share the subscription in-store via QR with a finder's fee (£0.30/meal of the CM – still £1.50 net direct); (c) let the channels fight and lose one. Channel conflict is not a pricing error; it is a design decision with a bill. The chosen answer: differentiate, and say so to the retailer with the map on the table.

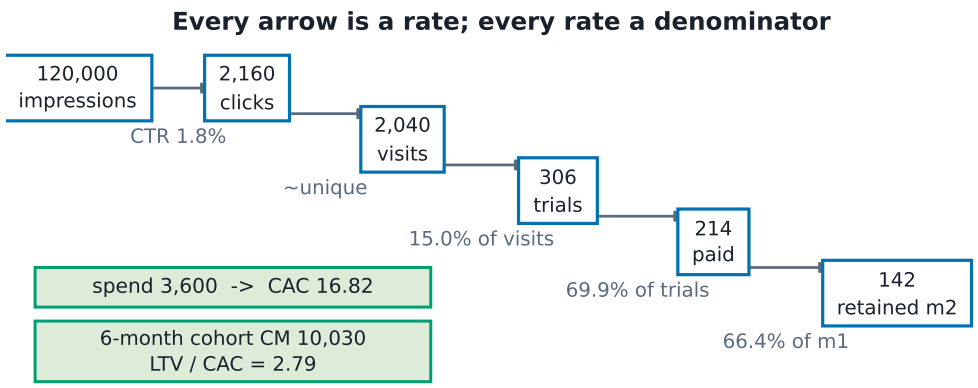
Common mistakes. Comparing channels on margin alone (each does different jobs). Letting direct price undercut the shelf partner without a differentiation story. And forgetting the channel's information flow: direct owns the data that makes unit 7.6's funnel measurable at all.

Summary. Retail: £4.55 in, £1.45 CM, discovery. Direct: £2.70 less logistics, plus the data. Conflict is designed, not stumbled into.

Promotion and digital measurement

Why it matters. £3,600 is about to be spent telling the segment the news. Digital marketing's gift is that everything is measurable; its curse is that everything measurable is not therefore meaningful. Honest denominators are the whole unit.

The idea, with a picture. The month-1 funnel: 120,000 impressions → 1.8% CTR = 2,160 clicks → 2,040 unique landing visits → 306 trial signups (15.0%) → 214 paid conversions (69.9%) → month-2 retention 142 (66.4%). Every arrow is a rate with a denominator, and every rate can be gamed.



Every arrow a rate, every rate a denominator - and CAC meets LTV at the bottom

The definitions. **CTR** = clicks/impressions; **conversion rate** = orders/visits (state which). **CAC** = spend ÷ new customers; **Lifetime value** (LTV) = cohort contribution over life. **Attribution** assigns credit across touches; last-click is the default and the confession. **ROAS** = ad revenue/spend.

The derivation. CAC = 3,600/214 = **£16.82** ✓. Six-month cohort contribution at 34% first-month churn then 15%: £10,030 total → **LTV/CAC = 2.79** ✓ (month-1 CM £2,889; cohorts 214, 142, 121, 103, 88, 75 subscriber-months × £13.50). The healthy range most operators cite: ≥ 3 to absorb overheads – at 2.79 the campaign is close but not home; the lever is retention, not CTR: +5 points of month-2 retention is worth ~£900 of cohort contribution, which no plausible CTR gain buys.

> **Worked example 1 – the denominator police.** "15% conversion!" – visits or clicks? 306/2,040 unique visits = 15.0% ✓, but 306/2,160 clicks = 14.2%, and against impressions it is 0.26%. Each is true; only one is the campaign's steering number. Same discipline on CAC: £3,600 of PAID spend only – the founder's cousin's graphic design is a cost of the decision and belongs in it. Honest metrics name their numerator and denominator in the same breath.

> **Worked example 2 – escalate to attribution.** A subscriber saw a podcast ad, searched, clicked a retargeting ad, bought. Last-click gives retargeting the credit – so retargeting looks like the best channel and the budget migrates there, every quarter, forever. Two repairs: a holdout test (turn retargeting off for a random city – did sales fall by the claimed amount?) and incrementality thinking: credit a channel only for orders that would NOT have happened without it. The funnel measures the pipes; experiments measure the pressure.

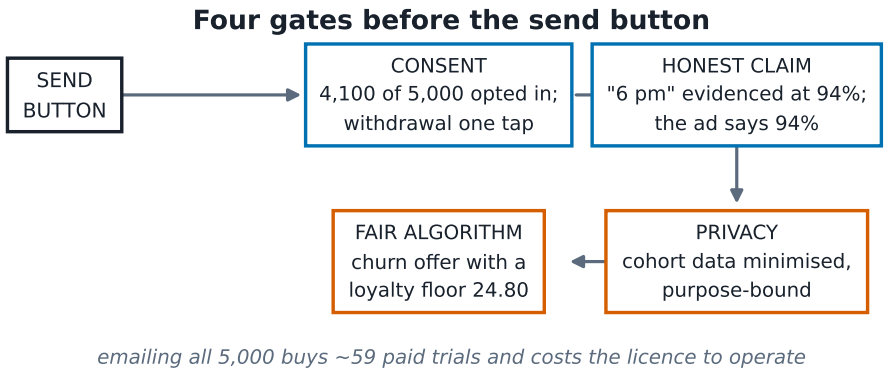
Common mistakes. Optimising CTR into clickbait (the rate rises as the denominator fills with the wrong people). Comparing CAC across periods with different channel mixes. And quoting LTV/CAC from churn assumptions nobody re-tests – the model is a hypothesis with a spreadsheet attached.

Summary. Rates are numerator-over-denominator, stated together. CAC £16.82 vs 6-month LTV/CAC 2.79 – retention is the lever. Attribution defaults lie; holdouts tell the truth.

Marketing ethics

Why it matters. The campaign's final asset is trust: the consent behind the email list, the honesty behind the ad claims, the fairness inside the algorithm. Every earlier number is only worth having if the data and the promises behind it were clean.

The idea, with a picture. The email list: 5,000 addresses, 4,100 with marketing consent. Contacting all 5,000 reaches 18% who never agreed – a small conversion gain bought with the campaign's licence to operate.



Four gates before the send button - each with its evidence line on the sheet

The definitions. Consent : opt-in, informed, withdrawable. **Honest advertising**: claims that are true, fair and substantiable – puffery ("delicious") is legal, measurable claims ("6 pm guaranteed") must be deliverable. **Dark patterns** : design that tricks (the hard cancel, the sneaky basket). **Privacy**: data minimisation and purpose limitation – collect what the purpose needs, use it for that purpose.

The derivation. The arithmetic of the gates: emailing all 5,000 adds 900 × (62% × 15% trial ≈ 84) ≈ 84 trials and ~59 paid (at 69.9%) – call it £1,000 of cohort contribution. The cost: regulator exposure, list poisoning, and the one-word review that outlives the campaign. Trust compounds; the sheet's final line prices it in retention, not fines.

> **Worked example 1 – the substantiable claim.** The ad draft says "always delivered by 6 pm". Operations says 94% by 6 pm. Three options: run it (an unsubstantiated absolute – one 6:15 delivery makes the ad false), soften to "delivered by 6 pm, 94% of the time" (true, checkable, and the number itself signals honesty), or fix operations first. The honest claim is not the cautious one – it is the one the evidence carries.

> **Worked example 2 – escalate to the algorithm.** Churn prediction flags the 72 subscribers likeliest to leave; the retention offer is £10 off. Should it go to (a) the 72 flagged, or (b) everyone? Option (a) is efficient – and punishes the loyal who do not churn, trains customers that threatening to leave pays, and if the model's error falls on the quiet, they never even know they were passed over. Fairness fix: a loyalty floor (never a worse offer than a new customer gets – unit 7.4's reversion priced this), an offer budget open to all at some level, and a human read of the model's misses. Algorithms inherit the objectives they are given; ethics is the objective statement.

Common mistakes. Treating consent as a tick-box (it is withdrawable – and the withdrawal must be as easy as the sign-up). Confusing legal with right. And measuring only the campaign that ran – the people the gate PROTECTED are in the denominator of trust.

Summary. Consent is the licence; claims must be substantiable. Fairness is designed into offers and algorithms, not audited in afterwards. Trust compounds in the retention line.

The campaign sheet (reproducible)

Month 1, Nimbus Kitchen subscription launch:

- Email: 5,000 sent, 408 responses (8.2%), 62% intent, 95% CI [57.3, 66.7]; A/B offer to 200: 27 accept (13.5%).
- Funnel: 120,000 impressions; CTR 1.8% → 2,160 clicks; 2,040 unique visits; 306 trials (15.0% of visits); 214 paid (69.9%); month-2 retention 142 (66.4%).
- Economics: spend £3,600 → CAC £16.82; bundle £29/5 meals → £5.80 effective, CM £2.70/meal, delivery £0.90 → £1.80 net; 6-month cohort CM £10,030 → LTV/CAC 2.79. Shelf: £6.50, retailer pays £4.55, CM £1.45. Launch offer: first 100 at £24.80 (CM £1.86), reverts month 2.
- Segments of 408: professionals 180 (44.1%), fitness 110 (27%), families 69 (17%), other 49 (12.0%).
- Ethics gates: consent 4,100/5,000; "6 pm" claim evidenced at 94%; retention-offer fairness floor £24.80 for all.

Every figure in this handout recomputes from these values.

Topic 08

University Business

What this course is, and how to use it

This is the eighth course in University Business: **organisational behaviour and HRM** – the study of how people actually behave at work, and how to build people processes that are fair, measurable and honest about their evidence.

The running case is **Nimbus Kitchen's front-of-house crew**: eight people, one Saturday rush, and the company's first hires as it grows from kitchen to brand. Every unit works the same crew forward – perception, motivation, teams, leadership, selection, performance, inclusion – with the survey, funnel and meeting numbers recomputable from the crew sheet at the end.

- Units 8.1–8.3 take the individual into the team: perception, motivation, and the design of working groups.
- Units 8.4–8.6 build the system: leadership, culture, selection, development, performance and change.
- Unit 8.7 audits it all: bias, inclusion and ethics with evidence.

How to read each unit: the question first, then the theory with its evidence honestly graded, then two escalating worked examples on the crew, the mistakes people actually make, and a three-line summary.

Individual behaviour and perception

Why it matters. The Saturday rush looks different from the pass and from the oven – and the manager judges what she sees. Perception is not a camera; attribution is not a court. Both can be audited.

The idea, with a picture. A customer films a chef snapping at a waiter; the clip goes round the office. Asked "why?", 25 of 40 staff (62%) say **personality** – "he's abrasive"; 15 (38%) say **circumstances** – "the ticket printer jammed during a double-booked hour". Same clip, two attributions.



Three questions decide an attribution - and the fundamental attribution error is skipping two of them

The definitions. **Selective attention** and **perceptual bias** (halo, recency, contrast) shape what is noticed. **Attribution theory** : behaviour is judged on **consistency** (does the person act this way over time?), **distinctiveness** (only here, or everywhere?), **consensus** (does everyone act this way here?). Low distinctiveness + low consensus → person; the reverse → situation. The **fundamental attribution error** is over-weighting personality and skipping the situation.

The derivation. The chef's ledger: consistency high (short on three of four Saturdays), distinctiveness LOW (also short at the farmers' market stall), consensus LOW (the calm colleague on the same tickets) – the cube points at the person ✓. But add the withheld fact – the printer jammed on all three occasions – and the situation returns. Attribution arithmetic is only as good as its evidence list; the audit question is always "what would I need to see to flip the verdict?"

> **Worked example 1 – the halo audit.** The crew rates the new waiter "brilliant" after one flawless Saturday. The manager checks the record: great with tables, poor at stock (three stock-outs). The halo effect – one vivid strength grading the whole person – had already priced him. Repair: rate the criteria separately (unit 8.6's appraisal is built for this), and the halo loses its seat at the table.

> **Worked example 2 – escalate to stereotype risk.** The oldest crew member is "slow with the tablet" – one instance, one busy day. Calling it an age fact commits the ultimate attribution error (a stereotype acting as a person-verdict). The evidence: two orders mis-keyed in month one, none in month two after a five-minute training. Track *trends within the person* before invoking *categories about the group* – and note the cost of the alternative: stereotyped staff disengage precisely when you need them learning.

Common mistakes. Judging on the vivid instance (recency). Grading the whole

person from one trait (halo). Explaining others by personality and yourself by circumstances (the actor–observer gap).

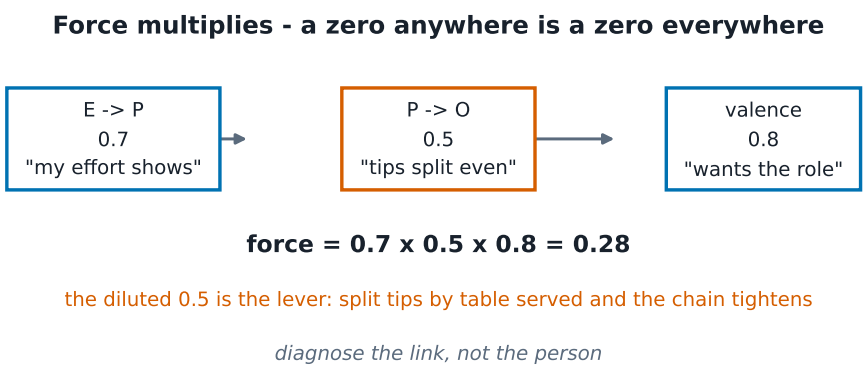
Summary. Perception selects; attribution explains. Consistency, distinctiveness, consensus – ask all three or the verdict is a mood.

Motivation

This unit is about **motivation** .

Why it matters. The Saturday rush is run on discretionary effort – the ten percent nobody can check. Motivation theory is the design language for that effort: what starts it, what directs it, and what sustains it.

The idea, with a picture. Two families: **content** theories ask WHAT motivates (Maslow's ladder, Herzberg's two factors – hygiene prevents dissatisfaction, motivators create satisfaction); **process** theories ask HOW motivation happens. **Expectancy theory** : force = $E \to P \times P \to O \times$ valence.



Force multiplies: any link near zero kills the drive - find the broken link, not the lazy worker

The definitions. Maslow: physiological → safety → belonging → esteem → self-actualisation (a lens, weakly evidenced). **Herzberg's two-factor theory** : pay and conditions are hygiene – their absence demotivates, their presence does not motivate; achievement, recognition, growth are motivators. **Goal-setting theory** : specific, difficult, accepted goals outperform "do your best" – the most replicated finding in OB.

The derivation. The tip-share scheme: if the crew believes effort leads to performance ($E \to P = 0.7$ – staffing makes it reachable), performance leads to the outcome ($P \to O = 0.5$ – tips are split evenly, so the link is diluted), and the outcome is valued (valence = 0.8), then force = $0.7 \times 0.5 \times 0.8 = \mathbf{0.28}$ ✓. Multiplication is the theory's discipline: a zero anywhere is a zero everywhere. The diluted 0.5 is the design lever – split tips by table served and the chain tightens.

> **Worked example 1 – find the broken link.** A quiet, skilled cook stops volunteering for specials. The manager assumes laziness; the chain says measure: $E \to P$ fine (he can cook them), $P \to O$ broken (specials earn team tips, not recognition), valence fine (he wants the sous-chef role). The repair is recognition with a name on it – a specials menu credit – not a pep talk. Process theory turns "motivate them" into "fix the link".

> **Worked example 2 – escalate to goals.** Two Saturday crews: one set "do your best on covers", the other "95 covers, no order over 12 minutes, checked at 2 pm and 6 pm". Stretch-goal crew: 14 of 20 shifts hit the targets; do-your-best crew: 8 of 20 hit the (unspecified) standard. Specific + difficult + feedback – and the acceptance condition is where it earns its honesty: a goal the crew did not accept is an instruction, and instructions do not carry Saturday energy.

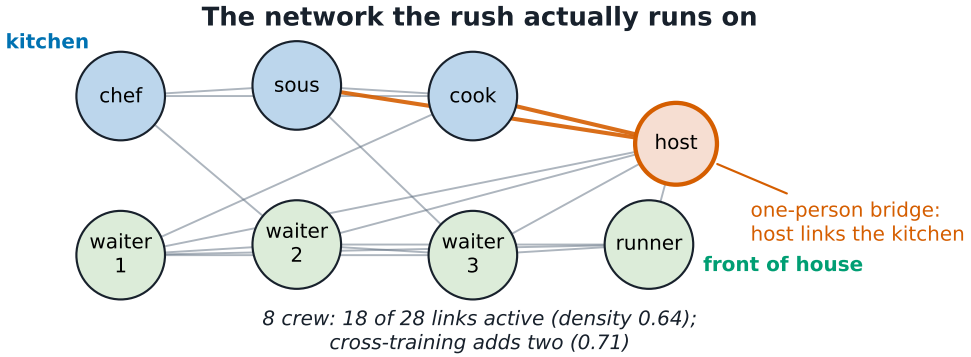
Common mistakes. Treating pay as a motivator when the evidence says hygiene. Setting vague goals ("improve service"). And diagnosing motivation problems as people problems before auditing the expectancy chain.

Summary. Content says what, process says how. Force = $E \to P \times P \to O \times V$ – find the broken link. Specific, difficult, accepted goals win.

Teams

Why it matters. Eight people staff the Saturday rush; the rush is won or lost as a system, not as eight individuals. Team design decides before the first ticket prints.

The idea, with a picture. The crew's sociogram: 18 of the 28 possible communication links are active (density 0.64) – and while a few waiter-to-cook shortcuts exist, both pass-critical links (host–cook and host–sous) run through one person. That bridge is a single point of failure wearing a smile.



Density 0.64 with a one-person bridge - the network the rush actually runs on

The definitions. Tuckman's forming–storming–norming–performing (a map, not a

law). **Role theory** (Belbin): teams need complementary roles, not clones. **Psychological safety** : the shared belief that raising problems is safe – the strongest single predictor of team effectiveness in Google's Project Aristotle. **Social loafing** : effort drops when output is unidentifiable.

The derivation. The meeting audit: in twelve pre-service meetings, three distinct voices spoke – 3 of the 8 crew (37.5%). After the round-robin change ("each name, one blocker, thirty seconds"), voices per meeting rose to seven of the eight and stock-outs fell from three to one a month. The mechanism: the loafing fix (identifiable contribution) IS the safety intervention – a structure that makes speaking routine.

> **Worked example 1 – design before personality.** Two strong personalities collide on plating. The resolution is not mediation first; it is role clarity: one owns tickets, one owns pass quality, the interface is a checklist. Belbin's point: conflict often signals two people cast for the same role – redesign the cast, not the characters.

> **Worked example 2 – escalate to the bridge.** The bridge person takes a festival Saturday off; the crew mis-routes nine orders. The network fix: cross-train one waiter and one cook as a second bridge (adds two links, density to $20/28 = 0.71$), and the checklist makes the hand-off explicit. Networks are infrastructure – they need redundancy like any other infrastructure.

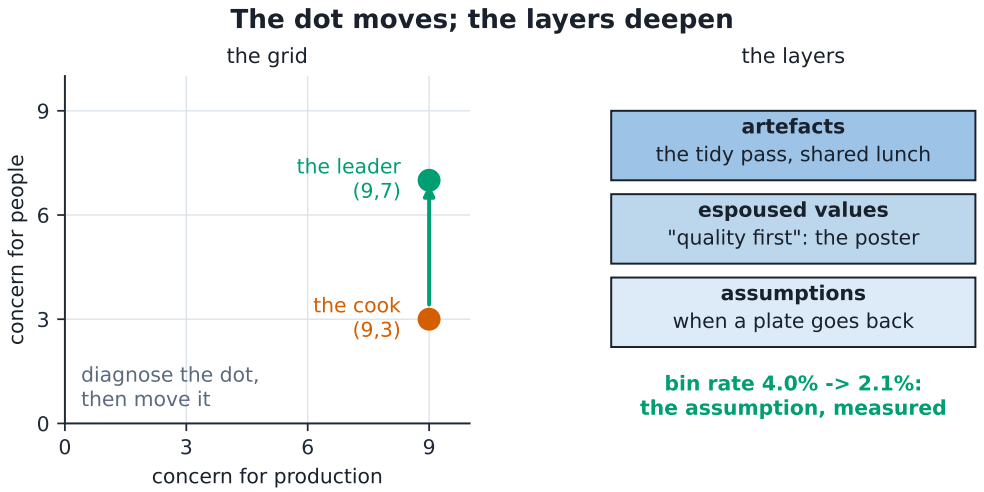
Common mistakes. Building teams of clones ("all stars") and losing the complements. Mistaking harmony for safety – safe teams disagree out loud. Ignoring social loafing until the annual review discovers the invisible shirker.

Summary. Design roles complementary, links redundant. Make contribution identifiable – loafing and silence share a cure. Safety is measured in voices, not vibes.

Leadership and culture

Why it matters. The founder can cook, sell and schedule – and at twelve staff, the job changes from doing to enabling. Leadership theory is the map of that transition, honestly graded by evidence.

The idea, with a picture. Theory families as a grid: task versus relationship behaviours (the managerial grid), then the modern evidence: **transformational** leadership (vision, individual consideration, intellectual stimulation) outperforms transactional bookkeeping – with effect sizes that are real but modest, and the caveat that most studies are correlational.



The dot moves; the layers deepen - culture is what the crew does when the founder is off

The definitions. Managerial grid: concern for production × concern for people; the evidence favours high–high, situational leadership says match the follower's readiness. **Culture** in three layers: **artefacts** (the tidy pass, the shared lunch), **espoused values** ("quality first"), **basic assumptions** (what the crew assumes without asking). Culture is read at the deepest layer that shows.

The derivation. The founder's audit: she scores 9,3 on the grid (nine task, three relationship) – the cook's score. The crew's engagement survey says the strongest driver is "someone notices my development": the dot must move toward 9,7. The behaviour priced concretely: one development conversation a fortnight, diarised like a supplier order – because what gets scheduled gets led.

> **Worked example 1 – situational honesty.** The newest waiter needs directing (new task, low readiness); the sous-chef needs delegating (seasoned, high readiness). Applying the founder's 9,3 everywhere over-directs the sous-chef into quiet disengagement and under-directs the new waiter into anxiety. Situational leadership's prescription (getting the readiness diagnosis right) is common sense made explicit – its evidence base is thinner than its popularity, and the honest sentence says so.

> **Worked example 2 – escalate to culture arithmetic.** Espoused value: "quality first". Artefact: waste bin audit shows 4% of mains binned for defects; assumption test – what happens when a waiter sends back a late plate? If the chef blames, the assumption is "speed hides errors" whatever the poster says; if the chef thanks, the assumption is catching. Culture change starts at assumptions by changing what gets praised – the bin rate is a culture metric, and its trend (4.0% → 2.1% over a quarter) is the change measured.

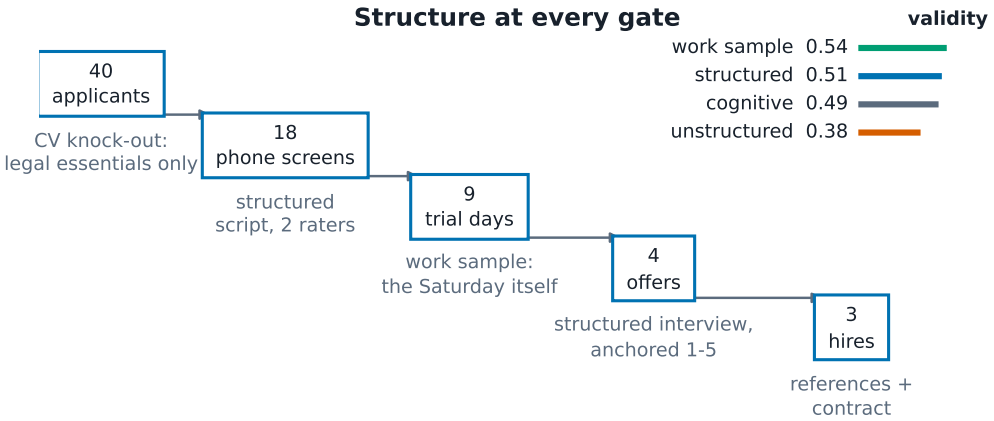
Common mistakes. Anointing one theory as the answer (each explains a slice). Reading culture from the poster (artefacts) instead of the behaviour under pressure. And leadership training without scheduled behaviour change – inspiration decays like milk.

Summary. Grid for the diagnosis, situational for the match, transformational for the evidence (modest, correlational). Culture = artefacts, values, assumptions – change what is praised.

Recruitment and development

Why it matters. Every people problem downstream is cheaper to solve at selection than at correction. Nimbus's first structured hire builds the funnel that the whole course can audit.

The idea, with a picture. The funnel: 40 applicants → 18 phone screens → 9 trial days → 4 offers → 3 hires. Each arrow is a rate, and the funnel's honesty depends on the criteria at each gate – structure beats instinct by half again in every meta-analysis.



Structure beats instinct at every gate - and the work sample is the strongest single predictor

The definitions. Predictor validity (correlation with job performance): work samples ≈ 0.54, structured interviews ≈ 0.51, cognitive tests ≈ 0.49, unstructured interviews ≈ 0.38 (the illusion of insight). **Base rate**: the share of strong performers in the pool (0.2 here). **Reliability** (consistency) precedes validity (a ruler that changes length measures nothing).

The derivation. With base rate 0.20 and a test of validity $r = 0.4$, the top-quartile selection lifts the expected share of strong performers to about 0.35 – the maths of why a mediocre test plus a cut still beats instinct. The funnel's design choice: work-sample trial day (validity ≈ 0.5+) as the widest expensive gate, CVs only as a knock-out for legal requirements. Cost arithmetic: trial day costs £90 a candidate × 9 = £810; one mis-hire costs a season of Saturday chaos (recruitment rule of thumb: 30–50% of salary) – the funnel is an investment with a number.

> **Worked example 1 – structure the interview.** Old question: "Are you good under pressure?" (everyone says yes; answers legally risky). A **structured interview** asks: "Tell me about a Saturday that went wrong – what did you do first?" – same behaviour probe for every candidate, scored on a 1–5 anchored scale by two interviewers independently. The validity jump (0.38 → 0.51) is bought with a question sheet, not a personality.

> **Worked example 2 – escalate to development ROI.** The stock-out problem traces to a training gap. Design: two hours on the ordering system, then measure the gap closed – stock-outs fall from three a month to one, each worth about £40 of lost margin: £80 a month, £960 a year against £180 of training time. ROI honestly stated: the transfer conditions (practice on real Saturdays, a checklist that survives the rush) are where training dies; the evaluation asks whether the BEHAVIOUR changed, not whether the room was warm.

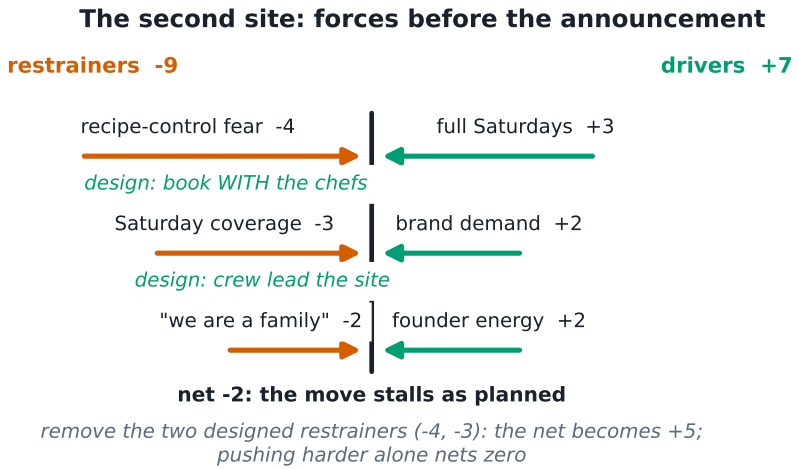
Common mistakes. Unstructured interviews defended by gut ("I know in five minutes" – the five-minute illusion). Validity claimed without a criterion to correlate against. And training evaluated by satisfaction sheets while the behaviour never moves.

Summary. Funnel with structure at every gate; work samples lead validity. Base rate × validity decides the cut's value. Development ROI = measured behaviour change.

Performance and change

Why it matters. Performance appraisal is where fairness becomes visible – or not. And the move to a second site is the founder's first real change programme: resistance guaranteed, information guaranteed too.

The idea, with a picture. Appraisal with measurable criteria: each crew member rated on 3–6 agreed measures (covers served, defect rate, development goal) by two raters independently – the design's fairness is in its redundancy. For change: **force-field analysis** – driving forces against restraining forces, listed, weighted and drawn.



Equilibrium is a balance of forces - the change plan either adds drivers or removes restrainers, never both by accident

The definitions. Rating **errors**: halo, central tendency, recency, leniency – controlled by anchored scales and multiple raters. **BARS**: behaviours anchored to numbers. Change: Lewin's **unfreeze–change–refreeze** ; resistance as information about forces the plan missed, not as insubordination.

The derivation. The second-site force-field: drivers – full Saturday demand (+3), brand pull (+2), founder energy (+2); total +7. Restrainers – recipe-control fear (−4), Saturday coverage worry (−3), "we're a family" identity (−2); total −9. Net −2: the move stalls unless the plan removes restrainers. The designed move: write the recipe book WITH the chefs (removes −4 by converting control to authorship), promise the founding crew the second site's leads (removes −3). Net becomes +5 – and the arithmetic says why "push harder" alone fails: adding +2 of founder pressure against unchanged fears still nets zero.

> **Worked example 1 – the recency rescue.** The waiter's year: strong spring, midsummer wobble, strong autumn. Recency-biased raters score the autumn; the anchored diary (monthly one-line ratings) scores the year. The difference changed one rating band – and one pay decision. Measures defend the rater from themselves.

> **Worked example 2 – escalate to resistance as data.** The sous-chef's resistance to the recipe book turns out to be knowledge: unwritten variations the book would erase. The design changes – the book gains a "house variations" page with his name on it – and the strongest restrainer becomes a driver. Resistance audited is free consulting; resistance crushed is the same information purchased later at the price of trust.

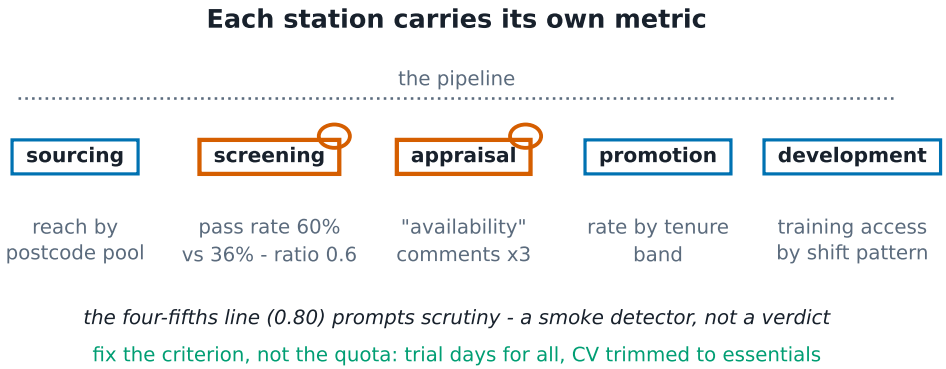
Common mistakes. One rater, one memory, one band. Rating traits instead of behaviours. Leading change with communication alone while the force-field is net negative – inspiration does not alter forces.

Summary. Anchored criteria, multiple raters, the year not the fortnight. Force-field before the announcement: remove restrainers, then push. Resistance is the plan's missing information.

Inclusion and ethics

Why it matters. People processes decide who gets hired, praised, paid and promoted – and every one of them can carry bias invisibly. The audit is what makes inclusion a practice instead of a poster.

The idea, with a picture. The process audit map: each people process (sourcing, screening, appraisal, promotion) checked against its own data for group differences that the criteria cannot explain.



Each station carries its own metric - the audit flags what the criteria cannot explain

The definitions. Adverse impact : a selection rate for a group less than four-fifths of the top group's rate – the standard's trip line, a prompt for scrutiny rather than a verdict by itself. **Bias audit** : comparing process outcomes across groups BEFORE adjusting for legitimate criteria. Ethical frames: utilitarian (greatest good), rights, justice – named, because most HR dilemmas are frames colliding.

The derivation. The screening audit: of applicants from two postcodes, one group passes the CV knock-out at 60%, the other at 36% – a ratio of 0.60, below the four-fifths line ($0.80 \times 0.60 = 0.48$; $36\% < 48\%$). Investigation: the knock-out listed "fine-dining experience", which the one postcode's casual-dining workforce could never show – a criterion sounding neutral, functioning as a gate. Fix: trial-day work samples for all, CV reduced to legal essentials – and the pass-rate gap closes to 52% vs 46%. The audit's arithmetic found what the intent never wrote down.

Worked example 1 – the appraisal distribution. Ratings cluster: the two crew members with caring responsibilities receive "availability concerns" comments at three times the rate of others – with identical measured covers and defect rates. The comment field is carrying the bias the numbers do not hold. Fix: comments must cite a measured criterion or be deleted – making the appraisal's prose as auditable as its scores.

Worked example 2 – escalate to the ethical collision. A key Saturday chef has a second job that sometimes collides. Utilitarian: flex him (the rush needs him); rights: the contract is the contract; justice: flexibility offered to one must be offerable to all. The designed answer: a swap-board with transparent rules – the flexibility is real, the justice condition is met by access, not by exception. HR ethics is mostly the design of systems where the frames stop colliding.

Common mistakes. Declaring processes unbiased because intent was clean – bias lives in outcomes, not intentions. Auditing once and framing it. Treating the four-fifths rule as either proof or permission – it is a smoke detector.

Summary. Audit outcomes by group; fix the criterion, not the quota. The four-fifths line prompts scrutiny. Ethical frames named = dilemmas designed, not ducked.

The crew sheet (reproducible)

The eight-member front-of-house crew, month 12:

- Attribution clip: 40 staff, 25 person (62%) / 15 situation (38%); the chef's evidence ledger: consistency high, distinctiveness low, consensus low – with the printer-jam fact withheld.
- Motivation: $E \to P$ 0.7, $P \to O$ 0.5, valence 0.8 $\to$ force 0.28; goal crews 14/20 versus 8/20.
- Team: 8 members, 18/28 links (density 0.64), one bridge; voices 3/12 meetings $\to$ 7/8 after the round-robin; stock-outs 3 $\to$ 1 a month.
- Leadership: grid 9,3 $\to$ target 9,7; bin rate 4.0% $\to$ 2.1%.
- Selection: 40 $\to$ 18 $\to$ 9 $\to$ 4 $\to$ 3; base rate 0.20; screen pass rates 60% vs 36% (ratio 0.60) $\to$ 52% vs 46% after the fix.
- Change: drivers +7, restrainers −9, net −2 $\to$ +5 after the two designs; training ROI £960 a year on £180.

Every figure in this handout recomputes from these values.

Topic 09

University Business

What this course is, and how to use it

This is the ninth course in University Business: **corporate finance** – the discipline of deciding which cash to take, which projects to fund, and what that cash costs. Accounting (course 2) recorded the past; finance prices the future.

The running case is **Nimbus Kitchen's second-site raise**: £60,000 of new capital, one five-year project, and every number machine-verified. The course's one habit: **a pound today is not a pound tomorrow – so every comparison happens at one point in time.**

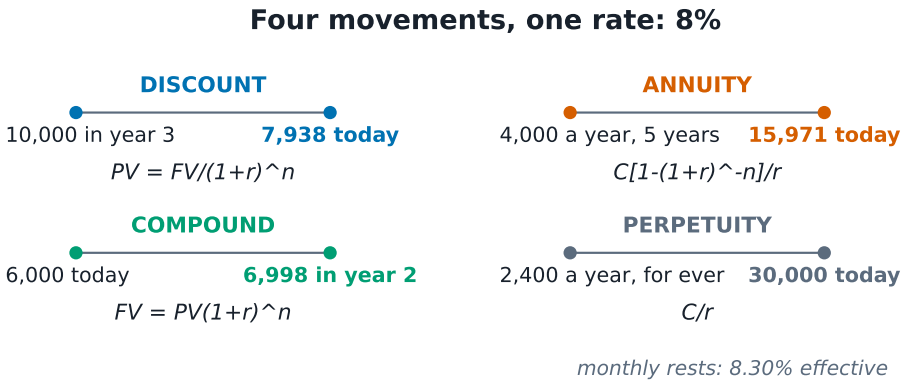
- Units 9.1–9.3 build the machine: time value, incremental cash flows, and the appraisal rules (NPV, IRR, payback).
- Units 9.4–9.5 price the risk and the capital: expected return, diversification, WACC.
- Unit 9.6 closes the loop: financing choices and working capital as the cash the growth eats.

How to read each unit: the question first, then the machinery, then two escalating worked examples on the raise, the mistakes people actually make, and a three-line summary.

Time value of money

Why it matters. The second site costs £60,000 now and pays back over five years. Until those pounds are moved to one common date, "it pays back more than it costs" is not even a sentence.

The idea, with a picture. At 8%, a pound today grows to £1.08 in a year; a pound next year is worth 1/1.08 today. Everything in the course is one of four movements: **compounding** forward, **discounting** back, an **annuity** (equal payments), or a **perpetuity** (equal payments for ever).



Every comparison happens at one point in time - usually today

The definitions. $PV = FV/(1+r)^n$; $FV = PV(1+r)^n$. Annuity of C for n years: $C[1-(1+r)^{-n}]/r$. Perpetuity: C/r. An APR of 8% compounded monthly is $\frac{8\%}{12}$ a month – an effective annual $(1+0.08/12)^{12}-1 = 8.30\%$ ✓.

The derivation. The four checks: £10,000 in 3 years at 8% → $10\,000/1.08^3 = £7\,938$ ✓; £6,000 today in 2 years → $6\,000 \times 1.08^2 = £6\,998$ ✓; £4,000 a year for 5 years → $4\,000 \times [1-1.08^{-5}]/0.08 = £15\,971$ ✓; £2,400 for ever → $2\,400/0.08 = £30\,000$ ✓. Four formulas, one idea.

> **Worked example 1 – which is worth more?** £10,000 in 3 or £6,000 today plus £4,000 in year 2? Move everything to today: £7,938 against $6\,000 + 4\,000/1.166 = 6\,000 + 3\,430 = £9\,430$. The second wins by £1,492 – and the comparison took one line because both sides sit at the same date. Never compare a future pound with a present pound; move one of them.

> **Worked example 2 – escalate to the mortgage.** The site's fit-out is financed at 8% nominal, monthly rests: the monthly rate is 0.667%, and the effective annual rate is 8.30% – not 8%. On £60,000 over 10 years, quoting 8% flat understates the true rate's cost by roughly £180 of extra interest a year. The convention (nominal ÷ rests) is law in some markets and a trap in all of them: always ask "how many rests, and which rate did you quote?"

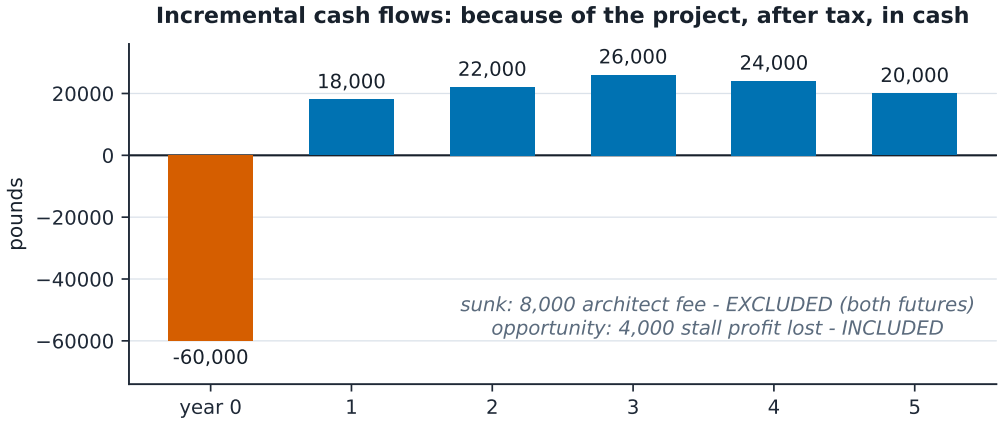
Common mistakes. Comparing cash at different dates. Using the nominal rate as if it were effective. And discounting at a rate that does not match the payment frequency (monthly cash needs the monthly rate).

Summary. $PV = FV/(1+r)^n$; annuities and perpetuities are discounting pre-assembled. Nominal ≠ effective. One date, then compare.

Cash flows and valuation basics

Why it matters. NPV is arithmetic on a list; the craft is the LIST. Most valuation errors are cash-flow errors: sunk costs included, opportunity costs missed, cannibalisation ignored.

The idea, with a picture. The second site's INCREMENTAL cash flows exclude the £8,000 already spent on the architect (sunk), include the £4,000 a year the market stall loses (opportunity), and add back depreciation's tax shield after tax.



Incremental means: because of the project, after tax, in cash

The definitions. Incremental cash flow : the difference the project makes. **Sunk cost** : already spent, never enters. **Opportunity cost** : the best alternative forgone, always enters. **Cannibalisation**: sales stolen from your own products, entered net. Financing costs are in the DISCOUNT rate (unit 9.5), never in the flows.

The derivation. The raise's flows: −60,000; then +18,000, +22,000, +26,000, +24,000, +20,000 – constructed by: revenue 78,000 rising to 92,000; cash costs 54,000 rising to 62,000; depreciation 12,000 a year (non-cash) with tax at 25% on (revenue − costs − depreciation), then depreciation added back; plus the working-capital injection of 4,000 in year 0 recovered in year 5. Every line differences to zero against "do nothing".

> **Worked example 1 – the architect's 8,000.** The founder argues "we've spent 8,000 on drawings, so the project must go ahead". The 8,000 is in BOTH futures – with and without the site – so it differences to zero. Include it and the NPV falls by exactly 0. The rule is not heartless; it is bookkeeping: the money is gone at the decision, and grief is not a cash flow.

Worked example 2 – escalate to the stolen Saturday crew. The second site will use the existing Saturday crew on festival days: £3,000 of contribution that would otherwise be earned ANYWAY is not incremental – but the £1,200 of lost festival orders (capacity genuinely diverted) is. Net entry: −1,200. The test for every line: does the cash change BECAUSE of the project? If yes, in it goes at the moment it moves.

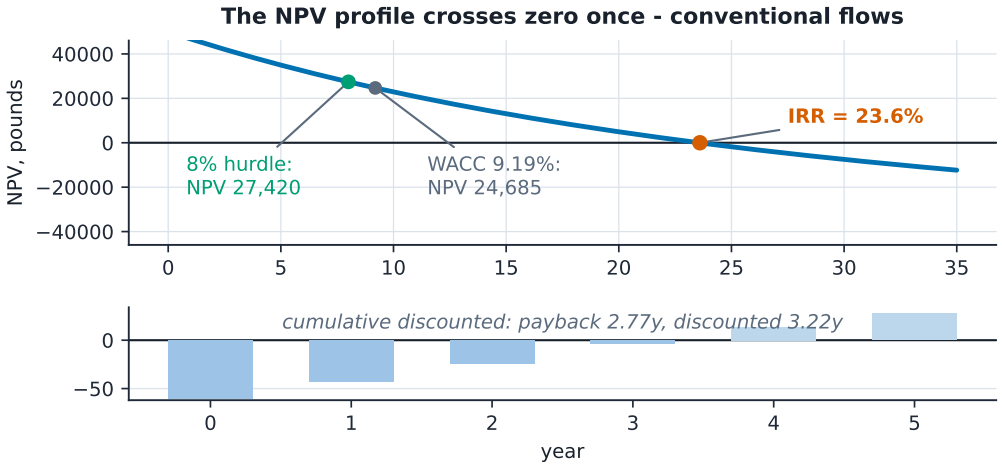
Common mistakes. Including sunk or financing flows (double count). Missing opportunity costs (the stall). And using accounting profit where cash moves: depreciation shelters tax but is not a cash out.

Summary. Incremental = because-of, after-tax, in-cash. Sunk never; opportunity always; financing in the rate. Depreciation's only entry is its tax shield.

Investment appraisal

Why it matters. The five flows are on the table; the decision needs a rule. Three are on the table in every boardroom – NPV, IRR, payback – and they do not always agree. Finance's answer: NPV is the referee.

The idea, with a picture. NPV at 8% = **£27,420** ✓ (positive – create value). IRR = **23.6 per cent** ✓ (above the 8% hurdle → same verdict). Payback = **2.77 years** ✓; discounted payback 3.22 years.



The profile crosses zero once (conventional flows) - where IRR is honest

The definitions. NPV = Σ $CF_t/(1+r)^t$ − I: pounds of value created today. **IRR**: the rate where NPV = 0. **The hurdle rate** is the opportunity cost of capital (unit 9.5's WACC). Payback: years to recover the outlay – a liquidity measure, not a value measure.

The derivation. Year by year at 8 per cent: 18,000/1.08 = 16,667; 22,000/1.166 = 18,858; 26,000/1.260 = 20,635; 24,000/1.360 = 17,647; 20,000/1.469 = 13,612. The inflows sum to 87,420, so NPV = 87,420 − 60,000 = 27,420. The IRR solves the same sum for zero: at 23.6 per cent the discounted inflows equal 60,000. Payback: after year 2 the cumulative inflow is 40,000, and year 3 delivers 26,000, so payback is 2 + 20,000/26,000 = 2.77 years (discounted: 3.22).

Worked example 1 – the three rules on one project. NPV says +27,420: do it. IRR says 23.6% > 8%: do it. Payback says 2.77 years: comfortable. All three agree here because the flows are conventional (one sign change) and the project is independent. The discipline: quote NPV as the decision (pounds of value), IRR as a summary (a rate – useful for comparing scale-free), payback as a risk note (how long the money is exposed). Each answers a different question; only NPV answers "how much richer?".

Worked example 2 – escalate to where IRR lies. Alongside site A, a quick refit project C needs only £5,000 (flows −5,000; +4,200; +2,700 over two years): its IRR is 26.6% – HIGHER than A's 23.6% – yet its NPV is £1,204, a twenty-third of A's £27,420. Ranking by IRR prefers C; NPV prefers A, and NPV is right: the company gets 27,420 richer with A. IRR ignores SCALE – the higher rate on £5,000 adds £1,204 of value; the slightly lower rate on £60,000 adds £27,420. With non-conventional flows (a later outflow), IRR can exist twice or not at all – the profile crosses zero twice. The referee stays NPV; the others are commentary.

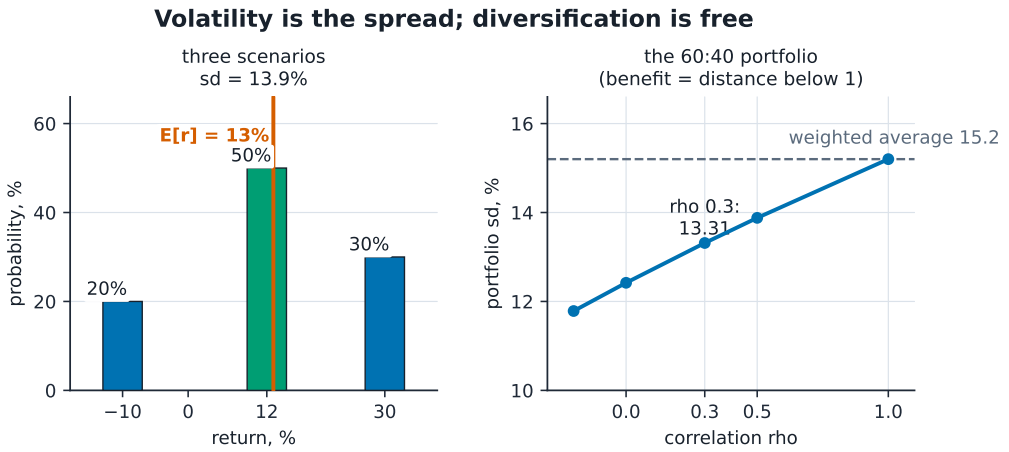
Common mistakes. Ranking mutually exclusive projects by IRR. Accepting on payback alone (it ignores everything after recovery). And quoting IRR without the profile when flows change sign more than once.

Summary. NPV = value in pounds – the decision rule. IRR = where the profile crosses zero; misleads on scale and multiple crossings. Payback = exposure time. Quote all three; decide on NPV.

Risk and return

Why it matters. The 8% hurdle is not given by nature – it is the return investors demand for the risk they carry. This unit builds the demand curve: expected return, volatility, and the one free lunch in finance.

The idea, with a picture. The second site's return distribution: boom 30% (p = 0.3), base 12% (p = 0.5), bust −10% (p = 0.2). **Expected return** = 0.3×30 + 0.5×12 + 0.2×(−10) = **13.0%** ✓; **volatility** (sd) = **13.9%** ✓ – the spread around the mean.



Volatility is the spread; diversification is the only free lunch

The definitions. $E[r] = \sum p_i r_i$; $\sigma^2 = \sum p_i (r_i - E[r])^2$. **Risk premium** : the excess over the 4% risk-free rate. **Diversification** : combining imperfectly correlated assets lowers portfolio volatility below the weighted average – nothing else in finance is free.

The derivation. Variance: $0.3(0.30-0.13)^2+0.5(0.12-0.13)^2+0.2(-0.10-0.13)^2 = 0.0193 \rightarrow \sigma = 13.89\%$ ✓. A 60:40 portfolio of a stock ($\sigma = 20\%$) and a bond ($\sigma = 8\%$) at correlation 0.3: $\sigma_p^2 = 0.36(0.04) + 0.16(0.0064) + 2(0.6)(0.4)(0.3)(0.2)(0.08) = 0.01771 \rightarrow \sigma_p = 13.31\%$ ✓ – below the 15.2% weighted average, because the third term shrinks with ρ. Risk that diversification cannot remove is **systematic** ; risk it can is **specific** . Only the systematic kind is priced.

> **Worked example 1 – price the site's risk.** The site's 13% expected return against a 4% risk-free rate is a 9-point premium. If the site's volatility is mostly specific (a single location), investors holding many sites bear little of it – and should not demand much premium. The finance answer: the 8% hurdle was never about this site alone; it is the WACC of unit 9.5, where only systematic risk survives the portfolio. A single-site founder FEELS idiosyncratic risk; a diversified investor does not price it.

> **Worked example 2 – escalate to the correlation knife.** The same 60:40 portfolio at $\rho = 1.0$: $\sigma_p = 15.2\%$ (no benefit); at $\rho = 0$: 12.4%; at $\rho = -0.2$: 11.8%. Diversification's power is exactly the distance of ρ below 1 – which is why "add another bakery" ($\rho \approx 0.9$ with the first) barely helps, while "add the bond fund" ($\rho \approx 0.3$) does. In the 2008 crisis correlations rose toward 1 and the free lunch briefly vanished; the honest course says so.

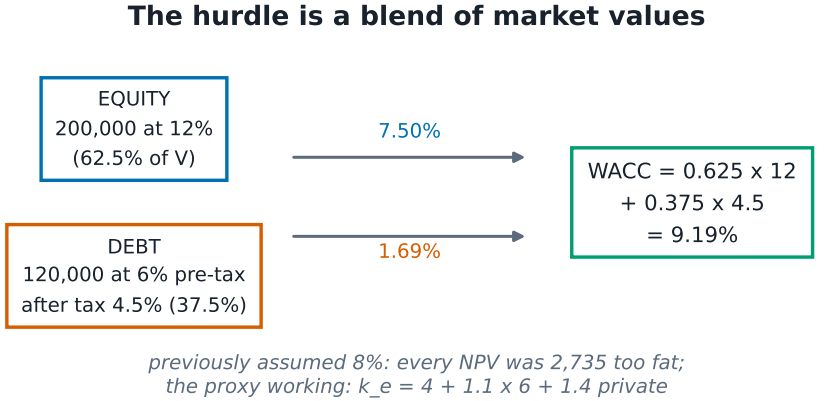
Common mistakes. Using past average return as the expectation without a scenario or a model. Adding volatilities (variances add, with the covariance term). And treating specific risk as priced – the market pays for systematic risk only.

Summary. $E[r] = \sum p_i r_i$; σ = the spread. Portfolio σ falls with the covariance term; only systematic risk earns a premium.

Cost of capital

Why it matters. Every appraisal rule needs the hurdle rate, and the hurdle rate is the blended cost of the money actually raised: £200,000 of equity and £120,000 of debt. Get the blend wrong and every NPV the company ever computes is wrong in the same direction.

The idea, with a picture. WACC with market-value weights: equity £200,000 at 12%, debt £120,000 at 6% pre-tax with tax at 25% → **9.19%** ✓. The weights are what the capital is worth NOW, not the letterhead.



The hurdle rate is a blend of market values - and debt is cheaper only because the taxman pays part of it

The definitions. $WACC = \frac{E}{V}k_e + \frac{D}{V}k_d(1-t)$. Cost of equity k_e: what equity investors require (CAPM's $k_e = r_f + \beta(E[r_m] - r_f)$ when a beta exists). The **tax shield** : interest is deductible, so debt's effective cost is $k_d(1-t)$. **Book vs market weights**: the market values, always.

The derivation. $E = 80\,000$ shares × £2.50 = £200,000; D = £120 000; V = 320 000. Weights 62.5% / 37.5%. $WACC = 0.625(12\%) + 0.375(6\% \times 0.75) = 7.5\% + 1.69\% = 9.19\%$ ✓. Recompute the site's NPV at 9.19%: the inflows discount to £84,685 → NPV £24,685 – still comfortably positive, but £2,735 thinner than the 8% guess. The hurdle is not a rounding detail.

Worked example 1 – the beta-less private firm. Nimbus is private: no beta, no quoted price. Proxy: a listed food company with $\beta = 1.1$, $E[r_m] = 10\%$, $r_f = 4\%$ $\rightarrow$ $k_e = 4 + 1.1 \times 6 = 10.6\%$; add a private-risk premium of 1.4 points $\rightarrow$ 12% (the number above). Every step is written down because each is a judgement: the proxy, the premium, the market estimate. A WACC with no working shown is a number wearing a lab coat.

Worked example 2 – escalate to the debt illusion. Borrow more at 6% (after tax 4.5%) and WACC falls: at D = £240,000 the blend is 7.91%. Cheaper for ever? No – each tranche of debt raises the equity's risk (β rises toward 1.5, k_e toward 13%+) and eventually the 6% becomes 8% for the riskier borrower. The curve is U-shaped: the bottom is the right amount of debt, not the most the bank will lend. WACC is a snapshot of a balance sheet, not a dial to minimise by borrowing.

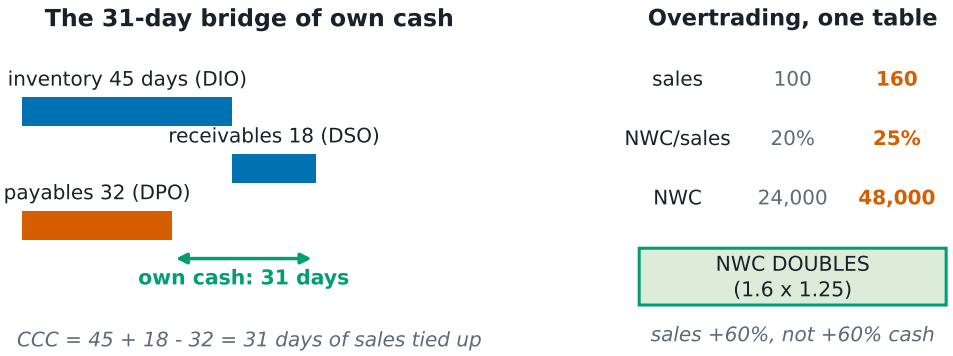
Common mistakes. Book-value weights (equity at par when the market says 2.50). Forgetting the tax shield (using k_d pre-tax) – or claiming it while forgetting interest must actually be paid. And recomputing WACC per project with project-level debt: the hurdle moves with the COMPANY's blend, adjusted for the project's RISK.

Summary. $WACC = \frac{E}{V}k_e + \frac{D}{V}k_d(1-t)$ at market weights; private k_e by proxy with the judgements written down. Debt helps to a point; the U-curve finds the point.

Financing and working capital

Why it matters. The raise succeeded; now the cash has to survive contact with growth. Fast-growing kitchens die of one thing: the cash cycle eating the working capital before the profit arrives.

The idea, with a picture. The **cash conversion cycle** : inventory 45 days + receivables 18 − payables 32 = **31 days** ✓ of the business's own cash tied up. Growth multiplies it.



Growth multiplies the cycle - and overtrading is the multiplier turning profit into a cash crisis

The definitions. $CCC = DIO + DSO - DPO$. **Overtrading**: sales grow faster than the working capital funding them – ratios worsen as the cycle is financed by stress (overdraft, late payments). **Debt vs equity financing**: debt is cheaper and tax-shielded but rigid; equity is permanent but dilutive. **Matching**: fund long assets with long money.

The derivation. Second site opens: sales rise 60%; with the same terms the NWC need rises 60% – but new-supplier terms are worse (NWC/sales 20% $\rightarrow$ 25%), so working capital must rise $1.6 \times 1.25 = 2.0\times$ ✓: DOUBLE, not 60%. At the old scale NWC was £24,000; the new need is £48,000 – £24,000 of extra cash the profit-and-loss will never show as a cost, and exactly the trap the overtrading table names.

Worked example 1 – choose the funding. The £60,000 raise: option A, all equity at 12%; option B, £24,000 debt at 6% (after tax 4.5%)

- £36,000 equity. Blended cost: A = 12.0%; B = 0.4×4.5 + 0.6×12 = 9.0% – cheaper, and the WACC curve (unit 9.5) said this much debt is left of the U's bottom. The matching principle checks the term: the fit-out is a 10-year asset, the debt is 5-year – mismatched; either lengthen the debt or accept the refinancing risk, in writing.

Worked example 2 – escalate to the rescue package. Six months in: sales +60% as planned, but the cash cycle stretched to 38 days and the overdraft is at its limit. The diagnosis in the overtrading table: NWC/sales has hit 26% – the business is profitable and insolvent at the same time. The levers, in order of speed: collect faster (DSO 18 $\rightarrow$ 12 releases six days of sales, about £3,200 on the new £192,000 a year), stretch suppliers honestly (DPO 32 $\rightarrow$ 35 releases three days of purchases), and only then the slow lever – permanent capital. The lesson of the whole course in one sentence: profit is an opinion; cash is a fact, and the cycle decides which fact.

Common mistakes. Reading profit as liquidity (the overtrading trap). Funding long assets with overdrafts. And comparing debt's 6% with equity's 12% without the shield and the risk difference.

Summary. $CCC = DIO + DSO - DPO$; growth multiplies it, and ratio-worsening doubles it. Match term to life. Profit is an opinion; cash is a fact.

The project sheet (reproducible)

Second-site raise, five years:

- Flows: −60,000; +18,000; +22,000; +26,000; +24,000; +20,000.
- NPV at 8% = 27,420; at 9.19% (WACC) = 24,685. IRR = 23.6%. Payback 2.77 years; discounted 3.22.
- Risk: E[r] = 13.0%, σ = 13.9% (boom 30%/0.3, base 12%/0.5, bust −10%/0.2). 60:40 portfolio σ 13.31% at $\rho = 0.3$.

- WACC: E 200,000 (80,000 × 2.50) at 12%; D 120,000 at 6% pre-tax, tax 25% → 9.19%. CAPM proxy: 4 + 1.1×6 = 10.6% + 1.4 private.
- Working capital: CCC = 45 + 18 − 32 = 31 days; sales +60% with NWC/sales 20% → 25% means NWC doubles (24,000 → 48,000).

Every figure in this handout recomputes from these values.

Topic 10

University Business

What this course is, and how to use it

This is the tenth course in University Business: **operations management** – the discipline of making the promise real. Finance (course 9) priced the second site; this course runs it: the line that bakes, the quality that holds, the stock that neither runs out nor drowns, the forecast that sizes it, the chain that feeds it, the project that built it, and the footprint it leaves.

The running case is **the second site's bake line on the festival contract**: 120-bun batches, 480 buns an hour, 19,200 a week, and every number machine-verified. The course's one habit: **find the constraint, measure the error, state the boundary.** An operations claim without a number behind it is an opinion.

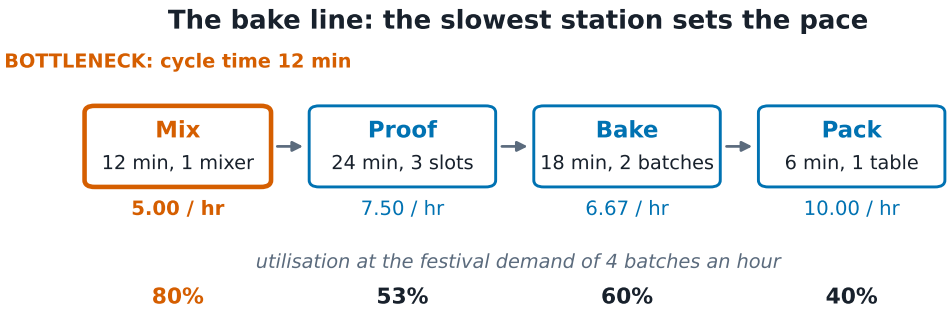
- Units 10.1–10.4 run the line: processes and capacity, quality, inventory, forecasting.
- Units 10.5–10.6 look outward and back: the supply chain and the project that built the site.
- Unit 10.7 closes the loop: the footprint, with the boundary stated.

How to read each unit: the question first, then the machinery, then two escalating worked examples, the mistakes people actually make, and a three-line summary.

Processes and capacity

Why it matters. The festival contract promises 3,840 buns a day. Whether the site can deliver is not a matter of effort – it is arithmetic on four stations, and the answer lives at one of them.

The idea, with a picture. The line: **mix** 12 min a batch, a three-slot **prover** at 24 min, a two-batch **oven** at 18 min, **pack** at 6 min. Capacities: mix 5.00 batches/hr, proof 7.50, bake 6.67, pack 10.00. The slowest – mix – is the **bottleneck** : the line runs at 5 batches = 600 buns an hour, and the **cycle time** is 12 minutes a batch.



The line runs at its slowest station - nothing else matters until Mix moves

The definitions. **Throughput** : the rate the line actually produces. **Capacity utilisation** = throughput ÷ capacity. **Little's law** : $WIP = TH \times FT$ – things in the system equal the rate through it times the time through it. **Work in progress** (WIP): batches inside the line.

The derivation. Capacity per station = $60 \times \text{slots}/t$: mix $60/12 = 5.00$; proof $60 \times 3/24 = 7.50$; bake $60 \times 2/18 = 6.67$; pack $60/6 = 10.00$. At the festival demand of 4 batches an hour, utilisations are $4/5 = 80\%$, $4/7.5 = 53.3\%$, $4/6.67 = 60.0\%$, $4/10 = 40\%$. Flow time $12 + 24 + 18 + 6 = 60$ min $= 1$ hr, so Little's law: $WIP = 4 \times 1 = 4$ batches ✓ – exactly what a walk down the line counts.

> **Worked example 1 – where an hour is won.** The head baker wants to speed up packing (40% utilised). Speeding pack to 3 min moves the line nowhere: mix still caps it at 5 batches an hour. The hour is won at mix: a second mixer lifts mix to 10.00 an hour – and the bottleneck MOVES to bake (6.67), so the line gains only $6.67 - 5 = 1.67$ batches an hour (200 buns). Every investment question in operations starts the same way: where is the constraint now, and where does it move?

> **Worked example 2 – escalate to the Saturday surge.** Festival Saturdays need 6 batches an hour for 8 hours. One mixer caps the line at 5.00: a shortfall of 1 batch an hour, 8 batches over the day. Two answers. Buy the second mixer (bake's 6.67 then clears the 6) – capital spent for one day a week. Or BAKE AHEAD: two extra hours on Friday evening at the line's 5 an hour build a buffer of $5 \times 2 = 10$ batches – inventory is rate × time, the same law – covering the 8 with two spare. Buffers are not a failure of the line; they are how a fixed line meets a moving demand, at the price of freshness, which is why the buffer is Friday's buns and not Saturday's promises.

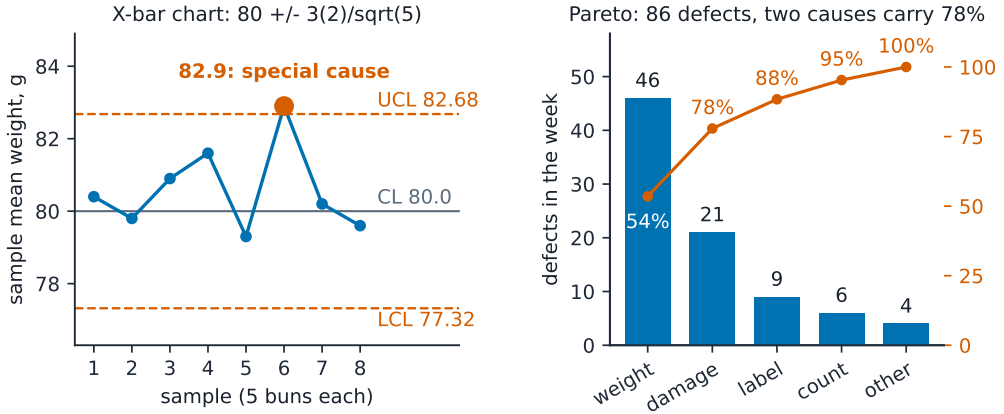
Common mistakes. Balancing the line to raise average utilisation (instead of raising the bottleneck). Quoting capacity in buns before checking the batch size. And confusing flow time (1 hr through the process) with cycle time (12 min between outputs).

Summary. Capacity = the slowest station; cycle time is its inverse. Little's law ties WIP, throughput and flow time. Fix the constraint or buffer around it – in that order.

Quality

Why it matters. The contract's tolerance is ±10% on an 80 g bun. Weighing every bun finds bad ones; the statistical process control discipline finds the batch going bad BEFORE the buns ship.

The idea, with a picture. Buns are sampled 5 at a time. With the process centred at 80.0 g and bun-to-bun sd 2.0 g, the sample-mean **control chart** runs $80 \pm 3 \times 2/\sqrt{5} = 80 \pm 2.68$ g: limits 77.32 and 82.68 ✓.



A point beyond the limits is a signal; within them, leave the process alone

The definitions. **Statistical process control** (SPC): steering the process, not inspecting the product. **Common cause** variation: the noise inside the limits. **Special cause**: a point beyond the limits or a run – evidence something changed. **Pareto analysis** : rank the defects; the vital few carry most of the count. **Cost of quality**: prevention + appraisal + internal + external failure.

The derivation. Limits $80 \pm 3\sigma/\sqrt{n} = 80 \pm 3 \times 2/\sqrt{5} = 80 \pm 2.68$ ✓. The same limits from the range table: $\bar{R} = d_2\sigma = 2.326 \times 2 = 4.65$ g, and $A_2\bar{R} = 0.577 \times 4.65 = 2.68$ ✓ – two routes, one answer. The R-chart's upper limit: $D_4\bar{R} = 2.114 \times 4.65 = 9.83$ g. The week's eight means: 80.4, 79.8, 80.9, 81.6, 79.3, **82.9**, 80.2, 79.6 – the sixth sits beyond 82.68 ✓ (the new flour lot), and the next week's eight means (79.4 to 79.9, all below centre for eight samples running) flag a drift by the run rule even though every point is inside the limits.

Worked example 1 – read the chart, don't nudge it. Sample 6 reads 82.9. The wrong response: retune the mixer by −0.5 g – that is tampering, adding the operator's hand to the noise (Deming's funnel). The right response: FIND the cause – the flour lot changed at that batch, its protein runs higher, water absorption shifted the dough. Fix the lot, and the chart returns to centre on its own. Within limits, leave the process alone; beyond them, look for a nameable cause.

Worked example 2 – escalate to the 19,000-pound week. The defect Pareto: underweight 46, damaged pack 21, mislabel 9, wrong count 6, other 4 (86 total; the first two carry 67/86 = 77.9% ✓). The cost ledger: prevention 1,200 + appraisal 2,400 + internal failure 5,600 + external failure 9,800 = 19,000 – 19% of the week's 100,000 revenue, and failure is 81% of it. The lever the Pareto names: underweight traces to one filler head – 800 spent rebuilding it (prevention) removes 46 defects and most of the 9,800 external bill. Quality is not inspected in; it is paid for at the front, where it is cheapest.

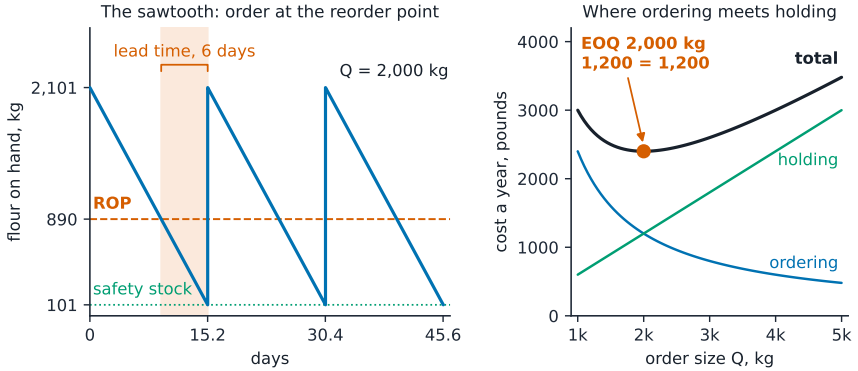
Common mistakes. Chasing every inside-limits wiggle (tampering inflates variance). Reading the R-chart with the X-bar's limits. And reporting the defect count without the Pareto – 86 means nothing; "46 from one filler head" means Monday morning.

Summary. Limits at $3\sigma/\sqrt{n}$; a point beyond or a run is a signal with a nameable cause. Pareto ranks what to fix; the cost of quality says fix it early. Steer, don't nudge.

Inventory

Why it matters. Flour is 60% of a bun and 48,000 kg a year of the site's cash. Too little stops the line; too much is money mouldering in bins. The economic order quantity is where the two costs balance.

The idea, with a picture. Ordering costs £50 an order; holding costs £1.20 per kg a year. $EOQ = \sqrt{2 \times 48{,}000 \times 50/1.20} = 2{,}000$ kg ✓ – 24 orders a year, one every 15.2 days, £1,200 of ordering and £1,200 of holding, exactly equal ✓.



Order when the sawtooth falls to the reorder point - not when the bin looks empty

The definitions. Economic order quantity (EOQ): the order size minimising ordering + holding. **Reorder point** (ROP): the level that triggers an order – lead-time demand plus safety stock. **Safety stock** : $z \times \sigma_{LT}$, the buffer bought for a service level. **Continuous review**: watch the level; **periodic review**: order on a clock.

The derivation. Daily demand 48,000/365 = 131.5 kg; over the 6-day lead time 131.5 × 6 = 789.0 kg leaves the bin while the order is on the road. Demand sd 25 kg/day gives $\sigma_{LT} = 25\sqrt{6} = 61.2$; at 95% service $z = 1.645$, so safety stock = 1.645 × 61.2 = 101 kg ✓ and $ROP = 789 + 101 = 890$ kg ✓. At the EOQ the annual costs tie: ordering 24 × 50 = 1,200, holding 1,000 × 1.20 = 1,200 – the balance the square root buys.

Worked example 1 – the discount that had to be checked. The mill offers 2% off the £2.00 kg price for orders of 4,000 kg or more. The tempting answer is "2% of 96,000 – take it". The check: at 4,000 kg, purchase 48,000 × 1.96 = 94,080, ordering 12×50 = 600, holding 2,000×1.20 = 2,400 – total £97,080 against £98,400 at the EOQ ✓, saving 1,320 a year. But the margin is honest: the extra logistics cost is exactly 600, so the breakeven discount is 600/96,000 = 0.625% – below that, "cheap" flour costs more than it saves. The EOQ is not always the order; it is the number every deviation must beat.

Worked example 2 – escalate to the shortage that was predictable. Festival week: demand runs 180 kg a day (37% above the 131.5 mean). The bin hits the 890 kg reorder point and an order goes out – but in the 6 days it takes to arrive the line eats 180 × 6 = 1,080 kg: a 190 kg stockout ✓, about a day of silence. The 101 kg of safety stock was sized for noise with sd 25 a day, not for an event 49 kg a day above the mean for a week. The honest reading: safety stock is the wrong tool for anything on a calendar. The operational answer is a scheduled pre-buy (one extra 1,100 kg order the week before) or a temporary periodic-review target computed from festival demand – planning, not buffering.

Common mistakes. Reordering "when the bin looks low" (no lead time). Treating the EOQ as eternal – demand moves, the square root moves. And buying a discount without the total-cost check (the 0.625% breakeven above).

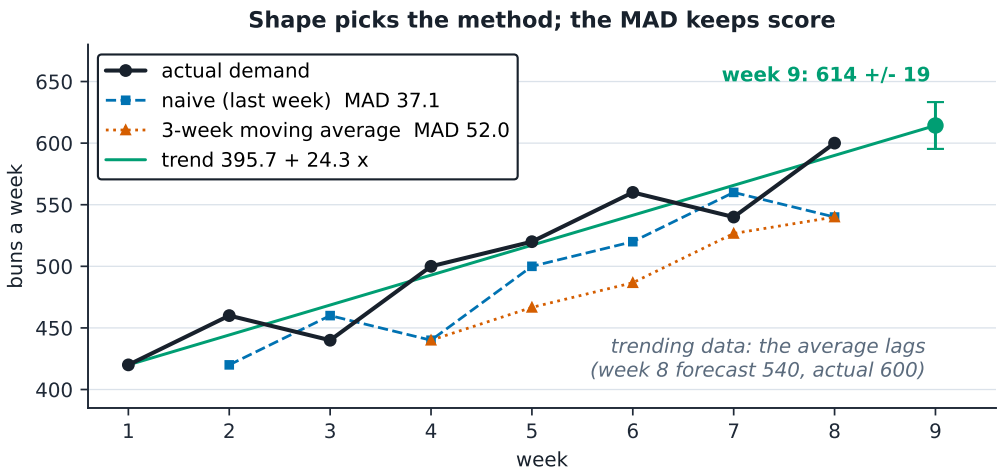
Summary. $EOQ = \sqrt{2DS/H}$, where ordering = holding. ROP = lead-time demand + $z\sigma_{LT}$. Safety stock is for noise; calendars need planning, not buffers.

Forecasting

Why it matters. Every number in this course – capacity, stock, staff – is a forecast wearing a uniform. The craft is picking the method that fits the shape, and then measuring how wrong it is.

The idea, with a picture. Eight weeks of bun demand: 420, 460, 440, 500, 520, 560, 540, 600 – a rising trend. Three methods, one scoreboard: the **mean absolute**

deviation (MAD).



On trending data the naive method wins - the moving average is built for flat demand

The definitions. **Naive forecast**: today's actual is tomorrow's forecast. **Moving average** : the mean of the last n weeks. **Exponential smoothing** : $F_{t+1} = F_t + \alpha e_t$ – a weighted past with fading memory. **MAD** = mean of $|e_t|$; divide demand by MAD to size buffers. **Tracking signal** = cumulative error ÷ MAD – drift past ±4 means the method has lost the process.

The derivation. Naive errors (weeks 2–8): 40, −20, 60, 20, 40, −20, 60 → MAD 260/7 = 37.1 ✓. Three-week **moving average** (from week 4): 440, 466.7, 486.7, 526.7, 540 against 500, 520, 560, 540, 600 → MAD 260/5 = 52.0 ✓. Exponential smoothing at $\alpha = 0.4$ (seed 420): forecasts 420, 436, 437.6, 462.6, 485.5, 515.3, 525.2 → MAD 48.3 ✓. The least-squares trend: $\hat{y} = 395.7 + 24.3x$; week 9 = **614** ✓ with residual sd 19.3 – so quote 614±19 as the two-thirds band, and buy flour for the band, not the point.

> **Worked example 1 – the method must match the shape.** The naive MAD (37.1) beats the moving average (52.0) here – not because naive is clever but because demand is TRENDING and the average drags it down (week 8: 540 forecast against 600 actual). The same moving average on flat demand would win. Method choice is not taste: fit the shape (flat → average; trend → trend model; new product → judgement), then let the MAD scoreboard arbitrate – and re-run it monthly.

> **Worked example 2 – escalate to the transparent adjustment.** Week 9 base forecast: 614. The sales lead hears the festival added a second stage and wants "add 20%". The honest protocol: base 614 × 1.10 = 675, RECORDED as "judgemental uplift +10%, second stage, sales lead, 2026-05-04" – and at week 9 the actual 700 lands 25 above the adjusted number (3.7%), while 614 alone would have been 86 low. The lesson is not that judgement won; it is that the adjustment is written down, so next year's 614 base gets a better prior. A forecast you cannot audit is a rumour with decimals.

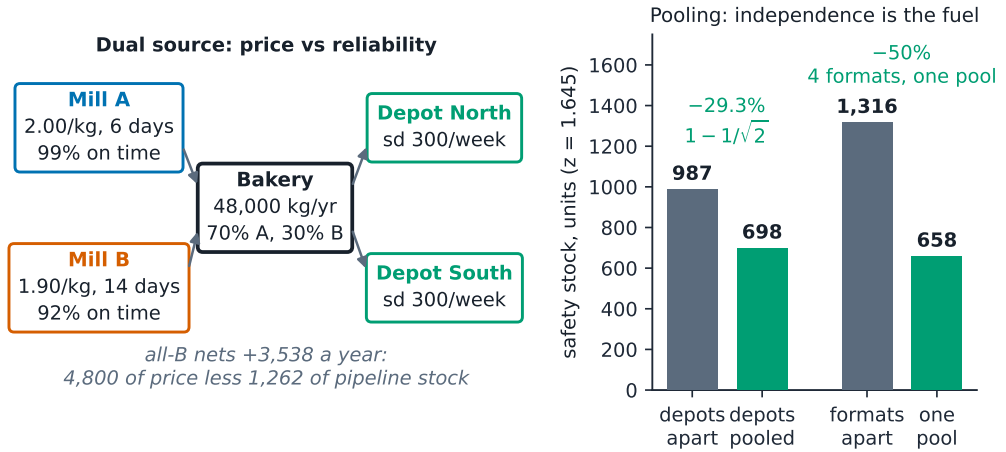
Common mistakes. Scoring accuracy on the wrong weeks (the moving average has no forecast for weeks 1–3 – score like with like). Letting α rise to chase noise (0.4 follows trend; 0.8 tracks weather). And hiding the judgemental uplift inside the number instead of beside it.

Summary. Shape picks the method; MAD scores it; trend data punishes averages. Forecast a band, not a point – and write every judgement beside the number it moved.

Supply chains

Why it matters. The site does not make buns; it makes buns out of a chain – mills, a bakery, two depots. Every chain choice trades cost against service against risk, and the trade is arithmetic, not slogan.

The idea, with a picture. The design question: one supplier or two, local or distant, stock at each depot or pooled centrally?



Pooling safety stock cuts it by one minus one over root two - independence is the fuel

The definitions. **Lead time** : order to arrival. **Risk pooling** : aggregate demand is less variable than its parts – $\sigma_{sum} = \sqrt{\sum \sigma_i^2}$ for independent demands (the **square-root law**). **Postponement** : hold the generic form, differentiate late. **Single sourcing**: cheaper, riskier; **dual sourcing**:

the reverse.

The derivation. Each depot faces weekly demand sd 300 units (independent), one-week lead time, 95% service: each holds $1.645 \times 300 = 494$; both, 987. Pooled at one central store: $\sigma = \sqrt{300^2 + 300^2} = 424.3$, safety stock $= 1.645 \times 424.3 = 698$ – saving $987 - 698 = 289$, exactly $1 - 1/\sqrt{2} = 29.3\%$ ✓. The same law prices postponement: four pack formats each needing 1,316 units of finished safety stock ($4 \times 1.645 \times 200$) can be held as one generic pool of $1.645 \times 400 = 658$ – half ✓.

> **Worked example 1 – the cheap mill that wasn't cheap.** Mill B bids £1.90 against A's £2.00 – saving $0.10 \times 48{,}000 =$ £4,800 a year. But B is a 14-day lead against A's 6, so 8 extra days of pipeline stock, $8 \times 131.5 = 1{,}052$ kg at £1.20 to hold = £1,262 – the net is still +£3,538 ✓. The arithmetic says yes; the RISK column says read on: B delivers on time 92% of the time against A's 99%, single-sourced, harvest-exposed. One missed 14-day delivery is a stopped line (the site holds 1,052 kg – 8 days; a second miss is closure). The decision: dual-source, 70% A for reliability, 30% B for price and leverage. The number never decides alone; it prices the menu.

> **Worked example 2 – escalate to the postponed bun.** Four pack formats (plain 40%, seeded 30%, branded festival 20%, gluten-free 10%) each carry their own forecast error. Label-and-pack at the depot instead of the bakery: the bakery ships generic buns (one forecast, sd 400) and differentiation happens after demand is visible – safety stock halves from 1,316 to 658 ✓, and the gluten-free format stops being baked to a two-week-old guess. The cost of postponement is real – a depot packing line at £18,000 – so it is checked like every other number in this course: write-offs fall with the safety stock, about 600 buns a week at £0.50 = £300 a week, £15,600 a year, so the line pays back in 14 months ✓. Postponement is the pooling law in time instead of space.

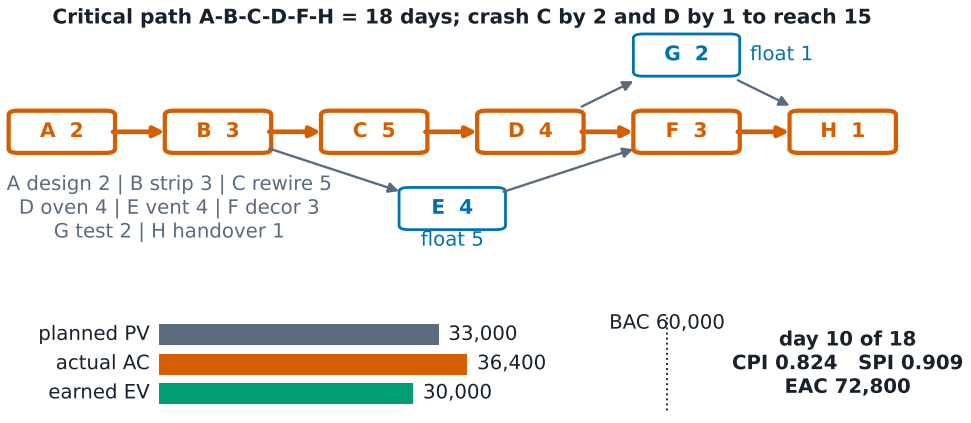
Common mistakes. Adding supplier standard deviations (variances add; the square root does the work). Sourcing on price alone after the pipeline-stock check says otherwise. And calling a chain "resilient" with no number – resilience is measured in days of coverage and per cent on-time.

Summary. $\sigma_{sum} = \sqrt{\sum \sigma_i^2}$: pooling and postponement both live on independence. Lead time is inventory in disguise. Cost, service, risk – price all three, then choose.

Project management

Why it matters. The second site was built in 18 days of planned work – and the plan, the crash and the cost tracking are three disciplines that keep a project from being a story told afterwards.

The idea, with a picture. The fit-out as activities: A design sign-off 2 d, B strip-out 3, C electrical rewire 5, D oven install 4, E ventilation 4, F decor 3, G equipment test 2, H handover 1. The **critical path** A-B-C-D-F-H totals **18 days** ✓; E floats 5 days, G floats 1.



The critical path is the project - a day lost on it is a day lost everywhere

The definitions. **Critical path** : the longest path; float zero on every activity. **Total float** : how late an activity can start without delaying the project. **Crashing** : buying time on the critical path at a cost per day. **Earned value analysis** : $CPI = EV/AC$, $SPI = EV/PV$ – cost and schedule health in two ratios.

The derivation. Forward pass: ES of H = 18; backward pass gives E float $14 - 9 = 5$ and G float $17 - 16 = 1$, everything on A-B-C-D-F-H zero ✓. Crashing: C can compress 2 days at £400/day (overtime), D one day at £600/day (crew); a day saved is worth £700 (the festival deadline). Crash C by 2: 16 days, cost 800 against 1,400 of value; the near-critical path A-B-C-D-G-H is now 15, so crashing D one more day pulls both to 15 – cost 600 against 700 ✓. Total: 18 → 15 days for £1,400 of crash against £2,100 of value, net +£700 ✓, and E's float absorbs the noise.

> **Worked example 1 – the float is not yours.** The ventilation contractor asks to start E five days late to fit another job – "it has 5 days of float, nothing moves". Arithmetically true today; the error is ownership. Float belongs to the PROJECT, not the trade: spend it on E and a later surprise elsewhere has nowhere to go. When the decor crew (F) is then held two days by late paint, F's slip lands straight on the end date, where before it would have been absorbed by moving F after an on-time E. The discipline: float is a shock absorber managed by the project manager, released to a trade only in writing, with the end date recomputed the same day.

Worked example 2 – escalate to the earned-value audit. Budget at completion £60,000 (the raise of course 9). Day 10 of 18: planned value 33,000, actual cost 36,400, earned value 30,000 (half the scope done). $CPI = 30{,}000/36{,}400 = 0.824$ ✓ and $SPI = 30{,}000/33{,}000 = 0.909$ ✓: 82p of value earned per pound spent, 9% behind schedule. $EAC = BAC/CPI = 60{,}000/0.824 =$ £72,800 ✓ – a £12,800 overrun flagged on day 10, while it is still a management problem, not a post-mortem. The naive reading ("we spent 36,400 of 60,000, 61% through the money") hides that only 50% of the VALUE is earned. The two-ratio discipline says what is wrong, where, and early.

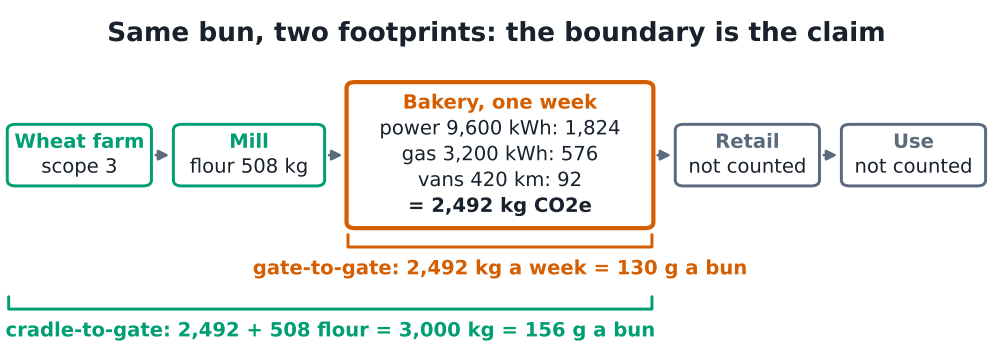
Common mistakes. Crashing the cheapest activity instead of the critical one (money spent off the path buys nothing). Letting trades consume float. And tracking cost without EV – spend is not progress.

Summary. The critical path is the project; float is a shared shock absorber. Crash where the path runs, at a price below the day's value. CPI and SPI flag overruns while they are still decisions.

Sustainability in operations

Why it matters. The festival's sponsor wants "a carbon-neutral bun". Whether the claim survives arithmetic depends entirely on where the boundary is drawn – and unstated boundaries are how greenwash is manufactured.

The idea, with a picture. The site's own operation (gate-to-gate) runs at 2,492 kg CO2e a week ✓: electricity 9,600 kWh × 0.19 = 1,824, gas 3,200 × 0.18 = 576, deliveries 420 km × 0.22 = 92. Draw the boundary one step wider and flour adds 508 kg – 20% more, from one input.



Same bun, two footprints - the boundary is the claim

The definitions. Carbon footprint : CO2e emitted within a stated **system boundary** . Gate-to-gate: the site's own operation. Cradle-to-gate: plus everything upstream. **Scope 1** (direct fuel), **Scope 2** (purchased electricity), **Scope 3** (everything else in the chain). **Circularity**: keeping materials in use – reuse beats recycle beats dispose.

The derivation. Weekly energy: 9,600 × 0.19 = 1,824; 3,200 × 0.18 = 576; 420 × 0.22 = 92.4 – total 2,492 kg ✓, which is 2,492/19,200 = 130 g per bun ✓. Upstream flour at 0.55 kg CO2e per kg and 923 kg a week adds 508 kg, so cradle-to-gate is 2,492 + 508 = 3,000 kg – and the boundary, not the bun, decided the 20% jump. On a renewable tariff at 0.06 kg/kWh, electricity falls 1,824 → 576 and the site's own footprint halves (2,492 → 1,244, a 50% fall; 42% cradle-to-gate) with no operational change at all – the same fact that makes the claim honest makes it cheap to game, which is why the tariff must be additional, not a certificate shuffle.

Worked example 1 – the tray that pays only if it comes back. Cardboard trays: 42,000 a year at £0.22 = £9,240, six tonnes of board, 4.2 t CO2e upstream. The circular alternative: returnable trays at £5.50, a two-week cycle, so a pool of 808 × 2 = 1,616 trays (£8,888). The whole decision is the return rate. At 98% returns the site loses 840 trays a year (£4,620) and washes 42,000 (£1,470): running cost £6,090, saving £3,150 a year, payback 2.8 years ✓. At 95% returns, losses alone cost £11,550 and the scheme LOSES £3,780 a year ✓. Breakeven return rate: 96.6% ✓. Emissions at 98%, amortised honestly: the pool spread over a five-year life (323 a year) plus 840 losses = 1,163 trays a year × 0.6 kg × 3.0 kg CO2e per kg = 2,093 kg, plus 300 kg of wash energy = 2.4 t against 4.2 t – a 43% cut ✓, but only with the pool's manufacture, the losses and the wash INSIDE the boundary. Circularity claims with no loss rate are decoration.

Worked example 2 – escalate to the neutral bun. The sponsor's claim, audited: cradle-to-gate 3,000 kg for the site's 19,200 buns a week = 156 g a bun ✓; add the retail stage and the customer's drive and it only climbs – no boundary makes the bun neutral. What survives scrutiny is a REDUCED claim with the boundary stated: "156 g a bun cradle-to-gate, week-41 audit, renewable tariff and tray pool in place, boundary excludes retail and use". The operations answer to greenwash is always the same shape: a number, a boundary, a date, and something the competition can verify. "Neutral" without those four is a word, not a fact.

Common mistakes. Comparing footprints with different boundaries (the 2,492 vs 3,000 trick is the standard one). Offsetting Scope 2 while Scope 3 dwarfs it. And treating circularity as free – every reuse loop has a wash, a reverse logistics leg, and a loss rate, all inside the boundary.

Summary. The boundary IS the claim: state it, then measure inside it. Scope 2 moves with the tariff; Scope 3 needs the chain. Circularity pays when the loop's costs stay inside it.

The project sheet (reproducible)

Second-site bake line, festival contract:

- Line: mix 12 min / proof 24 (×3) / bake 18 (×2) / pack 6 → capacities 5.00 / 7.50 / 6.67 / 10.00 batches an hour; bottleneck mix, cycle time 12 min. At 4 batches an hour: utilisations 80 / 53.3 / 60 / 40%. Little: WIP $4 = 4 \times 1$ hr. 480 buns an hour; 19,200 a week.
- Quality: X-bar $80 \pm 2.68 \to [77.32, 82.68]$; $\bar{R} = 4.65$, R-chart UCL 9.83; sample 6 at 82.9 = special cause; Pareto 46/21/ 9/6/4 (77.9% in two); COQ 19,000 = 19% of sales, 81% failure.
- Inventory: EOQ $\sqrt{2 \times 48{,}000 \times 50/1.20} = 2{,}000$ kg; 24 orders, 15.2-day cycle; £1,200 = £1,200. ROP = 789 + 101 = 890 kg ($z = 1.645$, $\sigma_{LT} = 61.2$). Discount check: 2% at 4,000 kg saves 1,320; breakeven 0.625%.
- Forecast: 420→600 over 8 weeks; MAD naive 37.1 / MA3 52.0 / ES(0.4) 48.3; trend $\hat{y} = 395.7 + 24.3x$; week $9 = 614 \pm 19$.
- Chain: mill B nets +3,538 but 92% on-time single-sourced → dual 70/30. Pooling: 987 → 698 ($29.3\% = 1 - 1/\sqrt{2}$); postponement 1,316 → 658 (half).
- Project: A-B 2-3, C-D 5-4, E 4, F 3, G 2, H 1 → critical path A-B-C-D-F-H = 18 d; floats E 5, G 1; crash C −2 and D −1 → 15 d, +£700 net. EVM day 10: PV 33,000, AC 36,400, EV 30,000 → CPI 0.824, SPI 0.909, EAC 72,800.
- Footprint: gate-to-gate 2,492 kg/wk (1,824 + 576 + 92); 130 g/bun; cradle-to-gate 3,000 kg (flour +508); renewable tariff: site 2,492 → 1,244 (−50%), cradle-to-gate −42%.

Every figure in this handout recomputes from these values.

Topic 11

University Business

What this course is, and how to use it

This is the eleventh course in University Business: **business law and ethics** – the rules the game is actually played under, and the ones it should be played under. Every rule in this course is labelled with its **jurisdiction**: the working law is **England and Wales**, with named comparisons where the contrast teaches. No rule is presented as universal, because none is.

The running case is **the Nimbus Kitchen legal file**: one allegation of a bad bun, one flour contract that arrived twenty days late, one crew member whose status nobody wrote down, one name the founder wants to protect, one sponsor's claim the founder should not make, and one public collapse studied at arm's length. Every number in the file is checked against the statutes and public record.

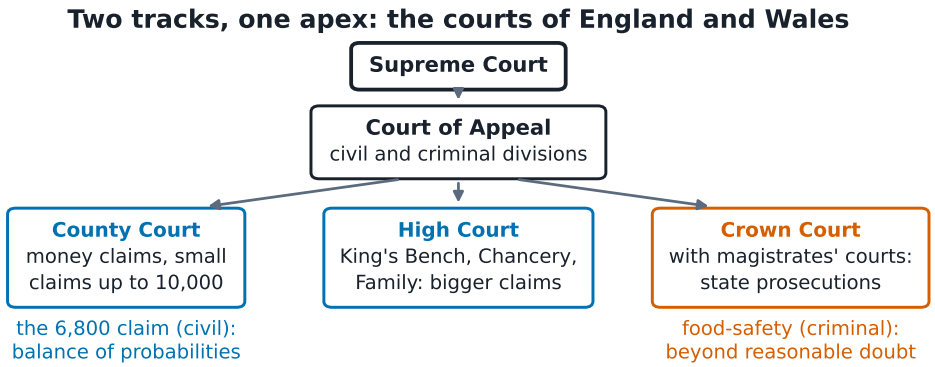
- Units 11.1–11.2 build the frame: which court, which standard of proof, and what a contract is before it breaks.
- Units 11.3–11.5 run the business through the law: forms and liability, employment, intellectual property and data.
- Units 11.6–11.7 turn from law to ethics: how to reason to a decision you can defend, and what governance failures look like from the outside.

How to read each unit: the situation first, then the rules, then two escalating worked analyses, the mistakes people actually make, and a three-line summary. This is a study course, not legal advice.

Legal systems and jurisdiction

Why it matters. A customer alleges the festival buns made her ill. Before anyone argues about facts, three questions decide everything: is this civil or criminal, which court hears it, and who must prove what to what standard. Getting these wrong loses the case before it starts.

The idea, with a picture. Civil liability is the claimant's claim to be made whole (the customer's injury claim); **criminal liability** is the state's claim that a rule of public law was broken (a food-safety prosecution by the authority). Same facts, different parties, different **standards of proof**: the civil **balance of probabilities** (more likely than not, just past 50%); the criminal **beyond reasonable doubt** (no numeric standard, deliberately).



Citations: [2015] UKSC 67 = year, court, judgment number; [1893] 1 QB 256 = volume, report series, page

Civil and criminal ride different tracks to the same apex - and money claims of 10,000 pounds or less usually stop at the small claims track

The definitions. **Jurisdiction** : the power of a court to hear a matter. **Precedent** : like cases decided alike – the Supreme Court binds every court below (and itself since 1966 only exceptionally); the Court of Appeal binds itself subject to narrow exceptions. A **case citation** `[2015] UKSC 67` reads: year 2015, court UKSC (Supreme Court), 67th judgment of that year; older cases cite a law report: `[1893] 1 QB 256` is volume 1 of the Queen's Bench reports for 1893, page 256.

The derivation. The customer's claim: a private wrong (negligence or breach of statutory duty), so civil, so she sues in the County Court; alleged loss £6,800 (two days off work £340, medical £160, injury £6,300 – her schedule, uncontested arithmetic for this exercise). At £6,800 the claim falls under the £10,000 small-claims threshold: small claims track, no cost recovery to speak of, her own time. The regulator's parallel power: if inspection finds the fridge out of temperature range, that is a criminal matter under food-safety law, prosecuted by the authority, proved beyond reasonable doubt, and her civil claim does not depend on its outcome – different parties, different standards, different questions.

> **Worked example 1 – read the citation before the case.** The founder is sent "Carlill v Carbolic Smoke Ball Co [1893] 1 QB 256" by a well-meaning friend. Reading it: 1893, volume 1 of the Queen's Bench law reports, page 256. "QB" names the report series, not the court: that series prints both High Court and Court of Appeal decisions, and Carlill was decided by the Court of Appeal, so its reasoning on unilateral **offers** binds the High Court, the county courts and the Court of Appeal itself. The lesson: the citation tells you the era and where to find the case, and the report's heading tells you the court and so the weight – and a case from 1893 still binding on offer law is the system working, not failing. Look up, never guess: the wrong level of court wastes a day.

> **Worked example 2 – escalate to the parallel proceedings.** Six weeks in, the authority prosecutes (temperature log falsified by a supervisor – criminal, strict-liability flavour, unlimited fine headroom), the customer's civil claim continues, and the festival sponsor threatens to sue for its fee under the contract. Three proceedings, one set of facts, three different proofs. The operational rule the founder writes on the office wall: never let one proceeding's defence contradict another's admissions – the civil statement "our fridge failed" is prosecution evidence. Answers to any proceeding get written once, with all three in view.

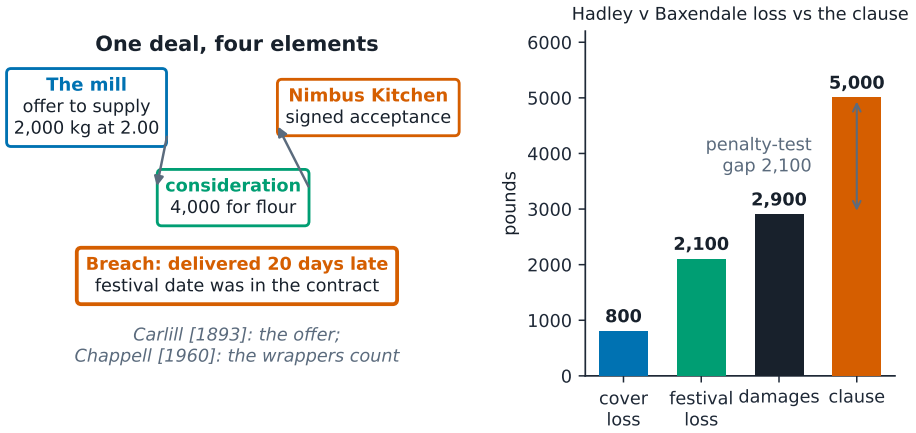
Common mistakes. Treating "not guilty" as "not liable" (the standards differ on purpose). Suing in the wrong court because it felt right. And reading a Scottish, US or EU ruling as if it were English law – every rule in this course is jurisdiction-labelled, and so is every one you will ever rely on.

Summary. Civil makes the claimant whole on the balance of probabilities; criminal punishes beyond reasonable doubt. Money claims of £10,000 or less usually take the small claims track. Read the citation, check the court, label the jurisdiction.

Contract formation and breach

Why it matters. Nimbus buys 2,000 kg of flour a fortnight at £2.00 a kilogram. When the mill delivers twenty days late, whether Nimbus recovers £2,900 or nothing depends on four words taught in every first-year law course: offer, acceptance, consideration, intention.

The idea, with a picture. The flour deal: the mill's order confirmation (offer), Nimbus's signed order (acceptance), £4,000 for 2,000 kg (consideration – the price of the promise), two businesses (intention presumed commercially). Formed. Then broken: delivered 20 days late against a contract that named the festival date.



Expectation damages put the buyer where performance would have - the clause must survive the penalty test on top

The definitions. Offer : an intention to be bound on specified terms, distinguished from an invitation to treat (the shop shelf). **Acceptance** : unqualified assent, communicated. **Consideration** : something of value moving each way – adequacy is not judged (the peppercorn), existence is (Chappell v Nestlé [1960] AC 87: the three wrappers were good consideration). **Intention**: presumed commercially, rebutted between family. **Breach of contract** : failure to perform; **expectation damages** put the innocent party where performance would have (Hadley v Baxendale (1854) 9 Ex 341: losses arising naturally, or within the parties' contemplation). A **liquidated damages** clause is a genuine pre-estimate; a **penalty clause** is unenforceable – the test is legitimate interest proportionate to loss, not deterrence for its own sake (Cavendish v Makdessi [2015] UKSC 67).

The derivation. The breach arithmetic: cover purchase 2,000 kg at £2.40 = £4,800 against the contract £4,000 → £800 price-difference loss ✓. The festival weekend the flour missed, known to and contemplated by both parties at signing: £2,100 profit lost ✓. Expectation damages = 800 + 2,100 = **£2,900** ✓. The contract's late-delivery clause, £250 a day × 20 days = **£5,000** ✓ – £2,100 more than the loss, so the penalty question is live, not automatic.

> **Worked example 1 – the four elements on one deal.** The mill's email quotes "2,000 kg, £2.00, delivery 30 May" – an offer (not a display or an advert). Nimbus replies "accepted" – unqualified assent, communicated. £4,000 moves for flour: consideration both ways. Both are companies: intention presumed. Four boxes, four ticks, one enforceable contract – and the founder's checklist for every order is exactly those four boxes, initialed on the confirmation.

> **Worked example 2 – escalate to the clause that failed its test.** Nimbus claims the clause: £5,000. The mill says: penalty, your loss is £2,900. Applying Cavendish: is £250 a day a legitimate protection of Nimbus's interest in timely delivery (perishable stock, festival dates), or extravagant deterrence? The gap (5,000 − 2,900 = £2,100, 72% above loss ✓) is not decisive; what decides the essay is evidence that the daily figure tracked a genuine daily cost of lateness. The drafting lesson costs less than the litigation: a clause that states the daily loss it protects (cold storage, lost slots) with a reasonable sum survives; a round number plucked for pain does not. There is no middle path: if the clause fails the test, the court does not cut it down to the genuine part; it strikes the clause out and Nimbus recovers only its proven loss, £2,900.

Common mistakes. Treating an advertisement or shelf price as an offer (Carlill is the exception that proves the carve-out for unilateral offers). Judging consideration by adequacy rather than existence. Claiming the clause and skipping the loss proof – the penalty test is argued on both legs.

Summary. Offer + acceptance + consideration + intention = enforceable. Damages: the loss that flowed within contemplation – here £2,900. Liquidated must protect, not punish: £250 a day lives or dies on its evidence.

Business forms and liability

Why it matters. The raise succeeded (course 9), the site is growing, and the founder must choose the legal skin everything lives in: her house on the line, or not; her control diluted, or not. The form decides who bears the bad-bun claim and who owes what to whom.

The idea, with a picture. Four forms: sole trader (unlimited liability – the bad-bun claim is the founder's house), ordinary partnership (joint and several, unlimited), LLP (limited, but profit-taxed like a partnership), Ltd (limited liability – Salomon v Salomon [1897] AC 22: the company is a separate person).

The form is a wall with a toll

	sole trader	partnership	LLP	Ltd
liability	**unlimited**	**joint, unlimited**	limited	limited
	the bad-bun claim →			Salomon [1897]
who pays tax	the trader	the partners	the members	the company
main filings	self-assessment	partnership return	accounts at CH	accounts + CS
capital rules	none	none	none	PLC 50,000

Ltd duties: CA 2006 s.172 promote success, s.174 care and skill; audit exemption while turnover stays at or under 10.2m

The form is a wall: unlimited for the trader and partners, limited for LLP and Ltd - and the wall costs filings, duties and, eventually, an audit

The definitions. Limited liability: shareholders lose only what they invested. **Incorporation**: a company exists when Companies House registers it – certificate, number, constitution; a PLC needs £50,000 allotted share capital ✓, a private Ltd needs none. Small companies escape audit while turnover stays ≤ £10.2m ✓. **Director duties** (Companies Act 2006): s.172 act to promote the success of the company; s.174 care, skill and diligence. **Agency** : a principal is bound by an agent acting with authority – **apparent authority** binds when the principal's holding-out made the deal look authorised (Freeman & Lockyer v Buckhurst Park [1964] 2 QB 480).

The derivation. Nimbus Kitchen Ltd: 80,000 shares at £2.50 = £200,000 (the course-9 equity), one director-founder, small-company audit exemption (turnover £100,000 « £10.2m ✓, so no audit – but annual accounts and a confirmation statement still filed). The bad-bun claim, if it comes, is against the company –

£6,800 of schedule loss against a company holding the raise, not the founder's house. The price of the wall: public filings, the duties in ss.172–174, and the no-conflict, no-profit rules (Regal (Hastings) [1967] 2 AC 134) that catch the founder's side-deal on packaging in unit 11.7's spirit.

> **Worked example 1 – the signature that bound the company.** The festival sponsor's manager recalls a "deal" with the Saturday supervisor. Agency analysis: did the supervisor have actual authority (companies-house-checked list of directors: no)? Apparent authority – did the company hold her out as able to sign (title "Head of Operations", email signature, past conduct)? If yes, the company is bound; if no, the sponsor's claim fails. The operational fix the case teaches: a one-page authority schedule (who may sign what, up to how much) published internally and honoured – apparent authority is manufactured by what the business lets people look like.

> **Worked example 2 – escalate to the duty at the raise.** Course 9's second site looked marginal at WACC 9.19% (NPV still positive). Now suppose a director had quietly taken a 5% interest in the fit-out contractor. Analysis: s.175 (avoid conflicts of interest) and s.177 (declare an interest in a proposed transaction) breached, s.172's duty to promote the company's success compromised, and the gain is held on trust for the company whatever the site's NPV – the duty is not a numbers test, it is a loyalty test. The comparison that teaches: a well-run Ltd does the same deal transparently – declares the interest, excludes itself from the vote, takes board approval – and keeps both the profit and the office. Disclosure converts a breach into governance.

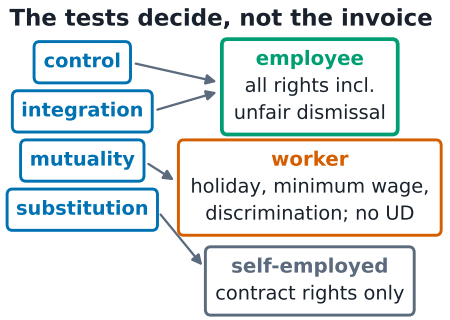
Common mistakes. Believing "Ltd" ends personal exposure (personal guarantees, wrongful trading, and s.214 IA 1986 all punch through). Trading as partners "informally" (two people carrying on business for profit ARE a partnership). And appointing directors without the duties register, training and insurance the duties imply.

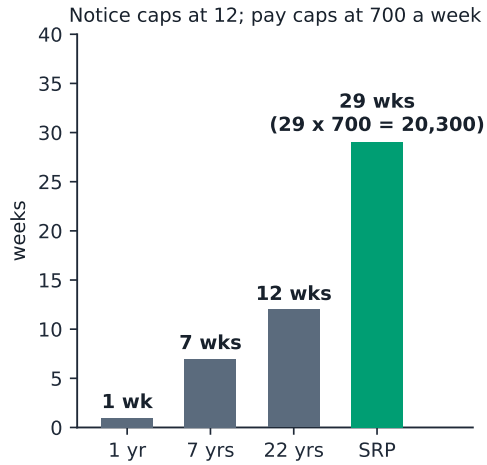
Summary. The form is a wall with a toll: Ltd/LLP stop the claim, at the price of filings and the Companies Act duties. Authority is what you appear to give, not what you meant. Declare interests; the no-profit rule is not a suggestion.

Employment

Why it matters. The Saturday crew has worked for three years on a handshake, one of them just asked about "redundancy if the festival ends", and nobody can find a contract. Whether she is an employee, a worker or self-employed decides what she is owed – and it is decided by law, not by what her invoices call her.

The idea, with a picture. Three statuses: **employee** (all rights: unfair dismissal, notice, redundancy pay), worker (holiday, minimum wage, discrimination – no unfair dismissal), genuinely self-employed **contractor** (contract rights only). The tests are about control, integration, mutuality and substitution – in that order of stubbornness.





The label on the invoice does not decide the status - the tests do, and the entitlements follow the status

The definitions. Employee/contractor tests: control (Ready Mixed Concrete [1968] 2 QB 497), integration, mutuality of obligations, actual obligations over written pretence (Autoclenz v Belcher [2011] UKSC 41; Pimlico Plumbers v Smith [2018] UKSC 29). **Unfair dismissal**: needs 2 years' service ✓, a fair reason and a fair process (ERA 1996). Notice: 1 week under 2 years, then a week a year capped at 12 weeks ✓. **Statutory redundancy pay**: a week's pay capped at £700 ✓, 0.5/1.0/1.5 weeks by age band (under 22 / 22–40 / 41+), 20 years counted. The Equality Act 2010 prohibits discrimination on nine protected characteristics; UK GDPR duties apply to the loyalty app's data.

The derivation. The Saturday baker: 22 years' continuous service part-time ending at age 59, weekly pay £800. Redundancy pay: only the last 20 years count, so years at 41+ = 18 × 1.5 + years 22–40 = 2 × 1.0 = 29 weeks ✓, at the capped £700 = **£20,300** ✓. Notice at 22 years: the 12-week cap ✓. (The basic award for unfair dismissal uses the same age-banded formula – also 20,300 here – which is why the arithmetic is worth drilling once.)

> **Worked example 1 – the handshake status.** Facts: fixed shifts she cannot refuse, wears the brand apron, no substitute ever sent, paid weekly through payroll for twenty-two years (she is the Saturday baker above). Analysis: control high, integration high, mutuality high, substitution nil – employee, on any current authority, whatever her early "freelance" invoices said. Her entitlements on the festival ending: consultation, notice 12 weeks (paid), redundancy £20,300, and no unfair-dismissal issue because redundancy is a fair reason run fairly. The cost of the handshake was never zero; it was a liability accruing at 1.5 weeks a year.

Worked example 2 – escalate to the app's data. The loyalty app holds 14,000 festival customers' emails, orders and one allergy field. The marketer wants to sell the list. UK GDPR: Nimbus is the **data controller** ; the allergy field is health data (special category); selling without a lawful basis is a breach. Exposure: the ICO's upper tier – up to £17.5m or 4% of worldwide turnover, whichever is higher ✓ (standard tier £8.7m / 2% ✓). On £100,000 turnover that reads absurd; on a growing group it reads existential – and the ICO publishes its fines, so it is brand risk too. The lawful alternative costs one afternoon: consent, purpose limitation, and a processor contract with the app vendor.

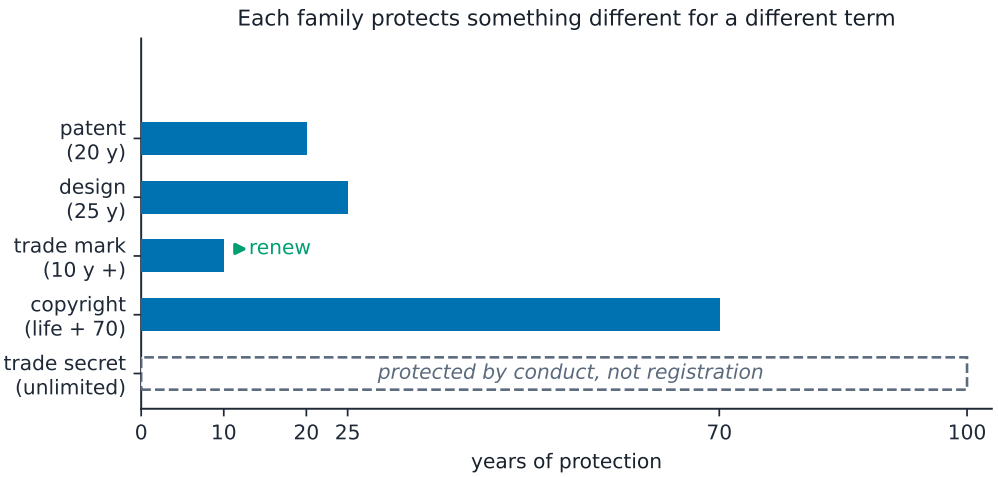
Common mistakes. Choosing the label before the facts (the tests decide, and sham terms are read against the draftsman). Forgetting the 2-year clock starts on day one, not at dismissal. Treating "we're small" as a data-law exemption – there isn't one.

Summary. Status by tests: control, integration, mutuality, substitution. Redundancy: capped £700, age-banded, 20-year limit – £20,300 in the worked case. Data duties scale with the fine tiers, not the company's size.

Intellectual property and data

Why it matters. The bun has a name ("Nimbus Bloomer"), the packaging has a design, the recipe is in a folder marked secret, and a rival has started selling a "Nimbus Blummer". Which of the four IP families protects what – and what it costs to find out – is a map, not a mood.

The idea, with a picture. The families: **patent** (inventions, 20 years ✓, new and inventive – the proving cabinet mechanism, if it were novel), **trade mark** (signs distinguishing goods – the name and logo, 10 years renewable for ever ✓), **copyright** (original works – the packaging artwork and the website, life + 70 years ✓, automatic), registered design (25 years ✓), and **trade secret** (the recipe – protected as long as it stays secret and protected).



Each family protects something different for a different term - the recipe's protection is conduct, not registration

The definitions. Trade mark: registerable signs with distinctiveness; infringement by identical/similar marks on identical/similar goods with a likelihood of confusion. **Copyright**: subsists automatically in original artistic, literary and musical works. **Patent**: full disclosure for a 20-year monopoly – the deal is publication, which is exactly why the recipe stays a trade secret. **Trade secret / confidentiality**: protection by reasonable endeavours (NDA, access control, exit process); once disclosed, gone. **Licensing**: permitting use while owning – the sponsor's festival use of the logo is a licence with terms (territory, term, quality control), not a transfer.

The derivation. The rival's "Blummer": identical mark on identical goods – if Nimbus's "Bloomer" is registered, this is the textbook s.10(1) case: infringement without even a confusion argument. The arithmetic of protection: UK trade mark registration ≈ £170 for one class online (the official fee – checked against the registry fee schedule at authoring) against a passing-off claim's five-figure legal bill with an uncapped outcome; the registration is the cheapest insurance the business will ever buy. The recipe's alternative: patent = 20 years of monopoly but published; secret = indefinite but one leak from worthless – the choice is strategic, and the folder marked "secret" that everyone can open is neither.

Worked example 1 – the four families on one product. The Bloomer: name and logo (trade mark, register now), packaging artwork (copyright, automatic, keep dated originals), proving cabinet mechanism (patent only if novel and non-obvious – it is not: standard racking), recipe (trade secret: NDA, two named holders, exit interviews). One product, four regimes, four different actions – and the mistake that costs the case is treating "IP" as one thing on a form.

Worked example 2 – escalate to the licence the sponsor wrote. The sponsor's contract "grants all rights in the Bloomer name for festival use". Analysis: that is an assignment in the draftsmanship of someone who did not mean it – but the words bind. The fix negotiated in the case: a licence, exclusive only as to festival sites, two seasons, quality control retained, reversion on breach, and the mark never registered by anyone but Nimbus. The general lesson: never let a counterparty register your mark (the registry checks ownership, not fairness), and never sign a rights clause you have not read upside down.

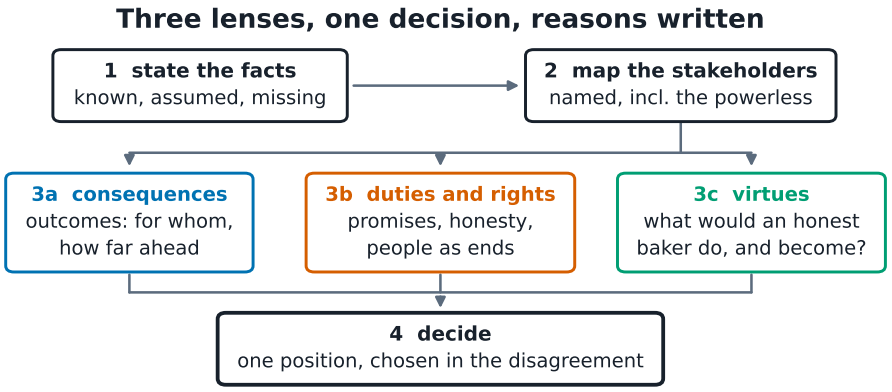
Common mistakes. Disclosing the invention before filing (novelty dies). Trusting "©" without dated originals. And the classic: the founder who patents the recipe, publishing it, to stop a rival who only ever needed an NDA.

Summary. Trade mark 10 + renewals; copyright life + 70; patent 20 for disclosure; secret = conduct. Registration beats litigation on price every time it is available.

Ethical reasoning in business

Why it matters. Course 10's sponsor wants a "carbon-neutral bun". The founder can probably get away with the claim; the question in this unit is whether it would be right – and that question has method, not just instinct.

The idea, with a picture. Three lenses, one decision: **consequentialism** (what outcomes does the claim cause – the measured 156 g of unit 10.7 against the buyer who believed zero), **deontology** (duties and rights – a promise of "neutral" you know to be unmeasured is a lie regardless of outcome), **virtue ethics** (what would an honest baker do – the character the business is building with every claim it makes).



Three lenses, one decision, reasons written - the tree ends in a defensible position, not a comfortable one

The definitions. Stakeholder : anyone the decision affects – mapped with power/interest and named, not averaged. **Consequentialism**: judge by outcomes (utilitarianism is the business default – and its trap). **Deontology**: some acts are wrong whatever the outcome (deception, breaking promises). **Virtue ethics**: judge the character the act builds. The decision protocol: facts → stakeholders → lenses → decision → written reasons → review date. **Greenwashing**: a claim without a number, boundary, date and verifiable method (unit 10.7's four-part standard) – an ethical defect before it is a legal one under consumer law.

The derivation. The neutral-bun claim through the tree: facts – cradle-to-gate 156 g measured, gate-to-gate 130, offsets unverified; stakeholders – sponsor (wants the word), customers (buy the feeling), staff (must answer questions), regulator (reads advertising law); lenses – consequentialist: brief sales bump, long trust damage when a journalist recomputes (they will; it is one division); deontological: the word says zero, the number says 156 – a false statement; virtue: the bakery's brand IS honesty at the counter. Decision: the reduced, numbered claim of unit 10.7 – "156 g a bun cradle-to-gate, measured, week-41 audit" – with the sponsor sold the stronger story: a number beats a word because a number can be defended.

Worked example 1 – the cheaper supplier's wage. Mill B (unit 10.5's 1.90) subcontracts harvesting to gang-labour rates below a living wage. Pure consequentialism says: 4,800 saved, no one sees. The tree: stakeholders include the harvest crew (affected, powerless to object – stakeholder theory's whole point is they still count); deontology: knowingly benefiting from a below-living wage treats people as means; virtue: the brand that audits its chain sleeps at night. The negotiated answer in the case: dual-source with a written wage clause in the contract (law gives the hook: supplier codes with audit rights), accepting +£400 a year of cost as the price of a claim Nimbus can make truthfully – an ethics decision with a number on it, which is the only kind that survives a board.

Worked example 2 – escalate to the deadline lie. A star engineer asks whether the refit will finish before the festival (course 10.6's day-10 SPI 0.909 says: probably not). The founder is tempted: "yes, on track". The tree compressed: consequences – the sponsor books around the lie; duties – truth-telling is not outcome-sensitive; virtue – organisations get the culture they rehearse. The honest version that still leads: "we are 9% behind schedule; here is the crash plan and the fallback" – deontology and management are not enemies; the lie is just lazier than the risk-adjusted truth.

Common mistakes. Averaging stakeholders into "the greater good" (the crew below the line is where ethics actually starts). Using a lens as a verdict-generator instead of a stress test – the lenses disagree, and the reasoning happens in the disagreement. And treating "legal" as the finish line (unit 11.1's standards are floors, not ceilings).

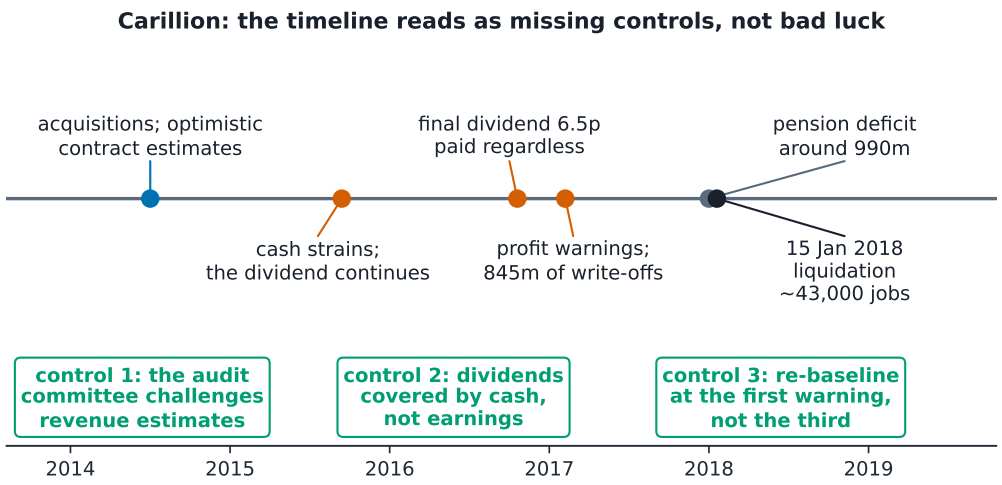
Summary. Facts, stakeholders, three lenses, decision, reasons, review. A claim you can defend has a number and a boundary. Ethics with a number on it survives the

board; ethics without one is a mood.

Governance cases

Why it matters. Governance failures are legible in public: dates, numbers, and the controls that were missing. One studied honestly teaches more than a shelf of codes – and the bakery's own controls are the same controls in miniature.

The idea, with a picture. Carillion plc: the construction and outsourcing giant that entered compulsory liquidation on 15 January 2018 with roughly 43,000 jobs worldwide (about 20,000 in the UK) ✓, a pension deficit reported around £990m ✓ – after years of aggressive accounting, dividends paid while the dividend was being earned only on paper, and a board that heard the warnings and kept the music playing.



Every failure on the timeline has a control that would have caught it - the controls, not the villain, are the lesson

The definitions. Corporate governance : the system by which companies are directed and controlled – board, risk, audit, remuneration. The trio this case makes concrete: board challenge (the non-execs' duty to ask), the audit relationship (challenge vs familiarity), and risk appetite said aloud (a written number, not a sentence). The UK Corporate Governance Code is comply-or-explain – which works only when the explaining is honest.

The derivation. The failure path, read as controls: (1) contract revenue recognised optimistically → control: revenue-recognition policy challenged by the audit committee (the write-offs that arrived in 2017 had lived in the estimates for years); (2) dividends paid while leverage rose → control: a distribution policy tied to cash cover, not earnings cover (the 2016 final dividend of 6.5p ✓ was paid into deteriorating liquidity); (3) warnings normalized (profit warnings in 2017 treated as the weather) → control: a board that re-baselines on the first warning, not the third. Each control was in the Code's letter; none was in the board's practice.

> **Worked example 1 – the audit committee's question.** The July 2017 write-off (£845m of contract charges across 2017 – public record ✓) did not fall from the sky; it reversed years of estimates. The committee question that bites: "show me the ten largest contracts' assumptions, and what last year's assumptions on those same contracts would say today". If the answer is uncomfortable, that is the answer. Nimbus's miniature: the festival contract's revenue is recognised when the buns ship, not when the sponsor grins – one sentence of policy, worth more than the audit fee.

> **Worked example 2 – escalate to the mini-Carillion checklist.** The founder writes three lines into the company's governance calendar: distributions only from cash cover (cash in bank after the dividend ≥ 3 months' payroll ✓ – at £8,000 monthly payroll, the dividend waits until £24,000 sits behind it); the first missed forecast triggers a written re-baseline, not a hopeful one; and every related-party deal (unit 11.3's packaging interest) goes to the minutes with the interest declared. None of this needs Carillion's scale – a company with £4,000 fortnightly flour orders and a £60,000 raise gets the same failure path in miniature, and the same three controls stop it.

Common mistakes. Studying the villain instead of the control (fraudsters are rare; missing controls are common). Treating the Code as a compliance form (comply-or-explain without honest explanation is comply-and-lie). And small businesses concluding governance is for giants – the failure path scales down; so does the cure.

Summary. Carillion: optimistic revenue, earnings-covered dividends, normalised warnings, £990m pension hole, 43,000 jobs, 15 January 2018. The controls: challenge estimates, cover dividends in cash, re-baseline at the first warning. Write the three lines; they scale down.

The project sheet (reproducible)

The Nimbus Kitchen legal file, one page:

- Jurisdiction: England and Wales throughout; civil balance of probabilities against criminal beyond reasonable doubt; small claims ≤ £10,000; limitation 6 years (12 by deed).
- Contract: offer + acceptance + consideration (the peppercorn counts, Chappell) + presumed intention. Breach arithmetic: cover £4,800 − contract £4,000 = £800 + £2,100 festival loss = £2,900 expectation damages; clause 20 × £250 = £5,000 survives only the Cavendish legitimate-interest test.
- Forms: Ltd wall (Salomon); PLC £50,000 capital; audit exemption ≤ £10.2m turnover; s.172 success, s.174 care; apparent authority is what you hold out (publish an authority schedule).

- Employment: tests – control, integration, mutuality, substitution (Autoclenz: real terms beat paper). Redundancy: £700-capped week, 0.5/1.0/1.5 by age, 20-year limit → 22 years ending at 59 = £20,300; notice caps at 12 weeks. Data: ICO tiers £8.7m/2% and £17.5m/4%.
- IP: trade mark 10 + renewals, copyright life + 70, patent 20, design 25, secret = conduct. Register before you litigate; license festival use, never assign.
- Ethics: facts → stakeholders → consequences/duties/virtues → decision with written reasons and a review date. A defensible claim has a number, a boundary, a date and a method.
- Governance: Carillion 15 Jan 2018 – 43,000 jobs, £990m pension deficit; controls = challenge revenue estimates, cash-cover dividends ($\geq$ 3 months' payroll behind: £24,000 at Nimbus scale), re-baseline at the first warning.

Every rule in this handout is England and Wales unless a named comparison says otherwise, and every number recomputes from these values or the cited public record. This is a study course, not legal advice.

Topic 12

University Business

What this course is, and how to use it

This is the twelfth course in University Business: **strategic management** – the discipline of choosing what the business will be, testing that the choice hangs together, and turning it into Monday morning. Eleven courses built the pieces; this course spends them.

The running case is **Nimbus Kitchen at the crossroads**: the festival contract is won, the second site runs (course 10), the legal file is tidy (course 11), and now the founder must choose between growing what works, adding products, or buying a café. Every framework is scored, every score is arithmetic, and every number recomputes from the project sheet.

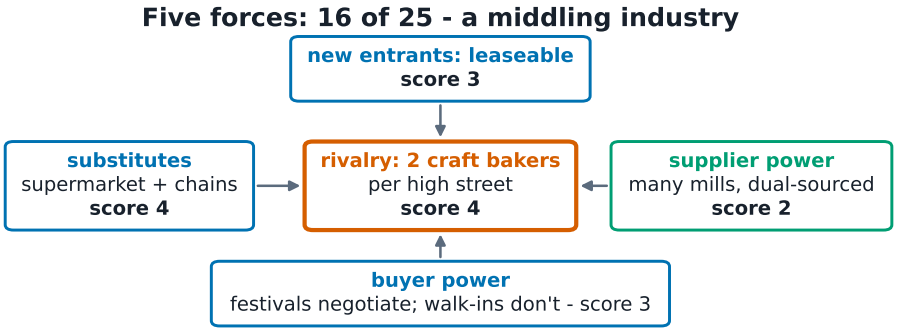
- Units 12.1–12.2 analyse: the industry's structure and the firm's right to win.
- Units 12.3–12.4 choose and implement: where to grow, and what has to be true for it to work.
- Units 12.5–12.6 decide under uncertainty and put the whole case on one page.

How to read each unit: the decision first, then the framework, then two escalating worked analyses, the mistakes people actually make, and a three-line summary.

Industry and internal analysis

Why it matters. The founder thinks the craft-bakery segment is "obviously attractive" because queues are long. Structure says otherwise half the time – and the difference between an opinion and an analysis is a score.

The idea, with a picture. Five **forces** , each scored 1–5 (higher = worse for us): rivalry 4 (two craft bakers per high street), new entrants 3 (a bakery is leaseable), substitutes 4 (supermarket in-store + coffee chains), supplier power 2 (many mills, course 10.5 dual-sourced), buyer power 3 (festivals negotiate, walk-ins don't). Total **16 of 25** ✓ – a structurally middling industry where the STRATEGY, not the industry, must make the money.



total 16 of 25: the profit must come from position; industry average 6%, top quartile 11%

A middling 16 of 25 - so the profit must come from position, not from the industry

The definitions. Five forces: rivalry, entrants, substitutes, supplier power, buyer power – industry profit is their residue. **VRIO** VRIO : is the resource Valuable, Rare, difficult to Imitate, and Organised for capture – the last question separates competitive parity from sustained advantage. **Value chain** : margin decomposed along the activities.

The derivation. The value chain of the 2.50 bun: materials 0.95, labour 0.70, energy 0.20, distribution 0.15, overhead 0.30 → cost 2.30, margin 0.20 = **8.0%** ✓ – above the 6% industry average and below the 11% top quartile, which is exactly where a middling-industry differentiator sits. VRIO on five resources: the ovens (valuable, not rare → parity), the crew (valuable, not rare → parity), the recipe (valuable, rare, but reverse-engineerable → temporary), the loyalty data (valuable, rare, partly organised → temporary), the festival brand (valuable, rare, built on nine years of history – causally ambiguous, organised through the contract → **sustained**). One resource carries the strategy; the other four are the tax.

> **Worked example 1 – score the force, then find the lever.** Buyer power scores 3: festivals negotiate hard (few, large, concentrated) but walk-ins don't. The lever hiding inside the score: contract design (course 11.2's licence terms) converts one negotiator into a relationship, and the loyalty app (course 10's 14,000 records) makes the walk-ins measurable – the score is a diagnosis, and every diagnosis names a treatment or it is decoration.

> **Worked example 2 – escalate to the margin bridge.** To reach the top quartile's 11%, Nimbus needs +0.08 a bun (0.03 price on the festival line, 0.03 materials through the pooling of course 10.5's safety stock, 0.02 labour through the second mixer's 1.67 batches an hour). Frameworks do not close that gap; they locate it – and the bridge above is three line items long, each owned by a unit of this course.

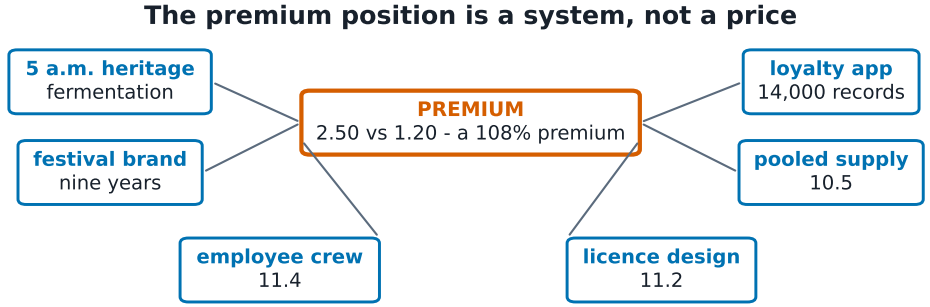
Common mistakes. Scoring forces with adjectives instead of evidence. Reading VRIO as a checklist (the I question is about imitation COST, including time). And treating industry attractiveness as destiny – 16 of 25 industries contain 11% quartiles.

Summary. Forces total 16 of 25: the industry is middling; the margin bridge from 8 to 11% is three line items. VRIO: one sustained resource (the brand), two temporary, two parity. Analysis names levers or it is decoration.

Competitive strategy

Why it matters. "Premium quality at a value price" is not a strategy; it is a wish with two adjectives. A strategy is a POSITION – premium or value, not both – and a system of activities that only makes sense in that position.

The idea, with a picture. Nimbus's choice against the supermarket in-store bakery: its cost is 1.60 a bun at 200,000 a week; Nimbus's is 2.30 at 19,200. Nimbus cannot win cost; it can win willingness to pay – 2.50 against 1.20, a premium of 108% ✓ carried by the festival brand.



the 2.50 bun: materials 0.95 + labour 0.70 + energy 0.20 + distribution 0.15 + overhead 0.30

= cost 2.30, margin 0.20 (8%) - against the supermarket cost line of 1.60

Fit is the moat: each activity makes sense alone, together they are hard to copy

The definitions. Cost leadership : win on the cost curve. **Differentiation** : win on willingness to pay. **Stuck in the middle** : neither – the profitless place. **Activity system** : the map of reinforcing activities that delivers the position; fit makes the system hard to imitate even when every activity is visible. **Dominant strategy** : a move best whatever the rival does.

The derivation. The game (weekly contribution, Nimbus / rival, in thousands): Nimbus Premium against rival Ignore **22 / 14**; Premium against Match 12 / 10; Standard against Ignore 15 / 12; Standard against Match 13 / 9. Rival: Ignore beats Match against both of Nimbus's moves ($14 > 10$, $12 > 9$) – Ignore is dominant. Nimbus: Premium is NOT dominant – it wins against Ignore ($22 > 15$) but loses against Match ($12 < 13$) – so Nimbus's answer must wait on the rival's. The rival's dominance resolves it: Ignore is coming, and against a rival who Ignores, Premium 22

beats Standard 15. The unique equilibrium: **(Premium, Ignore), 22 / 14** ✓.

> **Worked example 1 – the value line that loses money.** The founder floats a "Nimbus Value" bun at 1.80 to catch the supermarket shopper. Arithmetic: the value line's cost is 1.95 (smaller batches, separate packaging, no festival subsidy) → margin **−0.15** ✓ per bun. Worse than unprofitable: it spends the brand's scarcity (course 12.1's only sustained resource) to attack the one rival whose cost curve Nimbus cannot beat. This is the stuck-in-the-middle trap wearing a growth costume – the analysis kills it in one subtraction.

> **Worked example 2 – escalate to the fit that defends.** Suppose the rival MATCHES the premium: weekend festival stalls of its own. The activity map is the defence: the rival can copy the buns (activity one) but not profitably carry the rest – the 5 a.m. heritage fermentation needs the crew (course 11.4's employee tests found loyalty), the contract terms (course 11.2's licence design), the pooled supply (course 10.5), the brand's nine years. Matching one activity buys the cost without the willingness to pay – which is exactly why the payoff matrix says Match pays the rival 10 against Ignore's 14. Fit, not secrecy, is the moat.

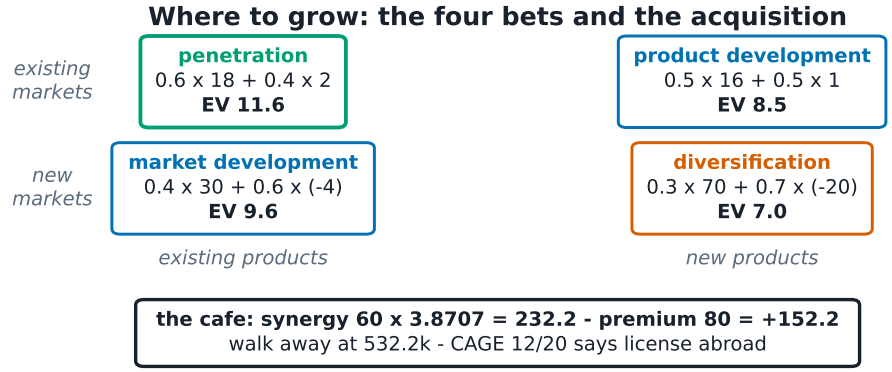
Common mistakes. Reading the activity map as a list (the ARROWS are the strategy). Playing the game without asking what the rival's dominant move is. And pricing a value line before subtracting its own cost.

Summary. Position first: premium at a 108% premium, carried by fit. The equilibrium is (Premium, Ignore) – 22 against 14. A value line at −0.15 is the middle, and the middle is where margins die.

Corporate and international strategy

Why it matters. Growth is not a strategy; it is four different bets with different risks (Ansoff), and one of them – buying the café – is a number, not a feeling. The same for the export idea: distance is measurable, and measured distance usually says license, don't ship.

The idea, with a picture. The four bets, each with an expected value: penetration **11.6k**, market development **9.6k**, product development **8.5k**, diversification **7.0k** ✓ – and the twist the table forces you to see: the highest-EV bet is the safest one, and the most exciting one has the worst EV and the biggest tail.



The boring quadrant wins the arithmetic - and the exciting one is an option, not a bet

The definitions. Ansoff matrix : market/product, old and new – penetration, development, diversification, each step out raising risk. **CAGE distance** CAGE : cultural, administrative, geographic, economic distance – the corrective to "the world is flat". **Synergy** : value only if $2 + 2 > 5$ after the premium is paid. **Parenting advantage**: HQ must add more value than it costs.

The derivation. The EVs: penetration $0.6 \times 18 + 0.4 \times 2 = 11.6$k ✓; market development $0.4 \times 30 + 0.6 \times (-4) = 9.6$k ✓; product development $0.5 \times 16 + 0.5 \times 1 = 8.5$k ✓; diversification $0.3 \times 70 + 0.7 \times (-20) = 7.0$k ✓. The acquisition: the café standalone is worth 300k; Nimbus offers 380k – an 80k premium. Synergies (festival outlet, pooled flour, the loyalty list) 60k a year for five years at the course-9 WACC 9.19%: annuity factor 3.8707, **PV 232.2k** ✓. The deal creates value for Nimbus if the premium is below the synergy PV: $232.2 - 80 = +\mathbf{152.2}$k ✓ – the arithmetic says buy, and the CAGE table on the export idea says license (cultural 3 + administrative 4 + geographic 3 + economic 2 = **12 of 20** ✓: far; export dies in the admin, licensing keeps the recipe home).

> **Worked example 1 – the premium is the price of the synergy.** Two boards argue: "the synergies are 232k, buy at any price" against "acquisitions destroy value". Both are wrong by the same arithmetic: value created = synergy PV − premium, and the premium is the negotiation. At 380k the deal is worth +152.2k; at 550k it is worth −17.8k (232.2 − 250) ✓; the walk-away number is 300 + 232.2 = **532.2k** ✓. The board that does not know its walk-away number is not negotiating; it is hoping.

> **Worked example 2 – escalate to the parent's own test.** If Nimbus buys the café, HQ (the founder) must run two businesses. The parenting test: what does HQ actually add – the flour pooling (real, course 10.5's numbers), the brand licence (real, priced in the synergy), the founder's sixteen-hour days (a COST, and the scarcest resource in the file). The honest answer in the case: buy only with the Saturday supervisor promoted to café manager (course 11.4's status tests already classify her as an employee with the skills), because a parent that adds management hours it does not have is subtracting.

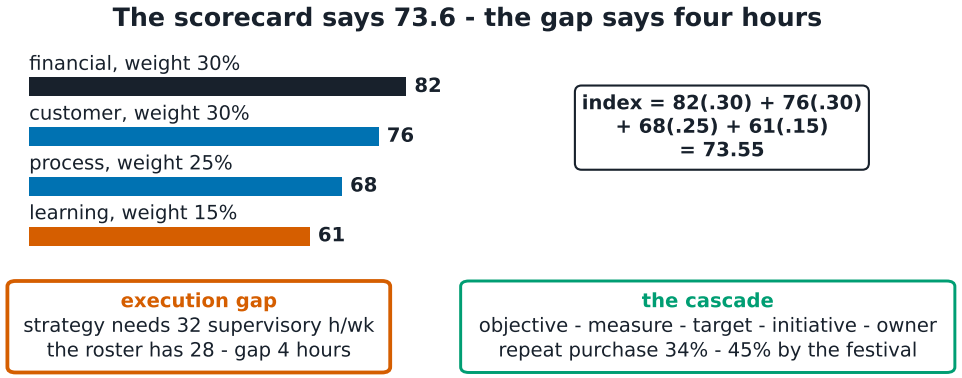
Common mistakes. Reading Ansoff as a menu of equal squares (risk rises left to right and top to bottom). Valuing synergies gross and premiums net. And calling licensing a defeat – with CAGE at 12 of 20 it is the highest-EV form of the international bet.

Summary. EVs: 11.6 / 9.6 / 8.5 / 7.0 – penetration funds the future. The café: value = 232.2 − premium; walk away at 532.2k. CAGE 12 of 20: license the brand, don't ship the buns.

Implementation and organisation

Why it matters. The best strategy in this file fails at 6 a.m. if the crew, the measures and the money point somewhere else. Implementation is not the soft part; it is the part with the scorecard and the arithmetic.

The idea, with a picture. The **balanced scorecard** : financial 82, customer 76, process 68, learning 61 – weighted 30/30/ 25/15 (past results get weight; the future gets some) → index **73.6** ✓. And the gap the index hides: the strategy needs 32 supervisory hours a week; the roster has 28. Four hours a week is the whole execution plan.



The index is the average - the 4-hour gap is the diagnosis

The definitions. Balanced scorecard: financial, customer, process, learning – measures chained by cause (learning → process → customer → financial). **Strategy map**: those chains drawn. The **cascade**: objective → measure → target → initiative → owner, down to the Saturday roster. **Execution gap** : resources the strategy needs minus resources the organisation has – stated in hours and pounds.

The derivation. The index: $82(0.30)+76(0.30)+68(0.25)+61(0.15) = 24.6+22.8+17.0+9.15 = 73.55$ ✓. The weakest perspective (learning 61) is the leading indicator: the process score 68 follows it by a quarter, the customer score 76 by two. The cascade's arithmetic: premium position needs 32 supervisory hours (5 a.m. fermentation checks, festival logistics, the café walk-through); the roster supplies 28; the gap of **4 hours** ✓ is bought three ways in the case – promote and train (+3, the course-11 supervisor), cut one low-value activity (−1, the value line is dead anyway), and measure the rest.

> **Worked example 1 – find the gap before the plan.** A strategy review that never states the gap produces calendars, not change. The case writes it as one sentence: the premium position requires 32 hours of supervision and the organisation has 28 – every initiative that does not close, buy or shrink those four hours is theatre. Compare: the balanced scorecard's 73.6 looks healthy and predicts nothing; the 4-hour gap predicts the August collapse that follows in the case, which is why the gap, not the index, goes on the wall.

> **Worked example 2 – escalate to the measure that changes behaviour.** The cascade's customer objective: "the festival brand means the best bun" → measure: repeat purchase in the loyalty app (currently 34%) → target 45% by the festival → initiative: the 5 a.m. fermentation log, owned by the head baker. The test of a good measure: the owner can move it this month. "Brand equity" fails the test; "34% repeat" passes – and the scorecard's customer 76 becomes auditable one number at a time.

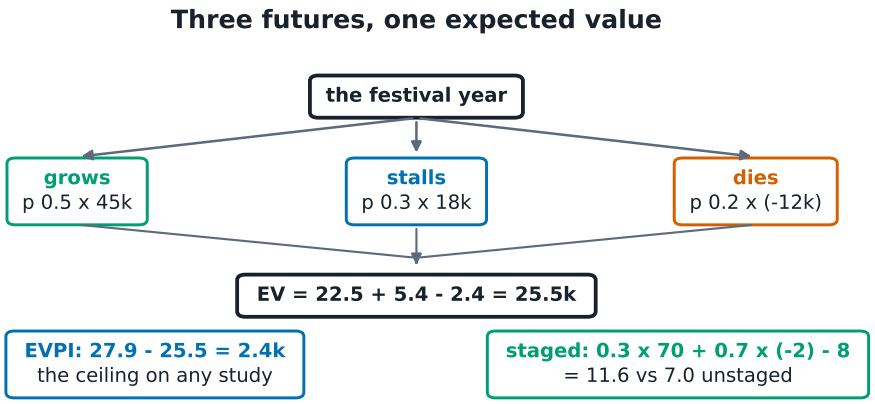
Common mistakes. Weighting the scorecard to flatter (financial at 70% is last year wearing a scorecard). Cascading to measures nobody owns. And writing the implementation plan before finding the execution gap.

Summary. Scorecard 73.6 – and the 4-hour supervisory gap is the real headline. Cascade every objective to an owner who can move the number this month. Learning 61 is the leading indicator; fund it first.

Strategy under uncertainty

Why it matters. The founder does not know whether the festival doubles, stalls or dies – and "we'll adapt" is not a plan. Uncertainty has arithmetic: expected values, the value of information, and the price of staging.

The idea, with a picture. Three scenarios for the festival year: grows (p 0.5, payoff 45k), stalls (p 0.3, 18k), dies (p 0.2, −12k). **Expected value** = $22.5 + 5.4 - 2.4 = \mathbf{25.5}$k ✓. **EVPI**: with perfect information you act only on the first two → $0.5 \times 45 + 0.3 \times 18 = 27.9$k; the information is worth $27.9 - 25.5 = \mathbf{2.4}$k ✓ – that is the most any research study may cost.



Three futures, one expected value - and the pilot is worth more than the forecast

The definitions. **Scenario** : a coherent future, not a forecast. **EV**: probability-weighted payoff. **EVPI** : the ceiling value of information – what perfect foresight would add. **Real option** : paying to keep a decision open; the pilot is a call option on expansion. **Trigger** : the pre-committed number that forces the review.

The derivation. The staged café: pilot the takeover with an 8k festival-weekend trial. Staged: $0.3 \times 70 + 0.7 \times (-2) - 8 = \mathbf{11.6}$k ✓ against the unstaged 7.0 – staging adds +4.6 by cutting the tail, not by raising the head. The scenario arithmetic above gives EV 25.5 and EVPI 2.4; the trigger written into the case: if festival pre-orders fall 20% below plan by week 4, the expansion option lapses unexercised – a decision made in daylight, revisited by rule.

> **Worked example 1 – price the study before commissioning it.** A consultant offers a festival-demand study for 6k. The EVPI arithmetic caps ALL information about the scenarios at 2.4k – the study is unbuyable at any price above that, however good. The reframe that is actually worth 2.4k: a week of loyalty-app data (course 10's 14,000 records) is nearly free and moves the probabilities. Most "research" dies at the EVPI line, and the discipline is one subtraction.

> **Worked example 2 – escalate to the option portfolio.** The case's three options: the café pilot (cost 8, value added 4.6), the school contract pilot (cost 2, adds 1.1 by the same staging arithmetic), the second-oven expansion (already committed – course 10's capacity maths). The portfolio view: options are bought when cheap and uncertainty is high, exercised when the trigger hits, abandoned without grief when it doesn't – the opposite of the plan that is defended because it was written. The founder's one-line summary: spend 10k on options, 2.4k maximum on knowing, and nothing on pretending.

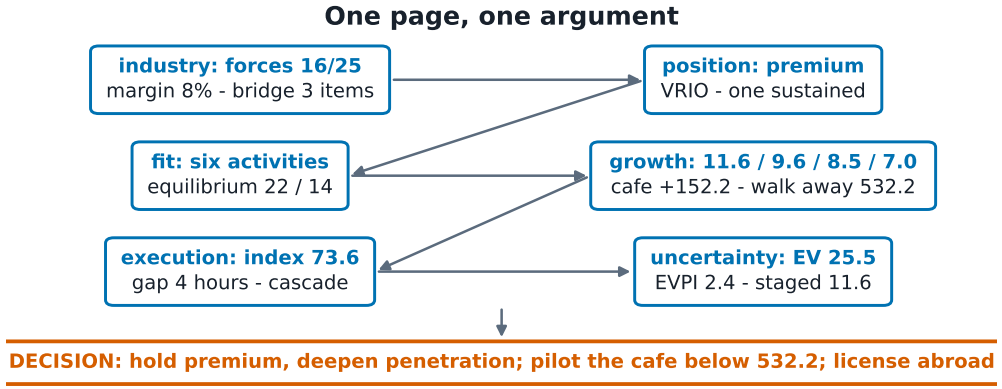
Common mistakes. Treating the EV as a forecast (it is a weighted average of futures that will not happen). Buying information above its EVPI. And confusing flexibility with indecision – the option has a trigger, a price and an expiry or it is procrastination.

Summary. EV 25.5; EVPI 2.4 – the ceiling on research. Staging the café: 7.0 → 11.6 for an 8k pilot; the tail is where options pay. Write the trigger before the uncertainty resolves.

Integrated case

Why it matters. Twelve courses built one bakery. The case pulls the file together – and the test of everything above is whether one page can carry it.

The idea, with a picture. The logic map: the industry (16/25) → the position (premium, one sustained resource) → the fit (six activities) → the growth bets (11.6 / 9.6 / 8.5 / 7.0) → the execution (73.6 index, 4-hour gap) → the uncertainty (EV 25.5, staged to 11.6) → the decision.



One page, one argument - each box recomputes from its unit

The definitions. **A case analysis** : situation → frameworks (chosen, scored) → options (numbered) → decision → risks and review. The discipline from course 1 still holds: evidence for every claim, a number for every argument.

The derivation. The page's spine, recomputed: forces 16 of 25 and margin 8% against a bridge to 11%; VRIO one sustained; the equilibrium (Premium, Ignore) at 22/14 and the value line at −0.15; growth EVs 11.6/9.6/8.5/7.0 with the café at +152.2 and a 532.2 walk-away; the scorecard 73.6 with the 4-hour gap; uncertainty at EV 25.5, EVPI 2.4, staged 11.6. Every number traces to a unit and a line of arithmetic – that is what makes it a case and not a story.

> **Worked example 1 – the decision memo.** Forty words: "We hold the premium position and deepen it (penetration EV 11.6): the second stage, the pooled supply, the fermentation log. We buy the café only at or below 532.2k with the supervisor as manager, piloted first (staged 11.6). We license abroad (CAGE 12). We do not ship buns, and we do not launch the value line." Every clause carries a number from the file – which is why it survives a board.

> **Worked example 2 – escalate to the review.** The memo ends with its own falsification: repeat purchase 34% → 45% by the festival (owner: head baker); supervisory gap 4 → 0 hours by June (owner: founder); festival pre-orders −20% by week 4 triggers the option's lapse (rule). A strategy that cannot say what would prove it wrong is a belief; this one names three numbers and dates.

Common mistakes. Framework tourism – five models, no decision. Numbers without owners, owners without dates. And the retrospective rewrite: the case is written before the festival, not after.

Summary. One page, one argument, every number traceable. The memo holds forty words and three review triggers. Analysis ends in a decision or it never started.

The project sheet (reproducible)

The Nimbus Kitchen strategic file, one page:

- Industry: forces $4 + 3 + 4 + 2 + 3 = \mathbf{16/25}$; bun economics $0.95 + 0.70 + 0.20 + 0.15 + 0.30 = 2.30$ against $2.50 \rightarrow$ margin $0.20 = 8\%$; bridge to the 11% quartile = price 0.03 + materials 0.03 + labour 0.02.
- Position: VRIO – brand sustained; recipe, data temporary; ovens, crew parity. Premium 2.50 vs 1.20 (108%); supermarket cost 1.60; the value line 1.80 vs cost $1.95 = \mathbf{-0.15}$.
- Game: (Premium, Ignore) = 22/14 unique – Ignore dominates for the rival ($14 > 10$, $12 > 9$); Premium wins against Ignore ($22 > 15$).
- Growth: EVs penetration 11.6 / market 9.6 / product 8.5 / diversify 7.0 (thousands). Café: synergy 60k × 3.8707 = 232.2k; premium $80 \rightarrow \mathbf{+152.2}$; walk-away 532.2k. CAGE 3+4+3+2 = 12/20 → license.
- Execution: scorecard 82/76/68/61 at 30/30/25/15 → **73.55**; supervisory need 32 vs have 28 → **gap 4 hours**.
- Uncertainty: scenarios $0.5\times45 + 0.3\times18 + 0.2\times(-12) = \mathbf{EV\ 25.5}$; EVPI $27.9 - 25.5 = \mathbf{2.4}$; staged café $0.3\times70 + 0.7\times(-2) - 8 = \mathbf{11.6}$ vs 7.0 unstaged. Triggers: repeat 34→45; gap→0; pre-orders −20% by week 4.

Every number in this handout recomputes from these values.