

ACCG2000

Cost Analysis and Managerial Decision Making

Where a product's cost actually comes from, why three defensible methods give three different answers, and which of those numbers should be allowed anywhere near a decision.

Contents

PART ONE *Costs, and how they behave*

GROUNDWORK	How this book is organised	4
CHAPTER ONE	Management accounting and cost concepts	7
CHAPTER TWO	Cost-volume-profit analysis	16

PART TWO *Costing the product*

CHAPTER THREE	Job costing	23
CHAPTER FOUR	Process costing	28
CHAPTER FIVE	A closer look at overhead costs	34
CHAPTER SIX	Activity-based costing	39
CHAPTER SEVEN	Absorption and variable costing; service costing	44

PART THREE *Costs for decisions*

CHAPTER EIGHT	Information for decisions	50
CHAPTER NINE	Product mix decisions	57

PART FOUR *Planning and control*

CHAPTER TEN	Budgeting	62
CHAPTER ELEVEN	Standard costing for control	69

PART FIVE *The tool*

CHAPTER TWELVE	Excel for management accounting	77
----------------	---------------------------------	----

Each chapter gives the idea first, then the calculation, then the mistakes that are easy to make. Worked examples are stepped through in full — cover the working and attempt them before reading on.

PART ONE

Costs, and how they behave

Before anything can be costed, decided or budgeted, the word “cost” has to be pinned down — because it means six different things depending on why you are asking. These two chapters settle the vocabulary and then put it to work on the single most useful model in the subject.

How this book is organised

This is a book about cost: where a product's cost comes from, why several defensible methods give several different answers for the same product, and which of those numbers should be allowed near a decision. Unlike financial accounting, almost nothing here is prescribed by a standard. Every number is a choice, and the choices have consequences.

The single idea underneath the subject

A financial accountant reports one figure because the standards require one figure. A management accountant produces whatever figure answers the question in front of them — and different questions need different figures for the same object. The full cost of a product, its variable cost, its avoidable cost and its contribution margin are four different numbers, all correct, each useful for exactly one kind of decision and misleading for the others.

THE QUESTION BEHIND EVERY CALCULATION IN THIS BOOK

What decision is this number for? Costing a product for the inventory figure in the accounts is a different exercise from costing it to decide whether to keep making it. Get the purpose right and the method follows; get it wrong and a perfectly executed calculation gives the wrong answer.

What you should be able to do by the end

1. **Classify** any cost on the six bases that matter, and say which classification the question needs.
2. **Cost a product** under job costing, process costing, conventional overhead absorption and activity-based costing — and explain why the four disagree.
3. **Isolate the relevant costs** for a decision, discarding sunk and unavoidable costs however large they are.
4. **Build a budget** from a sales forecast through to a cash position, and then explain the variances between that budget and what happened.
5. **Do all of it in a spreadsheet** that someone else could audit.

The five parts

PART	WHAT IT SETTLES
One — Costs, and how they behave	The six classification bases; product against period cost; how cost behaviour is estimated; and cost–volume–profit analysis, which is the model everything later leans on.
Two — Costing the product	The five costing systems: job, process, conventional overhead absorption with plantwide and departmental rates, activity-based costing, and absorption against variable costing. Plus costing a service, which has no inventory to hide in.
Three — Costs for decisions	Relevance. Which costs matter to a decision, which are sunk, and how to rank products when a resource runs short.
Four — Planning and control	The master budget as a linked set of schedules, and standard costing as the machinery for explaining why actual results differ from it.
Five — The tool	The spreadsheet skills the work actually requires, and how to build a model that can be checked.

The topics, week by week

Chapters are grouped by theme rather than by calendar — the five costing systems belong together, and are far easier to learn side by side than a fortnight apart. The week each topic is taught is printed at the head of its chapter, and the table below maps one order onto the other.

WK	TOPIC	CHAPTER HERE
1	Introduction, and basic cost concepts and terms	One
2	Cost–volume–profit analysis	Two
3	Product costing systems — job costing	Three
4	Process costing	Four
5	A closer look at overhead costs	Five
6	Activity-based costing	Six
7	Excel	Twelve
8	Absorption and variable costing; service costing	Seven
9	Information for decisions	Eight
10	Product mix decisions	Nine
11	Budgeting	Ten

WK	TOPIC	CHAPTER HERE
12	Standard costing for control	Eleven

How each chapter is built

The idea first, then the calculation, then the traps. Formulas are boxed where the exact form matters. Worked examples are stepped through in full — cover the working and attempt them before reading on, because in this subject the arithmetic is easy and the setup is everything. Each chapter closes with a paragraph carrying the argument into the next.

FURTHER READING

The text covering this ground in full is Langfield-Smith, Smith, Andon, Hilton and Thorne, *Management Accounting: Information for Creating and Managing Value*, 9th edition (McGraw-Hill, 2021). Where this book compresses a topic to its working essentials, that is where the long version lives.

Start with the word itself. “Cost” looks like a simple noun and is in fact a question with six different answers, and choosing the wrong one is the most common way this subject goes wrong.

Management accounting and cost concepts

This chapter does two jobs. It separates management accounting from the financial accounting you already know, and it settles the vocabulary of cost. The second job matters more than it looks: almost every error later in the subject is a classification error made here and carried forward.

What management accounting is

DEFINITION

Management accounting is the set of processes and techniques focused on the effective and efficient use of an organisation's resources, to support managers in their decision making and in improving customer and shareholder value.

Note what is absent from that definition: any external user, any standard, and any prescribed report. Management accounting exists to serve a decision maker inside the organisation, and it is constrained only by usefulness.

MANAGEMENT AGAINST FINANCIAL ACCOUNTING

	MANAGEMENT ACCOUNTING	FINANCIAL ACCOUNTING
Users	Internal — managers, supervisors, executives	External — shareholders, lenders, the ATO, regulators
Governed by	Nothing but usefulness to the decision	Accounting standards and the Corporations Act
Time orientation	Future-oriented as much as historical — forecasts, budgets, projections	Historical — what happened in the period
Frequency	As often as the need arises, up to real time	Half-yearly and annually
Level of detail	Whatever the decision needs — by product, customer, region, activity	Aggregated to the entity
Scope of data	Financial and non-financial; internal and external; estimates welcome	Financial, verifiable, transaction-based
Audited?	No	Yes, annually and independently

THE CONSEQUENCE OF “GOVERNED BY NOTHING”

Because no standard prescribes a management accounting number, two competent accountants can produce different product costs for the same product and both be right. That is not a defect. It is why the subject spends so much time on *which* method to use for *which* purpose, and why a bare cost figure with no stated basis is close to useless.

The three words that have to be exact

DEFINITIONS

Cost — resources given up to achieve a particular objective, measured in monetary terms.

Cost object — any item for which management wants a separate measure of cost: a product, a job, a batch, a process, a department, a customer, a distribution channel.

Cost driver — a factor or activity that causes a cost to be incurred. The plain-language test: *what drives the cost of this?*

A fourth term links them: the **level of activity** is the amount of work performed in the organisation — units produced, machine hours run, invoices processed, patients treated. Cost behaviour is always behaviour *with respect to* a stated level of activity, and a cost that is fixed against one measure of activity may be variable against another.

The six bases of classification

Costs are classified differently for different purposes, because they are used in different ways. There is no single correct classification of a cost — only a correct classification for the question being asked.

BASIS	CLASSIFICATIONS	USED FOR
Behaviour	Variable and fixed (and mixed, or semi-variable)	CVP analysis, budgeting, short-run decisions
Traceability	Direct and indirect	Product costing — deciding what can be traced and what must be allocated
Timing of expense	Product and period cost	Inventory valuation and the income statement
Manufacturing function	Direct material, direct labour, manufacturing overhead	The three elements of manufacturing cost
Value chain	Upstream, manufacturing, downstream	Strategic cost management — where in the chain cost is created
Controllability	Controllable and uncontrollable	Performance evaluation — holding a manager responsible only for what they can affect

BEHAVIOUR: VARIABLE, FIXED AND MIXED

VARIABLE COST	FIXED COST
Changes <i>in total</i> in proportion to the level of activity. Constant <i>per unit</i> .	Does not change in total as activity changes within the relevant range. Falls <i>per unit</i> as activity rises, because the same total is spread over more units.
Direct materials, piece-rate labour, sales commission, packaging.	Factory rent, supervisor salary, straight-line depreciation, insurance.

TWO STATEMENTS THAT ARE CONSTANTLY MISREAD

“Variable costs do not change” is true of the cost **per unit** and false of the **total**. “Fixed costs fall as production rises” is true **per unit** and false in **total**. Always say which you mean, because the two halves of each sentence point in opposite directions.

And no cost is fixed in the long run. Rent is fixed until the lease is renegotiated; the supervisor’s salary is fixed until a second shift needs a second supervisor. “Fixed” is shorthand for *fixed within the relevant range, over the planning horizon*.

A **mixed** or **semi-variable** cost has both a fixed and a variable component — a phone plan with a monthly charge plus usage, or maintenance with a standing crew plus consumables. Splitting the two is what cost estimation is for, below.

Step-fixed costs are fixed over a range and then jump: one supervisor covers up to 20 staff, so the 21st employee adds a whole salary. Treating a step-fixed cost as variable, or as flatly fixed across a big change in volume, is a standard modelling error.

TRACEABILITY: DIRECT AND INDIRECT

Direct costs can be traced to a cost object *in an economic manner* — the timber in a music box, the labour hours booked to a job. **Indirect costs** cannot: the factory supervisor’s salary, the electricity, the glue.

“IN AN ECONOMIC MANNER” IS DOING REAL WORK IN THAT DEFINITION

The glue in a music box is physically part of the product, so it *could* be traced — weigh the glue on each box. Nobody does, because the tracing would cost more than the information is worth. Glue is therefore an indirect material and goes into overhead, even though it is physically direct. Direct and indirect is a judgement about the economics of measurement, not about physics.

The same cost can also be direct to one object and indirect to another. A department supervisor’s salary is *direct* to the department and *indirect* to each product passing through it.

TIMING: PRODUCT AND PERIOD COST

PRODUCT COST	PERIOD COST
Assigned to goods manufactured or purchased for resale. Held in inventory as an asset, and expensed as cost of sales only when the goods are sold.	Expensed in the accounting period in which it is incurred, rather than attached to units.
Direct material, direct labour, manufacturing overhead.	Selling, administrative and finance costs — sales commission, advertising, office salaries, audit fee, interest.

Two consequences worth holding onto. A service business has no inventory, so it has no product costs in this sense — everything is a period cost. And a manufacturing cost incurred this period may not hit this period's profit at all, if the units it attached to are still sitting in the warehouse. That timing difference is the whole of Chapter Seven.

THE THREE MANUFACTURING COST ELEMENTS

Direct material	+ Direct labour	= Prime cost
Direct labour	+ Manufacturing overhead	= Conversion cost
Direct material	+ Direct labour + Overhead	= Total manufacturing cost

Prime cost is the traceable core of the product. **Conversion cost** is what it takes to turn material into product — and it is the pairing that matters in process costing, where materials and conversion are tracked separately because they enter production at different rates.

TWO ITEMS THAT CATCH NEARLY EVERYBODY

Overtime premium. A technician earns \$10 an hour and works 45 hours in a week instead of 40, with overtime paid at 200% of the regular rate. The overtime hours cost \$20 each, of which \$10 is the ordinary rate and \$10 is the *premium*. Only the premium — 5 hours × \$10 = **\$50** — is treated as manufacturing overhead. The other \$10 per hour is ordinary direct labour.

The reasoning: the job that happened to be on the bench at 5pm did not cause the overtime; the total workload did. Charging the premium to that one job would make its cost an accident of scheduling. So the premium goes to overhead and is spread across everything.

Idle time. Time paid for but not worked — a machine breakdown, waiting for materials. It is not direct labour, because no product was being worked on. It is manufacturing overhead.

Depreciation is classified by *where the asset is*, not by the method used to calculate it: depreciation of factory equipment is manufacturing overhead and therefore a product cost; depreciation of the sales office is a period cost.

VALUE CHAIN AND CONTROLLABILITY

The **value chain** classification splits costs into *upstream* (research and development, design), *manufacturing*, and *downstream* (marketing, distribution, customer service). It matters strategically

because most of a product's cost is committed upstream, at design, long before manufacturing begins — which is why cutting factory costs late in a product's life achieves so little.

Controllability asks whether a particular manager can influence the cost within a given period. It exists for performance evaluation: holding a production supervisor responsible for a materials price rise they did not negotiate is both unfair and useless as a control. Note that controllability depends on *who* and *over what horizon* — almost everything is controllable by somebody, eventually.

The manufacturing cost flow

Three inventory accounts in sequence, and a cost sheet for each. Every figure in the schedules below is one of these flows.



```
Opening raw materials
+ Purchases
- Closing raw materials
= Direct material used

Direct material used + Direct labour + Manufacturing overhead applied
= Total manufacturing costs

Opening work in process + Total manufacturing costs - Closing work in process
= Cost of goods manufactured

Opening finished goods + Cost of goods manufactured
= Cost of goods available for sale
- Closing finished goods
= Cost of sales
```

WORKED • COST OF GOODS MANUFACTURED

Materials purchased \$18,000. Direct labour \$14,000. Manufacturing overhead applied \$20,000. Raw materials \$12,000 opening and \$14,000 closing. Work in process \$12,000 opening and \$8,000 closing.

```
Direct material used      = 12,000 + 18,000 - 14,000 = 16,000
Total manufacturing cost = 16,000 + 14,000 + 20,000 = 50,000
Cost of goods manufactured = 12,000 + 50,000 - 8,000 = $54,000
```

The trap is using *purchases* as the material figure. Purchases are what arrived; what belongs in the calculation is what was *issued to production*, and the two differ by the movement in the raw materials store.

WORKED • COST OF GOODS AVAILABLE FOR SALE

Total manufacturing costs \$5,390. Inventories: materials \$100 opening, \$120 closing; work in process \$60 and \$30; finished goods \$500 and \$600.

$$\begin{aligned}\text{Cost of goods manufactured} &= 5,390 + 60 - 30 = 5,420 \\ \text{Cost of goods available for sale} &= 5,420 + 500 = \mathbf{\$5,920}\end{aligned}$$

Two things to notice. The *materials* movement is already inside the \$5,390 — total manufacturing cost is stated after material usage, so adjusting for it again would double-count. And “available for sale” stops at opening finished goods; closing finished goods is deducted only when you go on to cost of sales. Reading past the words “available for sale” and subtracting the \$600 gives \$5,320, which is the planted wrong answer.

WORKED • CLASSIFYING A MONTH'S COSTS

The Casual Furniture Company incurred the following in January. Identify prime cost and period cost.

COST	\$	CLASSIFICATION
Timber	25,000	Direct material
Paint	5,000	Direct material
Glue	500	Indirect material → overhead
Wages — assembly personnel	20,000	Direct labour
Wages — factory supervisor	3,500	Overhead
Factory cleaner's wages	2,000	Overhead
Sales commissions	10,000	Period
Administrative staff salaries	4,000	Period
Depreciation — factory equipment	3,000	Overhead
Depreciation — sales office equipment	1,000	Period
Utilities and insurance — factory	6,000	Overhead
Utilities and insurance — sales office	2,000	Period
Advertising	8,000	Period
Total	90,000	

Prime cost	= 25,000 + 5,000 + 20,000	= \$50,000
Period cost	= 10,000 + 4,000 + 1,000 + 2,000 + 8,000	= \$25,000
Overhead	= 500 + 3,500 + 2,000 + 3,000 + 6,000	= \$15,000
<i>check: 50,000 + 25,000 + 15,000 = 90,000</i>		

The two deliberate traps: **glue** looks like direct material and is not, so it stays out of prime cost; and the paired **factory** / **sales office** items exist so that reading “depreciation” or “insurance” without reading the location gets it wrong. Always total all three groups and check they reconcile — it catches both errors at once.

Estimating cost behaviour: the high–low method

A mixed cost has to be split into its fixed and variable parts before it can be used in any model. The high–low method does it with two observations.

THE METHOD

$$\text{Variable cost per unit} = (\text{Cost at highest activity} - \text{Cost at lowest activity}) \div (\text{Highest activity} - \text{Lowest activity})$$

$$\text{Fixed cost} = \text{Total cost at either point} - (\text{Variable cost per unit} \times \text{activity at that point})$$

Highest and lowest activity — not highest and lowest cost. This is the single most common error in the topic, and questions are built to punish it.

WORKED • MAINTENANCE COST AT TIERCE CORPORATION

MONTH	MACHINE HOURS	MAINTENANCE COST
March	3,000	45,000
April	3,400	44,900
May	4,340	56,800
June	6,500	67,000
July	6,400	75,900

The highest *activity* is June at 6,500 hours, and the lowest is March at 3,000 hours. July has the highest *cost* at \$75,900 and is not used; April has the lowest cost at \$44,900 and is not used either.

$$\begin{aligned}\text{Variable} &= (67,000 - 45,000) \div (6,500 - 3,000) \\ &= 22,000 \div 3,500 = \mathbf{\$6.29 \text{ per machine hour}}\end{aligned}$$

$$\begin{aligned}\text{Fixed} &= 45,000 - (6.286 \times 3,000) = 45,000 - 18,857 = \mathbf{\$26,143 \text{ per month}} \\ &\text{check against June: } 26,143 + 6.286 \times 6,500 = 67,002 \checkmark\end{aligned}$$

Had you taken July and April as the extremes, the variable rate would have come out at \$10.33 an hour and the fixed cost at just over \$9,800 — a completely different cost model built on the same data. The distractors in a question of this type are usually exactly those figures.

WHY HIGH-LOW IS TAUGHT, AND WHY IT IS WEAK

It is taught because it is transparent: two points, one line, no software. It is weak because it uses only two of the observations and both are extremes, so an unusual month at either end distorts the whole model. Regression uses every observation and is what a real analysis would use — Chapter Twelve does it in a spreadsheet. High-low is the hand calculation that shows what regression is doing.

WORKED • PREDICTING TOTAL COST AT A NEW VOLUME

Total costs are \$140,000 when 10,000 units are made, and variable costs are \$4 per unit. What are total costs at 8,000 units?

Variable at 10,000	=	10,000 × 4	=	40,000	
Fixed	=	140,000 – 40,000	=	100,000	<i>fixed in total, so unchanged</i>
Total at 8,000	=	100,000 + (8,000 × 4)	=	\$132,000	

The wrong answer is \$112,000 — scaling the whole \$140,000 down by 20% as though every cost were variable. The distinction between total and per-unit behaviour is the entire content of the question.

Where the management accountant sits

Worth knowing for the discussion questions. In most large Australian organisations the accounting activities sit in a **finance function** headed by a chief financial officer, who is usually responsible for both management and financial accounting and acts as an adviser to managers rather than merely a scorekeeper. Accounting staff are often placed physically close to operations — some inside the factory — precisely so they work alongside the managers they serve.

Two points follow. First, management accounting is not the exclusive property of accountants: performance measurement systems are often designed with human resources, and costing systems are sometimes initiated by production engineers. Second, because the information shapes how people behave, the **behavioural** consequences of a system are part of its design, not an afterthought — a measure that can be gamed will be gamed. That theme returns in Chapter Ten with budgetary slack, and in Chapter Eleven with variances used to blame rather than to explain.

With costs sorted by behaviour, the most useful model in the subject becomes available: what happens to profit when volume changes. That is the next chapter, and everything in Part Three depends on it.

Cost–volume–profit analysis

One model, a handful of formulas, and a great deal of mileage. CVP analysis asks what happens to costs, revenue and profit when sales volume changes. It is the natural tool for short-run decisions, and its central quantity — contribution margin — reappears in every decision chapter later in the book.

DEFINITION

Cost–volume–profit analysis is a technique for determining the effects of changes in an organisation's sales volume on its costs, revenue and profit.

Despite the word “profit”, the model applies just as well to a not-for-profit: a charity needs to know how many places it can fund before its fixed costs are covered, which is the same arithmetic with a different name on the answer.

Contribution margin

THE FOUR FORMS

Unit contribution margin (UCM) = Unit selling price – Unit variable cost

Total contribution margin = Total sales revenue – Total variable cost
= UCM × units sold

Contribution margin ratio (CMR) = UCM ÷ Unit selling price
= Total contribution margin ÷ Total sales revenue

Contribution margin percentage = CMR × 100

What contribution margin *means*: the amount each unit contributes towards covering fixed costs, and — once fixed costs are covered — towards profit. It is not profit per unit. Below break-even every dollar of contribution is going to fixed costs; above it, every dollar goes to profit.

WHERE THE BREAK-EVEN FORMULA COMES FROM

Worth deriving once rather than memorising, because the derivation is what lets you handle the variants.

Profit = Total revenue – Total costs
= (Units × Price) – (Fixed costs + Units × Unit variable cost)
= Units × (Price – Unit variable cost) – Fixed costs
= (Units × UCM) – Fixed costs

At break-even, Profit = 0, so Units × UCM = Fixed costs

THE CORE FORMULAS

Break-even (units) = Total fixed costs ÷ Unit contribution margin
Break-even (dollars) = Total fixed costs ÷ Contribution margin ratio

Target profit (units) = (Fixed costs + Target profit) ÷ UCM
Target profit (dollars) = (Fixed costs + Target profit) ÷ CMR

After-tax target (units) = (Fixed costs + Target profit ÷ (1 - tax rate)) ÷ UCM

Safety margin = Budgeted sales revenue - Break-even sales revenue

The two cross-checks: break-even dollars = break-even units × unit price, and break-even units = break-even dollars ÷ unit price. If they disagree, something is wrong.

THE AFTER-TAX STEP PEOPLE SKIP

Tax is charged on profit, so a target stated *after* tax has to be grossed up to a before-tax figure before it joins fixed costs. If the target is \$147,000 after tax and the rate is 30%, the before-tax requirement is $147,000 \div 0.70 = \$210,000$, not $147,000 \times 1.30 = \$191,100$. Dividing by (1 - rate) and multiplying by (1 + rate) are different operations, and only the first is right.

WORKED • MILT BEAUTY CORP

Selling price \$1,000 per pack. Variable costs per pack: direct material \$250, direct labour \$100, manufacturing overhead \$60, selling costs \$80 — total \$490. Annual fixed costs: manufacturing overhead \$1,000,000, selling and administrative \$2,500,000 — total \$3,500,000. Forecast sales 20,000 packs. Tax rate 30%. Round part-units up.

$$\text{UCM} = 1,000 - 490 = \$510$$

$$\text{CMR} = 510 \div 1,000 = 0.51$$

1 Break-even in units

$$3,500,000 \div 510 = 6,862.7 \rightarrow 6,863 \text{ packs}$$

2 Break-even in sales dollars

$$3,500,000 \div 0.51 = \$6,862,745$$

check: $6,863 \text{ packs} \times \$1,000 = \$6,863,000$, the rounding-up difference

3 Packs for a profit of \$147,000

$$(3,500,000 + 147,000) \div 510 = 3,647,000 \div 510 = 7,151.0 \rightarrow 7,151 \text{ packs}$$

4 Safety margin

$$\text{Budgeted revenue} = 20,000 \times 1,000 = 20,000,000$$

$$\text{Break-even revenue} = 6,862,745$$

$$\text{Safety margin} = \$13,137,255$$

5 Packs for \$147,000 after tax at 30%

$$\text{Before-tax target} = 147,000 \div (1 - 0.30) = 210,000$$

$$(3,500,000 + 210,000) \div 510 = 3,710,000 \div 510 = 7,274.5 \rightarrow 7,275 \text{ packs}$$

6 Break-even if variable manufacturing overhead rises 20%

Variable manufacturing overhead 60 → 72, so unit variable cost 490 → 502

$$\text{New UCM} = 1,000 - 502 = 498$$

$$3,500,000 \div 498 = 7,028.1 \rightarrow 7,029 \text{ packs}$$

Note part 6: only the *variable* manufacturing overhead moves. The \$1,000,000 of fixed manufacturing overhead is a different line and is untouched, and mixing the two is the trap. A 20% rise in one \$60 component pushes break-even up by 166 packs — small inputs move break-even more than people expect, which is the real lesson of a sensitivity question.

Safety margin

The safety margin indicates the extent to which sales can decline before profit turns to zero. It can be expressed three ways, and a question may ask for any of them:

In dollars = Budgeted sales revenue – Break-even sales revenue

In units = Budgeted units – Break-even units

As a ratio = Safety margin in dollars ÷ Budgeted sales revenue

The ratio is the most informative: it says what proportion of forecast sales could evaporate before the business starts losing money. A safety margin of 66% is a comfortable business; 4% is a business betting on its forecast.

Operating leverage — why cost structure matters

Two businesses with identical sales and identical profit can respond completely differently to a change in volume, depending on how their costs are split between fixed and variable.

	COMPANY A MOSTLY VARIABLE	COMPANY B MOSTLY FIXED
Sales	500,000	500,000
Variable costs	350,000	100,000
Contribution margin	150,000	400,000
Fixed costs	100,000	350,000
Profit	50,000	50,000
<i>Contribution margin ratio</i>	<i>30%</i>	<i>80%</i>

Now raise sales 20%, to \$600,000. Company A gains $20\% \times 30\% \times 500,000 = \$30,000$ of extra contribution, so profit goes to \$80,000 — up 60%. Company B gains $20\% \times 80\% \times 500,000 = \$80,000$, so profit goes to \$130,000 — up 160%.

THE POINT, IN BOTH DIRECTIONS

The business with the higher proportion of fixed costs has the higher contribution margin ratio, and therefore gains more profit from each extra sales dollar once break-even is passed. The same structure works against it on the way down: a 20% fall in sales costs Company B \$80,000 of profit and pushes it to a \$30,000 loss, while Company A merely drops to \$20,000 profit.

High fixed costs mean high reward and high risk. That is the whole answer to “which company benefits most from a sales increase?” — and a complete answer names the downside too.

Multiple products: the sales mix

With more than one product there is no single unit contribution margin, so CVP works on a **weighted average** based on an assumed sales mix that stays constant.

$$\text{Weighted average UCM} = \sum (\text{UCM of each product} \times \text{its share of unit sales})$$

WORKED • MAXIE PTY LTD

Two products, Plain and Fancy. Plain sells for \$20.00 with variable costs of \$12.00; Fancy sells for \$35.00 with variable costs of \$24.50. Sixty per cent of unit sales are Plain. Annual fixed expenses \$45,000. The mix remains constant.

$$\text{UCM Plain} = 20.00 - 12.00 = 8.00 \qquad \text{UCM Fancy} = 35.00 - 24.50 = 10.50$$

$$\begin{aligned} \text{Weighted average UCM} &= (0.60 \times 8.00) + (0.40 \times 10.50) \\ &= 4.80 + 4.20 = \mathbf{\$9.00} \end{aligned}$$

How many units of Fancy must be sold for a target profit of \$31,500?

$$\begin{aligned} \text{Total units required} &= (45,000 + 31,500) \div 9.00 = 76,500 \div 9.00 = 8,500 \text{ units} \\ \text{Fancy} &= 40\% \times 8,500 = \mathbf{3,400 \text{ units}} \quad \text{Plain} = 5,100 \text{ units} \end{aligned}$$

The trap is answering **8,500** — the total — when the question asked for units of Fancy. The weighted average calculation always produces a total, and the final step is to apply the mix percentage to get back to the individual product. Read which one is wanted.

WEIGHTING BY UNITS, NOT BY REVENUE

The mix here is 60/40 by *units*. If a question gives the mix by *sales revenue* instead, you cannot use those percentages against unit contribution margins — weight the contribution margin *ratios* by revenue share instead, and work in dollars. Mixing a unit-based weight with a revenue-based mix is a silent error that produces a plausible wrong number.

The assumptions, and why they are the answer to half the discussion questions

ASSUMPTION	WHAT BREAKS WHEN IT FAILS
Selling price is constant	Volume discounts and price cuts to shift extra units make revenue non-linear. The model overstates the profit from big volume increases.
Costs are linear, and neatly split into fixed and variable	Step-fixed costs jump; overtime raises the variable rate at high volumes; bulk buying lowers material cost. Linearity holds only inside the relevant range.
Sales mix is constant	The weighted average UCM changes, so break-even changes even though nothing about any individual product moved.
Production equals sales — no inventory change	With inventory movement, absorption costing profit and variable costing profit diverge (Chapter Seven), and CVP tracks the variable costing figure.

ASSUMPTION	WHAT BREAKS WHEN IT FAILS
Efficiency and productivity are unchanged	Learning effects lower unit cost as volume accumulates, so the variable rate is not really constant over a long horizon.
Volume is the only cost driver	Batch-level and product-level costs move with the number of setups or the number of products, not with units. That objection is the whole case for activity-based costing in Chapter Six.

USING THE ASSUMPTIONS WELL

The assumptions are not a disclaimer to be recited. They are a list of the circumstances in which the model's answer should not be trusted. A strong answer names the assumption that is *actually* under pressure in the scenario given — a business planning a 300% volume increase has a relevant-range problem; a business changing its product emphasis has a sales-mix problem — rather than listing all six.

CVP needs a unit variable cost and a total fixed cost, and it takes both as given. Part Two asks the harder question underneath: where those numbers come from, and why the answer depends on which system you use to build them.

PART TWO

Costing the product

Five systems, one question: what did this thing cost? Direct materials and direct labour are the easy part. The whole difficulty — and every disagreement between the five systems — is overhead, which by definition cannot be traced and must therefore be allocated on some basis somebody chose.

Job costing

Job costing is the system for businesses that make distinguishable things: a house, an audit, a print run, a bridge. Each job collects its own costs. The chapter's real subject is the machinery for getting overhead onto a job before anyone knows what overhead actually was — and what to do at year end when the estimate turns out wrong.

Job costing or process costing?

JOB COSTING	PROCESS COSTING
Products are distinguishable . Each job is different, so each job needs its own cost record.	Products are essentially identical and flow continuously through a sequence of processes.
Costs are accumulated <i>per job</i> .	Costs are accumulated <i>per process per period</i> , then averaged across the units that passed through.
Construction, custom furniture, professional services, printing, shipbuilding, film production.	Petroleum refining, cement, flour, soft drink, paint, chemicals, sugar.

WHAT DECIDES WHICH SYSTEM A BUSINESS USES

Uniformity of output. Process costing assumes each unit produced is relatively uniform, or essentially identical — which implies the inputs of material, labour and overhead are also relatively uniform and the production process is consistent. Where that holds, averaging is not a compromise; it is the correct answer, because the units genuinely are interchangeable.

Where output is heterogeneous, averaging would be actively misleading: it would tell you the average cost of a kitchen renovation and a garage extension, a figure describing neither.

WHICH IS "MORE ACCURATE"? — THE WRONG QUESTION, ANSWERED PROPERLY

Job costing traces more cost to a specific output, so in a business with distinguishable jobs it is more accurate. But in a business producing 400,000 identical litres of paint, job costing is not more accurate — it is impossible, and the effort of pretending otherwise buys nothing, since every litre really did cost the same.

The honest answer names the condition: *accuracy depends on whether the units differ*. Neither system is superior in the abstract, and hybrid systems (operation costing) exist for businesses whose output is uniform in process and varied in materials, such as clothing manufacture.

The job cost sheet

The subsidiary record for each job, and the document the whole system is built around. It accumulates three things:

- **Direct material**, from materials requisitions charged to the job;
- **Direct labour**, from time sheets or job tickets booked to the job;
- **Manufacturing overhead applied**, at a predetermined rate $\times$ the job's actual usage of the cost driver.

The sum of all open job cost sheets equals the balance of Work in Process. The sum of completed-but-unsold job sheets equals Finished Goods. That reconciliation is the control over the whole system.

The predetermined overhead rate

Overhead cannot be charged to a job as it is incurred, because most of it is not known until the period ends — the electricity bill arrives in arrears, the annual insurance is one payment, the supervisor's salary covers everything. But a job finished in July needs a cost in July. The answer is to estimate in advance.

THE TWO FORMULAS

Predetermined overhead rate = **Budgeted** overhead cost $\div$ **Budgeted** level of cost driver

Overhead applied = Predetermined rate $\times$ **Actual** level of cost driver used

Budgeted over budgeted to set the rate; the rate times actual to apply it. Getting that pairing the wrong way round is the most common error in the topic, and it produces an answer that looks entirely reasonable.

The cost driver is whatever best explains overhead consumption: direct labour hours, machine hours, direct labour cost. In a labour-intensive workshop, labour hours; in an automated plant, machine hours. Chapter Five deals with what happens when a single driver is a poor explanation, and Chapter Six with what happens when no single driver will do.

WORKED • KELSTON MANUFACTURING

Kelston applies manufacturing overhead using a predetermined rate based on machine hours. At the beginning of the year it estimated total manufacturing overhead at \$550,000 and machine hours at 100,000. Actual overhead incurred during the year was \$150,000 and actual machine hours worked on jobs were 25,000.

Predetermined rate = $550,000 \div 100,000 = \$5.50$ per machine hour

Overhead applied = $5.50 \times 25,000 = \$137,500$

Actual overhead = $\$150,000$

Applied 137,500 – Actual 150,000 = **\$12,500 underapplied**

The distractor in a question like this is \$400,000 — the difference between the *budgeted* \$550,000 and the actual \$150,000. That comparison is meaningless: it compares a full year's plan against a quarter of the activity. Over- and underapplied overhead is always *applied against actual*, never *budget against actual*.

Over- and underapplied overhead

WHICH IS WHICH

- Applied < Actual → **UNDERAPPLIED** → not enough charged to jobs
 → debit balance left in the overhead account
 → product costs understated
- Applied > Actual → **OVERAPPLIED** → too much charged to jobs
 → credit balance left in the overhead account
 → product costs overstated

The clearest way to hold it is the control account. Overhead is *debited* as it is actually incurred and *credited* as it is applied to production:

MANUFACTURING OVERHEAD – DEBIT SIDE	CREDIT SIDE
Actual overhead cost incurred during the period	Overhead applied (charged) to production
The closing entry for overapplied overhead	The closing entry for underapplied overhead

WHY THE CLOSING ENTRIES LAND WHERE THEY DO

Overapplied overhead means the credits (applied) exceeded the debits (actual), leaving a *credit* balance. Clearing a credit balance requires a *debit*. So both actual costs incurred and the overapplied closing entry appear on the debit side – which is exactly the point a question about “what appears on the debit side” is testing. Underapplied overhead is the mirror image and is closed with a credit.

DISPOSING OF THE BALANCE

Two treatments, and the choice depends on materiality:

CLOSE TO COST OF SALES	PRORATE ACROSS INVENTORIES AND COST OF SALES
The whole balance is written off against cost of sales. Simple, and appropriate where the amount is immaterial. <i>Underapplied:</i> Dr Cost of Sales, Cr Overhead. <i>Overapplied:</i> Dr Overhead, Cr Cost of Sales.	The balance is apportioned between Work in Process, Finished Goods and Cost of Sales in proportion to the overhead already applied in each. More accurate, because the misestimate attached to units still in stock as much as to units sold. Used where the amount is material.

The journal entries, in flow order

1	Raw Materials <i>purchase of materials</i>	✓
	Accounts Payable	✓
2	Work in Process <i>direct material requisitioned to jobs</i>	✓
	Manufacturing Overhead <i>indirect material</i>	✓
	Raw Materials	✓
3	Work in Process <i>direct labour</i>	✓
	Manufacturing Overhead <i>indirect labour, idle time, overtime premium</i>	✓
	Wages Payable	✓
4	Manufacturing Overhead <i>other actual overhead incurred</i>	✓
	Cash / Payables / Accumulated Depreciation	✓
5	Work in Process <i>overhead applied at the predetermined rate</i>	✓
	Manufacturing Overhead	✓
6	Finished Goods <i>jobs completed</i>	✓
	Work in Process	✓
7	Cost of Sales <i>jobs sold</i>	✓
	Finished Goods	✓

THE PATTERN WORTH MEMORISING

Entries 2, 3 and 4 all **debit Manufacturing Overhead with the indirect portion**. Entry 5 is the only one that credits it. So overhead accumulates actual costs from three directions and is relieved from one — which is why a balance is left at period end unless the estimate was perfect.

Costing a service with job costing

Professional services — audit, law, architecture, consulting — are job costing with different words. Each client engagement is a job, and the procedure is the same four steps:

- 1 **Identify the direct costs** of the job — chiefly professional labour, traced by time sheet, plus any direct disbursements.
- 2 **Aggregate the overhead costs** into cost pools.
- 3 **Identify an overhead cost driver** and calculate a predetermined overhead rate.
- 4 **Apply overhead** to each job using that rate and the job's actual driver usage.

Note the classification point: in an accounting practice, wages paid to audit clerks are **direct labour**. The reflex is to call professional salaries an overhead because there is no factory, but the clerk's hours are

traced to a client engagement, and the engagement is the cost object. Chapter Seven takes up the other kind of service business, where averaging is the right approach.

Job costing works because the units are distinguishable. Turn that assumption off — 300,000 identical units flowing through a plant, some of them half-finished at midnight on the last day of the month — and a different machinery is needed.

Process costing

When output is uniform and continuous, cost per unit is just total cost divided by units. The complication is the units that are only partly finished at period end: they consumed some cost, so they must absorb some cost, and deciding how much is what the **equivalent unit** is for. Learn the four-step schedule as a fixed template and this becomes mechanical.

DEFINITION

Process costing accumulates the costs of each process for a period and then averages them across all the units that passed through that process, expressed in **equivalent units**.

An **equivalent unit** is the amount of work that would produce one completed unit. Four units 25% complete represent one equivalent unit of output.

EQUIVALENT UNITS ARE NOT PHYSICAL THINGS

Nobody can point at the equivalent unit on the factory floor. It is a measure of work done, not of objects, and its purpose is to make partly finished output commensurable with finished output so that a single average cost can be calculated. That is the whole idea, and it is worth being comfortable with the abstraction before touching the arithmetic.

Why material and conversion are tracked separately

Materials are typically added at the start of a process, while labour and overhead are consumed steadily throughout it. So a unit that is 25% through the process may be 100% complete for materials and only 25% complete for conversion. Averaging the two together would misstate both.

Conversion cost — direct labour plus manufacturing overhead — is therefore treated as a single stream, and every schedule has two columns: *direct material* and *conversion*.

The four steps — weighted average method

THE TEMPLATE

- Step 1** Analyse the physical flow of units
Beginning WIP + units started = units to be accounted for
Units completed and transferred out + ending WIP = units accounted for
- Step 2** Calculate equivalent units, separately for DM and conversion
Units completed and transferred out (always 100%)
+ Ending WIP × its percentage of completion
= Total equivalent units
- Step 3** Calculate cost per equivalent unit, separately for DM and conversion
(Cost in beginning WIP + cost added this period) ÷ equivalent units
- Step 4** Assign the total cost
Cost of units completed and transferred out
+ Cost of ending WIP
= Total cost accounted for — *must equal total cost in process*

THE DEFINING FEATURE OF WEIGHTED AVERAGE

Step 3 adds the cost already in beginning work in process to the cost added this period, and divides by equivalent units that count the completed output at 100%. In other words it blends last period's costs with this period's and strikes one average.

The consequence catches almost everyone: **the percentage completion of *beginning* WIP is irrelevant under weighted average**. A question will give it to you — “beginning WIP, direct material 75% complete and conversion 40% complete” — and under weighted average you do not use those figures anywhere. They matter under FIFO, which separates the two periods. If you find yourself reaching for them, stop.

WORKED • KESTREL INSULATION, FIRST PROCESS, JUNE

Work in process at 1 June: 60,000 units (direct material 75% complete, conversion 40% complete), carrying direct material cost of \$56,600 and conversion cost of \$14,900. Units started into production 280,000. Units completed and transferred out 300,000. Work in process at 30 June: direct material 50% complete, conversion 25% complete. Costs added during June: direct material \$385,000, conversion \$214,500. Use the weighted average method.

STEP 1 – PHYSICAL FLOW

Beginning WIP	60,000
+ Units started	280,000
= Units to account for	340,000
- Completed and transferred (300,000)	
= Ending WIP	40,000 units

STEP 2 – EQUIVALENT UNITS

	PHYSICAL UNITS	DIRECT MATERIAL	CONVERSION
Completed and transferred out	300,000	300,000	300,000
Ending WIP — DM 50%, conversion 25%	40,000	20,000	10,000
Total equivalent units	340,000	320,000	310,000

STEP 3 – COST PER EQUIVALENT UNIT

	DIRECT MATERIAL	CONVERSION	TOTAL
Cost in beginning WIP	56,600	14,900	71,500
Cost added during June	385,000	214,500	599,500
Total cost in process	441,600	229,400	671,000
Equivalent units (step 2)	320,000	310,000	
Cost per equivalent unit	\$1.38	\$0.74	\$2.12

STEP 4 – ASSIGN THE COST

Completed and transferred out		
300,000 × \$2.12	=	\$636,000
Ending work in process		
Direct material 20,000 × \$1.38 =	27,600	
Conversion 10,000 × \$0.74 =	7,400	
	=	\$35,000
Total cost accounted for	=	\$671,000
check against total cost in process: 441,600 + 229,400 = 671,000 ✓		

ALWAYS FINISH WITH THE RECONCILIATION

Total cost accounted for in step 4 must equal total cost in process from step 3. It is the only self-check the method has, it takes ten seconds, and it catches nearly every arithmetic slip — a wrong equivalent-unit figure, a cost column added into the wrong row, a completion percentage applied to the wrong stream.

Note also that completed units are costed at the *combined* rate of \$2.12, because they are 100% complete for everything. Ending WIP must be costed *stream by stream*, because its two percentages differ. Applying \$2.12 to ending WIP is a frequent and expensive shortcut.

WORKED • A SECOND PASS, WITH DIFFERENT NUMBERS

Beginning WIP 30,000 units, 60% complete for conversion; direct material \$12,000 and conversion \$15,800. Units started 320,000. Completed and transferred out 310,000. Ending WIP is 100% complete for direct material and 40% complete for conversion. Costs added: direct material \$163,000, conversion \$310,200.

Step 1 Ending WIP = 30,000 + 320,000 - 310,000 = **40,000 units**

Step 2 Equivalent units

DM 310,000 + (40,000 × 100%) = 350,000

Conversion 310,000 + (40,000 × 40%) = 326,000

Step 3 Cost per equivalent unit

DM (12,000 + 163,000) ÷ 350,000 = 175,000 ÷ 350,000 = **\$0.50**

Conversion (15,800 + 310,200) ÷ 326,000 = 326,000 ÷ 326,000 = **\$1.00**

Combined **\$1.50**

Step 4 Transferred out 310,000 × 1.50 = 465,000

Ending WIP DM 40,000 × 0.50 = 20,000

Conv 16,000 × 1.00 = 16,000 = 36,000

Total accounted for = **501,000**

check: 175,000 + 326,000 = 501,000 ✓

Notice that direct material in ending WIP is 100% complete here, so all 40,000 units count as 40,000 equivalent units of material. Material added at the *start* of a process is always fully complete in any unit that has entered the process at all — a detail that turns a hard-looking question into an easy one.

Costs transferred in from a previous process

Most process industries have several sequential processes, and the output of the first becomes the input of the second. In the second department the cost arriving with those units is a third cost stream, usually called **transferred-in cost**.

HOW TO HANDLE IT

Treat transferred-in cost as a **third column** alongside direct material and conversion, and run the identical four steps on all three. Transferred-in cost behaves exactly like a material added at the very start of the process: any unit present in the department at all is **100% complete** with respect to it, whatever its stage of processing otherwise.

So ending WIP in the second department is 100% complete for transferred-in cost, and some lower percentage for that department's own material and conversion.

The weighted average method against FIFO

WEIGHTED AVERAGE	FIFO
Blends the cost in beginning WIP with the cost added this period and strikes one average.	Keeps the two periods separate: beginning WIP is finished first at last period's cost, and only then are new units started.
Beginning WIP completion percentages are not used .	Beginning WIP completion percentages are used, to work out the effort needed to finish those units.
Simpler, and the method used throughout this book.	Gives a cost per unit that reflects only current-period cost, so it is better for cost control.

The two give the same answer when there is no beginning work in process, and diverge more the larger the opening inventory and the faster costs are moving.

Where this goes wrong

ERROR	HOW TO AVOID IT
Using the beginning WIP completion percentages	Under weighted average they are a distractor. The only percentages that enter the calculation are <i>ending</i> WIP's.
Costing ending WIP at the combined rate	Ending WIP is at different stages for material and conversion, so cost it stream by stream and add.
Forgetting beginning WIP cost in step 3	The numerator is beginning cost <i>plus</i> cost added. Dividing only this period's cost by weighted-average equivalent units mixes two methods and is wrong under both.
Deriving ending WIP wrongly	Beginning + started – completed. If the answer is negative, the “units started” and “units completed” figures have been swapped — a question will sometimes list completed output above started output to invite it.
Skipping the reconciliation	Ten seconds, and it catches almost everything.

Both systems so far have taken the overhead rate as given: budgeted overhead over budgeted driver, one rate for the plant. That single number conceals a great deal, and the next two chapters take it apart.

A closer look at overhead costs

One overhead rate for a whole factory assumes every product consumes overhead in proportion to the same single driver. That is rarely true. This chapter refines the rate twice: first by department, then by dealing properly with the support departments whose costs have to reach a product through another department first.

The problem with a plantwide rate

DEFINITION

A **plantwide rate** is a single overhead rate calculated for the entire production plant.

$$\begin{aligned} \text{Predetermined overhead rate} &= \text{Budgeted overhead cost} \div \text{Budgeted level of cost driver} \\ \text{Overhead applied} &= \text{Predetermined rate} \times \text{actual level of cost driver} \end{aligned}$$

It is defensible when the plant is homogeneous: one process, one pattern of overhead consumption. It distorts as soon as departments differ. Consider a plant with a highly automated machining department and a labour-intensive assembly department. A single rate based on direct labour hours charges almost no overhead to machining — which has few labour hours and most of the depreciation, power and maintenance — and loads it onto assembly instead. Products that pass mainly through machining are systematically undercosted, and products that pass mainly through assembly subsidise them.

Departmental rates

The refinement is to recognise overhead separately in each production department. Three steps:

- 1 **Assign overhead costs to production departments.** Costs incurred within a department are traced to it; costs shared across departments are allocated on a sensible basis (floor area for rent, headcount for canteen, machine value for insurance).
- 2 **Calculate a predetermined overhead rate for each department,** using a driver appropriate to *that* department — machine hours in machining, labour hours in assembly.
- 3 **Apply overhead to each product** at each department's rate, based on the product's actual driver usage in that department.

$$\text{Total overhead applied to a product} = \Sigma (\text{departmental rate} \times \text{driver used in that department})$$

WHAT DEPARTMENTAL RATES FIX, AND WHAT THEY DO NOT

They fix the distortion caused by *departments* consuming overhead differently. They do not fix the distortion caused by *products* consuming overhead differently within a department — a product needing forty machine setups and a product needing one still get overhead in proportion to their machine hours. That residual problem is what Chapter Six is about.

Support departments

Production departments make the product. **Support** (or service) departments do not: maintenance, the canteen, quality control, the tool room, IT. Their costs are genuine manufacturing overhead and must reach the product, but they cannot be applied directly — no product spends machine hours in the canteen. So support department costs are first allocated to production departments, and then flow to products through those departments' rates.

The complication that generates three methods: support departments serve *each other*. Maintenance repairs the canteen's dishwasher; the canteen feeds the maintenance crew.

METHOD	HOW IT TREATS SERVICES BETWEEN SUPPORT DEPARTMENTS	CONSEQUENCE
Direct	Ignores them completely. Every support department's cost goes straight to production departments, reallocating its percentages to add to 100%.	Simplest; least accurate. The reciprocal services are simply pretended away.
Step-down	Recognises them in one direction only. Support departments are allocated in sequence; each can pass cost to support departments not yet allocated, but nothing flows back to a department already closed.	Better. In simple cases it can capture all inter-support flows; in complex ones it recognises only about half of them.
Reciprocal	Recognises all of them, in both directions, by solving simultaneous equations.	Conceptually superior and the most accurate; more work, and needs algebra.

CHOOSING THE STEP-DOWN ORDER

Two rules, applied in order:

1. Allocate first the support department that **provides the highest proportion of its service to other support departments**.
2. If that is tied, allocate first the department with the **highest overhead cost**.

The logic: get the department with the most inter-support involvement out of the way first, so that the maximum number of reciprocal flows is actually captured before departments start being closed off.

WORKED • HARLOW FABRICATION, ALL THREE METHODS

Two support departments and two production departments. Overhead before allocation:

Maintenance \$180,000, Canteen \$90,000. Service provided:

PROVIDED BY	MAINTENANCE	CANTEEN	MACHINING	ASSEMBLY
Maintenance	—	20%	30%	50%
Canteen	10%	—	60%	30%

DIRECT METHOD

Inter-support services are ignored, so each support department's percentages are re-based on the production departments alone.

Maintenance 30 : 50, so 37.5% and 62.5%

Machining $180,000 \times 30/80 = \mathbf{67,500}$

Assembly $180,000 \times 50/80 = \mathbf{112,500}$

Canteen 60 : 30, so 66⅔% and 33⅓%

Machining $90,000 \times 60/90 = \mathbf{60,000}$

Assembly $90,000 \times 30/90 = \mathbf{30,000}$

	MACHINING	ASSEMBLY	TOTAL
From Maintenance	67,500	112,500	180,000
From Canteen	60,000	30,000	90,000
Total allocated	127,500	142,500	270,000

STEP-DOWN METHOD

Maintenance provides 20% of its service to another support department; the Canteen provides 10%. By rule 1, **Maintenance is allocated first.**

Maintenance \$180,000 allocated on its own percentages

Canteen 20% = 36,000 *the Canteen is not yet closed, so it can receive*

Machining 30% = 54,000

Assembly 50% = 90,000

Canteen now holds $90,000 + 36,000 = \mathbf{126,000}$

Allocated to production only 60:30 – nothing flows back to Maintenance

Machining $126,000 \times 60/90 = 84,000$

Assembly $126,000 \times 30/90 = 42,000$

	MAINTENANCE	CANTEEN	MACHINING	ASSEMBLY
Overhead before allocation	180,000	90,000	—	—
Allocate Maintenance	(180,000)	36,000	54,000	90,000
Subtotal	—	126,000	54,000	90,000
Allocate Canteen	—	(126,000)	84,000	42,000
Total allocated	—	—	138,000	132,000

RECIPROCAL METHOD

Let M be the total cost of Maintenance to be allocated, and C the total cost of the Canteen, each including the service received from the other. Then:

$$M = 180,000 + 0.10C$$

$$C = 90,000 + 0.20M$$

$$\begin{aligned} \text{Substituting: } M &= 180,000 + 0.10(90,000 + 0.20M) \\ &= 180,000 + 9,000 + 0.02M \end{aligned}$$

$$M(0.98) = 189,000$$

$$M = \mathbf{192,857}$$

$$C = 90,000 + 0.20(192,857) = \mathbf{128,571}$$

Now allocate each department's full reciprocal cost on its own percentages:

	MACHINING	ASSEMBLY
From Maintenance 30% and 50% of 192,857	57,857	96,429
From Canteen 60% and 30% of 128,571	77,143	38,571
Total allocated	135,000	135,000

Check: $135,000 + 135,000 = \$270,000$, the original $\$180,000 + \$90,000$. The reciprocal allocations to each other cancel out, as they must — no cost is created or destroyed by the method.

THE THREE ANSWERS SIDE BY SIDE

METHOD	MACHINING	ASSEMBLY
Direct	127,500	142,500
Step-down	138,000	132,000
Reciprocal	135,000	135,000

Same costs, same service percentages, three answers, a \$15,000 spread between the extremes for Machining. Nothing about the factory changed — only the allocation convention. If those departmental totals then drive product costs, and product costs drive pricing, the choice of method has just moved the price of a product.

WHY THE RECIPROCAL METHOD IS CONCEPTUALLY SUPERIOR

Three sentences, and they are the complete answer:

The **reciprocal** method fully recognises all services provided by support departments to other support departments. The **direct** method completely ignores all such services. The **step-down** method can, in certain situations, recognise all such inter-support cost flows, but in complex situations it recognises only one half of them — because once a department has been allocated and closed, nothing can flow back to it.

Superiority here is conceptual, not practical. Direct is defensible when inter-support services are trivial, and the cost of the extra precision has to be worth having — which is the cost-benefit judgement that runs through the whole subject.

Departmental rates and properly allocated support costs still leave one assumption standing: that within a department, overhead is consumed in proportion to volume. Break that assumption and you get the most-discussed costing system of the last forty years.

Activity-based costing

Conventional costing allocates overhead on a volume measure — labour hours, machine hours, units. Much overhead is not caused by volume at all: a setup costs the same whether the batch is ten units or ten thousand. Activity-based costing traces overhead to the activities that cause it, and then to the products that consume those activities.

Why volume-based allocation distorts

Overhead costs are caused at different levels of the organisation, and only the first level moves with units.

LEVEL	COST IS DRIVEN BY	EXAMPLES
Unit level	The number of units produced	Power to run a machine, consumables, some quality inspection
Batch level	The number of batches, regardless of batch size	Machine setups, first-item inspection, materials handling, purchase order processing
Product level	The existence of a product line	Product design and engineering, holding a part number, product-specific tooling and testing
Facility level	Nothing product-specific — the cost of having a plant at all	Factory rent, plant management salaries, security, insurance

THE DISTORTION, IN ONE SENTENCE

Where overhead is largely batch- and product-level but is allocated on a unit-level driver, **high-volume simple products are overcosted and low-volume complex products are undercosted** — because the high-volume product absorbs overhead in proportion to its units while the complex product actually consumes the setups, the engineering hours and the special handling.

Facility-level costs have no genuine driver at all. ABC allocates them on some arbitrary basis, or leaves them out of product cost and treats them as a period cost. That honesty about the arbitrariness is a point in ABC's favour, not against it.

The procedure

- 1 **Identify the activities** that consume resources — setups, machining, engineering, packing, inspection.
- 2 **Assign overhead cost to each activity**, forming an activity cost pool.
- 3 **Identify an activity driver** for each pool: the factor that causes that activity's cost.
- 4 **Calculate an activity rate** — budgeted cost of the pool ÷ budgeted level of its driver.

Apply cost to products — activity rate $\times$ the product's actual consumption of the driver, summed across all activities.

Activity rate = Budgeted overhead cost of the activity $\div$ Budgeted level of the activity driver

Overhead to a product = Σ (activity rate $\times$ driver units consumed by that product)

WORKED • WAVERLEY MOWERS, CONVENTIONAL AGAINST ABC

Two models of petrol lawnmower. Overhead has been allocated on direct labour hours.

	STANDARD	DELUXE
Annual sales (units)	25,000	26,750
Selling price	\$450	\$650
Direct material per unit	\$80	\$120
Direct labour per unit	\$50	\$40
Direct labour hours per unit	5	4

OVERHEAD ACTIVITY	COST	DRIVER	STANDARD	DELUXE
Setups	100,000	Number of setups	15	10
Machining	1,000,000	Machine hours	60,000	40,000
Engineering	800,000	Engineering hours	25,000	15,000
Packing	420,000	Packing orders	20,000	5,000
Total overhead	2,320,000			

CONVENTIONAL COSTING

Total direct labour hours = $(25,000 \times 5) + (26,750 \times 4)$
 $= 125,000 + 107,000 = 232,000$ hours

Predetermined rate = $2,320,000 \div 232,000 = \text{\$10 per direct labour hour}$

UNIT COST — CONVENTIONAL	STANDARD	DELUXE
Direct material	80.00	120.00
Direct labour	50.00	40.00
Overhead $\$10 \times 5 \text{ hrs} / \$10 \times 4 \text{ hrs}$	50.00	40.00
Total unit cost	180.00	200.00

ACTIVITY-BASED COSTING

Setups	$100,000 \div (15 + 10)$	= \$4,000 per setup
Machining	$1,000,000 \div (60,000 + 40,000)$	= \$10 per machine hour
Engineering	$800,000 \div (25,000 + 15,000)$	= \$20 per engineering hour
Packing	$420,000 \div (20,000 + 5,000)$	= \$16.80 per packing order

OVERHEAD APPLIED	STANDARD	DELUXE
Setups × \$4,000	60,000	40,000
Machining × \$10	600,000	400,000
Engineering × \$20	500,000	300,000
Packing × \$16.80	336,000	84,000
Total overhead	1,496,000	824,000
÷ units	25,000	26,750
Overhead per unit	59.84	30.80

Check: $1,496,000 + 824,000 = \$2,320,000$, the total overhead. ABC redistributes overhead; it never changes the total.

UNIT COST – ABC	STANDARD	DELUXE
Direct material	80.00	120.00
Direct labour	50.00	40.00
Overhead	59.84	30.80
Total unit cost	189.84	190.80

The cost of distortion

Distortion = Conventional cost – ABC cost

Standard	$180.00 - 189.84 = -9.84$	→ UNDERCOSTED by \$9.84 under conventional
Deluxe	$200.00 - 190.80 = +9.20$	→ OVERCOSTED by \$9.20 under conventional

UNDERCOSTED $ABC > CONVENTIONAL$

OVERCOSTED $CONVENTIONAL > ABC$

The selling price set from the conventional cost is **too low to cover the product properly**, so the product earns less than management believes and may be sold at a loss. Every extra unit sold makes the position worse, and the loss

The selling price is **inflated**. In a competitive market that is actively dangerous: customers go elsewhere, volume falls, and the company may abandon a product that was in fact profitable.

is concealed by the products subsidising it.

READ THIS EXAMPLE CAREFULLY – IT DOES NOT DO WHAT THE TEXTBOOK STORY USUALLY DOES

The stock narrative is that ABC exposes a low-volume product as a hidden loss-maker. Here it does the opposite. Standard and Deluxe sell in similar volumes, but Standard consumes **four times** the packing orders and considerably more engineering. Standard, not Deluxe, is the complexity hog, and conventional costing had been quietly subsidising it out of Deluxe.

So a marketing manager arguing for more emphasis on Deluxe was right, and more right than the conventional numbers showed: Deluxe's unit cost falls from \$200 to \$190.80 once overhead follows the activities.

The general rule still holds — volume-based systems undercost the products that consume disproportionate support activity. The lesson is that you have to *look at the drivers* to know which product that is, rather than assuming it is the low-volume one.

MARGINS, BEFORE AND AFTER

	Conventional		ABC	
Standard	450	180 = 270 (60.0%)	450	189.84 = 260.16 (57.8%)
Deluxe	650	200 = 450 (69.2%)	650	190.80 = 459.20 (70.6%)

The conclusion did not reverse, but its size changed and its basis became defensible. That is the ordinary outcome of an ABC exercise, and a more honest thing to expect than a dramatic reversal.

What ABC is good for, and what it costs

BENEFITS

More accurate product costs where overhead is large and diverse, and where products differ in complexity.

Better pricing, product mix and discontinuation decisions.

It makes the *cost of complexity* visible — small batches, many variants, short runs — which is management information in its own right.

Identifying activities is the first step to eliminating the ones that add no value.

LIMITATIONS

Expensive to design and maintain: activities have to be identified, drivers chosen, and consumption measured.

Facility-level costs still have to be allocated arbitrarily, or excluded.

Driver measurement is itself an estimate, so the extra precision is partly illusory.

It is not required for external reporting, so it is a second system running alongside the statutory one.

Most of the reported cost may still be unavoidable in the short run — which matters enormously for Chapter Eight.

WHEN ABC IS WORTH DOING

Overhead must be a large proportion of total cost; products must differ materially in how they consume support activities; and the decisions at stake must be big enough to justify the system. A single-product plant with 5% overhead gains nothing from ABC. A plant making forty variants where overhead is 60% of cost gains a great deal.

Every system so far has treated fixed manufacturing overhead as a cost of the product. The next chapter asks whether it should be — and shows what happens to reported profit when the answer changes.

Absorption and variable costing; service costing

Two topics, joined by a single theme: what belongs in the cost of a unit. The first half turns on the treatment of fixed manufacturing overhead, and shows how the same business reports two different profits depending on the answer. The second half deals with businesses whose output cannot be stored at all.

The one difference that generates everything else

ABSORPTION COSTINGFULL COSTING	VARIABLE COSTINGDIRECT OR MARGINAL COSTING
Fixed manufacturing overhead is a product cost . It attaches to units and sits in inventory until they are sold.	Fixed manufacturing overhead is a period cost . It is expensed in full in the period incurred, whatever happens to the units.
Unit product cost = DM + DL + variable overhead + fixed overhead	Unit product cost = DM + DL + variable overhead

Everything else is identical. Variable and fixed selling and administrative costs are period costs under both. Direct materials, direct labour and variable manufacturing overhead are product costs under both.

THE CONSEQUENCE

If fixed manufacturing overhead is buried in inventory, then **producing more than you sell moves fixed cost out of this period's expenses and onto the balance sheet**. Reported profit rises without a single extra sale. Variable costing removes that possibility, because fixed overhead is expensed whatever the production volume.

The two income statement formats

ABSORPTION – THE STATUTORY FORMAT	VARIABLE – THE CONTRIBUTION FORMAT
Sales – Cost of sales <i>at full unit cost</i> = Gross profit – Selling and administrative expenses = Profit	Sales – Variable manufacturing cost of sales – Variable selling and administrative = Contribution margin – Fixed manufacturing overhead – Fixed selling and administrative = Profit

The variable format is the one CVP analysis needs, because it separates costs by behaviour and produces a contribution margin. The absorption format separates costs by *function* and produces a gross profit, which is what external reporting requires but which no decision can be made from directly.

WORKED • BRINDABELLA TOOLS, THREE YEARS

Selling price \$50. Variable manufacturing cost \$20 per unit. Fixed manufacturing overhead \$120,000 a year, applied at \$12 per unit on normal capacity of 10,000 units. Variable selling and administrative \$3 per unit sold. Fixed selling and administrative \$60,000. Production is 10,000 units every year. Sales are 8,000 in year 1, 10,000 in year 2 and 12,000 in year 3. No opening inventory in year 1.

Absorption unit product cost = $20 + 12 = \$32$

Variable unit product cost = **\$20**

YEAR 1 IN FULL – PRODUCTION 10,000, SALES 8,000

ABSORPTION COSTING	\$	VARIABLE COSTING	\$
Sales $8,000 \times 50$	400,000	Sales	400,000
Cost of sales $8,000 \times 32$	(256,000)	Variable manufacturing $8,000 \times 20$	(160,000)
		Variable selling $8,000 \times 3$	(24,000)
Gross profit	144,000	Contribution margin	216,000
Selling and admin $24,000 + 60,000$	(84,000)	Fixed manufacturing overhead	(120,000)
		Fixed selling and admin	(60,000)
Profit	60,000	Profit	36,000

Absorption profit is \$24,000 higher. The reason: 2,000 unsold units carry \$12 of fixed overhead each into closing inventory instead of into this year's expenses.

Closing inventory, absorption $2,000 \times 32 = 64,000$
 Closing inventory, variable $2,000 \times 20 = 40,000$
 Difference = **24,000** = the profit difference

ALL THREE YEARS

	YEAR 1	YEAR 2	YEAR 3
Production (units)	10,000	10,000	10,000
Sales (units)	8,000	10,000	12,000
Change in inventory (units)	+2,000	—	-2,000
Absorption profit	60,000	90,000	120,000
Variable costing profit	36,000	90,000	144,000
Difference	+24,000	nil	-24,000

THE RECONCILIATION RULE

$$\begin{aligned} &\text{Absorption profit} - \text{Variable costing profit} \\ &= \text{Fixed manufacturing overhead per unit} \times \text{Change in inventory in units} \end{aligned}$$

IF	THEN	BECAUSE
Production > sales <i>inventory rises</i>	Absorption profit is higher	Fixed overhead is deferred into inventory
Production = sales <i>inventory unchanged</i>	The two are equal	All fixed overhead incurred is also expensed
Production < sales <i>inventory falls</i>	Absorption profit is lower	Fixed overhead deferred in earlier periods is released into this period's cost of sales

Over the life of the business the two totals are identical — \$216,000 across the three years under either method. The methods disagree about *timing*, never about total profit.

Which to use

ABSORPTION COSTING	VARIABLE COSTING
Required for external financial reporting and for tax, because accounting standards require inventory to be measured at full production cost. Consistent with the matching principle in the sense that manufacturing costs follow the units they made.	Better for decisions. Contribution margin is visible, so CVP analysis, short-run pricing, product emphasis and special-order decisions all work directly from it. Profit moves with sales, not with production, so it is harder to manipulate and easier to interpret. Fixed overhead is reported as one lump, where management can see and control it, rather than spread invisibly across units.

THE BEHAVIOURAL PROBLEM WITH ABSORPTION COSTING

If a manager's bonus depends on reported profit, absorption costing gives them a lever that has nothing to do with selling anything: **produce more**. Each extra unit made and not sold parks \$12 of fixed overhead in inventory and lifts reported profit, while consuming cash, filling the warehouse and risking obsolescence.

This is the clearest example in the subject of a measurement system creating an incentive nobody intended, and it is the argument for reporting internally on a variable costing basis even though the statutory accounts must use absorption. It returns in Chapter Ten: any target that can be met by doing the wrong thing eventually will be.

Costing a service

Services differ from manufacturing in ways that matter for costing: they are **intangible**, they cannot be **stored** (so there is no inventory and no absorption-versus-variable problem at all), production and consumption are **simultaneous**, and the customer is often part of the process. Labour is usually the dominant cost, and much of what would be overhead in a factory is here the main event.

THE SERVICE PROCESS TYPES

TYPE	CHARACTER	COSTING APPROACH
Professional services	Low volume, highly customised, high contact time with the customer, considerable employee discretion. People are the key input. Audit, law, architecture, management consulting, specialist medicine.	Job costing. Each engagement is a job.
Service shops	Between the two — moderate volume, some customisation, a mix of front and back office. Retail banking branches, hotels, vehicle servicing.	Hybrid, leaning on either approach depending on how distinguishable the jobs are.
Mass services	High volume of standardised services, mostly repetitive processes, equipment is a key input with an emphasis on minimising employee discretion, and little contact time — the main production happens in the back office. Transaction processing, telecommunications, airlines, mail sorting.	Average cost. Accumulate total cost and divide by services delivered — process costing in all but name.

THE CLASSIFICATION TRAP

Asked which statement is *not* a description of a mass service provider, the answer is “*it has some customisation of services, with mostly repetitive processes*”. Customisation is the marker of a professional service or a service shop; mass services are defined by *standardisation*. The other three descriptions — high volume of standardised services, equipment as the key input with minimal employee discretion, and little customer contact with production in the back office — all describe mass service correctly.

THE TWO PROCEDURES

PROFESSIONAL SERVICES	MASS SERVICES
1. Identify the direct costs of the job — chiefly professional labour by time sheet. 2. Allocate overhead to the job: a. aggregate overhead into cost pools; b. identify an overhead cost driver; c. calculate a predetermined overhead rate; d. apply overhead to the service.	1. Accumulate the total cost of delivering the service for the period. 2. Calculate the average cost per service: <i>Total cost ÷ number of services delivered</i>

Two classification points that recur. In a professional practice, wages paid to audit clerks are **direct labour**, because the hours are traced to a client engagement — not overhead, despite the absence of a factory. And because there is no inventory, every cost in a pure service business is a **period cost**; the product-cost machinery of Part Two has nothing to attach to.

Part Two has produced product costs by five different routes. Part Three asks the awkward question none of those systems answer: of all the cost now attached to a product, how much of it should influence a decision?

PART THREE

Costs for decisions

Part Two built full product costs. This part largely throws them away. A decision is affected only by cash flows that change because of it, and a full product cost is stuffed with allocations that will not change whatever is decided. Learning which costs to ignore is harder than learning to calculate them.

Information for decisions

One principle, applied to five recurring decisions. The principle is relevance: a cost or revenue matters only if it lies in the future and differs between the alternatives. Everything else — however large, however carefully allocated, however painful — is noise.

What makes a cost relevant

THE TWO TESTS, BOTH OF WHICH MUST BE MET

1. **It must be a future cost or revenue.** Anything already incurred is gone and cannot be changed by a decision made now.
2. **It must differ between the alternatives.** A cost that is identical under every option cannot influence the choice, so carrying it through the analysis only adds arithmetic and risk.

TERM	MEANING, AND HOW TO TREAT IT
Sunk cost	A past cost that will remain the same no matter which alternative is chosen. Always irrelevant. The research already spent, the machine already bought, the training already delivered.
Opportunity cost	The potential benefit given up when choosing one alternative requires sacrificing another. Always relevant , and never in the accounting records — which is exactly why it gets missed.
Avoidable cost	A cost that will not be incurred if an alternative is chosen. Relevant.
Unavoidable cost	A cost that continues regardless. Irrelevant, including most allocated overhead.
Differential (incremental) cost	The difference in total cost between two alternatives. This is what the analysis is trying to find.
Committed cost	A future cost the organisation is already contractually bound to. Future, but does not differ — so irrelevant.

WHY SUNK COSTS ARE SO HARD TO IGNORE

Because ignoring them feels like admitting a mistake. A business that has spent \$2 million developing a product will resist abandoning it, reasoning that the \$2 million would be “wasted”. But the \$2 million is already gone under both options; the only question is whether the *future* cash flows of continuing beat those of stopping. The \$2 million is a fact about the past, not an argument about the future.

The formal answer, worth having ready: *a sunk cost is a past cost that will remain the same no matter which of the alternatives under consideration is chosen; sunk costs should be ignored when making decisions, because one cannot change what has already happened.*

Decision 1 — Make or buy

The question is never “is the supplier’s price below our unit cost?” It is “is the supplier’s price below the cost we would *avoid* by not making it, adjusted for anything better we could do with the freed capacity?”

THE METHOD

1. Strip the unit cost down to its **avoidable** components. Variable manufacturing costs are avoidable. Fixed costs are avoidable only to the extent stated — allocated overhead that merely moves to other products is not.
2. Add any **opportunity cost** — the contribution the freed capacity could earn elsewhere.
3. Compare that total with the purchase price.

WORKED • XEBEX PTY LTD – THE MAXIMUM PRICE

Xebex uses 100,000 components a year. The annual cost of producing them: direct variable manufacturing costs \$300,000; direct fixed manufacturing costs \$100,000; allocated overhead \$50,000. If production ceased, direct fixed manufacturing costs would fall by 80%. Below what purchase price should Xebex buy?

Avoidable costs		
Direct variable manufacturing	300,000	
Direct fixed manufacturing 100,000 × 80%	80,000	
Allocated overhead	0	<i>continues regardless</i>

	380,000	

Avoidable cost per unit = $380,000 \div 100,000 = \$3.80$

Xebex should buy if the purchase price is below \$3.80. Note the two deliberate traps. The **allocated overhead** of \$50,000 is not avoided — it is reallocated to other products, so total company cost is unchanged. And only **80%** of the direct fixed cost goes: the remaining \$20,000 continues, so it belongs with the unavoidable costs.

The full unit cost of \$4.50 ($450,000 \div 100,000$) is the number a costing system would report, and using it would lead Xebex to buy at any price under \$4.50 — losing money on every unit priced between \$3.80 and \$4.50.

WORKED • FOTO COMPANY — WITH AN OPPORTUNITY COST

Foto makes 50,000 units a year of a part. The reported unit cost:

	\$ PER UNIT
Direct materials	12.00
Direct labour	10.10
Variable manufacturing overhead	2.00
Fixed manufacturing overhead	14.10
Unit product cost	38.20

An outside supplier offers the parts at \$37.30 each. If Foto buys, the facilities now used could make more of a product in high demand, adding \$310,000 a year of contribution margin. If the parts are bought, \$9.70 of the fixed manufacturing overhead per part would continue to be incurred and would be reallocated to other products.

(a) How much of the \$38.20 is relevant?

Direct materials	12.00	relevant
Direct labour	10.10	relevant
Variable manufacturing overhead	2.00	relevant
Fixed manufacturing overhead	14.10 – 9.70	4.40 relevant portion only

Relevant cost of making, per unit	28.50	

(b) The net advantage or disadvantage of buying

Cost to buy	50,000 × 37.30	= (1,865,000)
Contribution from freeing the facility		= 310,000

Net cost of buying		= (1,555,000)
Cost to make	50,000 × 28.50	= (1,425,000)
Net disadvantage of purchasing		= (130,000)

(c) The maximum Foto should pay per unit

Relevant cost of making	28.50
+ Opportunity cost	310,000 ÷ 50,000
	6.20

Maximum acceptable price	\$34.70

So Foto should keep making the part: the offer of \$37.30 exceeds the \$34.70 ceiling. Notice that the offer *is* below the reported unit cost of \$38.20 — which is precisely why the reported cost must

not be used. The \$9.70 of fixed overhead that continues anyway is the difference between a right answer and a \$130,000 mistake.

Decision 2 — Special orders

A one-off order at below the normal price. If there is **spare capacity**, accept it if the price exceeds the incremental cost — usually the variable cost of the order, plus any specific setup or delivery cost it causes. Fixed overhead is unaffected and irrelevant.

If the order requires **displacing existing sales**, the lost contribution margin is an opportunity cost and must be added to the incremental cost.

THE QUALITATIVE HALF OF THE ANSWER

A special order can be quantitatively attractive and strategically foolish. Regular customers may learn of the lower price and demand it. The order may commit capacity that a better customer needs later. In a market with any transparency, a price accepted once becomes a price expected. A complete answer notes the arithmetic *and* the precedent.

Decision 3 — Add or drop a product or segment

The test is whether the segment's contribution margin covers the fixed costs that would actually **disappear** if it were dropped. Fixed costs split into two kinds, and only one of them counts:

AVOIDABLE (TRACEABLE) FIXED COSTS	UNAVOIDABLE (COMMON) FIXED COSTS
Arise because the segment exists and end when it ends — a dedicated supervisor, product-specific advertising, a leased machine used only for it. Relevant.	Continue and are merely reallocated to the survivors — head office, factory rent, the general manager's salary. Irrelevant.
Segment contribution margin – Avoidable fixed costs = Segment margin – drop the segment only if this is negative	

A segment showing a loss after a share of common fixed costs may still be worth keeping: dropping it removes its contribution while leaving the common costs to be absorbed by fewer products, which usually makes total profit worse, not better.

Decision 4 — Sell now or process further

A **joint process** produces two or more products simultaneously from common inputs. The point at which they become separately identifiable is the **split-off point**, and the cost incurred up to it is the **joint cost**.

JOINT COSTS ARE ALWAYS IRRELEVANT AFTER SPLIT-OFF

Generally, joint costs are not relevant to a decision made after split-off **because they do not change regardless of any decision**. They have been incurred to produce the whole set of products, they cannot be avoided by declining to process one of them further, and any allocation of them between products is arbitrary.

The only relevant figures are the **incremental revenue** from further processing and the **incremental cost** of doing it. Decide each product separately.

WORKED • LIDO PRODUCTS

Products A and B come from a joint process costing \$80,000. Product A: 5,000 units, sellable at split-off for \$20 each, or processed further at an additional cost of \$20,000 and sold for \$25. Product B: 10,000 units, sellable at split-off for \$15 each, or processed further at an additional cost of \$20,000 and sold for \$16.

Product A

Incremental revenue	$5,000 \times (25 - 20) =$	25,000
Incremental cost		= (20,000)
Effect on profit		= +5,000 – process further

Product B

Incremental revenue	$10,000 \times (16 - 15) =$	10,000
Incremental cost		= (20,000)
Effect on profit		= -10,000 – sell at split-off

The \$80,000 joint cost appears nowhere in either calculation, and no allocation of it between A and B would change either answer. Any question that offers you a “share of joint cost” is offering you a distractor.

Decision 5 — Keep or replace equipment

The book value of the old machine is a **sunk cost** and irrelevant, as is any accounting loss on disposal. Relevant: the disposal proceeds now, the cost of the new machine, and the difference in future operating costs. This is the decision where sunk-cost reasoning is most tempting, because writing off a large book value looks like destroying value — when in fact the value was destroyed when the machine was bought.

Qualitative factors

Every decision in this chapter has a quantitative answer and a set of considerations the arithmetic cannot hold. A complete analysis states both.

- **Quality and reliability** of an outside supplier, and the cost of failure.
- **Loss of control** over a process, and of the internal expertise to resume it — a make-or-buy decision is easy to reverse on paper and hard in practice.
- **Employee and community effects** of closing a line, and the effect on the morale of those who remain.

- **Customer reaction** and reputation.
- **Long-run capacity**: capacity given up cheaply today may be expensive to buy back.
- **Supplier dependence**: the price is \$37.30 this year, and the supplier knows what your alternative costs once you have dismantled the line.

Relevance decides whether an alternative is worth taking. It does not decide what to do when every alternative is worth taking and there is not enough of something to do them all.

Product mix decisions

A short chapter with one idea in it, and the idea is easy to state and easy to get wrong. When a resource is scarce, the products to make are not those with the highest contribution margin per unit, but those with the highest contribution margin per unit *of the scarce resource*.

The constrained resource

A **constraint** or **limiting factor** is any resource in shorter supply than the demand for the products that need it: machine hours, direct labour hours, a scarce material, floor space, a licence, a bottleneck process. Where nothing is scarce, make everything with a positive contribution margin and the problem disappears.

THE RULE

Contribution margin per unit of constrained resource
= Unit contribution margin ÷ Units of the constrained resource per unit of product

Rank the products by that figure, highest first. Allocate the scarce resource to the highest-ranked product up to the limit of its demand, then to the next, and so on until the resource is exhausted.

Demand matters as much as ranking. There is no point allocating the whole constraint to the best product if only 3,000 units of it can be sold.

WORKED • THREE PRODUCTS, ONE BOTTLENECK

A firm produces X, Y and Z with unit contribution margins of \$4, \$5 and \$10. Only 5,000 machine hours are available. Machine hours required per unit are 1, 1 and 5 respectively. Demand is 1,000 units of X, 3,000 of Y and 2,000 of Z.

	X	Y	Z
Unit contribution margin	\$4	\$5	\$10
Machine hours per unit	1	1	5
Contribution per machine hour	\$4	\$5	\$2
Ranking	2nd	1st	3rd
Maximum demand (units)	1,000	3,000	2,000

Allocate the 5,000 machine hours in rank order:

Y	3,000 units × 1 hr = 3,000 hrs	full demand met
X	1,000 units × 1 hr = 1,000 hrs	full demand met

	4,000 hrs used, 1,000 hrs remaining	
Z	1,000 hrs ÷ 5 hrs per unit = 200 units	of the 2,000 demanded

Optimal mix: 3,000 Y, 1,000 X, 200 Z

Total contribution = (3,000 × 5) + (1,000 × 4) + (200 × 10)
 = 15,000 + 4,000 + 2,000 = **\$21,000**

Compare the naive answer of making Z first, because its \$10 unit contribution margin is the largest: 2,000 units of Z would consume 10,000 hours — twice the capacity available. Even 1,000 units of Z would eat the entire constraint for a contribution of \$10,000, less than half the optimum. Z looks best per unit and is worst per hour, because each unit devours five hours to earn its \$10.

The discussion version

WORKED • LAWNMATE COMPANY – THREE MANAGERS, ONE RIGHT ANSWER

LawnMate produces two lawnmowers, A and B. The company has spare machine hours but is always short of direct labour hours. Three managers offer views.

The sales manager: “There is a larger contribution margin on Product A, so we should manufacture only A.”

The factory manager: “Both products sell well and both take the same machining time, so we should produce equal quantities to use up factory capacity.”

The finishing manager: “Product B takes only half as long in direct labour as A, and its contribution margin is not much less. So concentrate on B, produce the maximum possible, and give any surplus capacity to A.”

The principle. To determine the optimal product mix, managers should maximise the profit generated from the limited resource. Here direct labour is the limiting factor, so the product with the higher **contribution margin per direct labour hour** takes priority.

MANAGER	VERDICT
Sales manager — <i>wrong</i>	A larger contribution margin <i>per unit</i> does not guarantee a larger contribution margin <i>per direct labour hour</i> . Since labour is the constraint, per-unit ranking is the wrong ranking. It is also wrong to make only A even if A did rank first: A's demand may not absorb the whole constraint.
Factory manager — <i>wrong</i>	The firm has spare machine hours, so machine time is not a constraint at all and comparing the products' machining requirements is pointless. Equal quantities is an arbitrary rule that optimises nothing, and “using up factory capacity” is not an objective — profit is.
Finishing manager — <i>right</i>	The reasoning is correct because it is based on contribution margin per direct labour hour : B uses half the labour for nearly the same contribution, so its contribution per labour hour is close to double A's. Producing the maximum B that can be sold and giving the residual capacity to A is exactly the ranking method.

WHY THE FINISHING MANAGER'S ANSWER IS STILL INCOMPLETE

“Produce the maximum number possible” must mean the maximum that can be *sold*, not the maximum that can be *made*. Ignoring the demand ceiling is how the ranking method goes wrong in practice: a firm fills the constraint with its best-ranked product and then cannot sell it, having declined orders for the second-ranked one.

More than one constraint

The ranking method works for a **single** constraint. With two or more scarce resources, ranking can give conflicting answers — the best product per labour hour may be the worst per kilogram of material — and the problem becomes one of **linear programming**: maximise total contribution subject to each constraint, solved graphically for two products or by computer for more. Chapter Twelve does exactly this with a spreadsheet solver.

Beyond the arithmetic

- **Relax the constraint instead of optimising within it.** Overtime, a second shift, subcontracting, or a small capital purchase may lift the limit. The value of one more hour of the bottleneck is its contribution margin per hour — which is exactly what you should be willing to pay to obtain it.
- **Customer relationships.** Dropping a low-ranked product may cost sales of a high-ranked one if customers buy the range.
- **Market position.** A product with a poor contribution per hour may be the entry point to an account, or may keep a competitor out.
- **The constraint moves.** Once labour is relieved, something else becomes the bottleneck, and the ranking changes with it. The analysis has to be redone, not inherited.

Decisions of this kind are made one at a time. Part Four takes the whole organisation for a year at once, sets out what it intends to do, and then explains what happened instead.

PART FOUR

Planning and control

The budget says what the organisation intends to do; standard costing explains what it did instead. The two chapters are one system: the second is meaningless without the first, and the first achieves nothing without the second.

Budgeting

A budget is a plan expressed in numbers. The master budget is a linked set of schedules in which each one feeds the next, so a single change in the sales forecast propagates through production, materials, labour, cash and the budgeted statements. The mechanics are simple arithmetic; the difficulty is the sequence and the behaviour of the people producing the numbers.

DEFINITION

Budgeting is the process of preparing detailed plans showing the financial consequences of an organisation's operating activities for a future period. Departmental budgets are aggregated to obtain a budget for the organisation as a whole — the **master budget**.

What budgeting is for


PURPOSE	WHAT IT MEANS IN PRACTICE
Planning	Forcing management to think ahead about volumes, prices, capacity and cash, before events force the thinking on them.
Coordination	Making the functions consistent. Production cannot plan for 10,000 units while sales plans for 8,000 and purchasing buys for 12,000.
Communication	Telling each manager what is expected of them, in numbers rather than adjectives.
Motivation	A target that is agreed and attainable improves effort. One that is imposed and unattainable destroys it.
Control	Providing the benchmark against which actual results are compared — which is the whole of Chapter Eleven.
Performance evaluation	Judging managers against what was agreed, for the things they control.
Resource allocation	Deciding, in advance and explicitly, who gets what.

TWO PURPOSES IN TENSION

Motivation wants a target that is challenging but achievable. **Planning** wants the most likely outcome, so that cash and capacity are sized correctly. These are different numbers, and an organisation that uses one budget for both is compromising one of them. Where a budget is also the basis of a bonus, the tension becomes an incentive to distort — see budgetary slack below.

The master budget, in sequence

Everything begins with sales, because sales volume determines production, production determines materials and labour, and all of it determines cash. Get the order wrong and the schedules cannot be built.



SCHEDULE	WHAT IT PRODUCES, AND WHAT IT NEEDS
1 · Sales budget	Units and revenue by period. The foundation, and the least reliable schedule in the set, because everything else is built on a forecast of other people's behaviour.
2 · Production budget <i>or purchases budget, for a retailer</i>	Units to be made (or bought). Driven by sales plus the desired closing inventory policy.
3a · Direct materials budget	Quantity of material needed for production, then quantity to <i>purchase</i> after allowing for the material inventory policy, then cost.
3b · Direct labour budget	Hours required from the production volume, then cost at the standard rate. Feeds capacity planning — whether the hours exist.
3c · Manufacturing overhead budget	Variable overhead from the activity level, fixed overhead as a period total. Produces the predetermined overhead rate of Chapter Three.
3d · Selling and administrative budget	The period costs, split into variable and fixed.
4 · Cash budget	Receipts, payments, and the resulting balance by period. The schedule that most often changes a plan, because a profitable plan can be unfundable.
5 · Budgeted income statement and balance sheet	The whole plan expressed as financial statements, so that the plan can be judged the way the results will be.

The two formulas that do most of the work

PRODUCTION AND PURCHASES

Production budget (units)

- Budgeted sales in units
- + Desired closing inventory of finished goods
- Opening inventory of finished goods
- = **Units to be produced**

Purchases budget (units) – for a retailer, or for raw materials

- Units required for sale or production
- + Desired closing inventory
- Opening inventory
- = **Units to be purchased**

Both are the same identity rearranged: *what you need, plus what you want left, less what you already have*. The opening inventory of one period is the closing inventory of the previous one, which is what links the periods together and what makes a multi-period question solvable at all.

WORKED • BEACTIVE SPORTING GOODS

BeActive sells tandem bicycles. From the quarterly sales forecast:

	EXPECTED SALES (UNITS)	TARGET CLOSING INVENTORY (UNITS)
April	1,400	315
May	1,575	412
June	1,650	425

Cost of one bicycle is \$125.

How many units should be purchased in May?

Sales in May	1,575
+ Target closing inventory, 31 May	412
- Opening inventory, 1 May = April's closing	(315)

Purchases in May	1,672 units

What dollar amount should be budgeted for June purchases?

Sales in June	1,650
+ Target closing inventory, 30 June	425
- Opening inventory, 1 June = May's closing	(412)

Purchases in June	1,663 units
× \$125	\$207,875

The trap in both parts is using the *same month's* target inventory as the opening figure, or forgetting to deduct opening inventory at all — which for June would give 2,075 units and \$259,375. The opening inventory is always the previous period's closing inventory, and the first period of a question is the only one where it has to be given to you.

WORKED • FROM PRODUCTION UNITS TO MATERIAL PURCHASES

A manufacturer plans to produce 12,000 units in July. Each unit requires 4 kg of material at \$6 per kg. Opening material inventory is 9,000 kg, and the policy is to hold closing material inventory equal to 20% of the following month's production requirement. August production is planned at 14,000 units.

Material required for July production	12,000 × 4 kg	= 48,000 kg
+ Desired closing inventory	20% × (14,000 × 4 kg)	= 11,200 kg
– Opening inventory		= (9,000) kg

Material to purchase in July		50,200 kg
× \$6 per kg		\$301,200

Note the two-stage structure: the production budget gives *units of product*, which convert to *units of material*, and only then does the inventory policy apply. Applying the 20% to units of product rather than kilograms of material is the standard error, and it is out by a factor of four.

The cash budget

The schedule that most often sends a plan back for revision, because profit and cash are different questions.

Opening cash balance	
+ Receipts	<i>cash sales; collections of receivables per the collection pattern; asset sales; new borrowings; capital contributions</i>
– Payments	<i>payments to suppliers per the payment pattern; wages; overheads in cash; capital expenditure; interest; tax; dividends</i>
= Closing cash balance	– compared against any minimum balance required

TWO THINGS THAT SEPARATE A CORRECT CASH BUDGET FROM A WRONG ONE

Timing, not accrual. A sale in May collected in July is a May revenue and a July receipt. Depreciation is an expense and never a payment. The cash budget is built from the collection and payment patterns, not from the income statement.

Non-cash items are excluded entirely. Depreciation, doubtful debts expense and accruals appear in the budgeted income statement and nowhere in the cash budget. Including depreciation as a payment is the most common single error in the schedule.

Approaches to preparing a budget

APPROACH	HOW IT WORKS, AND WHAT IT RISKS
Imposed (top-down)	Senior management sets the targets and passes them down. Fast, consistent with strategy, and poorly accepted by the managers

APPROACH	HOW IT WORKS, AND WHAT IT RISKS
	expected to deliver it — who did not agree to it and may not believe it.
Participative (bottom-up)	Managers who will be held to the budget help set it. Better information, because they know their operations; better acceptance, because they own the number. Slower, and it opens the door to slack.
Incremental	Last year's budget plus an adjustment. Quick and simple; perpetuates whatever inefficiency is already embedded, because nothing is ever re-justified.
Zero-based	Every item must be justified from nothing each cycle. Exposes inherited waste; expensive and exhausting to run every year.
Rolling (continuous)	A twelve-month horizon maintained by adding a period as each one closes. Always current; consumes budgeting effort continuously rather than annually.

Budgetary slack

DEFINITION

Budgetary slack is the difference between the revenue or cost projection a person provides and a realistic estimate of that revenue or cost.

A departmental manager who deliberately underestimates the level of sales activity by 10,000 units is creating budgetary slack. It is not an allowance for uncertainty, not appropriate caution, and not budgetary participation — it is a deliberate distortion of the estimate, and it has the precise name given above.

HOW MANAGERS CREATE IT	WHY THEY DO IT
Underestimating revenue or overestimating costs , so that the budget target is easier to achieve.	Performance looks better if you can “beat the budget”, particularly where a bonus depends on it. To cope with uncertainty — a buffer against events outside the manager's control. To compete for limited resources — a larger cost budget secures a larger share of what is going.

WHY SLACK IS A SYSTEM PROBLEM, NOT A CHARACTER PROBLEM

Slack is the predictable response to a system that uses one number for both planning and reward. If beating the budget is what gets rewarded, then making the budget easy to beat is rational behaviour, and the organisation has asked for it.

The remedies are structural: separate the planning forecast from the performance target; reward forecast *accuracy* as well as outcome; use rolling forecasts so a single annual number matters less; and review estimates against external benchmarks rather than against last year's budget. Exhortation does not work.

Other behavioural problems worth naming

- **Spending to the budget.** A department with money left in December spends it, because an underspend gets the budget cut next year. The remedy is to stop punishing underspends.
- **Short-termism.** Meeting this year's target by deferring maintenance, training or advertising — all of which show up as this year's saving and next year's problem.
- **Gaming the timing.** Pulling sales forward into the budget period or pushing costs out of it, changing the record without changing the business.
- **Budget as a weapon.** Where variances are used to allocate blame rather than to understand causes, managers stop reporting problems early — and the information system quietly stops working.

A budget without a comparison is a wish. The next chapter sets the standards precisely enough that the difference between plan and outcome can be split into its causes and handed to the person who can explain it.

Standard costing for control

A standard cost is what a unit *should* cost. Comparing it with what the unit did cost gives a total difference, and the point of the chapter is to split that difference into pieces small enough to be somebody's responsibility. The arithmetic is four or six formulas; the skill is in the direction of each variance and in what it actually tells you.

DEFINITION

A **standard cost** is a budgeted cost, based on estimates of the material, labour and overhead resources that *should* be used to make one unit of product. It serves as the benchmark in the budgetary control system.

A standard has two components, and both can be wrong independently: a **price** or rate standard (what an input should cost per unit) and a **quantity** standard (how much of the input a unit of product should need). That split is what makes the variance analysis possible — each variance isolates one component.

The flexible budget

COMPARE AGAINST OUTPUT ACTUALLY ACHIEVED

A **static** budget is set for one planned volume. Comparing actual costs at 4,800 units against a static budget for 5,000 units tells you almost nothing, because part of every difference is simply that fewer units were made.

A **flexible budget** restates the budget at the volume actually achieved: standard cost per unit × actual units produced. Every variance in this chapter is a comparison against the flexible budget, and that is why each formula contains a “standard quantity *allowed for actual output*” term.

The variance formulas

DIRECT MATERIAL AND DIRECT LABOUR

Material price variance purchased	$= (\text{Actual price} - \text{Standard price}) \times \text{Actual quantity}$
Material quantity variance	$= (\text{Actual quantity used} - \text{Standard quantity allowed}) \times \text{Standard price}$
Labour rate variance	$= (\text{Actual rate} - \text{Standard rate}) \times \text{Actual hours}$
Labour efficiency variance	$= (\text{Actual hours} - \text{Standard hours allowed}) \times \text{Standard rate}$

The pattern is consistent: a **price** variance holds quantity constant and prices the difference in rate; a **quantity** variance holds price constant at standard and values the difference in quantity. Never use actual price in a quantity variance — that would mix the two effects and double-count the price difference.

Note the asymmetry in the material formulas. The price variance uses quantity **purchased**, because that is when the price was incurred; the quantity variance uses quantity **used**, because that is when the usage happened. Where purchases and usage are equal the distinction disappears; where a question gives both figures, it is testing exactly this.

OVERHEAD

Variable overhead spending variance	$= \text{Actual variable overhead} - (\text{Actual hours} \times \text{Standard variable rate})$
Variable overhead efficiency variance	$= (\text{Actual hours} - \text{Standard hours allowed}) \times \text{Standard variable rate}$
Fixed overhead budget variance	$= \text{Actual fixed overhead} - \text{Budgeted fixed overhead}$
Fixed overhead volume variance	$= \text{Budgeted fixed overhead} - \text{Applied fixed overhead}$ $= (\text{Normal hours} - \text{Standard hours allowed}) \times \text{Standard fixed rate}$

The **volume variance** is the odd one out and worth understanding rather than memorising. It is not a spending difference at all: it arises purely because output differed from the level used to set the fixed overhead rate, so fixed overhead was applied at a rate that assumed a volume that did not happen. It measures capacity utilisation, not efficiency, and no one can be held responsible for it except whoever decided the production volume.

Favourable and unfavourable

THE CONVENTION

FAVOURABLE (F) — actual cost is **less** than standard. Profit is higher than the flexible budget predicted.

UNFAVOURABLE (U) — actual cost is **more** than standard. Profit is lower.

For every cost variance above, a **positive** result from the formula as written is unfavourable, because each formula is arranged as *actual minus standard*. Check the direction by asking whether the business spent more or less than it should have, rather than trusting the sign.

FAVOURABLE DOES NOT MEAN GOOD

A favourable material price variance obtained by buying inferior material causes unfavourable usage and labour efficiency variances later, and possibly warranty claims and lost customers after that. A favourable labour efficiency variance obtained by rushing produces scrap. A favourable variance is a difference from standard in a particular direction, and nothing more; interpreting it requires knowing how it was achieved.

WORKED • RIDGEWAY COMPONENTS, MONTH OF MARCH

Standard cost per unit:

	STANDARD	\$ PER UNIT
Direct material	3 kg @ \$8.00 per kg	24.00
Direct labour	2 hrs @ \$22.00 per hr	44.00
Variable overhead	2 hrs @ \$6.00 per hr	12.00
Fixed overhead	2 hrs @ \$9.00 per hr	18.00
Total standard cost		98.00

The fixed overhead rate is set on normal capacity of 5,000 units, that is 10,000 standard hours, giving a budgeted fixed overhead of \$90,000.

Actual results for March: **4,800 units produced**. Material purchased and used 15,000 kg costing \$114,000. Direct labour 9,400 hours costing \$211,500. Actual variable overhead \$58,000. Actual fixed overhead \$92,000.

SET UP THE STANDARDS ALLOWED FOR ACTUAL OUTPUT

Standard quantity of material allowed	= 4,800 × 3 kg = 14,400 kg
Standard hours allowed	= 4,800 × 2 hrs = 9,600 hours
Actual material price	= 114,000 ÷ 15,000 = \$7.60 per kg
Actual labour rate	= 211,500 ÷ 9,400 = \$22.50 per hour

DIRECT MATERIAL

Price variance	= (7.60 - 8.00) × 15,000 = -0.40 × 15,000 = 6,000 F
Quantity variance	= (15,000 - 14,400) × 8.00 = 600 × 8.00 = 4,800 U
	Net material 1,200 F
check: actual 114,000 against standard for output 14,400 × 8 = 115,200 → 1,200 F ✓	

DIRECT LABOUR

Rate variance	= (22.50 - 22.00) × 9,400 = 0.50 × 9,400 = 4,700 U
Efficiency variance	= (9,400 - 9,600) × 22.00 = -200 × 22.00 = 4,400 F
	Net labour 300 U
check: actual 211,500 against standard 9,600 × 22 = 211,200 → 300 U ✓	

VARIABLE OVERHEAD

Spending variance	= 58,000 - (9,400 × 6) = 58,000 - 56,400 = 1,600 U
Efficiency variance	= (9,400 - 9,600) × 6 = -200 × 6 = 1,200 F
	Net variable OH 400 U
check: actual 58,000 against applied 9,600 × 6 = 57,600 → 400 U ✓	

FIXED OVERHEAD

Budget variance	= 92,000 - 90,000	= 2,000 U
Volume variance	= 90,000 - (9,600 × 9 = 86,400)	= 3,600 U
	or (10,000 - 9,600) × 9 = 400 × 9 = 3,600 U	
	Net fixed OH	5,600 U
check: actual 92,000 against applied 86,400 → 5,600 U ✓		

THE TOTAL, RECONCILED

VARIANCE	\$	
Material price	6,000	F
Material quantity	4,800	U
Labour rate	4,700	U
Labour efficiency	4,400	F
Variable overhead spending	1,600	U
Variable overhead efficiency	1,200	F
Fixed overhead budget	2,000	U
Fixed overhead volume	3,600	U
Total variance	5,100	U

Actual total cost	114,000 + 211,500 + 58,000 + 92,000	= 475,500
Standard cost of actual output	4,800 × 98	= 470,400

		5,100 U ✓

ALWAYS RECONCILE TO THE TOTAL

The individual variances must sum to the difference between actual total cost and the standard cost of actual output. It is the only check the method offers, it catches sign errors and misapplied quantities, and a set of variances that does not reconcile is a set of variances you cannot defend.

Reading the variances together

Individual variances are rarely independent, and reading one on its own is how a variance report gets misinterpreted. The useful question is whether a pattern of variances has a single common cause.

WORKED • BOSAN PTY LTD – IS THERE A CONNECTION?

For the quarter, one manufacturing department reported: material price variance **unfavourable**; material quantity variance **favourable**; labour efficiency variance **favourable**; labour rate variance **negligible**. The favourable variances together exceeded the unfavourable price variance by a considerable amount. Output quality was as usual, and practically all the raw material used in the quarter was bought in the quarter.

The answer. A connection cannot be guaranteed, but one explanation is very plausible: **higher than standard quality material was purchased.**

- Paying above standard for better material produces the **unfavourable material price variance**.
- When that material was used, it was easier for labour and machines to process, so there was **less waste** — giving the **favourable material quantity variance**.
- Less rework and fewer stoppages meant fewer hours than standard for the output achieved — giving the **favourable labour efficiency variance**.
- Wage rates were unaffected, which is why the labour rate variance is negligible.

And the conclusion that matters: because the favourable variances exceed the unfavourable one, the purchasing decision was **net beneficial**, even though the buyer's own variance looks bad. A control system that judged the purchasing manager on the price variance alone would penalise them for a decision that made the company better off.

PATTERN	A PLAUSIBLE COMMON CAUSE
Material price F, material quantity U, labour efficiency U	Cheaper, poorer material was bought. It saved on price and cost more in waste and processing time. The mirror image of the BoSan case.
Labour rate U, labour efficiency F	More experienced or higher-graded staff were used: they cost more per hour and needed fewer hours.
Labour rate F, labour efficiency U	Cheaper or less experienced staff were used, taking longer.
Labour efficiency U and variable overhead efficiency U, of similar proportion	Not two problems. Variable overhead is applied on labour hours, so excess hours automatically create both. One cause, two variances.
Fixed overhead volume U with no other unfavourable variance	Production was below the normal capacity used to set the rate. Nothing was overspent; the plant was underused.

Investigating variances, and who is responsible

VARIANCE	USUALLY THE RESPONSIBILITY OF
Material price	Purchasing — subject to market movements outside its control
Material quantity	Production — unless caused by the material bought, as above
Labour rate	Human resources or production scheduling, through the grade of labour used
Labour efficiency	Production supervision — subject to material quality and machine condition
Variable overhead spending	The department incurring the overhead
Fixed overhead budget	Whoever authorised the spending
Fixed overhead volume	Whoever set the production volume — frequently sales, not production

Not every variance is worth investigating. **Management by exception** means setting a threshold — a dollar amount, a percentage of standard, or a statistical limit — and looking only at what breaches it. Investigation has a cost, and a variance smaller than the cost of investigating it should be left alone. Trends matter more than single periods: a small unfavourable variance repeating every month for a year is more informative than one large one.

The limits of standard costing

- **Standards go stale.** A standard set when material cost \$8 a kilogram measures nothing once the market price is \$11. Variances then report the obsolescence of the standard rather than the performance of the department.
- **It suits stable, repetitive production.** In a business with short product lives, high customisation or continuous process improvement, there may be no stable standard to set.
- **It focuses on cost, and only on manufacturing cost.** Quality, delivery reliability, customer satisfaction and design improvement do not appear in any variance.
- **It can drive the wrong behaviour.** Buying in bulk to secure a favourable price variance builds inventory nobody needs. Keeping a machine running to avoid an unfavourable volume variance produces goods nobody ordered. Any measure that can be met by doing the wrong thing eventually will be.
- **Favourable is not good and unfavourable is not bad** — the point is worth making twice. A variance is a question, not a verdict.

Every schedule, allocation and variance in this book is arithmetic that nobody does by hand any more. The last chapter is about doing it in the tool it is actually done in — and doing it so that somebody else could check it.

PART FIVE

The tool

None of the arithmetic in this book is done by hand in practice. It is done in a spreadsheet, and a spreadsheet that cannot be checked is worse than no spreadsheet at all — because it produces confident numbers nobody can defend.

Excel for management accounting

A NOTE ON THIS CHAPTER

This is the one chapter written from the subject's stated outcome — the ability to use Excel for managerial accounting planning, control and evaluation — rather than from a set of lecture materials. The functions and techniques below are the ones the calculations in the preceding eleven chapters actually require, and the model-building discipline is what makes such work reviewable. Check the specific functions against the demonstration videos, which will settle exactly which are expected.

Two things are being learned at once here, and only one of them is Excel. The first is a handful of functions. The second, and the more valuable, is how to build a model whose logic a reviewer can follow — which is the difference between an analysis and a number.

Structure before formulas

SEPARATE INPUTS, CALCULATIONS AND OUTPUTS

Every assumption goes in **one** place, in its own clearly labelled block, and appears **nowhere else**. Every other cell refers to it. Nothing is typed into a formula.

INPUTS	Selling price per unit	B4	50
	Variable cost per unit	B5	20
	Fixed costs	B6	120,000
	Units sold	B7	8,000
CALCULATIONS	Unit contribution margin	B10	=B4-B5
	Total contribution	B11	=B10*B7
	Profit	B12	=B11-B6
	Break-even units	B13	=B6/B10
OUTPUTS	the summary a reader sees, all of it referring back to the block above		

The test of a good model: **change one input and everything that depends on it changes correctly, with no other edits**. If updating the selling price means hunting through formulas for the number 50, the model is broken even if today's answer is right.

HARD-CODING IS THE SINGLE MOST COMMON FAILURE

=B7*50-120000 gives the same answer today as =B7*B4-B6 and is worthless tomorrow. The number 50 is invisible to a reviewer, will not update, and will be missed when the price changes. Any constant inside a formula — other than a genuine structural constant like the 12 in a monthly division — is a defect.

Absolute and relative references

The distinction that makes a schedule possible to fill by copying.

FORM	BEHAVIOUR WHEN COPIED	USE FOR
A1	Both column and row shift	Cell-by-cell arithmetic across a table
\$A\$1	Neither shifts — fully locked	A single assumption used everywhere: a tax rate, an overhead rate
\$A1	Column locked, row shifts	A column of rates applied across months
A\$1	Row locked, column shifts	A row of headings or periods applied down products

F4 cycles through the four forms while the reference is selected. The mixed forms are what let a single formula fill an entire two-dimensional grid — an allocation table with departments down the side and cost drivers across the top needs exactly one formula, written once and copied.

Named ranges go further: name cell B4 Price and the formula becomes =Units*Price-FixedCosts, which is self-documenting and cannot break when a row is inserted. For a model anyone else will read, names beat coordinates.

The functions the work actually needs

FUNCTION	WHAT IT DOES IN THIS SUBJECT
SUM, SUMPRODUCT	SUMPRODUCT is the one worth knowing: total cost of a mix in a single cell. =SUMPRODUCT(quantities, unit_costs) replaces a column of multiplications.
SUMIF, SUMIFS	Totalling cost by category from a transaction listing — overhead by department, cost by activity, sales by product. =SUMIFS(\$D\$2:\$D\$500,\$B\$2:\$B\$500,"Machining",\$C\$2:\$C\$500,"Setups"). The workhorse of any ABC exercise.
IF, IFS	Conditional logic: a volume discount above a threshold, a stepped fixed cost, a bonus that applies only above target.

FUNCTION	WHAT IT DOES IN THIS SUBJECT
XLOOKUP or VLOOKUP / INDEX+MATCH	Pulling a standard cost, an activity rate or a department rate out of a rate table rather than retyping it. =XLOOKUP(A2, rate_table[Activity], rate_table[Rate]).
ROUND	Rounding <i>where the business rounds</i> — a unit price to cents, a break-even to whole units with ROUNDUP. Do not confuse formatting with rounding: formatting changes the display, ROUND changes the value.
SLOPE, INTERCEPT, RSQ	Cost estimation by regression — the proper version of the high–low method from Chapter One. =SLOPE(costs, activity) gives the variable cost per unit and =INTERCEPT(costs, activity) the fixed cost, using every observation instead of two.
NPV, IRR	Evaluating a proposal over several years, where the timing of cash flows matters.

REGRESSION AGAINST HIGH-LOW, IN ONE SCREEN

Machine hours in A2:A13, maintenance cost in B2:B13

Variable cost per hour	=SLOPE(\$B\$2:\$B\$13,\$A\$2:\$A\$13)
Fixed cost per month	=INTERCEPT(\$B\$2:\$B\$13,\$A\$2:\$A\$13)
Goodness of fit	=RSQ(\$B\$2:\$B\$13,\$A\$2:\$A\$13)

Three cells, twelve observations, and an RSQ that tells you whether machine hours explain maintenance cost at all — which high–low cannot tell you, because two points always fit a line perfectly. An R^2 of 0.31 means the cost driver is the wrong one, and that is worth knowing before building a budget on it.

Pivot tables

For cost analysis from a transaction listing, a pivot table does in thirty seconds what a page of SUMIFS does slowly: drag the cost category to rows, the department or product to columns, the amount to values. Use it to summarise, then *reference the summary* from the model rather than retyping figures out of it.

Two disciplines. A pivot table does not refresh itself — if the underlying data changes, it must be refreshed, and a stale pivot is a silent error. And build it on a proper Excel **Table** (Ctrl+T) so that added rows are picked up automatically.

Sensitivity analysis: the three tools

TOOL	ANSWERS	USE IT FOR
Data Table <i>What-If Analysis</i>	“What is profit at every price from \$45 to \$60, and every volume from 6,000 to 12,000?”	A grid of outcomes across one or two varying inputs. The proper way to present a CVP sensitivity.

TOOL	ANSWERS	USE IT FOR
Goal Seek	“What volume makes profit exactly zero?” or “what price gives a profit of \$147,000?”	Working backwards to a single input that produces a target output — break-even, target profit, the maximum price in a make-or-buy.
Scenario Manager	“How does the whole model look under best, base and worst case?”	Several inputs changing together, saved and compared side by side.

GOAL SEEK AND BREAK-EVEN

Break-even has a formula, so Goal Seek is not needed for it — but it is a good way to prove the model is wired correctly. Set the profit cell as the target, ask for zero, and let it change the units cell. If the answer matches $\text{=FixedCosts/UnitCM}$, the model works. If it does not, there is a broken reference to find, and better to find it now.

Goal Seek earns its place where there is no closed-form formula — a price that produces a target return when volume itself responds to price, or the maximum acceptable purchase price in a make-or-buy with several interacting adjustments.

A TWO-VARIABLE DATA TABLE FOR CVP

Put the profit formula in the top-left corner cell of the grid.
 Put the varying unit prices down the left column.
 Put the varying volumes across the top row.
 Select the whole block, then Data → What-If Analysis → Data Table
 Row input cell = the cell holding volume
 Column input cell = the cell holding price

The result is the entire profit surface, which shows something a single break-even figure cannot: how *fast* profit deteriorates as assumptions move, and which of the two assumptions matters more. That is what a sensitivity analysis is for.

Solver, for product mix with several constraints

Chapter Nine ranked products by contribution per unit of *one* scarce resource. With two or more constraints, ranking breaks down and the problem becomes linear programming. Solver does it directly.

Decision variables the cells holding units to produce of each product
 Objective total contribution margin $\text{=SUMPRODUCT(units, unit_CMs)}$ → Maximise
 Constraints machine hours used ≤ hours available
 labour hours used ≤ hours available
 material used ≤ material available
 units of each product ≤ its maximum demand
 all units ≥ 0 and integer, if part-units are meaningless

Set the model up so each constraint is a SUMPRODUCT of the decision variables against the per-unit requirements, sitting next to the available quantity. Solver then reports the optimal mix, and its sensitivity report gives the **shadow price** of each constraint — the extra contribution one more hour would earn, which is precisely the maximum worth paying for overtime or subcontracting.

Charts that say something

- **The CVP chart.** Volume on the horizontal axis; total revenue and total cost as two lines; the crossing point is break-even and the widening gap beyond it is profit. Plot fixed cost as a horizontal line to show why the cost line starts above zero.
- **A scatter plot with a trend line** for cost estimation, which makes an outlier visible — something no regression statistic will point at for you.
- **A variance waterfall** from standard cost to actual cost, one bar per variance. It shows at a glance that the material price variance was favourable and everything else was not.

Label the axes, state the units, and put the finding in the chart title rather than the variable name.

Making the model auditable

WHAT A REVIEWER SHOULD BE ABLE TO DO

- **Find every assumption in one place**, labelled, with its source or basis noted beside it.
- **Trace any output back to inputs** — Formulas → Trace Precedents draws the arrows, and a formula that fills the formula bar is a formula that should be broken into steps.
- **See the checks pass.** Build them in: does the allocation total equal the cost pool? Does cost accounted for equal cost in process? Do the variances sum to the total variance? A row of =IF (ABS (a-b) <0.01, "OK", "CHECK") cells at the top of the sheet is the cheapest quality control available.
- **Distinguish inputs from calculations at a glance** — one fill colour for cells to be typed in, another for cells never to be touched.

ERROR	WHAT TO DO INSTEAD
Hard-coded constants inside formulas	Every number in its own input cell, referenced.
A formula copied across a range that needed an absolute reference	Lock the reference with \$ before copying, then check the last cell of the range as well as the first.
Ranges that stop short of the data	Use Excel Tables, or whole-column references, so added rows are included.
Rounding by formatting	ROUND the value where the business genuinely rounds; otherwise leave full precision and let the display format do the presenting. Cells that look like they sum and do not are usually this.
Manual overrides typed over a calculated cell	Never. Add an explicit adjustment input, labelled, so the override is visible.

ERROR	WHAT TO DO INSTEAD
No documentation	A notes area stating purpose, author, date, data source and known limitations. It takes two minutes and is the difference between a model someone can use next year and one they rebuild.

Where each chapter meets the spreadsheet

CHAPTER	THE EXCEL WORK IT IMPLIES
One — cost concepts	Classifying a cost listing with SUMIFS; splitting a mixed cost with SLOPE and INTERCEPT, and checking the fit with RSQ.
Two — CVP	A contribution-format model driven from an input block; break-even by formula and confirmed by Goal Seek; a two-variable data table over price and volume; a CVP chart.
Three and Four — job and process costing	The four-step process costing schedule as a template, with built-in reconciliation checks; overhead applied by referencing a single rate cell.
Five — overhead allocation	Direct and step-down allocation as a grid filled by one mixed-reference formula. The reciprocal method solved with simultaneous equations, or with Solver by driving the residual to zero.
Six — ABC	An activity rate table, then XLOOKUP to pull each rate and SUMPRODUCT to apply it. Conventional and ABC costs side by side, with the distortion as a calculated column.
Seven — absorption and variable	Both income statements from one input block, with the reconciliation as a check row that must equal fixed overhead per unit times the inventory movement.
Eight and Nine — decisions	Relevant-cost comparisons as two columns with a difference column. Product mix by ranking, then by Solver once a second constraint appears.
Ten — budgeting	The master budget as linked sheets, where each period's opening inventory references the previous period's closing figure. Change the sales forecast and watch the cash budget move.
Eleven — standard costing	A variance schedule computing all eight variances from actual results and a standard cost card, with the reconciliation to total variance as a built-in check.

And that is the subject. The through-line, if there is one, is the question the first chapter opened with: what decision is this number for? A cost is not a fact waiting to be discovered but a construction built for a purpose, and the whole discipline is knowing which construction answers which question — and being able to show your working when someone asks.