

ACCG2050

Information Systems and Data Analytics for Accountants

How information systems shape the business processes that accounting measures — the cycles that move money in and out, the controls that keep them honest, the data that comes out the other end, and the ethics of running the whole machine.

Contents

PART ONE *Systems, strategy and the enterprise*

GROUNDWORK	How this book is organised	4
CHAPTER ONE	Introduction to accounting information systems	7
CHAPTER TWO	Strategy, competitive advantage and information systems	12
CHAPTER THREE	Enterprise systems — ERP, CRM and SCM	18

PART TWO *Risk and control*

CHAPTER FOUR	Internal controls	23
--------------	-------------------	----

PART THREE *Data and analytics*

CHAPTER FIVE	Data and analytics in the accounting profession	34
CHAPTER SIX	Analytics skills and tools for accountants	42

PART FOUR *The business cycles*

CHAPTER SEVEN	The revenue cycle	48
CHAPTER EIGHT	The expenditure cycle	58
CHAPTER NINE	HR management and the payroll cycle	64
CHAPTER TEN	Documentation, the general ledger and reporting	69

PART FIVE *Building and governing systems*

CHAPTER ELEVEN	Systems development and acquisition	78
CHAPTER TWELVE	Ethics, cybersecurity and privacy	89

Chapters follow the unit's themes rather than its calendar; the week each was taught is printed at its head. Worked answers are set out in full — cover them and attempt the question first.

PART ONE

Systems, strategy and the enterprise

What an accounting information system is, why a business buys one, and what the enterprise-wide packages — ERP, CRM, SCM — are actually for. Nothing in the later chapters makes sense without the vocabulary settled here: business process, value chain, transaction cycle, competitive advantage.

How this book is organised

This is a book about the machinery around the accounts — the processes that generate transactions, the systems that record them, the controls that make the record trustworthy, and the analytics that turn the record back into a decision. There are no journal entries in it and almost no arithmetic. Read this page first; it explains the shape of everything that follows.

The single idea underneath the subject

An accountant does not observe transactions. What an accountant sees is whatever the system recorded. The reliability of every number in a set of financial statements is therefore the reliability of a business process and of the controls sitting inside it — which is why a subject called “information systems” belongs in an accounting degree rather than beside it.

THE QUESTION BEHIND EVERY CASE IN THIS BOOK

Given this process, **what could go wrong**, **what would it cost the business**, and **what control would stop it**? Almost every worked problem here is a version of that question in a different costume — an online retailer’s order process, a hospital’s billing department, a payroll supervisor with too many duties, a data-entry clerk called John.

What you should be able to do by the end

1. **Evaluate** the role an information system plays in supporting a particular organisation’s business processes — not describe the software, but say what the process would look like without it.
2. **Analyse** a described process for risk, and design a proportionate control plan for it.
3. **Document** a process precisely enough, in flowchart or BPMN notation, that someone who has never seen it can find its weaknesses.
4. **Judge** how an organisation should acquire or build a system, and reason through the ethical and privacy consequences of what it then does with the data.

The six parts

PART	WHAT IT SETTLES
One — Systems, strategy and the enterprise	The vocabulary: data against information, what an AIS is made of, the value chain, the transaction cycles, competitive advantage, and the three enterprise packages (ERP, CRM, SCM).
Two — Risk and control	The control framework the rest of the book applies: five control activities, segregation of duties, general against application controls, and the limits of the whole apparatus.
Three — Data and analytics	What analytics means in an accounting job, the process for running an analysis end to end, and the tools and techniques that do the work.
Four — The business cycles	Revenue, expenditure, payroll, and the general ledger — each pulled apart into activities, documents, data stores, threats and controls, plus the notation for drawing them.
Five — Building and governing systems	How a system is specified, acquired, implemented and maintained; and the ethics, security and privacy obligations attached to running it.

The topics, week by week

The chapters are grouped by theme rather than by calendar, because the four cycle chapters share one structure and are far easier to learn side by side. 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 to accounting information systems	One
2	Organisational strategy, competitive advantage and information systems	Two
3	Customer relationship management and supply chain management	Three
4	Internal controls	Four
5	Data and analytics in the accounting profession	Five
6	Data analytics skills and tools for accountants	Six
7	Revenue cycle	Seven
8	Expenditure cycle	Eight
9	Systems development and acquisition	Eleven
10	Ethics and cybersecurity	Twelve

WK	TOPIC	CHAPTER HERE
11	Documentation and financial reporting	Ten
12	HR management and the payroll cycle	Nine

How each chapter is built

The idea first, then the mechanics, then the traps. Definitions are boxed where the exact wording matters. Worked problems are set out in full — cover the answer and attempt them before reading on, because the analysis is the point and the answer is only the residue of it. Each chapter closes with a paragraph carrying the argument into the next.

Further reading

Four books cover this ground in more depth, and between them supply most of its standard vocabulary.

TEXT	WHERE IT GOES FURTHER
Rainer & Prince, <i>Introduction to Information Systems</i>, 9th ed.	The general information-systems ground — strategy and competitive advantage, ERP, CRM and SCM, business analytics, systems development, ethics and security. Written for business generally, not for accountants.
Parkes, Considine, Olesen & Blount, <i>Accounting Information Systems</i>, 5th ed.	The accounting ground — internal controls, the revenue and expenditure cycles, documentation and the general ledger, with detailed risk and control tables for each activity.
Dzuranin, Geerts & Lenk, <i>Data and Analytics in Accounting</i>, 1st ed.	Analytics as an accounting discipline, including the plan–analyse–report process and the critical-thinking framework used in Chapter Five.
Stone, Cox, Gavin & Carpini, <i>Human Resource Management</i>, 11th ed.	HR as a business function, which is the context the payroll cycle sits inside.

With the map in hand, start where the subject starts: with the difference between a fact and a piece of information, which sounds like definitional throat-clearing and is in fact the reason the discipline exists.

Introduction to accounting information systems

This chapter sets the vocabulary: data and information, what an AIS is made of, how it adds value, and the two organising frameworks — the value chain and the transaction cycles — that every later chapter maps onto. It is the most definitional chapter in the book and the easiest to underrate.

Data and information

DEFINITION

Data are facts collected, recorded and stored by the system. A fact is a number, a date, a name, a quantity. **Information** is data organised within a context so that it has meaning to someone making a decision.

The stock example: 22/02/26 ABC Company 123 99 3 20 60 is data. Arrange the same digits into a sales invoice — date, customer, invoice number, item, quantity, unit price, total — and it is information, because someone can now act on it.

The practical consequence is that *more* is not better. Information overload reduces decision quality; part of what information technology is for is to filter and condense, so that the person deciding sees less and knows more.

The seven characteristics of useful information

Expect these as a multiple-choice question, usually by giving you a scenario and asking which characteristic is missing.

CHARACTERISTIC	WHAT IT MEANS, AND WHAT ITS ABSENCE LOOKS LIKE
Relevant	Bears on the decision being made. A customer's aged receivables balance is relevant to extending credit; their favourite colour is not.
Reliable	Free from bias and error, so it can be depended on. Reliability usually costs time, and often depends on the independence of the source.
Complete	Omits nothing important about the event or activity. A sales report that silently excludes credit notes is incomplete, not merely optimistic.
Timely	Available in time to affect the decision. Perfectly accurate information delivered after the decision has been made has no value at all.
Understandable	Presented in a form the user can actually use.
Verifiable	Two knowledgeable, independent people working from the same source would reach the same conclusion.
Accessible	Available to the user when and where it is needed.

THE PAIRS THAT FIGHT EACH OTHER

Reliable and verifiable both take time; timely does not have time. Complete and understandable pull in opposite directions — the fuller the report, the harder to read. A good analysis notices the trade-off rather than listing all seven as though they were free.

What an accounting information system is

DEFINITION

An **accounting information system** is a system that collects, records, stores and processes data to produce information for decision makers.

It is not software. Software is one of its components. The standard decomposition:

- **People** who use the system;
- **Procedures and instructions** for collecting, processing and storing data;
- **Data** about the organisation and its activities;
- **Software** that processes the data;
- **Information technology infrastructure** — computers, peripherals, networks;
- **Internal controls and security measures** that safeguard the data.

The last one is not an optional extra. An AIS must include controls that give reasonable assurance of the **accuracy, security and availability** of the data. A system that produces information nobody can rely on has not done its job, however elegant the interface.

How an AIS adds value

A well-designed AIS earns its cost in six ways.

1. Improving the **quality** of products or services and **reducing their cost** — better information about defects, waste and process time.
2. Improving **efficiency** — timely information lets activities happen without idle time between them.
3. **Sharing knowledge** across the organisation, so expertise held in one unit is available to the others.
4. Improving the **efficiency and effectiveness of the supply chain** — giving suppliers visibility of demand reduces stock-outs and carrying costs alike.
5. Improving the **internal control structure**, and so the reliability of the record.
6. Improving **decision making**, by identifying situations that need action, providing alternatives, and providing feedback on results.

Isolated systems and integrated enterprise systems

A stand-alone accounting system does three things with accounting data: **stores** it, **analyses** it and **safeguards** it. Its limitation is the word “accounting”. It knows what was invoiced; it does not know what was promised to the customer, whether the goods left the warehouse, or whether the person who shipped them was rostered on.

An **integrated enterprise system** shares data across functional areas — accounting, HR, sales, production, purchasing — so that one event is recorded once and every function sees the same version of it. That is the ERP idea, and it is the subject of Chapter Three.

WHY THIS MATTERS FOR THE REST OF THE BOOK

Integration is what makes many controls possible. A three-way match between a purchase order, a receiving report and a supplier invoice only works if all three live in the same system. When you are asked why an ERP improves control, that is the answer — not “because it is modern”.

Business processes

DEFINITION

A **business process** is a set of related activities performed by an organisation to create value by transforming inputs into outputs.

To understand an enterprise system you must first understand the business process underneath it, because the system is only ever an implementation of the process. That order — process first, system second — is the single most useful habit in the unit, and it is what the professional practice report is marked on.

The value chain

Porter’s value chain is the organising framework for business processes. Activities divide into two kinds:

PRIMARY ACTIVITIES – DIRECT VALUE
TO THE CUSTOMER

Inbound logistics – receiving and storing inputs
Operations – transforming inputs into the product
Outbound logistics – distributing the product
Marketing and sales – getting the customer to buy
Service – post-sale support

SUPPORT ACTIVITIES – MAKE THE PRIMARY ONES POSSIBLE

Firm infrastructure – accounting, finance, general management
Human resource management
Technology development
Procurement

Note where accounting sits: in *support*, under firm infrastructure. Accounting does not itself add value to the customer; it makes the value-adding activities measurable and accountable.

The **supply chain** is the wider system: the organisation's own value chain plus its suppliers upstream and its distributors and customers downstream.

Transaction cycles

Business processes group into transaction cycles, and the cycles are the skeleton of Part Four of this book.

CYCLE	THE EXCHANGE	GIVE / GET
Revenue	Goods and services to customers, cash back	Give goods, get cash
Expenditure	Cash to suppliers, goods and services back	Give cash, get goods
Production	Raw materials and labour converted into finished goods	Internal transformation
Human resources / payroll	Cash to employees, labour back	Give cash, get labour
Financing	Cash from investors and lenders, returns paid back	Get cash, give returns
General ledger and reporting	Summarises every other cycle into the financial statements	Not an exchange – a synthesis

This book examines the revenue, expenditure, HR and payroll, and general ledger cycles. Production is referred to where it interacts with them but is not treated in its own right.

Modules and cycles in a baseline accounting system

Small-business packages organise the same cycles as modules. A standard week-one exercise asks you to match an event to its module and to its cycle; the answers, which are worth learning because they recur in the multiple choice, are these.

ACCOUNTING EVENT	MODULE	TRANSACTION CYCLE
Create purchase orders	Vendors	Expenditure (purchasing)
Enter bills	Vendors	Expenditure
Pay bills	Banking	Expenditure
Create invoices	Customers	Revenue (sales)
Receive customer payments	Customers	Revenue
Make deposits	Banking	Revenue
Enter time	Employees	Payroll
Calculate payroll	Employees	Payroll
Adjusting entries	Financial	General ledger and reporting
Tax and financial reports	Reports	General ledger and reporting

MODULE AND CYCLE ARE DIFFERENT QUESTIONS

Several events sit in two places at once, and which answer is right depends on which question is being asked. “Make deposits” is a *banking module* function but a *revenue cycle* event; “pay bills” is a banking module function but an expenditure cycle event. The module says which part of the software performs the task; the cycle says which economic exchange the task belongs to. Confusing the two is the most common error in reading a system’s structure.

Why an accountant studies this at all

Four roles, and they are worth naming because each puts the same body of knowledge to a different use:

- **Users** of information systems — every accountant operates one daily;
- **Auditors** of information systems — the large firms audit the system, not just the ledger, because the ledger is only as good as the system;
- **Consultants** — advising on selection, design and implementation, which for the Big Four is a larger revenue line than audit;
- **Designers and evaluators of controls** — the accountant is often the only person in the room who understands both what could go wrong financially and how the process works.

Professional certifications follow the same logic: CPA and CA ANZ both examine information systems, and specialist credentials such as CISA exist for auditing them.

So far the AIS has been described as a recording machine. The next chapter asks a harder question — whether a system can be a source of advantage rather than merely a cost of doing business.

Strategy, competitive advantage and information systems

This chapter asks what makes a business competitive, whether technology can be a source of that advantage or merely a cost of staying in the game, and how to improve a process once you have one. The frameworks — VRIO, five forces, the three generic strategies, BPI against BPR — all earn their keep by being applied to a particular firm, not recited.

Change is the premise

Begin with adoption curves: landline telephones took roughly seventy years to reach saturation, mobile phones perhaps twenty, smartphones and internet access faster still. The point is not the numbers. It is that the interval in which an advantage can be defended keeps shrinking, so a firm that is competitive today has no entitlement to be competitive in five years.

OBSERVATION, INFERENCE, ASSERTION

A chart lets you make an *observation* and offer a *suggestion*. It does not let you make an assertion about cause. “Smartphone uptake was faster because people had more disposable income” is not something an adoption curve can establish. Say what the data shows, then say what it might indicate, and label the second as inference. The discipline matters in every claim made about data, and it recurs in Part Three.

Competitive strategy and competitive advantage

DEFINITIONS

A **competitive strategy** is a statement identifying a business’s approach to competing, its goals, and the plans and policies required to carry them out (Porter, 1985). Because competitors are actively trying to stop you, a strategy must anticipate and counter their moves, not merely plan your own.

A **competitive advantage** is an asset or capability that gives the organisation an edge over its competitors: it allows superior margins and generates value for the firm and its shareholders.

The test of a genuine advantage is **durability**. If a rival can copy it cheaply and quickly, it is not an advantage; it is a feature. A price cut is not a competitive advantage. Neither is buying the same software everyone else has.

VRIO – THE FOUR QUESTIONS TO ASK OF ANY RESOURCE

QUESTION	IF THE ANSWER IS NO
Valuable? Does it let the firm exploit an opportunity or neutralise a threat?	Competitive disadvantage.
Rare? Do few competitors possess it?	Competitive parity — it keeps you in the game, nothing more.
Inimitable? Is it costly or impossible to copy or substitute?	Temporary advantage only.
Organised? Is the firm structured to actually exploit it?	Unused potential — the resource is wasted.

The fourth is the one most often skipped, and the one an accountant should care about most. A firm can hold excellent assets and excellent staff and still capture nothing, because its processes, incentives and reporting are not arranged to use them. VRIO is the natural tool whenever the question is whether a firm's advantage will last rather than whether it exists.

WHERE ADVANTAGE ACTUALLY COMES FROM

- **Brand loyalty** — customers who will not switch on price alone.
- **Scale** — volume that drives unit cost below what a smaller rival can reach.
- **Intellectual property** — patents, proprietary software, algorithms, process know-how.
- **Location** — access to inputs, or to customers, that a rival cannot replicate.
- **Network effects** — the product becomes more valuable as more people use it, and each additional participant costs almost nothing to add.
- **Exclusive relationships** — supply contracts and partnerships that lock a scarce input away from rivals.

THE UBER QUESTION

Asked what Uber's competitive advantage is — technology, business model, venture capital or price — the defensible answer is the **business model**, and more precisely the network effect it creates: more drivers means shorter waits, which attracts more riders, which attracts more drivers. The technology is imitable, capital is available to rivals, and price is the least durable advantage of all.

Porter's five forces

A framework to *apply* to a particular industry rather than recite. The five forces describe the profitability of an industry, not of a firm.

FORCE	WHAT TO LOOK FOR – AND WHAT IT DOES TO IT
Rivalry among existing competitors	Number and similarity of competitors, rate of new product introduction, price competition. The internet generally <i>intensifies</i> rivalry: price comparison is free and geography stops protecting anyone.

FORCE	WHAT TO LOOK FOR – AND WHAT IT DOES TO IT
Threat of substitutes	What else can the customer buy to meet the same need if your price rises? Digital goods increase substitution because switching costs approach zero.
Bargaining power of suppliers	Few suppliers means high supplier power. The more suppliers a firm can choose between, the greater its control. B2B marketplaces and EDI have generally <i>reduced</i> supplier power by widening the field.
Bargaining power of buyers	Can customers switch easily? Can they force competition on price alone? The internet has sharply <i>increased</i> buyer power through low search and switching costs.
Threat of new entrants	How high are the barriers to entry? Compare opening a café with opening a semiconductor fabrication plant. Digital distribution lowers barriers in many industries; regulation and capital intensity raise them.

THE NET EFFECT OF THE INTERNET

Intense rivalry, greater buyer power because switching costs fall, but also the opportunity to reach and enlarge a customer base far beyond a physical catchment. It is not uniformly good or bad for incumbents – which is exactly the nuance a strategy question is looking for.

The three generic strategies

STRATEGY	THE BET	EXAMPLE
Cost leadership	Produce at the lowest cost in the industry and win on price.	Kmart, Aldi
Differentiation	Offer something rivals cannot – features, quality, service – and charge for it.	Apple, premium airlines
Focus (niche)	Serve one narrow segment better than a generalist can, at either the luxury or the value end.	Luxury handbags; specialist medical software

Rainer adds two more that are worth knowing because they are often the honest description of what an IT investment achieves: an **innovation strategy** (new products or new ways of doing business) and an **operational effectiveness strategy** (doing the same things better than competitors), plus a **customer-orientation strategy** of concentrating on customer satisfaction.

STRATEGIC INFORMATION SYSTEM

A **strategic information system** is one that provides a competitive advantage by helping the organisation implement its strategic goals and improve its performance and productivity. The word doing the work is *strategic*: a payroll system that merely pays people correctly is operational, not strategic.

IT as risk and as opportunity

The **business environment** is the combination of social, legal, economic, physical and political factors in which an organisation operates. Significant change in any of them creates business pressure.

Globalisation — the integration and interdependence of economic, social, cultural and ecological life, made possible by information technology — is the largest such pressure, and its effect is to increase competition: organisations can now operate largely without regard to geography, time or distance.

Within a firm, IT changes processes in three ways worth naming:

- **Synergies** — the output of one unit becomes the input of another, lowering cost;
- **Knowledge sharing as a core competency** — expertise moves across business units instead of dying in one;
- **Network-based strategies** — exploiting network economics, where the marginal cost of adding a participant is near zero while the marginal gain is not; creating **virtual organisations** that link people, assets and ideas without traditional boundaries and share risk and return; and forming **business ecosystems**, in which one or a few keystone firms provide a platform that niche firms build on (Microsoft, Intel and the software industry around them).

Business–IT alignment

DEFINITION

Business–IT alignment is the tight integration of the IT function with the organisation's strategy, mission and goals.

The six characteristics of excellent alignment are a standard list:

1. IT is viewed as an engine of innovation that continually transforms the business and often creates new revenue streams.
2. The organisation treats internal and external customers, and the customer service function, as supremely important.
3. Business and IT professionals are rotated across departments and job functions.
4. Overarching goals are completely clear to every IT and every business employee.
5. IT employees understand how the company makes — or loses — money.
6. The organisation has a vibrant and inclusive culture.

Numbers 3 and 5 are the operative ones. Misalignment is usually not a disagreement about goals; it is that the two groups do not understand each other's work.

Improving a business process

A business process has three elements — **inputs** (materials, services and information transformed by the process), **resources** (people and equipment that perform the activities) and **outputs** (the product or service created) — and two objectives that are frequently confused:

OBJECTIVE	MEANING
Efficiency	Doing things <i>well</i> — moving from one activity to the next without delay or waste of money and resources.
Effectiveness	Doing the things <i>that matter</i> — creating outputs of genuine value to the process customer.

A process can be highly efficient at producing something nobody wants. Efficiency without effectiveness is the definition of well-run waste.

THE THREE IMPROVEMENT APPROACHES

Business process improvement (BPI) is an incremental approach that moves an organisation toward process-centred operations, focused on reducing variation in process outputs by identifying the underlying cause of that variation. Its five phases follow the DMAIC sequence:

- 1 **Define** the problem, the customer and the process requirements.
- 2 **Measure** — establish a data collection plan, determine defects, gather metrics on the current process.
- 3 **Analyse** — establish the root cause of the variation and the defects.
- 4 **Improve** — eliminate the root causes through innovative solutions.
- 5 **Control** — institutionalise the change so the organisation does not fall back into old habits.

Business process management (BPM) is the management system that supports continuous BPI across core processes over time — the discipline that keeps step 5 honest.

Business process reengineering (BPR), from Hammer and Champy's *Reengineering the Corporation* (1993), is a radical redesign of an organisation's processes to increase productivity and profitability, examining them with a clean-slate approach. A BPR project's leadership team may deliberately include influential external stakeholders, precisely because insiders find it hard to abandon the existing design.

BPR	BPI
High risk, high cost	Low risk, low cost
Radical redesign	Incremental change
Top-down	Bottom-up
Time consuming	Takes less time
Impacts can be overwhelming	Quantifiable results
High failure rate	All employees can be trained in it

Benchmarking supports both: comparing the efficiency and effectiveness of your processes against industry best practice or a defined standard, and measuring performance against it.

CHOOSING BETWEEN THEM

Faced with a firm in trouble, the choice between BPI and BPR turns on the diagnosis. If the process is fundamentally sound but produces variable output, that is BPI. If the process itself is obsolete — it was designed around paper, or around an organisational structure that no longer exists — incremental improvement optimises the wrong thing, and BPR is the honest answer, with the failure rate named as the cost of it.

Two of the three improvement routes end at the same place: a package that promises the best practice already built in. That promise, and what it costs, is the subject of the next chapter.

Enterprise systems — ERP, CRM and SCM

Three acronyms, one idea: take data that used to sit in one function's filing cabinet and make it available to every function that needs it. ERP looks inward at the enterprise, CRM outward at the customer, SCM outward at the supplier. The professional practice report asks you to identify which of the three your chosen organisation runs, so learn the boundaries, not only the definitions.

Enterprise resource planning

DEFINITION

An **ERP system** is a set of integrated computer program modules that attempts to unify the different functional areas of an organisation around a single shared database. It is designed on the basis of **best practice** — the vendor's view of the best way of performing a particular process.

Typical modules: financials and general ledger; sales and distribution; materials management and purchasing; production planning; inventory; human resources and payroll; project management; customer service. The historical core is manufacturing and production planning — the original question ERP answered was how to schedule production so that materials arrive when they are needed and no earlier.

WHAT INTEGRATION BUYS YOU

- One event is entered once. A sales order automatically checks inventory availability, reserves stock, triggers a production or purchase requisition if stock is short, and creates the receivable when it ships.
- Everyone sees the same numbers, so the argument about whose spreadsheet is right disappears.
- Controls become possible that no isolated system can perform — matching across documents, automatic credit limit checks, real-time exception reporting.
- Management reporting improves because data is consistent enough to consolidate.

WHAT IT COSTS YOU

- **Best practice is somebody else's practice.** Adopting the package means changing your process to fit the software, or paying heavily to customise the software — and customisation makes every future upgrade harder.
- Implementation is expensive, slow and disruptive; the change management usually costs more than the licence.
- Integration concentrates risk. One failure now affects every function at once, which is exactly the Woolworths story in Chapter Eleven.
- Vendor dependence — the firm's core processes now run on a product it does not control.

KEY TAKEAWAY

ERP stands for enterprise resource planning. Its goal is to **improve and streamline internal business processes** by integrating data across functions. If a question asks what ERP is for, that sentence is the answer; the module list is supporting detail.

Customer relationship management

DEFINITION

CRM refers to the strategies, practices and technologies an organisation uses to manage and analyse customer interactions and data across the whole customer lifecycle, with the aim of improving satisfaction and loyalty, increasing retention and driving sales growth.

Note that CRM is described as a *strategy* supported by technology, not as a piece of software. Operational CRM systems have two components, and the distinction between them matters:

CUSTOMER-FACING APPLICATIONS AN EMPLOYEE DEALS WITH THE CUSTOMER	CUSTOMER-TOUCHING APPLICATIONS THE CUSTOMER DEALS WITH THE TECHNOLOGY DIRECTLY
Customer service and support Sales force automation (SFA) Marketing Campaign management	Search and comparison capabilities Technical and other information and services Customised products and services Personalised web pages FAQs Email and automated response Loyalty programs

Operational CRM delivers three things: efficient and personalised marketing, sales and service; a 360-degree view of each customer; and access, for any sales or service employee, to the complete history of that customer's interaction with the organisation regardless of which touch point it came through.

Analytical CRM sits behind it, mining the accumulated data for segments, propensities and churn risk.

CRM BENEFITS MAPPED TO COMPETITIVE STRATEGY

The mapping below is the one to reach for whenever the question is how a system supports a strategy rather than what the system does.

CRM BENEFIT	STRATEGY IT SUPPORTS
Better customer knowledge and anticipation of needs	Customer intimacy
Segmentation	Differentiation and focus
Better communication	All strategies, especially differentiation and focus
Customer retention	All strategies
Protection of data privacy	Reputation-, brand- and CSR-based strategies

Supply chain management

DEFINITION

SCM is the planning, organising and coordination of all the activities of the supply chain — the flow of materials, information and money from raw material supplier through to the end customer, and back again.

Three segments and three flows:

SEGMENTS	FLows
Upstream — suppliers and their suppliers; sourcing and procurement Internal — the organisation's own manufacturing, packaging and in-house logistics Downstream — distributors, retailers, the end customer	Material flow — physical goods, including returns and disposal Information flow — orders, schedules, status, changes Financial flow — payments, credit terms, consignment arrangements

The five basic SCM processes: **plan**, **source**, **make** (execute), **deliver** and **return**.

WHAT GOOD SCM DELIVERS

- Selection of higher-quality suppliers and objective tracking of their performance;
- Feedback into internal product design and innovation;
- Analysis of historical data for insight into demand;
- Faster adaptation to changing markets and customer demand;
- Tracking of raw materials and finished inventory through the chain;
- Lower inventory and carrying costs without an increase in stock-outs.

THE BULLWHIP EFFECT – THE REASON INFORMATION FLOW MATTERS

Small fluctuations in end-customer demand amplify as they travel upstream: the retailer over-orders slightly to be safe, the distributor over-orders to cover the retailer, the manufacturer builds capacity for a spike that never existed. The cure is not better forecasting at each link but **shared demand information** along the chain — which is precisely what an integrated SCM system, EDI or vendor-managed inventory provides. If you are asked why supply chain visibility is worth paying for, this is the sharpest answer.

KEY TAKEAWAY

SCM stands for supply chain management. Its goal is to **improve coordination between the participants of the supply chain** — suppliers, manufacturers, distributors, customers. ERP looks inside the firm; SCM looks along the chain; CRM looks at the customer end of it.

Telling the three apart in a case

IF THE NARRATIVE EMPHASISES...	THE SYSTEM IS...	BECAUSE
One database shared by finance, HR, sales and production; entering a transaction once	ERP	Internal integration is the defining feature
Customer history visible at every touch point, campaigns, loyalty programs, self-service portal	CRM	The object of the data is the customer relationship
Supplier portals, demand forecasts shared upstream, EDI, RFID tracking, vendor-managed inventory	SCM	The object is coordination across organisations

They overlap. A large ERP suite will contain CRM and SCM modules, and saying so is a better answer than pretending the categories are watertight — provided you can say which function each module is performing.

Every system in this chapter multiplies the consequences of a single bad instruction: enter one transaction wrongly and it now propagates to every function at once. That is why the book turns next to control.

PART TWO

Risk and control

One chapter, and perhaps a third of everything that follows.

Each of the four cycle chapters in Part Four is an application of the vocabulary settled here, and “identify the weaknesses and recommend controls” is the instruction this whole book is built to answer.

Internal controls

Controls are policies and procedures adopted in response to identified risk. This chapter covers the five types of control activity, the three ways of classifying any control, general and application controls, disaster recovery, and the limits of the whole apparatus. It ends with a fully worked risk-and-control analysis.

Where control begins: risk

Errors can occur anywhere in the financial reporting process, and **not all risks are financial**. Auditing standards are the basis of external financial statement audits and their primary concern is the financial accuracy of the statements. The concern of an accountant working with an AIS is broader: reputational, operational, privacy and continuity risks all matter, and none of them appears as a line item.

The sequence is always the same, and it is worth setting out explicitly, because a control with no risk behind it is simply a cost:

- 1 **Identify the risk.** What could go wrong in this activity, given who does what?
- 2 **Evaluate its extent.** How likely is it, and how bad would it be? Controls are proportionate to this, not to the imagination.
- 3 **Implement policies and procedures to address it.** Those policies and procedures are what we call control activities.
- 4 **Monitor** whether they are operating, and revisit when the process or the environment changes.

Applying internal controls therefore involves an evaluation of *assertions* (that recorded sales occurred, that all sales were recorded, that they are accurate, that they fall in the right period) coupled with a risk assessment.

The five types of control activity

These are the categories used by auditing standard ASA 315. Learn them as a checklist to run over any narrative.

CONTROL ACTIVITY	WHAT IT IS
1. Authorisation	Procedures assuring that transactions and events are carried out by those with appropriate authority. Requires defined roles, responsibilities and adherence mechanisms — delegation limits, approval thresholds, credit limits.
2. Performance reviews	Review or analysis of performance, comparing actual outcomes with expected or planned ones. Budget variance analysis, comparison of gross margin by product to prior period, review of exception reports.
3. Information processing controls	Controls directed at the accuracy (data entered is correct and reflects the actual event), completeness (all events are recorded)

CONTROL ACTIVITY	WHAT IT IS
	and authorisation of transactions. Subdivided into general controls (policies and procedures supporting applications across the organisation) and application controls (manual or automated procedures at the business-process level, relating to the processing of transactions by an individual application).
4. Physical controls	Controls that physically protect the organisation's resources from theft or damage — locks, safes, restricted areas, cameras, fireproofing.
5. Segregation of duties	Certain key functions must not be performed by the same person. Applies to the organisation's IT systems as much as to its cash.

Segregation of duties — the CAR rule

THREE RESPONSIBILITIES THAT MUST NOT COMBINE

Custody — handling the assets themselves: cash, cheques, inventory, goods.

Authorisation — approving transactions: purchase orders, payments, discounts, credit.

Recording — keeping the records: journal entries, inventory logs, the customer ledger.

No one person should hold two or more of C, A and R for the same class of transaction. If they do, that is the weakness — and it is worth saying so in exactly those terms, which is clearer and shorter than a paragraph of narrative.

COMBINATION TO LOOK FOR	WHAT IT ENABLES
Custody + Recording <i>handles inventory and updates the inventory records</i>	Steal the goods and adjust the record to hide the loss.
Authorisation + Custody <i>approves payments and issues the cheques</i>	Approve a payment to yourself or a fictitious supplier and hand yourself the cheque.
Authorisation + Recording <i>records sales and approves discounts and credit notes</i>	Pocket a customer receipt and write it off as a discount or a return.

Three corollaries follow: sales staff should not bill customers; accounts receivable staff should not handle incoming payments; warehouse staff should not maintain the inventory records. Each is one of the three combinations above.

THE OBVIOUS LIMITATION

Segregation is defeated by **collusion**, and in a small business it may be unaffordable — there are only three people. The compensating answer is *not* to pretend the duties are segregated but to substitute controls that do not require numbers: owner review of the bank reconciliation, mandatory annual leave, surprise counts, dual signatures, and reliance on system-enforced access rights rather than staff structure.

Classifying controls, two ways

BY HOW THEY DEAL WITH THE RISK

TYPE	WHAT IT DOES	EXAMPLES
Preventive	Stops the error or irregularity before it happens — reduces the probability of occurrence.	Passwords; required fields; a firewall; an input validation rule; approval before a payment is released; segregation of duties.
Detective	Does not prevent anything, but alerts users to errors and anomalies after the fact.	A virus scan; a bank reconciliation; an exception report; a physical stocktake; a duplicate-payment query.
Corrective	Fixes the error or irregularity once it has occurred, and restores the position.	A disaster recovery plan; restoring from backup; antivirus software that quarantines and repairs; a correcting journal.

Some controls sit in more than one box, and saying so is correct rather than evasive. Antivirus software is detective when it scans and corrective when it removes. A reconciliation is detective, but the discipline of knowing it will be performed is preventive.

BY WHERE THE CONTROL ACTIVITY TAKES PLACE

STAGE	PURPOSE	EXAMPLES
Input	Operates as data enters the system; addresses accuracy, validity and completeness.	Standardised forms; field and range checks; pre-numbered documents; batch totals.
Processing	Ensures data is handled correctly once inside — that the right data stores are updated with the right values.	Run-to-run totals; reconciliations; sequence checks; hash totals.
Output	Protects the outputs the process generates, and who can see them.	Access privileges over reports; page numbering and end-of-report footers; controlled printing of confidential material.

KEEPING THE TWO AXES APART

Input, processing and output controls operate *within* a business process and are designed around the particular risks of that process. They are therefore **application controls**. The preventive / detective / corrective scheme is different in kind: it can be applied to *both* general and application controls. Do not confuse the two axes — one is about *where*, the other about *when relative to the event*.

What a computerised AIS is trying to achieve

PROPER AUTHORISATION	PROPER RECORDING	COMPLETENESS	TIMELINESS
User privileges, access rights, user restrictions, approval over defined thresholds	Records of the right type and format, reflecting the reality of the underlying transaction or event — input accuracy	Every transaction captured, at both the individual transaction level and the business process level	Processing need not be instantaneous; it must suit organisational needs. Batch, online real-time, or online gathering with batch processing

General controls

General controls relate across *all* the information systems in an organisation. They do not attach to a particular transaction; they create the environment in which the application controls can be trusted. Six headings:

1. PHYSICAL CONTROLS

Concerned with restricting access to physical resources; the major question is *who has physical access*. Locked computing premises; discreet premises that do not advertise what is inside; swipe card access; biometric access controls; on-site security; cameras recording access.

2. SEGREGATION OF DUTIES WITHIN IT

Separation between the users of IT, the maintainers of the systems, the system designers, the system testers, and those with access to the data. A developer who can also move code into production and also access live data can do anything at all and leave no trace.

3. USER ACCESS

Logical access of users to systems — unique identification codes and passwords, access rights matched to role, removal of access on termination. Note that *unique* identification is what makes the audit trail meaningful; shared logins destroy accountability even when the password is strong.

4. SYSTEMS DEVELOPMENT PROCEDURES

Policies, procedures and restrictions governing the maintenance and development of information systems, including differentiated privileges such as system administrator. The essential preventive control: assurance that untested, incompatible, unreviewed or unlicensed software is not placed on the system.

5. USER AWARENESS OF RISKS

Security education, training and awareness programs ensuring employees know the information system risks, the security threats and issues, the organisational security policies, and the policies for detecting fraud. The best-configured system in the world is defeated by one person who does not know what a phishing email looks like.

6. DATA STORAGE PROCEDURES

Customer, staff and intellectual property data is stored on servers; release would carry financial and non-financial consequences. Controls: data access logs, restriction of user privileges, encryption of stored data, and scheduled backups — batch or real time. Cloud storage brings its own risks:

- inability to audit and monitor at file level;
- dependence on internet access to reach your own data;
- unresolved security issues — data security, network security, data integrity, web application security;
- uncertainty about jurisdiction and about who else can compel access.

Sitting over all six: **security policies** documenting the organisation's approach to security, usually following a recognised framework or standard, and — the part that fails in practice — understood and used by all users.

Application controls

Application controls are built around the operation of a particular process and relate to the three system stages.

INPUT CONTROLS	WHAT IT PREVENTS
Standardised forms	Omission — the form asks for everything needed, in a fixed order.
Pre-numbered documents	Unrecorded and duplicated transactions; makes completeness testable.
Sequence checks	Gaps in a numbered series — the missing invoice that was never posted.
Turnaround documents	Re-keying error — output from the system returns as input, machine-readable.
Automated form completion	Transcription error; pulls the customer's stored details rather than retyping them.
Transaction authorisation procedures	Unauthorised transactions entering the system at all.
Batch totals	Lost or added records in a batch — compare the count and the value entered with the control total.
Independent review	Errors the preparer cannot see in their own work.
Field, range, limit, validity, reasonableness and check-digit tests	Impossible or implausible values: alphabetic characters in a quantity, a date in 1899, a customer code that does not exist, an order of

INPUT CONTROLS	WHAT IT PREVENTS
	10,000 units where 100 is normal.
PROCESSING CONTROLS	OUTPUT CONTROLS
Run-to-run totals — a total carried from one processing stage to the next and compared Reconciliations — subsidiary ledger to control account Batch totals Sequence checks Hash totals — a meaningless sum (of account numbers, say) used purely to prove nothing changed in transit	Access privileges over who may generate and view reports (segregation of duties) Page numbering and end-of-report footers, so a reader knows nothing has been removed Physical control — confidential material should not print to a device every staff member can reach Database queries to detect irregularities or anomalies, for example querying credit returns to detect cash or receivables fraud

Disaster recovery

DEFINITION

A **disaster recovery plan** is the strategy to be put into action in the event of a disaster that disrupts normal operations, to restore critical systems and services as quickly as possible and to recover the data and resources relating to the organisation's processes.

Disaster types: terrorism; natural disasters — fire, flood, cyclone; and failures of online operations such as a web server going down. To which 2020 onwards added pandemic and the loss of physical premises entirely.

A defensible plan contains, at minimum:

- **Identification of critical systems and data**, and the maximum tolerable outage for each — you cannot restore everything first;
- **Backup arrangements** — frequency, retention, off-site or geographically separate storage, and encryption;
- **Alternative processing facilities** — a hot site ready to run, a warm site partially equipped, or a cold site that is only space and power;
- **Documented recovery procedures**, in a form usable by someone other than the person who wrote them;
- **Named responsibilities and a communication plan** for staff, customers, suppliers and regulators;
- **Testing**, at defined intervals, because an untested plan is a document rather than a control.

Disaster recovery is a **corrective** control. Backups themselves are preventive of data loss but corrective of a disaster; either classification is defensible if reasoned.

The limitations of internal control

An internal control system does **not** provide 100% assurance that objectives will be achieved. It provides reasonable assurance, and the gap has names:

LIMITATION	WHAT IT MEANS
Judgement error	Controls are operated by people making decisions under time pressure with incomplete information.
Unexpected transactions	Designers cannot anticipate every possible outcome or future event; a novel transaction meets no control designed for it.
Collusion	Two or more people conspiring together defeat segregation of duties, which assumes they are adversaries.
Management override	Controls only work if management is ethical and promotes good behaviour. Those who design the controls can suspend them.
Weak internal controls	Controls that are poorly designed, or well designed and not operated.
Conflicting signals	Management says compliance matters and rewards only the sales number. Employees believe the reward.
Management incompetence	Incompetence at the top flows down through the whole organisation.
External factors	Events beyond the organisation's control with dramatic effects — natural disaster, supplier collapse.
Fraud	Deliberate circumvention of controls for personal gain.
Regulatory change	Changes in the regulatory environment alter how the organisation must operate, sometimes faster than its systems can follow.
Information technology	Constant flux — new platforms, new vulnerabilities, new threats.

Two more that always belong in a complete answer: **cost–benefit** — a control that costs more than the loss it prevents should not be implemented — and **obsolescence**, where a control still operates faithfully over a risk that has moved.

The COSO frame around all of this

The framework standing behind everything above has five interrelated components. Its vocabulary is the common language of internal control, in audit practice and in regulation, and the five components are worth holding in mind as the container the control activities sit inside.

COMPONENT	WHAT IT COVERS
Control environment	The attitude, emphasis and awareness of an organisation's management towards internal control and its operation — integrity, ethical values, competence, board oversight, assignment of authority. The foundation for everything else.
Risk assessment	Identifying and analysing the risks to achieving the organisation's objectives, as the basis for deciding how they should be managed.
Control activities	The policies and procedures that ensure management's risk responses are actually carried out — the five types above.
Information and communication	Capturing and exchanging the information people need to carry out their responsibilities, in a form and timeframe that lets them do so.
Monitoring	Assessing over time whether the components are present and functioning — ongoing supervision plus separate evaluations such as internal audit.

HOW RISK ASSESSMENT AND CONTROL ACTIVITIES FIT TOGETHER

Risk assessment and control activities are sequential and dependent. Risk assessment identifies and rates what could prevent the organisation meeting its objectives; control activities are the policies and procedures put in place *in response* to those rated risks. Without risk assessment, control activities are arbitrary — expensive where nothing threatens and absent where something does. Without control activities, risk assessment is a document that changes nothing. The link is also what makes controls proportionate: the extent of the control should reflect the assessed likelihood and consequence.

Worked answer — a data entry process

THE CASE

Users log on with a user ID (first initial plus the first six letters of the surname) and a password assigned when they join the firm, which cannot be changed. Users have internet access and several have installed messaging and chat programs on their machines. John's main task is data entry. Each day he receives a bundle of orders from the customer assistant and enters the details. He enters the customer name, address and contact number, then clicks "Next" to enter the items and quantities. If the customer name is not provided, the computer prompts him to go back and fill it in before proceeding. The computer will only accept numeric values for quantities. Once all orders are entered John clicks "Done" and the computer displays the number of orders entered. John usually ignores this, because by then it is lunchtime.

The task is to identify five risks, suggest a control for each, say whether that control is present or missing in the case, and classify it as general or application. Note the structure — four columns, one row per risk — because a table forces all four parts to be answered for every risk, which prose does not.

RISK IN THE PROCESS	CONTROL	PRESENT OR MISSING	GENERAL OR APPLICATION
Passwords are assigned and cannot be changed, and the user ID follows a published formula. A departing or curious employee can guess an ID and, once a password leaks, it leaks permanently — unauthorised access and repudiation of transactions.	Force a password change at first logon and at intervals; enforce complexity; lock out after failed attempts; add multi-factor authentication.	Missing — the case states passwords cannot be changed	General <i>user access control</i>
Unrestricted internet access and staff-installed chat software introduce malware and give an uncontrolled channel for confidential customer data to leave the firm.	Restrict local installation rights; whitelist permitted software; content filtering and endpoint protection; an acceptable-use policy with staff awareness training.	Missing	General <i>systems development / user awareness</i>
John ignores the count of orders entered, so a batch that is short — an order dropped, a page stuck to another — is never detected. Sales are understated and a customer never receives goods.	Batch total control: the customer assistant records the number of orders and a control total of quantities before handing over the bundle; John reconciles the system count to it and signs off.	Partly present — the system displays the count, but no comparison to an independently prepared total is required	Application <i>input control</i>
Only the customer name is mandatory and only quantity is validated as numeric. Prices, item codes, addresses and quantities can be wrong or implausible — wrong goods shipped, wrong amount billed.	Extend input validation: mandatory fields for all critical data, validity check of the item code against the item master, range and reasonableness checks on quantity, automated completion of customer details from the customer master.	Partly present — a required-field check and a numeric field check exist; the rest are missing	Application <i>input control</i>
There is no evidence that the orders themselves are authorised or that the bundle John receives is complete. Fictitious or altered orders could be entered, and there is no audit trail linking an entered order to a source document.	Pre-number the source order documents and perform a sequence check; require the customer assistant's authorisation on each order; retain the source documents and reconcile periodically to entered orders.	Missing	Application <i>input control</i> , supported by a general control over records retention

METHOD, SO IT TRANSFERS TO ANY CASE

Read the narrative twice. On the first pass highlight every **verb performed by a named person** and ask which of C, A and R it is. On the second pass highlight every **system behaviour** the narrative bothers to mention — a prompt, a check, a displayed total — because a detail of that kind is there either because it *is* the control or because it is the control that someone ignores. Then write the risk as a consequence, not as a description: not “John ignores the number” but “orders may be omitted from the system without detection, understating sales and leaving customers unserved”.

Controls make the record trustworthy. What the organisation then does with a trustworthy record — and the skills the profession now expects an accountant to bring to it — is the subject of the next two chapters.

PART THREE

Data and analytics

The newest material in the subject, and the reason accounting degrees now teach information systems rather than merely mentioning them. Two chapters: what analytics is and how an analysis is run from question to recommendation, then the skills and tools that do the work.

Data and analytics in the accounting profession

The accountant's traditional advantage is being the person who knows what the numbers mean. That advantage survives only if the accountant can also work with the data the numbers came from. This chapter covers what analytics is doing to the profession, where it is used in each branch of accounting, and the process for running an analysis end to end.

What is changing, and why now

Three things happened at once. Transaction data became complete and machine-readable, because everything runs through an enterprise system. Storage and processing became cheap enough that analysing an entire population became affordable. And the tools became usable by people who are not programmers.

The effect on accounting work falls into two categories:

EFFICIENCY – AUTOMATION OF ROUTINE TASKS

Journal entry preparation and posting; bank reconciliations; general ledger maintenance; tax filings; routine audit procedures; the assembly of recurring reports. Work that was manual, repetitive and slow becomes a scheduled process that runs unattended.

INSIGHT – ANALYSIS THAT WAS NOT PREVIOUSLY POSSIBLE

Testing an entire population rather than a sample; identifying anomalies across millions of transactions; monitoring performance continuously rather than monthly; forecasting from behaviour rather than from last year plus a percentage.

DEFINITION

Data analytics is the process of analysing raw data to answer questions or provide insights.

The habit that makes it useful is a two-step one: **observe** an area of interest or concern in a summary or visualisation of the data, then **investigate** it to understand the root cause. A dashboard that reports a decline without anyone asking why has produced description, not insight.

THE CLAIM WORTH BEING ABLE TO DEFEND

Analytics does not replace professional judgement; it relocates it. The judgement moves from “which forty invoices shall I look at?” to “which of these four thousand flagged exceptions actually matter, and what does the pattern tell me about the process that produced them?” The historical necessity of **sampling** was a response to cost, not a preference. When the whole population can be examined, sampling starts to look like an artefact of an era when it could not.

Where analytics is used, branch by branch

BRANCH	WHAT ANALYTICS IS USED FOR
Financial accounting	Automating routine tasks — bank reconciliations, tax filings, general ledger maintenance. Dashboards that let the firm, and its clients, monitor financial performance continuously. Anomaly detection in the ledger before the accounts are finalised; expected credit loss modelling; automated reconciliation across sources.
Management accounting	Dashboards monitoring projects and profitability; helping decision makers understand their own performance; KPI reporting built for the operational manager rather than the board.
Auditing	Automating routine audit procedures; risk assessment across the whole population; substantive procedures performed on every transaction rather than a sample; journal entry testing for unusual postings — round numbers, weekend entries, postings by unexpected users, amounts just below an approval threshold.
Tax accounting	Automating routine compliance work; analysing the tax efficiency of business units; identifying tax opportunities and evaluating them globally; dashboards that monitor tax positions in real time rather than at year end.
Forensic and fraud work	Benford's Law screening; matching employee bank accounts and addresses against the supplier master; detecting transactions split to stay under an approval limit; network analysis of related parties.

MANAGEMENT ACCOUNTING, IN MORE DETAIL

Five purposes account for most analytics work on the management side:

1. Using data more readily to **identify and manage risks**.
2. Improving **budgeting and forecasting**, by drawing on more data than a prior-year extrapolation.
3. **Automating internal reporting**, so the reporting cycle stops consuming the analyst.
4. Helping identify **operational improvements**.
5. Creating **dashboards of key performance indicators**.

The characteristic artefact is the **product performance dashboard**: operations or products separated into distinguishable categories, so that metrics and KPIs can be compared across them, trends made visible, and the comparison used as the basis for a recommendation. The value is in the categories chosen. A dashboard that groups everything into “sales” reveals nothing; one that splits sales by product line, channel and month makes the one deteriorating combination visible.

TAX, IN MORE DETAIL

Analytics assists tax work in three ways: maintaining **compliance**, by testing entire populations of transactions for correct treatment rather than sampling; **automating data collection** and routine tasks such as reconciliations, which leaves tax accountants time for analysis and planning; and building **tax dashboards** that let an organisation monitor changes in its tax position as they happen. A jurisdictional

apportionment dashboard is the standard example — a visualisation of the tax position by state or country, so a professional can see where exposure sits without assembling it by hand each quarter.

The four analytical methods

Every analysis belongs to one of four types, and naming the right one determines the data you need and the technique you use.

METHOD	OBJECTIVE	A QUESTION OF THAT TYPE	TYPICAL TECHNIQUES
Descriptive	Investigates <i>what</i> is happening now or has happened in the past.	What were gross sales by region for the past two years?	Sum, count, average, median, standard deviation, proportions
Diagnostic	Helps understand <i>why</i> something happened.	Why did sales decrease in Region 1 in the prior year?	Anomaly and outlier detection, trend analysis, pattern recognition
Predictive	Forecasts <i>what might happen</i> in the future.	What will sales be next year if we increase market share by 10%?	Forecasting, regression analysis, time-series analysis
Prescriptive	Helps understand <i>what should happen</i> to meet goals and objectives.	What is the most cost-effective way to ship our products?	Optimisation and what-if analysis

The four are cumulative in difficulty and in value, and most organisations that claim to do analytics are doing the first with better charts. An honest assessment of a firm’s analytics maturity says which of the four it actually reaches.

What the data looks like

STRUCTURED DATA	UNSTRUCTURED AND SEMI-STRUCTURED DATA
Organised into fixed fields and records; lives in relational tables. The general ledger, the customer master, the sales order table, a bank statement in CSV. Straightforward to query.	No predefined model: contracts, emails, board minutes, call recordings, images of invoices, social media, sensor output. Semi-structured formats such as JSON and XML sit between the two. Increasingly usable through text extraction and machine learning, and increasingly relevant as audit evidence.

Data is also either **internal** — the organisation’s own reports, transaction records and system data — or **external**: research, news, government statistics, geographic data, market and competitor information. Deciding which you need is part of planning the analysis, not part of doing it.

“Big data” is conventionally characterised by **volume**, **velocity** (the rate at which it arrives), **variety** (the mix of types above), **veracity** (how much of it can be trusted) and **value**. For an accountant the fourth is the interesting one: a very large dataset of uncertain provenance is a liability, not an asset.

Data quality

The phrase from the internal controls chapter — *garbage in, garbage out* — is the whole of data quality. The dimensions worth naming:


DIMENSION	THE FAILURE IT DESCRIBES
Accuracy	The value does not reflect the real-world event.
Completeness	Records or fields are missing; the population is not the population.
Consistency	The same entity is represented differently in different places — “BHP Ltd”, “B.H.P. Limited”, “bhp”.
Timeliness	The data is accurate as at a date that has passed.
Validity	Values fall outside the permitted domain — a date of 30 February, a negative quantity.
Uniqueness	The same transaction appears twice, inflating everything downstream.

THE CONNECTION BACK TO CHAPTER FOUR

Every one of those dimensions is the output of an input control that did or did not exist. Data quality is not a separate subject from internal control; it is internal control, observed at the far end. When the question is how to improve the reliability of a dataset, the strongest answers fix the process, not the spreadsheet.

The data analysis process

An analysis runs in three stages and six steps. The stages are **plan**, **analyse** and **report**; the six steps that fill them spell the mnemonic **MOSAIC** — motivation, objective, strategy, analysis, interpret, communicate.



STAGE 1 – PLAN

STEP	WHAT IT INVOLVES
Identify the motivation	Find the <i>reason</i> for performing the analysis. It may be external — a request from management, a regulator or another stakeholder — or internal : a desire to serve customers better, or to understand something affecting performance. An analysis with no motivation behind it produces a chart nobody uses.

STEP	WHAT IT INVOLVES
Determine the objective	Focus the analysis on a specific goal, and develop the specific questions that will guide it. “Are we collecting slowly?” is not an objective. “Has average days-to-collect for trade customers increased since the March credit policy change, and is the increase concentrated in particular customers or channels?” is.
Devise the strategy	Select the <i>type of data</i> and the <i>analysis method</i> . Which internal and external sources; at what granularity; which of the four methods above; and — the step most often skipped — what result would count as an answer either way, and what the data cannot tell you.

STAGE 2 – ANALYSE

STEP	WHAT IT INVOLVES
Prepare the data	Extract — retrieve the data from its source, such as a spreadsheet or database. Transform — clean, restructure and integrate it with other data so that it is valid. Load — import the transformed data into the software that will perform the analysis. Together, ETL , and in practice most of the work.
Model and visualise	Create the information — the statistics — the analysis needs, starting from the data collected. Calculating efficiency or profitability ratios from sales and expense data is modelling; so is building the comparison that a chart will display.
Explore	Interrogate the modelled data to generate insight: look for the anomaly, the trend, the group that behaves differently from the others.

STAGE 3 – REPORT

STEP	WHAT IT INVOLVES
Interpret	Review the analysis to be sure it makes sense against the objective, and that the results are valid and reliable . Investigate the causes behind the movements and trends found rather than reporting their existence. State the limitations honestly — missing periods, excluded records, assumptions made during cleaning.
Communicate	Deliver the finding to the decision maker in the form that lets them act. Orally, in writing, or visually — often a memo or presentation supported by visualisations and dashboards. The recommendation comes first, the evidence behind it second, the caveats attached.

THE TWO STEPS MOST OFTEN SKIPPED

Objective and **interpret**. Weak analytics work goes straight from “here is a dataset” to “here is a dashboard”, and then presents description as insight. If a piece of analysis produces a chart, the sentence beneath it should say what ought to be *done* differently, and what further evidence would make that recommendation more confident.

Worked example — LocalPics

THE SITUATION

LocalPics is a busy cinema chain with hundreds of sales each day and thousands on weekends. It has been tipped off by an employee that some workers are stealing cash from ticket sales, but the anonymous informant did not identify which of the seven cinemas the theft is occurring in. An accounting firm has been engaged to perform a fraud risk assessment. Set out the steps of the data analysis process for that engagement.

STAGE	STEP	WHAT IT MEANS HERE
Plan	Motivation	Discover whether theft is occurring. The motivation is external — a client engagement prompted by an allegation.
	Objective	Determine whether any employees are stealing cash, and if so, at which site and by whom. The informant’s inability to name the cinema is what the analysis has to overcome.
	Strategy	Collect ticket sales data and compare it with bank deposits, by cinema, by day and by shift. The method is diagnostic: what is happening, and why does one site differ from the others?
Analyse	Prepare the data	Extract sales and deposit records; determine whether they need cleaning or transformation — matching date conventions, aligning deposit dates to trading dates, removing card settlements from the cash comparison.
	Build the information model	Calculate what cash deposits <i>should</i> be, given recorded sales, for each cinema and each period. The gap between expected and actual is the statistic the whole engagement rests on.
	Explore the data	Examine the modelled differences for anomalies — a site with a persistent shortfall, shifts where the shortfall clusters, a pattern that appears only on the busiest days.
Report	Interpret	Review the analysis to be sure the differences are real and not artefacts — timing differences in banking, refunds, complimentary tickets, till float changes. Only what survives that review is a finding.
	Communicate	Prepare a memo or presentation summarising the findings, the sites and periods affected, the limitations of the analysis, and the controls that would prevent recurrence.

WHAT THE EXAMPLE TEACHES

The analysis never proves that a particular person stole anything. It narrows seven cinemas to one, and a month to a set of shifts, so that an investigation can be aimed. That is the honest description of what analytics contributes to a fraud engagement, and understating it is a better habit than overstating it.

The data analytics mindset

DEFINITION

A **data analytics mindset** is the professional habit of thinking critically throughout the data analysis process, before making and communicating a professional choice or decision.

SKILL	WHAT IT MEANS
Critical thinking	Disciplined reasoning used to investigate, understand and evaluate an event, opportunity or issue. Its central discipline in analytics is that correlation cannot imply causation : you must be aware of, and actively search for, hidden variables that explain both things you have observed.
Data literacy	The ability to understand and communicate data — to know what a figure represents, how it was produced and what it cannot support.
Technological agility	Awareness of current technological developments and a willingness to try new things, rather than forcing every problem into the tool you already know.
Reasoning	The process of logically forming conclusions, judgements and inferences from facts.
Communication	Conveying the result to someone who did not perform the analysis, in terms that let them act on it — the skill that determines whether any of the preceding work has an effect.

SPARKS – APPLYING THE MINDSET

Where MOSAIC describes the *steps* of an analysis, SPARKS is the critical-thinking check applied *at* each step: six things to interrogate before committing to a choice.

	ELEMENT	DEFINITION	WHY IT MATTERS
S	Stakeholders	The parties affected by the results of the analysis.	The potential impacts guide both the analysis and the communication choices.
P	Purpose	The reason driving the questions being asked.	Focusing on the purpose helps select the best data, analysis, interpretation and communication options.

	ELEMENT	DEFINITION	WHY IT MATTERS
A	Alternatives	Every choice should be the best option available.	Generating and ranking alternatives is what makes a choice a decision rather than a default.
R	Risks	All risks related to the data, the strategy, the biases and the assumptions.	Anticipating risk improves the data and analysis choices before they are locked in.
K	Knowledge	Proper strategy and interpretation depend on knowledge of accounting, economics, the industry, statistics and technology.	Domain knowledge is what gives the data and the results meaning; without it a result is a number.
S	Self-reflection	Reviewing each choice made across the process.	Reflecting on what worked improves the next analysis. It is also the only mechanism by which an analyst gets better.

HOW THE TWO FRAMEWORKS SIT TOGETHER

MOSAIC is the sequence: what you do, and in what order. SPARKS is the interrogation you run at each stage of that sequence — who is affected, what is the purpose, what were the alternatives, what are the risks, what knowledge does this require, and what would I do differently. Used together they answer the two questions an accountant is paid to answer about any analysis: *is it right*, and *should anyone act on it*.

Analytics and professional obligations

Three constraints distinguish an accountant's analytics from a data scientist's:


- **Privacy and consent.** Personal information collected for one purpose may not simply be repurposed for analysis. The Privacy Act 1988 and the Australian Privacy Principles apply to organisations with turnover above \$3 million, and to health service providers regardless of turnover — see Chapter Twelve.
- **Bias and fairness.** A model trained on past decisions reproduces the past, including its errors. A credit-scoring model built on historical approvals learns who was historically approved, not who was creditworthy.
- **Explainability.** An auditor cannot rely on evidence they cannot explain. A model whose output cannot be traced back to the transactions supporting it is not audit evidence, however accurate it happens to be.

That is the shape of the work. The next chapter is the toolbox — what you actually do to a file between receiving it and having something to say.

Analytics skills and tools for accountants

The practical chapter: getting data out of a system, cleaning it, analysing it and showing the result. It is best learned by doing it once on a real file, so the emphasis here is on the decisions that are easy to get wrong rather than on the keystrokes.

The workflow



The first three are **ETL**, and they consume most of the time in any real engagement. Expect roughly four-fifths of the effort to be here, and budget for it rather than being surprised by it.

Cleaning: the recurring problems

PROBLEM	WHAT TO DO, AND WHAT TO BE CAREFUL OF
Duplicate records	Deduplicate on a genuine key, not on the whole row — two legitimate identical transactions on the same day are not duplicates. Record how many you removed.
Missing values	Decide, and document, whether a blank means zero, unknown, or not applicable. These are three different things and treating them alike will change the answer.
Inconsistent formats	Dates stored as text, thousands separators, currency symbols inside numeric fields, trailing spaces, mixed case in names. Standardise before joining anything.
Inconsistent entity naming	The same supplier under four spellings. Match to a master file rather than to itself where possible.
Outliers	Investigate before excluding. In accounting data the outlier is very often the finding — the one payment fifty times the median is the interesting record, not the noise.
Wrong granularity	Header-level data joined to line-level data will double-count values. Check row counts before and after every join.

THE DISCIPLINE THAT MAKES AN ANALYSIS DEFENSIBLE

Keep the raw extract untouched and perform every transformation in a repeatable, documented way — a query, a script, a recorded list of transformation steps — never by typing over cells. Then state, with the result, the reconciliation: rows extracted, rows excluded and why, total value in the source against total value analysed. That paragraph is what makes an analysis auditable, and auditability is the accountant's distinctive contribution to it.

The tools, and what each is for

TOOL	GOOD FOR	WATCH OUT FOR
Excel	Everything small. Pivot tables for aggregation, XLOOKUP or INDEX-MATCH for joins, SUMIFS and COUNTIFS for conditional totals, Power Query for repeatable cleaning steps.	Row limits, silent formula errors, and no record of what was done by hand. Fine as a tool, dangerous as a process.
SQL	Getting exactly the data you want out of a database, and joining tables the size of a real ledger.	You must understand join types — an inner join silently drops unmatched records, which is how a population quietly stops being complete.
Power BI and Tableau	Interactive dashboards, drill-down, distribution to non-specialists, scheduled refresh.	A dashboard is a communication tool. It validates nothing, and a beautiful dashboard over unreconciled data is worse than no dashboard.
Python and R	Statistical modelling, large-scale text handling, anything repeated at scale.	More than most accounting work requires; know what they are for and when to ask for help.
Audit-specific tools <i>IDEA, ACL, CaseWare</i>	Population testing with a built-in audit trail of every command run.	Their advantage over Excel is precisely that log — the reason an auditor uses one is reproducibility, not capability.

THE SQL AN ACCOUNTANT SHOULD BE ABLE TO READ

```
SELECT c.customer_name,  
       SUM(i.amount) AS total_billed,  
       COUNT(*)     AS invoice_count  
FROM   invoices i  
JOIN   customers c ON c.customer_id = i.customer_id  
WHERE  i.invoice_date BETWEEN '2026-01-01' AND '2026-06-30'  
GROUP BY c.customer_name  
HAVING SUM(i.amount) > 100000  
ORDER BY total_billed DESC;
```

Six clauses, and they are all you need. **SELECT** which columns; **FROM** which table; **JOIN** to bring in a related table on a matching key; **WHERE** to filter rows before grouping; **GROUP BY** to aggregate; **HAVING** to filter the aggregated result; **ORDER BY** to sort. The join type matters: an **inner join** returns only matched rows, while a **left join** keeps every row from the first table even where no match exists — which is the one you want when the unmatched records are the point, such as invoices with no corresponding receiving report.

Techniques, by method

The four methods from Chapter Five map onto a specific and fairly short list of techniques. Choosing the technique is a consequence of having named the method, which is why the naming matters.

METHOD	TECHNIQUES
Descriptive	Sum, count, average, median, standard deviation, proportions. Aggregation and comparison — by period, by category, by entity.
Diagnostic	Anomaly and outlier detection, trend analysis, pattern recognition, drill-down from a summary figure into its components.
Predictive	Forecasting, regression analysis, time-series analysis, classification.
Prescriptive	Optimisation and what-if analysis — varying an input to find the combination that best meets an objective under constraints.

Accounting-specific tests worth knowing by name

TECHNIQUE	WHAT IT FINDS
Full-population testing	Every exception, rather than an estimate of the exception rate. Replaces sampling wherever the data supports it.
Gap and sequence analysis	Missing document numbers in a pre-numbered series — the deleted invoice, the unrecorded receipt.
Duplicate testing	The same invoice paid twice, or split into two payments to avoid an approval limit.
Threshold and clustering analysis	Transactions bunched just below an authorisation limit — a strong indicator of deliberate structuring.
Benford's Law	In many naturally occurring financial populations the leading digit 1 occurs about 30% of the time, falling away to about 4.6% for 9. A population departing materially from that distribution warrants investigation — a screening test that suggests where to look, never proof of anything.
Matching across files	Employee bank accounts or addresses appearing in the supplier master; supplier records created and paid by the same user.
Trend and ratio analysis	Movement against expectation — the classic analytical procedure, now performed at transaction level rather than on summarised

TECHNIQUE	WHAT IT FINDS
	accounts.
Correlation and regression	Relationships between variables, for forecasting and for cost behaviour analysis. Correlation is not causation — and the hidden variable that explains both is the thing to look for before concluding.

Business intelligence: the supporting infrastructure

TERM	MEANING
OLTP — online transaction processing	The operational systems that record transactions as they happen, optimised for writing many small transactions quickly and reliably. Your ERP.
OLAP — online analytical processing	Optimised for reading and aggregating large volumes across many dimensions. What analysis runs on.
Data warehouse	A repository of historical data drawn from multiple source systems and organised for analysis rather than for transacting — subject-oriented, integrated, time-variant and non-volatile.
Data mart	A subset of a warehouse serving one function or department: smaller, cheaper, faster to build.
Data lake	Raw data of all types stored in its native format and structured later, on use.
Data mining	Searching large datasets for previously unknown patterns and relationships — classification, clustering, association (“customers who buy X also buy Y”), regression.
Dashboard	A single screen presenting the indicators a particular role needs, usually with drill-down into the detail behind each.

Separating analysis from the transaction system is not only a performance decision. Running heavy queries against a live enterprise system degrades the system everyone else is working in, and it gives analysts read access to production data — a segregation-of-duties problem as much as a technical one.

Communicating the result

Choose the chart from the question, not from the software’s gallery:

IF THE QUESTION IS...	USE
Comparison between categories	Bar chart, sorted by value rather than alphabetically
Change over time	Line chart
Composition of a whole	Stacked bar, or a small table. Pie charts fail beyond about four segments
Relationship between two variables	Scatter plot
Distribution of one variable	Histogram or box plot
A single number against a target	A single number, large, with the target beside it

Then three rules that will not let you down: start value axes at zero for bar charts; label the axes and state the units; and put the finding in the title — “Collections have slowed by 11 days since March” rather than “Days sales outstanding”.

THE SENTENCE THAT HAS TO BE THERE

Every visualisation should be accompanied by a sentence saying what should be done differently as a result, and a second saying what would have to be true for that recommendation to be wrong. Without the first the analysis has no consequence; without the second it has no honesty.

Analytics needs a population to work on, and populations come from processes. Part Four takes the four processes that generate almost every transaction in an organisation and pulls each one apart — activities, documents, data stores, threats and controls.

PART FOUR

The business cycles

Four processes, one structure. Each chapter asks the same six questions — what is the cycle for, what technologies underpin it, what data does it hold, what activities does it perform, what can go wrong, and how is performance measured. Learn the structure once and the four chapters become variations rather than four separate memorisation jobs.

The revenue cycle

Selling something and getting paid for it. Four activities, a well-defined set of documents and data stores, and a set of threats that recur in almost every business that sells anything. This chapter and the next are mirror images of each other, and are best learned side by side.

DEFINITION

The **revenue cycle** is the recurring set of business activities and related information processing operations associated with providing goods and services to customers and collecting cash in payment. Its primary objective: **the right product, in the right place, at the right time, for the right price.**

How it connects to everything else

- To the **expenditure and production cycles** — sales data tells them when to buy or make more inventory.
- To the **HR and payroll cycle** — sales data drives commissions and bonuses.
- To the **general ledger and reporting cycle** — revenue, receivables and cost of goods sold flow through to the financial statements and the performance reports.

Every activity within the cycle performs the same three functions: capturing and processing data; storing and organising it for decisions; and providing controls to safeguard the assets involved.


The decisions the cycle exists to support

Before the mechanics, note what management is actually deciding here, because the data stores and controls that follow are all in service of these choices:

- **Product customisation** — how far the offering is tailored, and what that does to lead times and to inventory;
- **Inventory levels** — how much stock to carry against the risk of a stock-out;
- **Delivery methods** — own fleet or carrier, speed against cost;
- **Pricing**, including discount and rebate structures;
- **Credit policy** — to whom, how much, and on what terms;
- **Payment processing** — which methods to accept, and how quickly funds become available.

The common thread is **cash flow**. Each decision trades margin, service level or risk against the speed at which a sale becomes money in the bank.

The four activities



Data stores

The cycle's data divides into three kinds, and the distinction matters when you come to draw the process: master tables are referenced constantly and maintained occasionally, transaction files are created and updated as the work proceeds, and reference and history files support decisions without themselves being part of the transaction.

KIND	STORES
Master tables <i>standing data</i>	Customer table — a record of every authorised, regular customer: billing and delivery addresses, terms, credit limit, tax status. Inventory or items table — a record of each item stocked in the warehouse or regularly ordered from a vendor, with description, location, cost and price. Accounts receivable table — what each customer owes, aged, and the receipts applied against it.
Transaction and open document files <i>“open” means still editable, not yet finalised</i>	Sales order transaction table , with its sales order line table ; sales invoice transaction table ; cash receipts transaction table . Each is created when the transaction begins and updated as it progresses — a sales order is written on acceptance, updated when the goods ship, and updated again when the invoice is raised.
Reference and history files <i>support decisions and reporting</i>	Shipping and price data reference file; credit reference file; salesperson file; marketing data, being sales-oriented data from presales and sales activity; sales history file; cash receipts history file; accounts receivable reports file; undeposited funds table.

HEADER TABLE AND LINE TABLE

A **sales order table** holds one row per order — customer, date, order number, total. The **sales order line table** holds the line-by-line breakdown: if an order totalled \$100, which items were bought, the price of each, and any discount given. Almost every transaction file in an accounting system is a pair of this kind, and the pairing is the reason a careless join double-counts values (Chapter Six). When drawing a process, be explicit about which of the two an activity touches — confirming an order writes to both.

Documents, and what each one proves

DOCUMENT	PURPOSE — AND THE CONTROL IT SUPPORTS
Sales order	Records what the customer asked for. Everything downstream is checked back against it.
Customer acknowledgement	Notifies the customer that the order is accepted and gives the expected shipment date. A detective control — the customer will complain if it is wrong.
Picking ticket	Authorises the warehouse to pick the goods from the shelf and send them to shipping. Identifies the goods and usually the warehouse

DOCUMENT	PURPOSE – AND THE CONTROL IT SUPPORTS
	location. Without it, warehouse staff have no authority to release inventory.
Packing slip	Travels with the package, identifying the customer, destination and contents.
Bill of lading	The contract between the shipper and the carrier under which the carrier agrees to transport the goods to the customer. Evidence that the goods left, and on what terms.
Sales order notification	Message to accounting that a shipment is pending, so that billing knows to expect it.
Invoice	Notifies the customer of the obligation to pay for the merchandise ordered and shipped. Validated by comparing the sales order notification with shipping's notification.
Remittance advice	Accompanies the customer's payment and says which invoices it settles. Often a turnaround document detached from the invoice.
Credit memo	Authorises a reduction to the customer's account for a return or allowance. A high-risk document — it reduces a receivable without cash arriving.
GL inventory and sales update	Notifies the general ledger that inventory has been sold and cost of goods sold has increased.

Activity 1 — Sales order entry

The steps: the customer places the order; inventory availability is verified; credit approval is obtained; the accepted order is released to billing, shipping and the warehouse.

THREAT	CONTROLS
Incomplete or inaccurate customer order	Data entry edit checks; a verified customer list holding billing and delivery addresses; automated completion of stored customer details; mandatory fields.
Sales to customers with poor credit	A credit manager independent of the sales function sets credit policy and approves extensions of credit to new and existing customers. Sales staff may approve additional credit sales to existing customers only within the approved credit limit. Automatic system check of the order against the limit and the ageing. Critically, credit approval must occur before goods are released from inventory , not afterwards.
Orders that are not legitimate	Accept orders only against entries already in the customer table; credit sales to new customers carefully checked and approved by the credit manager; require business customers to supply a physical address; check for duplicated addresses and phone numbers between the customer file and the employee file, which is how a fictitious customer set up by a

THREAT	CONTROLS
	member of staff is caught; digital signatures or written orders for large accounts.
Losses from stock-outs, carrying costs and markdowns	Sales planning used to estimate and schedule purchasing and production; perpetual inventory system; just-in-time purchasing to reduce holding costs; vendor-managed inventory, giving the supplier direct access to sales data so that stock replenishes automatically; deliberate safety stock where the cost of a stock-out exceeds the cost of holding it; periodic physical counts.

THE SEGREGATION POINT THAT MATTERS MOST

Credit approval must be independent of sales. A salesperson paid on volume has an obvious incentive to approve credit that should be refused. This is an authorisation control, and its absence is the single most commonly planted weakness in a revenue cycle narrative.

Activity 2 — Shipping

The warehouse picks the goods against the picking ticket and transfers them to shipping; shipping compares what arrived with the sales order and the picking ticket; the goods go to the customer with a packing slip and bill of lading; shipping notifies billing.

THREAT	CONTROLS
Shipping errors — wrong goods, wrong quantity, wrong address	Compare the picking ticket to the goods picked; independent authorisation of shipment; compare the shipment to the sales order and picking ticket; review the file of open sales orders — the tickler file — to detect shipments that should have been made and were not, and to identify duplicate or unauthorised shipments; bar-code or RFID scanning at pick and at despatch.
Theft of inventory	Inventory kept in a secure location with access limited to responsible personnel; all inventory transfers documented; inventory released to shipping only against an approved sales order; both warehouse and shipping employees sign the transfer document; periodic physical counts reconciled to recorded quantities.

Activity 3 — Billing

The invoice is validated by comparing the details on the sales order notification with those on shipping's notification — the revenue-cycle equivalent of the three-way match.

THREAT	CONTROLS
Failure to bill the customer	Segregate shipping from billing; regularly review the tickler file of sales orders shipped but not yet billed; pre-number shipping documents and account for the sequence; send a copy of the invoice in the same package as the goods, so the customer sees at once what they are being charged for.
Billing errors — wrong quantity, wrong price	Compare the input shipping notice to the sales order table so that invoice, receivable and revenue reflect what was ordered and shipped; reconcile the quantities on the packing slip to those on the sales order and bill only what is on the packing slip; do not bill for items not shipped or on back order; use only authorised prices, terms, freight and discounts, taken from an independently maintained price file; separate the people who approve unit prices from the selling function.
Errors in maintaining customer accounts	Reconcile the receivables subsidiary ledger to the control account; send customer statements regularly, so the customer performs the detective control for you; independent review of adjustments and write-offs.

Activity 4 — Cash collections

Cash is the asset most exposed to misappropriation, so the controls here are the tightest in the cycle.

THREAT	CONTROLS
Theft of cash and cheques	Restrictively endorse incoming cheques immediately on receipt; minimise the handling of cash; use a lockbox so payments go directly to a bank-operated facility and never enter the office; use EFT and BPAY in preference to cheques; segregate cash handling from recording.
Lapping — concealing a stolen receipt with a later customer's payment	Separate the remittance advice from the cheque immediately; segregate the person who opens the mail from the person who posts to the receivables ledger; independent bank reconciliation; regular customer statements; mandatory annual leave and job rotation.
Posting errors and unapplied cash	Turnaround documents so the remittance advice is machine-read rather than re-keyed; reconcile the bank account regularly; monitor open receivables and unapplied receipts.
Fictitious write-offs and credit notes	Independent authorisation of every credit memo and write-off, by someone with no access to cash; review of credit notes by customer and by employee raising them.

Technologies underpinning the cycle

- **ERP** — integration of order, inventory, shipping and receivables data, which is what makes the cross-document checks possible;
- **EDI** — electronic exchange of orders and invoices directly between the customer's and the seller's systems, removing re-keying;

- **E-commerce and self-service portals** — the customer enters their own order, which removes a transcription risk and creates an authentication one;
- **Point-of-sale systems and bar coding / RFID** — capture at the moment of the event;
- **CRM** — the customer history behind the order.

Measuring performance

Useful KPIs, and worth quoting because “how would you measure this cycle?” is a standard part b: days sales outstanding; percentage of orders shipped complete and on time (order fill rate); percentage of orders taken by e-commerce or EDI rather than manually; number of credit notes as a percentage of invoices; bad debts as a percentage of credit sales; cost per sales order processed; customer complaints and returns.

Worked answer — XYZ Clothing

THE NARRATIVE

XYZ Clothing is an online retailer. Customers browse, add to cart and check out, entering billing and shipping information, then credit card details, and submit the order. The sales order is *automatically processed* and a confirmation page shows the order number and estimated delivery date. The order goes to the warehouse, where staff pick and package the items, and the package ships with an automatically generated shipping label. Invoices are automatically generated and then checked by the accounts receivable clerk against the sales order. The AR clerk reconciles customer accounts at month end. Where there are discrepancies she can raise queries with the sales, warehouse or shipping teams, *all of whom have access to customer data*. Customer data is backed up half-yearly by an IT staff member, and a copy of the customer file (a spreadsheet) is placed on the company’s database, with changes to that file linked and synchronised to the live customer data.

WEAKNESS	RISK IT CREATES	CONTROL
Orders are processed automatically with no credit or payment authorisation step described — the narrative says the card details are entered but never that authorisation is obtained before the order is accepted.	Goods are picked and shipped against a declined or fraudulent card; the sale is never collected and the inventory is gone.	Obtain payment authorisation from the card acquirer <i>before</i> the order is confirmed and released to the warehouse; block release of the picking instruction until authorisation is received; use address verification and 3-D Secure for cardholder authentication.
No verification of customer identity or of the delivery address against the billing address.	Fraudulent orders placed with stolen card details and delivered to an unrelated address; chargebacks borne by XYZ.	Address verification service; flag and manually review orders where the delivery address differs from the billing address, is new, or exceeds a value threshold.
No inventory availability check is described before the order is	Orders accepted for stock that does not exist; unfulfillable	System check of the order against real-time inventory availability before

WEAKNESS	RISK IT CREATES	CONTROL
accepted and a delivery date promised.	orders, refunds, reputational damage from a promised date that cannot be met.	confirmation; reserve the stock on acceptance; display back-order status to the customer.
Warehouse staff pick, package and ship with no independent check of what was picked against the order.	Shipping errors — wrong or short quantities — and theft of inventory concealed as a picking discrepancy. Custody and effective recording sit with the same person.	Independent check of picked goods against the picking ticket and sales order before despatch, evidenced by signature or scan; segregate picking from despatch; bar-code scanning at both points; periodic physical counts reconciled to the perpetual records.
The invoice is checked against the <i>sales order</i> only, not against evidence of shipment.	Customers billed for goods never shipped, or shipped goods never billed. Revenue and receivables misstated; the completeness assertion fails.	Match the invoice to the shipping record as well as the sales order; run a periodic report of shipments not yet billed and orders billed but not shipped, and clear it.
The AR clerk both generates or checks invoices and reconciles customer accounts, with no independent review of that reconciliation.	Errors and manipulation in the receivables ledger go undetected; the person who could create a fictitious credit is the person who reviews the account.	Independent review and sign-off of the monthly reconciliation by a supervisor; independent authorisation of credit notes and write-offs; send customer statements so customers detect errors.
Sales, warehouse and shipping teams <i>all</i> have access to customer data.	Unauthorised amendment of customer records — changing a delivery address to divert goods, or a bank or contact detail — and breach of privacy obligations over personal information.	Role-based access, read-only for those who only need to view; restrict amendment of the customer master to a small authorised group; log all changes to the master file and review the log; comply with the Australian Privacy Principles.
Customer data is backed up only half-yearly.	Up to six months of customer data lost on a failure. Recovery would be impossible in any practical sense; the business could not invoice or deliver.	Automated daily (or continuous) backup, held off-site or in a separate region, encrypted, with restoration tested at defined intervals as part of a documented disaster recovery plan.
Customer data is held in a spreadsheet that is linked and synchronised back to the live customer data.	An uncontrolled file with no access control, no audit trail, no validation and no version control can overwrite the production database. A single careless edit corrupts live customer records.	Eliminate the spreadsheet as a source of truth; maintain customer data only in the controlled application, with validation, access control and change logging. If an extract is needed, make it strictly one-way and read-only.

HOW TO WRITE UP A WEAKNESS

Each finding has three moving parts, and all three have to be present. **The weakness** — what the process does, or fails to do, stated as a fact about the narrative. **The risk** — what could therefore happen, expressed as a consequence to the business, not as a restatement of the weakness. **The control** — one specific, implementable measure that addresses that risk. “The supervisor distributes the cheques” is a weakness; “the supervisor can retain and cash the pay of a ghost employee” is the risk; “pay by EFT to verified accounts” is the control. Where there are more weaknesses than you can address, take the ones with the largest consequence first.

Worked answer — healthcare billing and payment

THE NARRATIVE

After a patient receives services, the provider’s billing department prepares a bill listing the services rendered and their costs. The bill is sent to the patient’s insurer. If the patient has no insurance, or the insurance does not cover the full amount, the patient is billed directly. The insurer reviews and either approves or denies the claim; if denied, the billing department must appeal. If approved, the insurer sends payment to the billing department, which posts it to the patient’s account. Any residual is billed to the patient, who can pay in person, by mail, online or by phone.

WEAKNESS	RISK	CONTROL
The billing department prepares the bill with no independent verification that the services listed were actually rendered.	Over-billing or under-billing; billing for services not provided, which in a healthcare setting is insurance fraud and a regulatory matter, not merely an error.	Reconcile the bill to the clinical record and the treating practitioner’s service log before it is issued; independent review of high-value claims; system-enforced links between a service code and a documented encounter.
The same billing department both prepares the bill and receives and posts the insurer’s payment.	Custody and recording combined: a receipt can be misappropriated and concealed by adjusting the patient’s account. Classic CAR breach.	Segregate cash receipting from billing and from posting to patient accounts; direct insurer payments to a bank account controlled by finance, not by billing; independent bank reconciliation.
No approved fee schedule or pricing authority is mentioned; costs for each service appear to be determined within billing.	Inconsistent or unauthorised charging; disputes and write-offs; revenue leakage.	Maintain an authorised fee schedule, maintained by someone outside billing, and generate charges from it automatically rather than by manual entry.
No tracking of claims submitted, denied or under appeal.	Denied claims are never appealed and simply lapse; approved claims are never followed up when payment does not arrive. Revenue is lost silently.	Aged claims report by status, reviewed weekly; escalation timeframes for appeals; reconciliation of claims submitted to payments received.

WEAKNESS	RISK	CONTROL
Patients may pay in person, by mail, online or by phone, with no controls described over any of those channels.	Cash and card details received across four channels with no receipting discipline — theft of receipts, and mishandling of card data.	Sequentially numbered receipts for in-person payments and daily banking intact; mail opened by two people with a listing prepared before it reaches billing; card payments processed through a compliant gateway, with card details never recorded on paper or in the patient file.
Patient health and payment information moves between provider, insurer and patient with no security or privacy control described.	Unauthorised disclosure of sensitive health information — the most strictly regulated category of personal information under the Privacy Act, and notifiable if breached.	Encrypt data in transit and at rest; role-based access to patient records; secure portals rather than email for claim documents; privacy training and a documented data breach response plan.

Reading a narrative for decision quality — AB Hi-Fi

THE EXTRACT

At 8am the warehouse officer collects picking tickets from the warehouse printer and picks the goods. If any are unavailable, the goods already picked are returned to the shelf and the picking ticket goes into an “awaiting goods” folder on the desk. If all are available, the goods and ticket go to shipping. The shipping officer scans the sales order number, and the computer displays the order. The shipping officer manually checks each item picked against the original sales order and the picking ticket, and checks the name on the sales order against the names highlighted on the *credit card status report*. If the name matches a name on that report, the goods and ticket are returned to the warehouse with a note that they cannot be shipped due to non-payment. Once satisfied that the credit card payment is not invalid and the correct items are being shipped, the shipping officer ticks a box on the sales order to indicate shipping is complete. The computer updates the sales order and inventory data and prints a delivery slip. The goods go to the loading dock and a carrier collects them at 3pm.

(a) What data does the shipping officer draw on? The sales order displayed on screen; the physical picking ticket; the physical goods themselves; and the credit card status report, which is a printed exception list of names with payment problems.

(b) Could there be errors in that data? Yes, and this is the point of the question — the officer is making an authorisation decision on data of uneven quality:

- The **credit card status report is matched on customer name**. Names are not unique. Two customers named J. Nguyen will produce a false match, blocking a legitimate shipment; a customer who spells their name differently on a second order will produce a false negative, releasing goods that should be held. The match should be on a unique customer identifier or order number, not on a name.
- The report is a **printed snapshot**, so it is out of date the moment it is produced. A payment that failed after printing will not appear; one that has since been settled will still appear.

- Names are **highlighted by hand**, so the control depends on someone highlighting correctly and on the officer reading the highlight correctly.
- The check of items against the order is **manual and visual**, so it is subject to fatigue and to the natural tendency to confirm rather than verify — particularly at volume. Scanning would remove the judgement.
- The picking ticket only reflects what was *meant* to be picked. If the warehouse officer picked the wrong item, the ticket does not know that; the only independent evidence is the goods.
- The “awaiting goods” folder on a desk is an **uncontrolled queue** — no record in the system that an order is held, no ageing, nothing to prevent a ticket being lost or a customer waiting indefinitely.

Turn the cycle around — cash out instead of cash in, goods in instead of goods out — and most of the same reasoning applies in mirror image. That is the expenditure cycle.

The expenditure cycle

Buying goods and services and paying for them. Structurally the mirror of the revenue cycle, but the risks are different in character: in the revenue cycle the danger is that value leaves without cash arriving; here it is that cash leaves without value arriving.

Objectives of the purchasing process

- Tracking purchases of goods and services from vendors;
- Tracking amounts owed;
- Maintaining vendor records;
- Controlling inventory;
- Making timely and accurate vendor payments;
- Forecasting purchases and cash outflows.

The optimum supply problem

TOO LITTLE INVENTORY	TOO MUCH INVENTORY
Lost sales; customers who do not return because they were let down; production stoppages; expensive emergency purchasing.	Storage and warehousing costs; insurance; staff to receive, count and monitor it; obsolescence and markdowns; and the opportunity cost of the cash tied up in it — the interest lost on money that could have been used elsewhere.

Everything in the cycle’s design is a trade-off between those two columns. Just-in-time supply pushes hard toward the left: goods arrive as they are needed, so carrying cost approaches zero — but JIT is only appropriate where the supply chain is reliable. If delays are likely, JIT converts a carrying cost into a stoppage.

Strategic implications

Purchasing is a genuine source of competitive advantage, and strategic fit matters: a cost-leadership firm and a differentiation firm should not buy the same way. The risks of getting it wrong are systemic — failure to manage purchasing creates problems that propagate into revenue and production, and poor payment practices damage both cash flow and supplier relationships, which is the sort of damage that is slow to repair.

Technologies underpinning the cycle

TECHNOLOGY	WHAT IT DOES
ERP	Improves the integration of enterprise-wide data — and makes the three-way match automatic.
EDI	Enables the exchange of data between two separate computer systems, eliminating the printing, posting and re-keying of orders and invoices.
RFID tags	A small tag attached to an item that can be scanned without line of sight, expediting stock handling. Unlike a barcode it can be read in bulk and rewritten.
SCM software	Improves the planning and execution of orders through supplier and customer integration.
Procurement cards	Controlled cards for low-value purchases, with limits by merchant category and value, removing hundreds of small transactions from the full purchase order process.

Inputs, outputs and data stores

INPUTS (DOCUMENTS)	OUTPUTS
Purchase requisition — the internal request to buy Purchase order — the external commitment to the supplier Vendor invoice — the supplier's claim for payment Receiving report — what actually arrived Bill of lading — the carriage contract Packing slip — what the supplier says is in the carton	Financial statement information Vendor payments Payment register Discrepancy reports Cash requirements forecast

DATA STORE	CONTENTS
Inventory master data	A record of each item stocked or regularly ordered.
Vendor master data	A record of each vendor approved for use by the organisation. A high-risk file — whoever can add a vendor can create a payee.
Purchase requisition data	All requisitions raised, by date and requester.
Purchase order master data	Open purchase orders, including the status of items on order, so commitments can be tracked.
Purchase receipts data	A record of each receipt of goods and services — quantity, condition, supplier, date.
Accounts payable data	What is owed, to whom, and when it falls due.

Business decisions in the cycle

STRATEGIC LEVEL	OPERATIONAL LEVEL
Purchase consolidation across sites and divisions How IT can improve efficiency and accuracy — automating purchasing Where inventories should be held	Determining the optimal inventory level Supplier selection on quality, service and price — a poor supplier can damage your reputation, not just your margin Cash flow considerations, including when to take discounts

The four activities



ACTIVITY	KEY STEPS
Determine demand	Collect requests from the departments; check available stock; create the purchase requisition. Forecasting demand is the analytical heart of the cycle, and is exactly where the professional practice report expects you to say something about data analytics.
Order goods	Choose the supplier; create the purchase order.
Receive goods	Accept the delivery; record the goods received.
Pay for goods	Receive and record the invoice; approve the payment; make the payment.

Threats and controls — purchasing

THREAT	CONTROLS
1. Stock-out or excess inventory	Accurate inventory control and sales forecasting; perpetual inventory system; bar coding or RFID; periodic physical counts; JIT acquisition including vendor-managed inventory to minimise carrying cost; warehouse bin location systems; locked storerooms.
2. High purchasing costs	Blanket purchase order — a commitment to buy specified items at specified prices from a particular supplier for a set period; EDI to eliminate the cost of printing and posting orders; vendor-managed inventory, where the vendor accesses point-of-sale and inventory data and replenishes automatically.
3. Purchasing items that are not needed	Integrate the databases of the various divisions, so existing stock elsewhere is visible; double-check the requisition against the requesting department.
4. Purchasing at inflated prices	Price lists for frequently purchased items; catalogues for low-cost items; competitive bids for high-cost and specialised products.
5. Purchasing goods of inferior quality	Approved supplier list; tracking of supplier performance; require quality certification such as ISO 9000; hold purchasing accountable for the cost of rework and scrap.
6. Purchasing from unauthorised or unreliable suppliers	Review of purchase orders; restricted access to the supplier master file; periodic review of the supplier list; procurement cards for small purchases.
7. Kickbacks	A “no gift” policy for buyers; employee training on gift handling; job rotation and mandatory annual leave; audits of buyers and of vendors; conflict of interest declarations, reviewed rather than merely filed.

Threats and controls — receiving

THREAT	CONTROLS
Theft of inventory	Secure storage locations; restricted access to inventory; documentation of intra-company transfers; periodic physical counts; segregation of the inventory custody function from receiving.
Accepting unordered items	Accept goods only where there is an approved purchase order on the system.
Mistakes in counting	Bar coding of ordered goods; quantities blanked out on the receiving copy of the order, so the receiving clerk must actually count rather than copy; signature of the receiving clerk; independent re-count by inventory control.

WHY THE QUANTITY IS BLANKED OUT

If the receiving clerk can see how many units were ordered, human nature supplies the count. Blanking the quantity converts a rubber stamp into a genuine independent verification — a small design decision that shows exactly what a control is *for*.

Threats and controls — paying

THREAT	CONTROLS
Paying for items not received	The three-way match : purchase order, receiving report and supplier invoice must agree on item, quantity and price before payment is approved. Automate it in the ERP so no payment can be released without all three.
Failure to take prompt payment discounts	File and track invoices by due date; prepare short-term cash flow budgets so the cash is available when the discount is worth taking.
Duplicate payments	Cancel invoices once paid; pay only on original invoices, never on statements or copies; approve invoices only with a complete voucher package; system check for duplicate invoice numbers and amounts against the same vendor.
Theft of cash	Restrict access to cash, cheques and cheque-signing machines or EFT facilities; use sequentially numbered cheques and account for the sequence; segregate cheque writing from accounts payable; require two signatures above a threshold; cancel all supporting documents once paid; have the bank reconciliation performed by someone independent of cash disbursements; log, encrypt, number and monitor all EFT transactions.
Fictitious suppliers	Restrict access to the supplier master file; limit the number of employees able to create a one-time supplier; independent review of all new and amended supplier records, especially bank account changes; periodically match supplier bank accounts and addresses against employee records.
Travel expenses and petty cash abuse	Procurement cards with category and value limits; surprise audits of the petty cash float; receipts required and reviewed by someone other than the claimant's subordinate.

THE SINGLE MOST VALUABLE SENTENCE IN THIS CHAPTER

“A payment should not be released unless the purchase order, the receiving report and the supplier invoice agree.” That one control addresses paying for goods not received, paying the wrong quantity, paying the wrong price and paying a supplier nobody ordered from. If a case narrative omits it, that is your first weakness.

Measuring performance

Purchase price variance against standard or against the last price paid; percentage of purchase orders raised without a requisition; percentage of invoices matched automatically without intervention; supplier

on-time delivery rate and defect rate; days payable outstanding; discounts lost as a percentage of discounts available; number of active suppliers per spend category; cost per purchase order processed.

The third cycle spends more of most organisations' money than the second, and unlike inventory it can walk out of the building on its own.

HR management and the payroll cycle

Taught last, and easy to underrate for that reason. Payroll is usually the largest single expense in a service business, it is paid to people who are inside the control environment, and its errors are both expensive and legally consequential.

What the cycle covers

The HR management and payroll cycle consists of all events involved in:

- hiring employees and maintaining personnel records;
- preparing and paying employees; and
- generating the reports and statements required by internal and external users — payroll tax, superannuation, Single Touch Payroll reporting to the ATO, workers compensation.

THE PRIMARY RESPONSIBILITY

To ensure that **human capital is paid for accurately and on time**. Both words matter. An accurate payment made late is a breach; a timely payment of the wrong amount is a different breach.

Strategic implications

A well-controlled payroll cycle is a competitive advantage in the plain sense that people leave employers who pay them wrongly. The most important organisational issue is **alignment of the HR and payroll cycle with the organisation's vision, culture, goals and strategy** — a cost-leadership firm and a differentiation firm need different workforces, and therefore different recruitment, remuneration and retention practices. Technology has changed the function itself: online recruitment and gamified assessment, self-service portals, remote and hybrid work, and analytics applied to turnover and engagement.

Technologies

- **ERP** — integration of enterprise-wide data, so time recorded against a job flows to costing, payroll and the ledger without re-entry;
- **Time clock scanners, swipe cards and biometric readers** — capture at the point of the event;
- **Online payment** of wages and salaries by EFT, which removes the physical cheque and with it a large class of fraud;
- **Employee self-service** for leave, expenses and personal details — efficient, and a new authorisation risk over bank account changes.

Data the cycle requires

- **Employee master data** — identity, tax file number, bank account, pay rate, entitlements, superannuation fund;
- **Time data** — hours attended and hours worked;
- **Job order data** — which job or cost centre the time is charged to;

- **Current period production cycle data**, where pay depends on output;
- **Payroll master file and payroll register** — the calculation and its result;
- **Accounts payable data** — for the payment of payroll liabilities to the ATO, superannuation funds and others.

The HR management activities

ACTIVITY	KEY STEPS AND DOCUMENTS
Employee recruitment	Selection against key selection criteria, shortlisting, interview. Induction — new employee form, tax file number declaration, superannuation choice, exception report of incomplete records.
Performance management	Assessment against agreed measures, and the adjustments — pay, role, development — that follow from it.
Employee termination	Identify terminations; manage the exit — final pay, entitlements, return of assets, and <i>removal of system access</i> , which is the control most often forgotten.

The payroll process activities



Threats and controls

THREAT	CONTROLS
Fictitious or “ghost” employees on the payroll	Separate the authority to add an employee to the master file (HR) from the authority to process the pay run (payroll); require documented approval and identity verification for every new starter; independent periodic review of the master file against HR records; distribute pay by EFT to verified accounts and investigate duplicate bank accounts.
Unauthorised changes to pay rates or bank details	Restrict amendment rights; log every master file change and review the log independently of the person who made it; notify the employee whenever their bank details are changed.
Errors in time recorded	Automatic capture of sign-in and sign-out; system logging of time against a job ticket at the point the work is retrieved; supervisor approval of timesheets; reconciliation of total time attended against total time recorded on jobs , which detects both over-claiming and mischarging.
Errors in the calculation	Batch totals and hash totals over the pay run; comparison of this run's totals with the last; exception reports for pays outside a normal range; independent review of the payroll register before release.

THREAT	CONTROLS
Theft of pay — a supervisor collecting or diverting pay	Pay by EFT direct to the employee's account ; if cheques exist, distribute them by someone independent of the supervisor, allow no collection on another's behalf, and audit unclaimed cheques.
Theft or misuse of payroll funds	Maintain a separate imprest payroll bank account , reimbursed with the exact net payroll amount, so that any residual balance is immediately visible; reconcile the payroll account independently.
Failure to remit payroll liabilities	Automated calculation and scheduling of PAYG withholding and superannuation; reconciliation of liabilities accrued to amounts remitted; diary of statutory deadlines.
Unauthorised disclosure of employee data	Restricted access to the employee master file; encryption; privacy obligations under the Australian Privacy Principles apply to employee data held by most organisations.

Worked answer — the Golf Club Company

THE NARRATIVE

The Golf Club Company makes custom golf clubs. The manufacturing supervisor interviews people with specialised manufacturing skills and informs payroll when an employee is hired. Employees use a time clock to record the time spent on each order. The supervisor approves all time cards. The accountant analyses job tickets and prepares a labour distribution summary. Payroll prepares the payroll register and paycheques. The supervisor distributes the paycheques to the employees. Payroll informs cash disbursements of the funds required to cover the entire payroll amount. The cash disbursements clerk ensures there are adequate funds in the company's regular cheque account to cover the payroll.

The question asks for three internal control weaknesses and the improvements you would make. There are more than three; these are the strongest.

WEAKNESS	RISK	IMPROVEMENT
The supervisor interviews and hires alone, and personally informs payroll of the hire.	The supervisor can add a fictitious or unqualified employee to the payroll and, since they also distribute the pay, collect the money. Authorisation and custody in one pair of hands.	Conduct interviews jointly with an HR representative; require HR, not the supervisor, to notify payroll of a hire; complete a new employee form signed by both the employee and the supervisor; verify identity and qualifications independently.
The supervisor approves all time cards for the employees they hired and supervise, with no reconciliation of time recorded to time actually attended.	Inflated hours, whether by employee or by supervisor, are undetected. Job costs and payroll are both overstated.	Automatic capture of sign-in and sign-out times, reconciled to the time recorded against job orders, with the exceptions reviewed by someone above the supervisor.

WEAKNESS	RISK	IMPROVEMENT
The supervisor distributes the paycheques.	The supervisor can retain the cheque of a ghost employee, or of a departed employee not yet removed, and cash it. This is the classic ghost-employee fraud, and the distribution step is what completes it.	Pay by EFT directly to employees' verified bank accounts. If cheques must be used, have a cashier independent of the supervisor distribute them against identification, prohibit collection on behalf of a colleague, and audit unclaimed cheques.
The payroll is funded from the company's <i>regular</i> cheque account, and the cash disbursements clerk merely ensures funds are adequate.	Payroll transactions cannot be distinguished from other expenditure, so the payroll bank reconciliation cannot detect an unauthorised or altered payment. Errors and fraud hide in the general flow.	Establish a separate imprest payroll account , transferring exactly the net payroll each period, so any remaining balance signals a discrepancy; reconcile it independently against the payroll register; support it with short-term cash budgeting so the funds are there.

THE PATTERN TO CARRY INTO ANY PAYROLL CASE

Payroll fraud almost always requires two of the three CAR responsibilities, and in a payroll case the person who has them is nearly always **the supervisor**, because supervisors naturally sit close to hiring, to time approval and to distribution. Find the supervisor in the narrative and ask what else they touch.

Multiple-choice traps in payroll questions

A more useful question than “which of these is a control?” is “which of these is the *weakest* control?”, because in payroll several genuine controls are usually available and they are not equal. The weakest is generally the one that relies on **a person choosing to be diligent** rather than on a system preventing the act:

- Colleagues attesting to each other's recorded time is weaker than a system that records it automatically — peers have every incentive to cooperate.
- A background check of candidates is weaker at preventing a fictitious employee than requiring a second person in the hiring decision — a fictitious person passes no check because nobody runs one on them.
- “Careful cash budgeting to ensure sufficient funds” is not really a control over misappropriation at all; it addresses liquidity, not integrity.

Measuring performance

Payroll cost as a percentage of revenue; error rate and number of off-cycle corrections per run; cost per payslip processed; time from timesheet close to payment; staff turnover and time-to-fill for vacancies; percentage of employees using self-service; compliance measures such as superannuation remitted by the due date.

Every cycle so far ends in the same place: a set of postings that have to be summarised, adjusted and reported. Before that comes the skill this subject rewards most directly — describing a process precisely enough that its weaknesses become visible.

Documentation, the general ledger and reporting

Two topics in one week. The first — how to draw a business process — is the one part of this subject that is genuinely a skill rather than a body of knowledge, and the only way to acquire it is to draw one. The second closes the loop: the cycle that takes every other cycle's output and turns it into financial statements.

Why document a process at all

- To **understand** it — you cannot find a weakness in a process you cannot see;
- To **evaluate controls**, which is what an auditor is doing when they walk a transaction through;
- To **design or redesign** it, and to specify a system that implements it;
- To **train** new staff, and to survive the departure of the person who knew how it worked — the exact failure that made the Woolworths ERP project worse (Chapter Eleven);
- To provide an **agreed record** between business and IT of what the system is supposed to do.

DEFINITION

Business process model and notation (BPMN) is a universal set of graphical symbols and rules for producing clear, readable diagrams of business processes — universal in the sense that the same diagram is intelligible to the business people who own the process and to the technical people who will build the system for it. It uses **swimlanes** to show which person, department or system is responsible for each step, and depicts the sequence of a process from start to finish using connecting arrows and distinct symbols for events (what happens), activities (the work done) and gateways (decision points).

The three notations

METHOD	SYMBOLS	SHOWS / DOES NOT SHOW
Flowchart	Five basic symbols — terminator, process, decision (diamond), document, and flow line; plus data store and connector.	Shows a sequence of activities and decisions; easy for accountants to read. Does not show who performs each step, nor the flow of data into and out of data stores. Often lacks detail.
Data flow diagram	Only four — agent (square), process (rounded rectangle), data store (three-sided rectangle), data flow (arrow).	Shows where data comes from, what transforms it, and where it is stored. Does not show decisions at all — there is no gateway symbol — nor timing or sequence.
BPMN 2.0	Pools and swimlanes, activities, gateways, events, data stores, data objects, sequence flows and data (message) flows.	Shows who does what, the activities, the decisions, the data stores and objects, and the entities that start and end the process. Combines the advantages of the flowchart and the data flow diagram, which is why it has become the standard.

THE COMPARISON, IF IT IS ASKED AS A SHORT ANSWER

A flowchart shows *sequence and decisions* but not *who* or *what data*. A DFD shows *data and stores* but not *decisions*. BPMN 2.0 shows all of it — swimlanes give the actor, activities give the work, gateways give the decisions, data stores and objects give the information. That is the answer, in that order.

The BPMN 2.0 elements

ELEMENT	WHAT IT REPRESENTS, AND THE GRAMMATICAL CLUE IN A NARRATIVE
Pool and swimlane	The participant, and the roles within it. One pool per organisation; lanes within the pool for its departments. The subject of the sentence — “ <i>the warehouse officer picks the goods</i> ”.
Activity (rounded rectangle)	Work performed. The verb — picks, enters, prepares, ships, generates. Name it verb-plus-object: “Pick goods”, not “Picking”.
Gateway (diamond)	A decision or a divergence in flow. The verbs are <i>verifies, checks, assesses, approves, if</i> . A gateway must be a question with labelled outgoing paths — yes and no.
Event (circle)	Start, intermediate and end events. The thing that triggers the process and the thing that concludes it.
Data store (cylinder)	A persistent table or file the process reads from or writes to — the customer table, the inventory table, the sales order table.
Data object (page symbol)	A document or message produced or consumed by an activity — an invoice, a notification email, a picking ticket.
Sequence flow (solid arrow)	The order of activities within a pool.
Message / data flow (dashed arrow)	Information passing between pools, or between an activity and a data store.

READ THE LANES FOR SEGREGATION OF DUTIES

A swimlane diagram makes segregation of duties visible in a way that prose does not, because it puts every activity in a named lane. Scan the lanes and ask which of custody, authorisation and recording each one holds. If the picking ticket is created in the warehouse lane, the people who take custody of the goods are also authorising their own release — a control weakness that is obvious in a diagram and easy to miss in a paragraph. Reading a process map for control is the same exercise as Chapter Four, performed with a picture.

The method — narrative to diagram

Follow it mechanically. Every element missed is a fact about the process that the diagram fails to record.

- 1 **Read the narrative and annotate it.** Highlight in one colour every actor, in another every activity verb, in another every decision verb, and underline every document or file mentioned.
- 2 **Fill in a business process map organiser** — a four-column table of swimlane, activity, gateway and data store or data object. The table is where the thinking happens; the drawing that follows is transcription.
- 3 **Add what the narrative does not say.** A narrative never mentions every update and notification. If an order is accepted, the sales order table is updated. If inventory is committed, the inventory table is updated. If a customer is told anything, that is a data object.
- 4 **Draw the diagram** from the completed organiser, one lane per actor, in sequence.
- 5 **Check every arrow.** This is where diagrams quietly go wrong.

THE ARROW RULE

An arrow running from a data store to an activity or gateway means the data is being *referenced* to enable that step. The items table pointing at “accept order” says the system looks up stock before accepting.

An arrow running from an activity to a data store means the data is being *updated* to record what happened. “Record payment” pointing at the accounts receivable table says the receipt is written to the ledger.

Some relationships are **both**, and must be drawn both ways. Creating an invoice reads the sales order table to find what was ordered *and* writes back to it to record that the order has now been invoiced. Verifying a credit card sends details to the card company *and* receives the result. A reversed arrow states the opposite of what happens; a missing arrow hides a data flow altogether. A diagram with the right boxes and careless arrows can be substantially wrong while looking right.

WORKED EXAMPLE – A SALES ORDER NARRATIVE


NARRATIVE

“The customer chooses the desired products from the website and places the order. The system verifies the customer’s identity against the customer files and approves the new customer if required. Once the bank verifies the credit card details and stock availability is confirmed, the order is approved.”

SWIMLANE / POOL	ACTIVITY	GATEWAY	DATA STORE OR OBJECT
Customer	Choose products; place order	—	Order (data object) sent to the sales system
Sales system	Receive order; verify customer identity	Is the customer identified? If not, approve as a new customer	Reads and writes the Customer table

SWIMLANE / POOL	ACTIVITY	GATEWAY	DATA STORE OR OBJECT
Sales system	Request credit card verification	Are the card details valid?	Message flow to and from the credit card company pool — bidirectional
Sales system	Check stock availability	Is stock available?	Reads the Inventory / items table
Sales system	Approve order	—	Writes the sales order table and the sales order line table ; sends an order confirmation (data object) to the customer
Credit card company	Verify card details	—	Returns the verification result (data object)

Drawn out, the completed organiser becomes this. Notice that the diagram carries information the table only implied — which decision leads where, and which table is read as against written.



Sales order entry. Arrows from a data store are *reads*; arrows into one are *updates*. The card verification runs both ways, so it is drawn as a pair of message flows. Every rejected path produces a notification to the customer — none of which the narrative mentions.

Three things in that diagram are not in the narrative and had to be supplied: the **writes to the sales order and sales order line tables** when the order is confirmed, the **notifications** sent to the customer on every path including the rejections, and the fact that the credit card flow runs **both ways**. The general rule: update every table the activity affects, including inventory, so that the business does not believe it holds stock it has already committed to a customer.

WORKED EXAMPLE – MIAPIZZA

NARRATIVE

“MiaPizza customers place an order by calling the store. They place their order after selecting the type of pizza. Often, customers will call the store if their order is not delivered after 60 minutes. They pay for the pizza when it is delivered.” (Only the customer’s side of ordering and receiving; not the delivery process itself.)

SWIMLANE	ACTIVITY	GATEWAY	DATA OBJECT
Customer	Select type of pizza	—	—
Customer	Call store and place order	—	Order details (spoken)
Customer	Wait for delivery	Has 60 minutes elapsed without delivery? If yes, call the store to follow up	Follow-up call
Customer	Receive pizza and pay	—	Payment

Two features are worth noticing because they generalise. The 60-minute wait is a **timer event** in BPMN, and the follow-up call is the path that leaves the gateway when the timer fires. And the scope statement — only the customer’s side — settles which pools to draw: the kitchen is out of scope, so drawing it would describe a process nobody asked about.

MATCHING EXERCISE – TERM TO EXAMPLE

EXAMPLE	BPMN TERM
Customer; Credit card company; Warehouse	Swimlane (or pool)
CUSTOMER table	Data store
Verify customer credit; Generate shipping labels; Authorise shipping	Activity
Customer lack-of-credit notification email; Shipping confirmation notification	Data object
Is customer credit verified?	Gateway

THE DISTINCTION MOST OFTEN BLURRED

Data store versus data object. A data store is a persistent file or table the organisation keeps — the customer table, the inventory table. A data object is a single document or message produced or used by one activity — an invoice, a notification, a picking ticket. “CUSTOMER table” is a store; “email telling the customer their credit was declined” is an object. Think worksheet versus document.

Critiquing a diagram against a narrative

The other form the question takes is a completed BPMN diagram and the narrative it was drawn from, with the instruction to find what is wrong. Work through this checklist:

- Is every actor in the narrative given a lane, and is every lane's work actually theirs?
- Is every activity in the narrative on the diagram, and is every activity on the diagram in the narrative?
- Does every gateway have a question and *both* labelled outgoing paths? A decision with one exit is not a decision.
- Are the data stores that must be read shown as read, and those that must be updated shown as updated? **A missing update is a missing fact about the process.**
- Is every notification the process must send actually drawn as a data object — even where the narrative does not mention it?
- Do the arrows point the right way, and is anything that must be bidirectional drawn as bidirectional?
- Does the process have a start event and at least one end event?

The general ledger and financial reporting cycle

This cycle summarises, adjusts and reports on the data produced by all the operational cycles. It is not an exchange with an outside party; its objective is to **synthesise and report**, including comparison against budget.

It extracts transactional data from accounts receivable (revenue cycle), accounts payable (expenditure cycle), payroll data (HR cycle) and production cycle data. Key business decisions relate to **budgetary considerations** — planning and controlling what resources are available — and to **accounting policies and procedures**. The **chart of accounts** establishes the basis on which every report can be generated, which is why designing it badly is expensive for years.

TECHNOLOGIES

- **ERP** — integration of enterprise-wide data, so the ledger is updated from the cycles rather than re-keyed;
- **Online banking** — enables reconciliation between amounts sent and amounts recorded;
- **XBRL** — a data standard used when generating financial reports, so they flow to users with minimal rework.

ACTIVITIES, RISKS AND CONTROLS

ACTIVITY	KEY STEPS	RISKS	COMMON CONTROLS
Prepare budgets	Determine budget values	Under- or over-estimating revenue and expenditure	Independent approval of budget estimates; preparation by the operational managers accountable for them; tight linkage between budget values and the performance monitoring system
	Record budget details	Data entry errors	Edit checks on input; reasonableness checks; batch totals

ACTIVITY	KEY STEPS	RISKS	COMMON CONTROLS
Update the general ledger	Extract and validate data	Incomplete or inaccurate data	Batch totals and hash totals; use of ledger control accounts
	Post transactions	Inaccurate updating	Batch totals; regular control account reconciliations
Record GL adjustments	Prepare adjustment journals	Errors in journal entries; unsupported or fraudulent adjustments	Independent approval of journal entries by senior accounting staff; full working papers attached to support the calculation
	Post adjustment journals	Data entry errors	Batch and hash totals; one-for-one checking of general journal inputs
Produce reports	Management reports; financial reports	Reports that are inaccurate, late, or that give access to those who should not have it	Access controls over report generation; review before issue; reconciliation of report totals to the ledger

WHY MANUAL JOURNALS ARE THE HIGHEST-RISK ITEM IN THE CYCLE

Every other posting in the ledger arrives from a cycle with its own controls — an invoice matched, a payment authorised, a payroll run reconciled. A manual journal has none of that history: it is a person typing a number directly into the financial statements. That is why the controls are approval by someone senior, mandatory supporting workpapers, and why journal entry testing is the first thing a data analytics audit looks at.

XBRL

DEFINITION

XBRL — eXtensible Business Reporting Language — is a standard used for digital financial reporting worldwide. It adds electronic tags to each financial amount, describing what the amount is. The tags are **metadata**: data about data.

A useful comparison is the product barcode, which can be added after the can is manufactured or applied during manufacture. XBRL is the same:

BOLT-ON TAGGING	INTEGRATED TAGGING
Tags are added after the reports have been created. Simpler to adopt; the tagging is a separate step that can itself contain errors.	Data is tagged during information processing, so the tag travels with the amount from the transaction onwards. More work up front, far more reliable.

BENEFITS	LIMITATIONS
Accuracy and accessibility; no re-keying; instant user access; more frequent updates; automated consumption by analysts and regulators; software independent; easily searched; less effort to use once implemented	New users must learn and conform to the standards of the language; evolving standards require continuing conformity to changing specifications; there is no requirement for auditors to provide assurance over the tagging; and the costs of transition are real

The third limitation is the interesting one for an accountant: a mis-tagged amount is machine-readable and wrong, which is worse than being unreadable and wrong.

Integrated reporting

An **integrated report** is a single report containing both financial and non-financial information about an organisation's economic, social and environmental performance. It matters here because the AIS is increasingly asked to capture non-financial data — emissions, safety incidents, workforce composition — to the same standard of reliability as financial data, which means it needs the same controls. That is a live issue in Australian practice as climate-related financial disclosure becomes mandatory.

Measuring the cycle

Days to close the month and the year; number of manual journals and their value as a proportion of postings; number and value of post-close adjustments; reconciliations completed on time; audit adjustments raised; forecast accuracy against actual. The maxim holds: *you measure what you aim for, and you manage what you measure* — which is also its own warning, because a cycle measured only on speed of close will close quickly.

All of this assumes the system exists. Part Five deals with what happens when it does not, and with what happens when the people running it decide the rules do not apply to them.

PART FIVE

Building and governing systems

How an organisation gets a new system, and what it owes the people whose data ends up inside it. The first is where most of the money and most of the failures are; the second is where the reputational damage is.

Systems development and acquisition

Five phases, a long list of alternatives to those five phases, and a decision about whether to build the thing at all. The skill worth acquiring is judgement: given a particular organisation, which approach fits, and why.

DEFINITIONS

The **systems development life cycle (SDLC)** is the progression through the phases of the systems development process. An **SDLC methodology** is a formalised, standardised and documented set of activities used to manage a systems development project.

Why it goes wrong — Woolworths

In 2015 Woolworths moved from a thirty-year-old in-house ERP to SAP, at a cost of around \$200 million. Suppliers waited up to eighteen months for payment. Store managers lost the tailored profit and loss reports they had run the business on. In 2016 the company wrote down \$766 million and laid off hundreds of employees.

REASON FOR FAILURE	WHICH PHASE IT BELONGS TO
Functional errors — inability to send orders to suppliers	Design and implementation: requirements not properly specified or not tested
Decreased accounting oversight of performance — new data collection methods incompatible with the old systems, and loss of institutional knowledge	Analysis: the current system was never adequately understood before it was replaced
Process documentation lacking — knowledge lived in the minds of key personnel, not in the system, and changes in project management lost it	Analysis, and a failure of documentation generally
Lack of senior management buy-in	Planning and investigation
Lack of integration with business strategic goals	Planning: the IT strategic plan was not aligned with the business plan
Extended timeline — poor planning and scheduling	Planning and scheduling feasibility

IF ASKED “IN WHICH STEP DID THE PROJECT GO WRONG?”

The honest answer is **investigation and analysis**, not implementation. The visible failures happened at go-live, but their causes were the failure to understand the existing processes, to capture the knowledge held by people rather than systems, and to secure senior sponsorship — all of which belong to the first two phases. Implementation failures are usually analysis failures arriving late.

Phase 1 — Investigation

Goal: identify problems with or opportunities in the current systems, and establish the feasibility of responding to them. Feasibility has five dimensions, and a project that clears four of them still fails:

FEASIBILITY	THE QUESTION IT ASKS
Economic	Will the benefits justify the time, money and other resources required?
Technical	Can it be developed and implemented using existing technology?
Legal	Does it comply with all applicable laws, statutes, regulations and contractual obligations?
Scheduling	Can it be developed and implemented in the time allocated?
Operational	Does the organisation have access to people who can design, implement and operate it — and <i>will</i> people use it? The second half of that question is the one projects fail on.

Economic feasibility is assessed with the usual capital budgeting models: **payback period** (how long until the investment is recovered — simple, ignores everything after payback and ignores the time value of money); **net present value** (future cash flows discounted to today; accept if positive); and **internal rate of return** (the discount rate at which NPV is zero, compared against the organisation’s minimum acceptable rate).

Phase 2 — Systems analysis

Two parts, in this order: **understand what the current system does and how it operates**, then **specify what the new system will need to do**. Skipping the first is the Woolworths error.

The **systems survey** aims to gain a thorough understanding of company operations, policies and procedures; data and information flows; AIS strengths and weaknesses; and available hardware, software and personnel. It then makes preliminary assessments of current and future processing needs, determines the extent and nature of change required, collects data identifying user needs, conducts the feasibility analysis and makes recommendations to management.

DATA GATHERING METHOD	WHAT IT IS GOOD FOR – AND ITS WEAKNESS
Interviews	Depth — why is there a problem, and why does it need addressing? Time-consuming, and subject to what the interviewee wants you to hear.
Questionnaires	Gathering a small, well-defined amount of data from many people cheaply. Cannot probe, and low response rates bias the result.
Observation	Validates what interviews and questionnaires claimed. People work differently when watched, and observation is expensive.
System documentation	Describes how the AIS is <i>supposed</i> to work — process maps, logical and physical DFDs, systems flowcharts. Frequently out of date, which is itself a finding.

USE MORE THAN ONE

Each method's weakness is another's strength: documentation says what should happen, interviews say what people believe happens, observation says what happens. Where the three disagree, you have found something.

Phase 3 — Systems design

Goals: evaluate design alternatives, prepare design specifications, and select the appropriate system. **Structured systems design** works backwards from what is needed: *output* → *data storage* → *input* → *processing procedures and operations*. Deciding the outputs first is deliberate — the system exists to produce them.

REPORT TYPE	CONTENT AND FORMAT / TIMING
Scheduled reports	Pre-specified content and format, prepared regularly.
Special-purpose analysis reports	No pre-specified content or format, not prepared on a schedule.
Triggered exception reports	Pre-specified content and format, prepared only in response to abnormal conditions.
Demand reports	Pre-specified content and format, prepared only on request.

Design must build in controls — *garbage in, garbage out*. The key control concerns to specify at design time: validity, authorisation, accuracy, security, numerical control, availability, maintainability, integrity and audit trail. Controls retro-fitted after go-live are more expensive and less effective, which is the reason the accountant belongs in the design phase rather than the acceptance test.

Phase 4 — Systems implementation

DEFINITION

Systems implementation is the set of procedures performed to complete the design contained in the approved systems design document, and to test, install and begin using the new or revised information system.

- 1 Complete the systems design.
- 2 Write (or configure), test and document the required programs and procedures.
- 3 Complete the user manuals.
- 4 Educate and train personnel.
- 5 Test — and test again.
- 6 Plan, control and conduct an orderly installation of the new system.

TESTING

TEST	WHAT IT ESTABLISHES
System test	Verifies the new system against the original specifications. Performed first by the development team, then by users with the team's assistance.
Acceptance test	A user-directed test of the complete system in a test environment. The users decide whether it does what they asked for.
Operations or environmental test	Runs a subset of the system in the actual production environment — because a system that works in test can still fail against real volumes, real interfaces and real hardware.

WHY EXHAUSTIVE TESTING IS NOT A STANDARD

It is tempting to say a system should be tested against every possible permutation of valid and invalid input. It cannot be — it is impossible for any non-trivial system, since the input space is effectively unbounded. Testing is **risk-based**: representative valid cases, boundary values, and classes of invalid input, chosen so that each class is exercised at least once. Exhaustive testing is not a standard anyone can meet.

DOCUMENTATION

- **Designers and programmers** — comment lines in programs, system flowcharts, program flowcharts;
- **Operators** — run manuals;
- **Users** — instructions on how to use the system, tutorials, help features;
- **Accountants and auditors** — all of the above, plus document flowcharts showing where the audit trail runs.

CONVERSION

Converting data and processes, with controls ensuring the accurate, complete and authorised conversion of both data and programs. Users suggest control totals to test the accuracy and completeness of the data conversion; program change controls ensure that only authorised, tested and approved versions are promoted to production.

CONVERSION STRATEGY	TRADE - OFF
Direct (“big bang”)	Switch off the old, switch on the new. Cheapest and fastest; no fallback if it fails.
Parallel	Run both systems simultaneously and compare outputs. Safest, and the most expensive — every transaction is processed twice by staff who have one job’s worth of time.
Phased	Convert one module or function at a time. Spreads the risk and the effort; requires temporary interfaces between old and new, which are themselves a risk.
Pilot	Convert one site or business unit first, then roll out. Contains the damage and produces a tested playbook; slower to complete.

Phase 5 — Maintenance and review

A **post-implementation review** examines a working system soon after implementation to determine whether the users’ requirements have been satisfied and whether the development effort was efficient and conducted in accordance with the organisation’s systems development standards. It looks at:

- **User requirements** — is the user satisfied; how does achieved performance compare with the performance requirements; what is the quality of the documentation, training and data conversion;
- **Development effort** — was the project management framework and the SDLC followed; should the standards manual be improved; how did actual costs and benefits compare with the projections, so that future estimates are better.

Maintenance is classified as *corrective* (fixing defects), *adaptive* (responding to a changed environment — new legislation, a new operating system) and *perfective* (improving performance or usability). It should be handled quickly and efficiently, should not cause other problems, should be in the organisation’s overall best interest, should be necessary, should be documented, and should minimise risk and disruption. The procedural controls: analysis and cost-benefit study before the change; charging user departments for maintenance costs so demand is disciplined; formal change request procedures with review and prioritisation; feedback gathering; program change controls; and high-quality documentation maintained as the change is made.

Across all phases — planning and people

New systems are disruptive, which is why planning takes a broad view and why the study team is drawn from beyond IT.

WHO	ROLE
Management	Provide support, resources and the authority to make the change stick.
Steering committee	Plan and oversee the IS function; set IS policies; prioritise between competing projects.
Project development team	Guide, plan and monitor the project.
Systems analysts and programmers	Study the existing AIS and design the new one.
Accountants	User, analyst, purchaser, implementer, consultant, internal auditor, external auditor — several roles across the SDLC.

WHAT THE ACCOUNTANT CONTRIBUTES

Three things, and they are the answer to “why is an accountant involved in systems development?”: providing technical expertise on **financial reporting requirements**; specifying **documentation standards** adequate for audit; and **verifying the adequacy of controls** at design time rather than after go-live.

Sitting above all of it, the **IT strategic plan** must be aligned with the organisation’s strategic plan, must provide an IT architecture linking users, applications and databases, and must allocate IT resources effectively among competing projects.

Buy, build, or neither

THE SELECTION PROCESS

Prepare a **request for proposal (RFP)** — a document sent to vendors inviting submission of plans for providing hardware, software and related services. Then evaluate the proposals: validate what the vendors claim (site visits, reference customers, benchmark tests), consider criteria beyond price, and assess the resources each option requires. **Point-scoring analysis** weights each requirement and scores each vendor against it, which makes the decision defensible — and is only as good as the weightings, which is where the judgement actually sits.

For each design alternative the key questions are: how well does it meet organisational and system objectives; how well does it meet user needs; is it economically feasible; and what are its advantages and disadvantages?

ALTERNATIVES TO BUILDING IT YOURSELF

OPTION	WHAT IT IS, AND WHAT IT COSTS YOU
Canned software	Sold on the open market to a broad range of users with similar requirements. Cheap and proven; the major problem is that it often does not meet all of a company’s information or processing needs.

OPTION	WHAT IT IS, AND WHAT IT COSTS YOU
Turnkey system	Software and hardware combined and sold as a finished, tested package ready for implementation. Suits small and medium enterprises; little flexibility.
Backbone system	Provides a basic system structure on which the organisation builds its own specifics.
Vendor-supported system	Customised and maintained by the vendor for the customer — low internal capability required, high dependence.
Application service provider / SaaS	The provider hosts web-based software on its own computers and delivers it over the internet. No hardware, no in-house maintenance, predictable subscription cost; you depend on their availability, their security and their pricing.
Outsourcing	Hiring an outside company to handle all or part of the organisation's data processing. A <i>mainframe outsourcing agreement</i> can go so far as the outsourcer buying the client's computers and hiring most of its staff, then operating the system on the client's site or migrating it to its own. A <i>client/server or PC outsourcing agreement</i> outsources a particular service, segment, function or PC support.

WHEN OUTSOURCING IS THE WRONG ANSWER

Outsource when the organisation lacks skilled IT staff who can keep the technology current, when it cannot manage the technology as well as a specialist firm can, or when it is simply cheaper. Do **not** outsource when **IT is a core competency** — when the system is itself the competitive advantage. Handing a rival-neutral vendor the thing that makes you different is how an advantage stops being inimitable. That distinction, rather than the cost comparison, is what makes an outsourcing decision defensible or not.

RELYING ON A SINGLE VENDOR — THE CONSEQUENCES

- **Lock-in** — switching costs rise over time as data, customisation and staff skills accumulate in one product;
- **Pricing power** — licence, support and upgrade fees can be raised once you cannot leave;
- **Continuity risk** — if the vendor fails, is acquired, or discontinues the product, the organisation's core processes have no supplier;
- **Development on their timetable** — features you need arrive when their roadmap says so;
- **Loss of internal knowledge** — nobody inside the organisation understands the system well enough to challenge the vendor or to run it if the relationship ends.

Alternatives to the SDLC itself

APPROACH	WHAT IT IS, AND WHEN IT FITS
Prototyping	A simplified working model is built quickly and cheaply and given to users to test. It supports rather than replaces the SDLC. Suitable where there is high uncertainty about the AIS, where it is unclear what questions to ask, where the final system cannot be clearly visualised, or where the likelihood of failure is high. Advantages: better definition of user needs; higher user involvement and satisfaction; faster development; fewer errors; more opportunity for change; lower cost. Risks: an unfinished prototype being pressed into production, and endless iteration without a specification.
Joint application design (JAD)	A group-based tool for collecting user requirements and creating system designs, combining requirements gathering and design into facilitated group sessions.
Rapid application development (RAD)	Combines JAD, prototyping and integrated CASE tools to produce a high-quality system rapidly: JAD sessions to collect requirements, tools to structure them and build prototypes, then user review of the prototypes in further JAD sessions.
Agile development	Delivers functionality in rapid iterations, usually measured in weeks, with frequent communication, development, testing and delivery. Focuses on rapid development and continual user contact, and on meeting the most important and immediate needs first.
End-user development	End users build their own applications with little or no formal assistance from IT. Fast and responsive; the reason so many organisations run on undocumented, uncontrolled and unaudited spreadsheets.
Object-oriented development	Traditional approaches focus on the system or the software; OO focuses on objects — tangible real-world entities, each with properties, data values and operations — and models how they interact to meet the new system's objectives.

ARE THE METHODOLOGIES MUTUALLY EXCLUSIVE?

No. They are complementary and are routinely combined: prototyping *within* the design phase of an SDLC, JAD sessions as the requirements-gathering technique, agile iterations delivering a system whose overall governance is still an SDLC. The SDLC is a framework of phases; the others are techniques for performing phases within it.

Two SDLC problems prototyping addresses: (1) users cannot specify what they want in the abstract, so requirements captured on paper are wrong — a prototype lets them react to something concrete; and (2) the SDLC's long lead time means that by delivery the requirement has changed — a prototype shortens the loop and surfaces the change while it can still be accommodated.

Prototyping and RAD compared

Both use an iterative approach, both save time, and both emphasise speed of development. The distinguishing point is that **RAD relies on integrated CASE tools** as part of the method, whereas prototyping does not require them — a prototype can be built in anything. It is not true that both methods depend on CASE tooling.

Small business systems

Small business owners need two things above all: visibility of **cash flow** and the ability to **evaluate profitability**. Common packages: Xero and MYOB in Australia, QuickBooks and Sage internationally. **Specialised accounting information systems** exist for particular industries — schools, dental practices, veterinary and pet retail — either as complete systems or as add-on modules to an integrated package, with independent developers creating programs that interface with the major packages.

Worked answer — Encosta Memories

THE CASE

Encosta Memories is a small family photography business with a good reputation and long-standing school photography arrangements. It uses a small accounting package, MYBP. Administration staff complain that it has become slow, and the computer running it is near capacity, not having been updated for years, and is also handling digital imaging. The business is considering an upgrade but fears it would have to employ a programmer to develop a new system and a full-time IT specialist to keep it running. Mr De Lago, part owner, is troubled by the possibility of investing capital now and again in a couple of years as technology changes, though he does not mind a large sum now if the system is a long-term answer: “Once the system is acquired then the business can get back to normal and do what it does best — take photos — without having to spend money on IT.”

1. Is De Lago’s statement correct? No. Acquisition is the beginning of the cost, not the end of it. A system requires *corrective* maintenance as defects emerge, *adaptive* maintenance as the environment changes — operating system updates, tax and payroll rule changes, new hardware, new security patches — and *perfective* maintenance as the business grows. It also requires backups, security, user training and eventual replacement. Industry experience is that maintenance consumes the majority of a system’s total lifetime cost. Treating IT as a one-off capital purchase is precisely how Encosta arrived at its present position: the current machine has not been updated for years and is now failing under a load it was never sized for.

2. Which of ASP, SDLC and outsourcing best limits the cost of maintaining the system once in place? An **application service provider** (in modern terms, a cloud/SaaS accounting package). The provider hosts the software and is responsible for maintenance, upgrades, backup and security; Encosta pays a predictable subscription and needs no programmer and no full-time IT specialist. A full SDLC means building a bespoke system, which is the most expensive option to maintain and requires exactly the staff De Lago wants to avoid. Outsourcing transfers operation to a third party but still presupposes a system to operate, and for a business of this size the contract management overhead is disproportionate. The ASP

model also answers De Lago's capital concern: the cost becomes an operating subscription rather than periodic large capital outlays, and upgrades arrive as part of the service.

3. Make or buy? Buy. Encosta's processes — invoicing schools, managing bookings, recording receipts, basic reporting — are entirely standard, so a packaged solution will meet them. It has no IT staff and no software development capability. Building would require hiring a programmer, exposes the business to key-person risk when that person leaves, produces no documentation unless someone insists on it, and delivers no competitive advantage: the business competes on photography and school relationships, not on the novelty of its ledger. Buying gives a tested product, vendor support, a shorter timeline and lower total cost. The one thing to verify before buying is whether any genuinely distinctive requirement — large-volume image storage linked to individual student orders, for instance — can be met by the package or a specialised add-on.

4. Four consequences of relying on a single vendor. As set out above: lock-in and high switching costs; the vendor's pricing power over licence, support and upgrades; continuity risk if the vendor fails or discontinues the product; and dependence on the vendor's development timetable, with no internal knowledge to fall back on. For a small business the continuity risk is the sharpest: it has no capacity to migrate quickly.

5. Five typical systems development problems Encosta may face.

1. **Requirements are poorly specified**, because the owners have never had to articulate what the system must do; the result is a package that does not fit the school photography workflow.
2. **Cost and time overruns**, since neither owner has run a systems project and the estimates will be optimistic.
3. **Data conversion problems** — years of MYBP data that must be migrated accurately and completely, with control totals to prove it.
4. **User resistance and inadequate training** — administration staff who are already frustrated will not absorb a new system without time set aside for it, and operational feasibility fails if they will not use it.
5. **Underestimating the ongoing commitment** — the maintenance, security and backup obligations De Lago's statement assumes away, together with the absence of anyone internally responsible for them.
6. **Inadequate documentation**, leaving the knowledge with whoever configured the system, which is the Woolworths failure at small scale.

Project scheduling

Systems projects are tracked by comparing budgeted with actual time, activity by activity, and projecting the remainder. The arithmetic is trivial; the discipline is in doing it while the project is still running rather than afterwards.

ACTIVITY	BUDGETED (DAYS)	ACTUAL (DAYS)	VARIANCE
Design the logical plan for the software operation	8	7	1 under
Select a programming language	2	2	—
Write source code	15	25	10 over
Test the source code	24	not yet available	—
Integrate software with existing applications	4	not yet available	—
Train users	7	not yet available	—
Total	60	34 to date + 35 remaining = 69	9 days over

Will it finish on time? No. Budgeted total is 60 days. Completed activities have taken 34 days against a budget of 25, a net overrun of 9 days. If the remaining activities — testing, integration and training, budgeted at 35 days — are completed on budget, the project will finish in 69 days, **9 days late**. **Which activity is over schedule?** Writing the source code, at 25 days against 15. And the overrun matters beyond its own nine days: code that took two-thirds longer than planned is usually code written under pressure, which makes an overrun in testing the likely next event rather than a surprise.

A system can be delivered on time, on budget and to specification, and still be used to do something indefensible. That is the final chapter.

Ethics, cybersecurity and privacy

Fraud, the ethical frameworks used to reason about what an organisation should do with information, the security controls that protect it, and the Australian privacy law that constrains its use. This is the material where a structured analysis separates itself most visibly from an opinion.

Fraud

COSO'S DEFINITION

"Fraud is any **intentional** act or omission designed to **deceive** others, resulting in the victim suffering a **loss** and/or the perpetrator achieving a **gain**."

Four elements, and each has to be present: intent, deception, and loss or gain. An error is not fraud, however expensive; the difference is intent.

Fraud attracts high public interest and poor data, because most of it is never reported. The surveys usually relied on are the Computer Security Institute annual survey, KPMG surveys, the Association of Certified Fraud Examiners survey, and the BDO fraud survey of not-for-profit organisations.

TYPES

TYPE	CHARACTER
Misappropriation of assets <i>employee fraud</i>	Theft of cash, inventory or information by employees. Individually smaller, far more frequent, and usually concealed through a weakness in segregation of duties.
Fraudulent financial reporting <i>management fraud</i>	Deliberate misstatement of the financial statements to deceive investors or lenders. Rarer, vastly larger, and defeats ordinary controls through management override.
Corruption	Bribery, kickbacks, conflicts of interest, bid rigging — usually involving a party outside the organisation.
Computer fraud	Any fraud that requires computer technology to perpetrate: input fraud (altering data going in), processor fraud (unauthorised use of the system), computer instruction fraud (altering the software), data fraud (copying, altering or destroying data) and output fraud (stealing or misusing displayed or printed output).

THE FRAUD TRIANGLE

ELEMENT	WHAT IT IS	WHAT MANAGEMENT CAN DO ABOUT IT
Pressure	A motive — financial need, addiction, unrealistic performance targets, fear of losing a job.	Least controllable, but avoid creating it: unattainable targets and bonus structures manufacture pressure.
Opportunity	The ability to commit the act and conceal it — almost always a control weakness, and usually a segregation-of-duties weakness.	The element internal control actually addresses. This is why control questions and fraud questions are the same question.
Rationalisation	The justification — “I am only borrowing it”, “the company owes me”, “everyone does it”.	Addressed through the control environment: tone at the top, a code of conduct that is enforced, ethics training, and consistent consequences.

THE POINT THAT MATTERS

All three elements must be present for fraud to occur. An organisation cannot eliminate pressure and cannot fully control rationalisation, so **internal control works by removing opportunity**. That single sentence links this chapter to Chapter Four, and it is the answer whenever the question is how a particular fraud could have been prevented.

Ethical theories

Ethical theories provide an *independent framework* for analysing an issue — which is why an analysis built on one is more persuasive than one that reports how the writer feels. The primary division:

CONSEQUENCE-BASED (TELEOLOGICAL)	DUTY-BASED (DEONTOLOGICAL)
An action is right if its consequences are good. The dominant version is utilitarianism : the right action produces the greatest net balance of good over harm for all affected.	An action is right or wrong <i>in itself</i> , by its adherence to a rule or duty, regardless of consequences. Deontologists look at duties — tell the truth, keep promises, respect persons — and ask what the action <i>is</i> , not what it produces.

LIMITATIONS OF UTILITARIANISM

- Consequential analysis tempts people to favour **quantifiable** variables over unquantifiable ones — dollars beat dignity because dollars can be added up.
- Measuring consequences is **subjective**: who decides what counts as more important? In this model, the decision maker — which is convenient for the decision maker.
- It can license serious harm to a minority where the aggregate benefit is larger.
- It requires prediction of consequences that may be unknowable.

One point is commonly reversed. Utilitarianism does *not* license self-interest: any act promoting the actor's own interest at the expense of others is unethical under it. The theory counts everyone's welfare equally, including the people the decision maker would rather not count.

THE OTHER THREE APPROACHES

APPROACH	THE TEST IT APPLIES
Fairness / justice	Treat all human beings equally, or, if unequally, then fairly, on a defensible standard . Most people accept higher pay for harder work or greater contribution — that is a defensible standard. A very large gap between executive and employee pay is tested by asking whether any defensible standard supports it, and whether a power imbalance rather than a standard produced it.
Rights	The ethical action is the one that best protects and respects the moral rights of those affected — to make one's own choices, to be told the truth, not to be injured, and to a degree of privacy. Rights are widely debated; commonly accepted ones include freedom of speech and religion.
Common good	Respect and compassion for all others is the basis for ethical action, emphasising the conditions important to everyone's welfare — a legal system, effective emergency services, healthcare, public education, public space. The test is whether the action strengthens or erodes a specific element of the common good.

AN ETHICAL DECISION-MAKING FRAMEWORK

- 1 **Recognise the ethical issue.** Is this a decision about more than what is legal or profitable? Who could be harmed?
- 2 **Get the facts.** What do we actually know, what is assumed, and what is unknown? Who are the stakeholders and what does each have at stake?
- 3 **Evaluate the alternative options** against each approach: which produces the most good and least harm (*utilitarian*); which best respects the rights of all stakeholders (*rights*); which treats people equally or fairly (*fairness*); which best serves the community as a whole (*common good*)?
- 4 **Make a decision and test it.** Would I be comfortable if this were reported on the front page? Would I be comfortable if it were done to me? What would a person I respect think?
- 5 **Act, then reflect on the outcome** — and feed what you learn back into how the organisation decides next time.

USING THE FRAMEWORK HONESTLY

Apply *every* approach, not the one that agrees with you, and say where they diverge. The most valuable sentence in an ethical analysis is usually of the form: “On a utilitarian analysis the decision is defensible because the aggregate benefit exceeds the harm; on a rights analysis it is not, because those affected never consented and were never told. The disagreement between the two is the reason the decision requires justification rather than assertion.”

Ethical issues in information technology

Four general categories, with Mason's **PAPA** model behind them:

CATEGORY	THE QUESTIONS IT RAISES
Privacy	What information about oneself should an individual be required to reveal? What can people keep to themselves? What should be kept in databases, and how secure is it there? What surveillance can an employer use on employees?
Accuracy	Who is responsible for the authenticity, fidelity and accuracy of information collected? How do we ensure it is processed properly and presented accurately? How do we ensure errors in databases, transmissions and processing are accidental rather than intentional? Who is accountable for errors, and how are injured parties compensated?
Property	Who owns the information? What is a just price for its exchange? How should software piracy be handled? When may proprietary databases be used? Can corporate computers be used for private purposes? How should experts who contribute knowledge to expert systems be compensated? How should access to information channels be allocated?
Accessibility	Who is allowed to access information? What may organisations charge for access? How is access provided for employees with disabilities? Who is provided with the equipment needed? What information does a person or organisation have a right to obtain, under what conditions and with what safeguards?

The surveillance debate sits underneath all four: the balance between national security and information privacy, illustrated by the Snowden disclosures — secret orders requiring a telecommunications carrier to hand over millions of call records daily, and the claim that an analyst at a desk could reach almost anyone's communications.

Privacy law in Australia

THE PRIVACY ACT 1988 (CTH)

The **Australian Privacy Principles** apply to Commonwealth government agencies and to private sector organisations with an annual turnover of **\$3 million or more** (and to some others regardless of turnover, notably health service providers). The APPs are **principles-based** — protecting privacy without burdening organisations with inflexible prescriptive rules.

They deal with all stages of the handling of personal information, setting standards for its **collection, use, disclosure, quality and security**, and they impose obligations concerning an individual's **access to, and correction of**, their own personal information.

Points worth knowing beyond the bare statement: personal information collected for one purpose generally may not be used or disclosed for an unrelated secondary purpose without consent; organisations must take reasonable steps to protect information from misuse, interference, loss and unauthorised access; and under the Notifiable Data Breaches scheme an eligible data breach likely to result in serious harm must be notified to affected individuals and to the Office of the Australian Information Commissioner.

PRIVACY CONTROLS

Privacy work deals with protecting *customer* information as distinct from internal company information — the second is confidentiality, the first is privacy, and the obligations differ. Controls include identification and classification of personal information holdings, encryption, VPNs for remote access, access restriction, retention limits and secure destruction.

Cybersecurity

THE THREATS

THREAT	WHAT IT LOOKS LIKE
Unintentional acts	Human error — the mis-addressed email, the lost laptop, the misconfigured share. The largest single category by volume, and the reason awareness training is a control rather than a courtesy.
Social engineering and phishing	Deceiving a person into granting access or transferring funds. Business email compromise — a spoofed instruction to change a supplier's bank details — is the version that costs Australian businesses most, and it defeats technical controls entirely because no system is breached.
Malware and ransomware	Software that steals, destroys or encrypts data. Ransomware is a business continuity event as much as a security one, which is why backups are the control that matters.
Denial of service	Overwhelming a system so it cannot serve legitimate users — an availability attack rather than a confidentiality one.
Unauthorised access and privilege abuse	Outsiders exploiting weak authentication, or insiders using access they legitimately hold for a purpose they do not.
Insider threat	The employee who already has access, already knows the controls, and is trusted. The hardest to detect and the reason for logging and monitoring.
Data exfiltration through everyday tools	Employee use of email and instant messaging is among the greatest threats to the confidentiality of sensitive information — not because it is malicious, but because it is frictionless.

ENCRYPTION

Encryption is a **preventive** control: the process of transforming normal content, called *plaintext*, into unreadable ciphertext; decryption reverses it.

TYPE	HOW IT WORKS	TRADE - OFF
Symmetric	One key both encrypts and decrypts.	Fast; vulnerable, because the key must be shared and anyone holding it can read everything.
Asymmetric	Different keys encrypt and decrypt — a public key and a private key.	Very secure; very slow.
Hybrid	Symmetric encryption for the information itself; asymmetric encryption to protect the symmetric key.	Speed of the first with the key distribution safety of the second. This is what actually runs on the internet.

A **digital signature** is a hash of a document encrypted with the creator's private key. It proves two things: that the document has not been altered since signing, and who created it. A **virtual private network** is a private communication channel — a tunnel — accessible only to parties holding the appropriate encryption and decryption keys.

THE REST OF THE CONTROL SET

- **Authentication** — something you know (password), something you have (token, phone), something you are (biometric). Multi-factor authentication combines two categories and defeats most credential theft.
- **Authorisation** — least privilege: each user has the minimum access their role requires, and no more.
- **Accountability** — logging and monitoring, which only work if identifiers are unique and the logs are actually reviewed.
- **Defence in depth** — firewalls, intrusion detection, network segmentation, endpoint protection, patching. No single control is relied on.
- **SETA programs** — security education, training and awareness, addressing the largest category of threat.
- **Backup and business continuity** — the corrective layer, tested rather than assumed.

THE SIMPLEST ANSWER TO "HOW DO I PROTECT MYSELF?"

A strong, unique password for every service, held in a password manager, with multi-factor authentication enabled, and a periodic check of whether your credentials have appeared in a known breach. Unglamorous, and it defeats the overwhelming majority of real attacks.

Two points that are easy to get backwards

Controlling access to system output. Of three plausible-sounding measures — a privacy policy, well-designed source documents, and training employees not to leave reports containing sensitive information in plain view on their desks — only the third is an output control. A privacy policy is a general

organisational control; source document design is an *input* control. The stage at which a control operates is not decided by what it is about.

Access and correction are obligations, not courtesies. It is sometimes assumed that an organisation is entitled to withhold from an individual the personal information it holds about them. The opposite is true: the Australian Privacy Principles require organisations to give individuals access to their own personal information and the means to correct it.

Worked answer — ABCD Medical Centre

THE CASE

Andrew Martins is CEO of ABCD Medical Centre. In 2005 he purchased MedicFlex, CRM software, to gather patients' personal information in one place, process historical disease information more efficiently, and let doctors communicate better with patients. By July 2010 information for almost 20,000 patients was stored in it. In November 2010 Andrew met David Wong, CEO of a large health insurer, who offered one million dollars for the patient database to promote his company's health services, promising to maintain its security and not to share it. Andrew decided to sell, arguing that both the customers and the medical centre benefit: most customers would benefit from the insurance offers, and he could use the money to renovate the centre.

How does a teleological approach differ from a deontological approach? A teleological approach judges an action by its *consequences* — the action is right if it produces the best outcome, typically the greatest net good across all affected parties. A **deontological** approach judges the action by its *conformity to duty or rule*, independently of outcome: some acts are wrong in themselves, and a good result does not make them right. The practical difference is what counts as evidence. A teleologist asks what happened as a result; a deontologist asks what was done, and whether it breached a duty owed — here, the duty of confidence a medical practice owes its patients. Andrew's reasoning is explicitly teleological: he weighs benefits (insurance offers, renovation) against harms and concludes the balance is positive.

On the theory of rights, was the decision ethical? Unethical. The rights approach asks which action best protects and respects the moral rights of those affected. The patients hold a right to privacy — to control information about themselves — and a right to be told the truth about how information they surrendered will be used. Health information is the most sensitive category there is. The patients gave that information for a specific purpose: their own medical care, and communication with their doctors. They were not told it would be sold, were not asked to consent, and were given no opportunity to refuse. That the recipient promised to keep it secure is irrelevant to the rights analysis: the wrong is the disclosure without consent, not the risk of a subsequent leak. Nor does it matter that most patients were pleased with the offers they received — a right is not defeated by the fact that its breach turned out well, and the patients who would have objected were never given the chance. Under the Privacy Act 1988 the same conclusion follows in law: personal information collected for the purpose of providing health services may not be used or disclosed for an unrelated secondary purpose — insurance marketing — without consent, and health service providers are covered regardless of turnover.

Who are the stakeholders, and what is each one's role?

STAKEHOLDER	ROLE AND INTEREST
The 20,000 patients	The subjects of the information. They bear the entire harm — loss of privacy and control over sensitive health data — and gave nothing more than the consent to be treated.
Andrew Martins, CEO	The decision maker. He holds the information in trust for the practice and its patients, and has a direct financial interest in the sale.
The doctors and clinical staff	They collected the information under a professional duty of confidence and will bear the consequences for the doctor–patient relationship. They were not consulted.
ABCD Medical Centre as an organisation	Receives the one million dollars and the renovation; carries the legal, regulatory and reputational exposure if the sale becomes known.
David Wong and the insurer	The purchaser and the beneficiary of the marketing. It acquires data it could not have collected itself, and takes on its own privacy obligations over it.
The regulator and the wider public	The Information Commissioner enforces the privacy regime; the public interest is in a health system where people disclose information to doctors without fear it will be sold. Every such sale erodes that.

Applying the decision-making framework. The model runs: recognise the ethical issue; get the facts and identify the stakeholders; evaluate the alternatives against the utilitarian, rights, fairness and common good approaches; make a decision and test it; act and reflect. Applied here:

- **Recognise.** This is not merely a commercial decision. It disposes of information belonging, in a moral sense, to people who are not party to the transaction.
- **Facts.** 20,000 patients; sensitive health information; collected for treatment; no consent sought; purchaser intends to market to them; consideration of one million dollars; the promise of security is unenforceable in practice.
- **Utilitarian.** Andrew's own analysis, and it is weaker than he thinks. The benefits — some patients like the offers, the centre is renovated — are modest and accrue partly to the decision maker. The harms are large and unquantified: loss of privacy for 20,000 people, unwanted marketing, the risk of health status influencing insurance pricing, and the systemic harm to trust in medical confidentiality. On honest accounting the utilitarian calculus does not favour the sale either.
- **Rights.** Fails, for the reasons set out above.
- **Fairness.** Fails. The benefit goes to Andrew and the centre; the cost falls on patients. There is no defensible standard on which those bearing the cost were selected, and the power imbalance is total — patients had no voice in the decision.
- **Common good.** Fails. Medical confidentiality is a condition of a functioning health system. If patients believe their records are saleable, some will withhold information from their doctors, and the harm is public rather than private.
- **Decide and test.** All four approaches point the same way once the utilitarian analysis is done properly. The publicity test settles it: no version of this decision survives being reported.
- **The alternative.** If the centre needs capital, borrow it. If the insurer's offers are genuinely valuable to patients, the centre can invite patients to *opt in* to receiving them, disclosing that the centre is paid. Consent converts an indefensible act into a defensible one, and it is available at almost no cost.

And that is the subject. The through-line, if there is one, is the sentence the first chapter opened with: an accountant does not observe transactions, only what the system recorded. Every framework in the twelve chapters since — the cycles, the control classifications, the BPMN notation, the development methodologies, the ethical theories — exists to make that record something a person can reasonably rely on, and to make clear who is answerable when it cannot be.