

Partner, Don't Police: AI in the Business Classroom

Michael Borck

2026-05-10

Partner, Don't Police
AI in the Business Classroom

Copyright © 2025 Michael Borck. All rights reserved.

Published by Michael Borck
Perth, Western Australia

ISBN: 979-8-2542-3619-1

First edition, 2025.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the author, except for brief quotations in reviews and certain non-commercial uses permitted by copyright law.

This work is also available under a Creative Commons Attribution (CC BY) licence at the companion website. See below for details.

AI disclosure: This book was written using the methodology it describes. AI tools were used as thinking partners throughout the drafting, iterating, and refining process. The author reviewed, challenged, and took responsibility for every sentence.

Companion website:

<https://michael-borck.github.io/partner-dont-police>

Source: <https://github.com/michael-borck/partner-dont-police>

Table of contents

Preface	1
Acknowledgments	9
1 Introduction - Why AI Matters for Business Education	11
2 Understanding AI and Large Language Models	19
3 Getting Started — Your First AI Conversation	25
4 Seven Essential Prompt Techniques for Business Teaching	31
5 Managing LLM Context: Working Smarter with AI Conversations	39
6 The Flight Simulator - Conversation Simulations	61
7 The AI Critique Toolkit - Becoming a Smart Business Professional	77
8 Ethics, Data Governance & Integrity	89
9 Rethinking Assessment: Process Over Product	117
10 AI as Study Buddy - Self-Assessment Tool	125
11 Virtual Company Simulation (Future Horizon)	141
12 Designing an AI-Integrated Unit	157
13 Transforming Content with AI	175
14 Global Perspectives and Adaptation	181
15 Implementation Practicalities	201
16 Conclusion: Where Do We Go From Here?	207
17 Colophon	217
Appendices	219
A Aligning AI Integration with Institutional Learning Outcomes	219
B AI-Integrated Assessment Rubric System	229
C Assessment AI Stress Test: A Prompt Sequence	251
D Glossary	259
E Further Reading	263

About the Author 273

Preface

Why This Book Exists

The conversation around AI in education has split into two camps, and neither is helping.

On one side: enthusiasm without scrutiny. AI will transform everything, students should use it for everything, and anyone who hesitates is falling behind. On the other: prohibition without understanding. AI is cheating, it threatens academic integrity, and the safest response is to ban it and hope it goes away.

If you are an educator caught between these positions, this book is for you.

I wrote it because the educators I work with are not looking for hype or prohibition. They are looking for practical guidance from someone who understands both the technology and the classroom. They want to know: What can I actually do with this? How do I protect academic integrity without pretending AI does not exist? How do I teach students to use these tools the way a professional would, critically and transparently, rather than as a shortcut around thinking?

Those are the questions this book answers.

Who This Book Is For

You are a business educator. You teach undergraduates or postgraduates in marketing, management, accounting, economics, information systems, tourism, supply chain, human resources, or a related discipline. You are not a computer scientist. You did not sign up to become an AI expert. But AI has arrived in your classroom whether you invited it or not, and you need to make professional decisions about what to do with it.

You may also be a student navigating a landscape where AI tools are

everywhere but guidance on using them well is scarce. The frameworks in this book apply to both sides of the classroom.

What This Book Is Not

This is not a technology manual. It does not teach you how to use ChatGPT, Claude, Copilot, or any specific platform. Interfaces change constantly; the principles here do not.

It is not a book of copy-paste prompts. You will find prompts throughout, but they are starting points for conversation, not finished products. If you are looking for a recipe book that removes the need to think about your teaching, this is the wrong book.

It is not a defence of AI, and it is not an attack on it. It does not argue that AI will save education or that it will ruin it. It argues that AI is a tool, that tools require judgement, and that your job as an educator is to develop that judgement in yourself and your students.

And it is not a book that requires you to change everything at once. The most common piece of advice in these pages is: start with one thing, in one unit, this semester.

If You Are Feeling Overwhelmed

You are not alone and you are not behind. Most educators feel some version of the same anxiety: that everyone else has figured this out already, that the technology is moving faster than they can follow, that they might look foolish trying something new in front of students who grew up with technology.

That anxiety is normal. It does not mean you are unsuited to this work. It means you are paying attention.

You may also worry that your students already know more about AI than you do. Some of them might. That is not a problem to solve. It is a resource to draw on. An educator who says “show me how you did that” and learns alongside their students is modelling exactly the kind of professional curiosity that matters in every business discipline.

You do not need to be the AI expert in the room. You need to be the expert in your discipline who knows how to ask good questions, evaluate evidence, and design learning experiences. AI does not change that. It amplifies it.

How This Book Is Structured

The book follows a progressive structure designed to build confidence gradually:

Foundation (Chapters 1–3) builds your understanding of AI and LLMs, introduces structured prompting through the CRAFT framework, and gets you into your first productive AI conversation. No prior AI experience is assumed.

Core Techniques (Chapters 4–7) gives you seven essential prompt techniques, strategies for managing AI context, the flight simulator concept for student practice, and the critique toolkit for teaching critical evaluation.

Teaching with AI (Chapters 8–12) moves into ethics and data governance, process-based assessment, self-assessment tools, virtual company simulations, and full unit design with a phased implementation roadmap.

Putting It Together (Chapters 13–16) covers transforming content with AI, global perspectives and adaptation, implementation practicalities (technical and accessibility), and where to go from here.

Appendices provide institutional alignment frameworks, rubric templates, stress test sequences for validating your assessments, a glossary, and further reading.

You do not need to read it in order. The table below will point you to the chapters that matter most for your situation.

Your Situation	Start Here
Completely new to AI	Chapter 1 (motivation) and Chapter 3 (your first AI conversation)
Want to implement AI in teaching	Chapter 4 (essential techniques) then Chapter 12 (unit design with phased roadmap)
Designing assessments	Chapter 8 (assessment: process over product) and the Rubric System appendix
Concerned about integrity	Chapter 7 (ethics, data governance, and integrity)
Need institutional justification	Institutional Alignment appendix

Conventions Used in This Book

Throughout the book you will encounter coloured callout boxes. Each serves a different purpose.

Chapter-opening quotes appear as indented text like this. They frame the key idea or challenge each chapter addresses.

Discipline Example

Green boxes contain discipline-specific examples — prompts, scenarios, or exercises tailored to a particular business field. In the online edition, these appear as interactive tabs covering all eight disciplines. In the print and ebook editions, a single representative example is shown.

Key Idea

Blue boxes highlight important concepts, ready-to-use prompts, or ideas worth pausing on.

Watch Out

Yellow boxes flag common mistakes, weak examples, or things that look right but are not.

Critical Point

Red boxes mark habits or principles that are essential. Do not skip these.

Important Notes

Technology changes rapidly. AI tools evolve quickly. Specific platform names and capabilities described here reflect the state of technology in early 2025. The principles and pedagogical approaches remain relevant even as specific tools change.

Context matters. This book was developed with Australian business programs in mind. If you are teaching in a different institutional or national context, you may need to adapt examples, learning outcomes, and regulatory references. The frameworks are transferable; the specifics may not be.

This is a starting point. Consider this book a foundation, not a complete solution. You will discover what works for your students, your teaching style, and your context through experimentation.

Ways to Engage with This Book

This book is available in several formats. Pick whichever fits how you work and learn.

- **Read it online.** The full book is freely available at the companion website, with dark mode, search, navigation, and interactive tabbed examples across all eight business disciplines.
- **Read it on paper or e-reader.** Available as a paperback and ebook through Amazon KDP, for those who prefer to read offline or away from a screen.
- **Converse with it.** The online edition includes a chatbot grounded in the book's content. Ask it questions, challenge its answers, and practise the methodology on the methodology itself.
- **Feed it to your own AI.** The `11m.txt` file provides a clean text version of the entire book, ready to paste into ChatGPT, Claude, or any AI tool for a conversation about the ideas.
- **Explore the source.** The full source is on GitHub, including every chapter, the build system, and the revision history. DeepWiki provides an AI-navigable view of the repository.
- **Browse all books.** This book is part of a series. See all titles at books.borck.education.

The online version is always the most current. The printed and ebook editions are updated periodically.

The Companion Book

The methodology in this book draws on *Conversation, Not Delegation: How to Think With AI, Not Just Use It* (Borck, 2025), which covers the full framework in a discipline-neutral way. If you want the underlying rationale for why these approaches work, or you want to share a version with non-educator colleagues, that book is the place to start. This book is independently complete, but the two are stronger together.

All titles are available at books.borck.education.

Copyright

Partner, Don't Police: AI in the Business Classroom

Copyright © 2025 Michael Borck. All rights reserved.

Published by Michael Borck
Perth, Western Australia

ISBN: 979-8-2542-3619-1

First edition, 2025.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the author, except for brief quotations in reviews and certain non-commercial uses permitted by copyright law.

This work is also available under a Creative Commons Attribution (CC BY) licence at the companion website. See below for details.

AI disclosure: This book was written using the methodology described in its companion, *Conversation, Not Delegation*. AI tools were used as thinking partners throughout the drafting, iterating, and refining process. The author reviewed, challenged, and took responsibility for every sentence.

Companion website: <https://michael-borck.github.io/partner-dont-police>

Amazon: *(add your Amazon listing URL here)*

Source: <https://github.com/michael-borck/partner-dont-police>

Acknowledgments

The ideas in this book were developed in classrooms, not in the abstract. The students in my business education units were the first to encounter these approaches — the flight simulators, the critique exercises, the process-based assessments — and their responses shaped every chapter. Some approaches worked immediately. Others failed in ways that were far more instructive than the successes. Both made the book better.

Colleagues across business disciplines provided feedback, scepticism, and the occasional “have you tried this with accounting students?” that sent entire chapters in new directions. Teaching and learning teams provided the institutional context that kept the ideas grounded in what is actually possible within a university semester.

This book draws on the methodology developed in *Conversation, Not Delegation*, and the two books were written in parallel, each sharpening the other. The pedagogical frameworks it builds on — experiential learning, reflective practice, authentic assessment — are not new. What is new is applying them through AI, and that application was only possible because decades of educational research provided the foundation.

The technical infrastructure behind this book is entirely open source. Quarto, Python, GitHub, and GitHub Pages made it possible to write, build, and publish across multiple formats without a traditional publisher. The open source community deserves more credit than it typically receives.

AI tools were used throughout the writing process, as described in the preface of the companion book. Claude (Anthropic) served as a conversation partner for drafting, iterating, and refining. The process was exactly what both books advocate: conversation, not delegation. The author’s judgement shaped every page. The AI made the work faster. It did not make the decisions.

Chapter 1

Introduction - Why AI Matters for Business Education

The question is not whether your students will use AI. It is whether they will use it well, and that depends on what you teach them.

1.1 The Problem We're Solving

As a business educator, you face a persistent challenge: **how do you prepare students for the messy, high-stakes reality of professional business work when your classroom is safe, controlled, and hypothetical?**

The specific challenge depends on your discipline. Here are just a few examples:

Example: Supply Chain & Logistics

You can teach supply chain theory and optimisation models, but you can't easily let students experience demand disruptions, supplier failures, or make real-time logistics decisions with uncertain information. You can't scale practice in crisis management across complex networks.

The common challenge across all these disciplines: You can't easily

give every student practice in high-stakes, complex, realistic scenarios with immediate feedback and the freedom to fail safely.

Until now.

1.2 The Flight Simulator Concept

Think about how pilots are trained. They don't learn to handle engine failure during a storm by reading a textbook. They don't practice emergency landings by watching videos. They use **flight simulators**, sophisticated environments where they can crash the plane, make terrible decisions, experience rare scenarios, and learn from catastrophic failures without anyone getting hurt.

That's what AI can do for business education.

AI conversation tools can create a **professional practice simulator** where your students can practice in their field:

Example: Human Resources

- Conduct recruitment interviews and evaluate candidate fit
- Handle sensitive employee relations and discrimination scenarios
- Navigate termination conversations with legal and emotional complexity
- Practice benefits negotiation and compensation discussions

And here's the remarkable part: **after the simulation, the AI can act as an expert supervisor**, reviewing the transcript of what happened and providing detailed critique based on discipline-specific theory, professional standards, and ethical principles.


1.3 What Makes This Different from Traditional Teaching?

Traditional business education often focuses on **product**, the final answer, the correct calculation, the right theoretical framework. Students write essays, complete exams, and submit reports that demonstrate they *know* things.

But professional business work is about **process**, the methodology of how you analyse financial data, the approach you take in strategic planning, the steps you follow to ensure market research validity, the communication style you use to negotiate deals.

AI allows us to assess and teach process, not just product.

Here's what this means in your discipline:

 Example: Accounting & Finance

Traditional: “Apply the going concern principle” (rules-based knowledge) **Process-focused:** “Audit this financial statement. What red flags did you identify? How did you investigate? What judgment calls did you make?”

This shift from testing knowledge recall to evaluating applied professional methodology is transformational for business education.

1.4 The Conversation, Not Delegation Framework

Most people use AI the same way: they give it a task and accept what comes back. Write this email. Summarise this report. Answer this question. This is **delegation**. It produces outputs. It does not produce understanding.

There is a different approach. Instead of handing tasks to AI, you think alongside it. You bring a question, not a task. You explore possibilities together. You push back on what it gives you. You refine, redirect, and challenge until the result reflects your judgement, not just the model's fluency.

This is **conversation**. And it changes what AI does for the person using it.

Delegation asks: “How do I get AI to do this for me?”

Conversation asks: “How do I use AI to think better with me?”

The difference matters because judgement cannot be delegated to something that has read everything but experienced nothing. AI has processed more text than any person could read in a lifetime. But it has never made a decision under pressure, never felt the weight of getting something wrong, never had a stake in the outcome. It can generate plausible answers. It cannot know which answer is right for your situation. That is your job.

1.4.1 The Four-Part Loop

The framework has four moving parts:

- **Brainstorm:** Arrive with a real question, not a task to outsource.

- **Ideate:** Go wide. Explore angles, alternatives, framings you had not considered.
- **Iterate:** Push back. Challenge what the AI gives you. Refine until it fits your context.
- **Amplify:** Take the best of what emerged and make it yours. You own the result.

Most good work passes through this loop more than once. The sign that you are done is not that the AI has stopped producing output. It is that your thinking has landed somewhere solid.

The core principle is simple: **your expertise + AI's breadth = amplified thinking**. The bottleneck is always your thinking, not the model.

Every technique in this book, from the flight simulator to the critique toolkit to the seven prompting techniques, is designed to keep you in conversation, not delegation. When you see a prompt that asks the AI to challenge your reasoning, or a simulation that requires you to respond in real time, or an assessment that grades your process rather than the AI's output, you are seeing this framework in action.

Are you conversing or delegating?

Here is a quick test. After working with AI on a teaching task, ask yourself: do I understand my pedagogical challenge more clearly than when I started? If the answer is yes, you were thinking alongside the AI. If the answer is no, you handed the task over and accepted what came back. The distinction matters because only one of those processes develops your professional judgement.

Skills compound or erode

Every interaction with AI is practice — but practice at what? If you routinely let AI do the thinking, your own capacity for that thinking weakens over time. If you use AI to challenge, extend, and refine your ideas, your expertise deepens. The same dynamic applies in your classroom: the habits your students build now will shape their professional capabilities for years.

1.4.2 Staying Critical: VET and the Cognitive Traps

Conversation only works if you stay critical. *Conversation, Not Delegation* introduces the **VET framework**, three questions to ask before acting on any AI output:

- **Verify:** Can I find this independently? Check sources, cross-reference claims.
- **Explain:** Can I explain this in my own words? If not, I do not understand it yet.
- **Test:** Does this hold up under scrutiny? Change a variable, try an edge case.

It also names three cognitive traps that undermine critical engagement:

- **Gell-Mann Amnesia:** You catch AI errors in your area of expertise, then trust it completely on topics you know less about.
- **The Sycophancy Trap:** AI is trained to agree with you. Ask “what do you think?” and you get flattery, not feedback. Ask “what are the three weakest points?” and you get something useful.
- **The AI Dismissal Fallacy:** Rejecting an idea solely because AI was involved. “That is just ChatGPT” is not a critique; it is a refusal to engage with the content.

These traps matter for teaching because students will fall into all three. Naming them makes them visible, and visibility is the first step to resisting them. The critique toolkit in this book and the VET framework from *Conversation, Not Delegation* reinforce each other: both teach the habit of evaluating AI output on its merits rather than accepting or rejecting it reflexively.


1.5 Three Core Principles of This Approach

As you read through this book and begin experimenting with AI in your teaching, keep these three principles in mind:

1.5.1 1. AI as Scaffolding, Not Replacement

AI is like a construction crew that can quickly build the framework for complex learning scenarios. But your role as the educator is irreplaceable: you design the learning objectives, you set the ethical boundaries, you guide students to inspect and refine their work, and you ensure the final structure is robust and professionally sound.

Examples of how this works across disciplines:

 Example: Tourism & Hospitality

AI creates demanding guests and service scenarios. You ensure cultural authenticity, teach service excellence principles, and guide

professional judgment about when to escalate.

1.5.2 2. Transparency Over Prohibition

Many educators worry about students using AI to cheat. This book takes the opposite approach: **give students the AI tools, teach them to use AI ethically, and grade them on their ability to critically evaluate and improve AI outputs.**

In the real world, professionals across all business disciplines will use AI tools:

Example: Information Systems

IT professionals use AI for code generation, systems analysis, and automation. Our job is to teach students to review AI outputs, maintain security and quality, and understand when to override automation.

Our job isn't to prevent AI use, it's to ensure students can use AI tools responsibly, identify their limitations, and maintain human judgment on ethical, legal, and disciplinary-specific matters.

1.5.3 3. Start Simple, Scale Gradually

You don't need to revolutionize your entire curriculum tomorrow. This book will show you how to start with a single prompt, try one simulation exercise, or enhance one assessment. Each chapter builds progressively, so you can adopt techniques at your own pace.

1.6 What You'll Learn in This Book

Chapters 1-2 introduce you to AI and walk you through your first successful AI interaction. No prior experience needed.

Chapter 3 gives you five proven prompt techniques specifically adapted for business education—tools you can use immediately.

Chapters 4-6 show you three powerful applications: conversation simulations, self-assessment tools, and virtual company scenarios. Each chapter includes complete worked examples across multiple business disciplines.

Chapters 7-8 reimagine assessment in business education and show you how to design complete AI-integrated units from scratch.

Chapters 9-10 address ethics, academic integrity, and practical implementation guidance directly.

The Appendices give you ready-to-use prompts, a workshop guide for colleagues, and a framework for aligning AI integration with institutional learning outcomes.

Let's begin.

Chapter 2

Understanding AI and Large Language Models

It has read everything and experienced nothing. That single fact explains both its usefulness and its limitations.

2.1 What AI Actually Is

Artificial Intelligence (AI) is software that can perform tasks that typically require human intelligence — recognising patterns, making decisions, understanding language, predicting outcomes, and generating content.

AI is not new. The term was coined in 1956. What changed is that three things converged: we now have massive amounts of data to learn from, computing power got cheap and fast, and better algorithms (particularly “deep learning”) were discovered. The result is AI that was science fiction ten years ago is now practical and affordable.

You do not need to understand how the engine works to drive the car. But you do need to know what the car can and cannot do. For business educators, the AI that matters most right now is the **large language model**.

2.2 How Large Language Models Work

A Large Language Model (LLM) is AI trained to predict the next word in a sentence. That sounds simple, but from that single task — predicting

the next word — these models learned to write essays, answer questions, translate languages, summarise documents, explain concepts, and hold conversations.

Examples you have probably heard of: ChatGPT (OpenAI), Claude (Anthropic), Gemini (Google), Copilot (Microsoft).

2.2.1 The Training Process

Imagine learning language by seeing billions of examples:

- “The cat sat on the _____” → You learn “mat” is likely
- “The capital of France is _____” → You learn “Paris”
- “If you drop a glass, it will _____” → You learn “break”

Now multiply that by billions of examples from Wikipedia, books, news articles, academic papers, code repositories, and conversations. By learning to predict the next word across all of that text, the model absorbed grammar, facts about the world, how arguments are structured, and how to write in different styles.

Think of it as autocomplete on your phone, but trained on trillions of words instead of just your texts, and able to predict entire paragraphs rather than a single word.

2.2.2 What Makes Them “Large”

“Large” means three things: huge amounts of training data (hundreds of billions of words), massive numbers of parameters (the patterns the model learned — GPT-4 has over a trillion), and enormous computing power to train. Bigger models learn more subtle patterns and handle more complex tasks, but are more expensive to run.


2.3 The Key Insight: Interpolation, Not Retrieval

If you remember one thing from this chapter, make it this: **LLMs interpolate, they do not retrieve.**

When you ask for a fact, the model is not looking it up in a database. It is predicting what a plausible answer would sound like, based on patterns in everything it was trained on. A convincing answer and a correct answer are produced by exactly the same process. The model cannot tell the difference.

This is why LLMs “hallucinate” — confidently stating false information. The model will fabricate a statistic, invent a citation, or state an incorrect

legal requirement with exactly the same tone and certainty as a verified fact. There is no built-in signal that distinguishes accurate output from hallucination.

 Confidence is not accuracy

An LLM will state a fabricated citation or an incorrect fact with exactly the same certainty as a verified one. If the information matters — and in business education it usually does — checking it is part of your workflow, not an optional extra.

This is not a flaw to be fixed in the next version. It is the nature of the technology. LLMs are extraordinarily capable pattern matchers trained on the written record of human thought. That makes them powerful tools. It also means they have real limits — limits that do not go away just because the outputs sound confident.

2.4 What LLMs Do Well and Where They Struggle

2.4.1 Strengths

LLMs excel at tasks where producing a convincing average is enough:

- **Writing and content generation** — emails, reports, summaries, marketing copy, lesson plans
- **Summarisation** — condensing long documents while maintaining meaning
- **Explanation** — breaking down complex topics for different audiences
- **Translation** — not just word-for-word, but with awareness of meaning and context
- **Brainstorming** — generating ideas, angles, and approaches you might not have considered

2.4.2 Limitations

- **Factual accuracy** — they predict plausible answers, not correct ones. Always verify.
- **Current information** — training data has a cutoff. They cannot browse the internet unless specifically connected to search.
- **Genuine reasoning** — they pattern-match, not reason. Complex logic, mathematics, and multi-step reasoning are unreliable.
- **Memory** — each conversation is isolated. They do not remember previous sessions.

- **Creativity** — they remix and recombine patterns, but do not have original insights or lived experience.

2.5 The Decision Framework: Average vs Precise, Small vs Large

There is a simple way to predict when an LLM will serve you well and when it will let you down. Think along two axes: **how precise does the output need to be?** and **how big is the task?**

	Small	Large
Average	Sweet spot. Drafts, summaries, brainstorming. Trust with light review.	Plausible but brittle. Looks right at first glance, falls apart on inspection.
Precise	Workable with verification. Facts, citations, specific details. Check before using.	Danger zone. Confident structure that is subtly wrong. Stay in the loop at every step.

Sweet spot tasks can often be handled with a light check. Everything else requires conversation, iteration, and human judgement proportional to where it sits on the grid.

This framework is covered in depth in *Conversation, Not Delegation*, including how it connects to the conversation loop and the two-chat workflow. For this book, the practical takeaway is: before using AI on any task, ask these two questions. The answers tell you how much oversight the task requires.

2.6 Why This Matters for Your Teaching

If you understand that LLMs are sophisticated prediction engines, not omniscient oracles, you will use them differently in your classroom.

You will not hand students a tool and tell them to trust the output. You will teach them to generate a draft and then apply their own judgement. You will not ask them to get the answer from AI. You will ask them to use AI to think through the problem. You will teach them to recognise when output is echoing a pattern rather than reflecting genuine reasoning, and to push back.

The difference between someone who uses AI well and someone who uses it poorly is rarely about technical skill. It is about understanding what the tool actually is — and what it is not. That understanding starts here, and it runs through every chapter that follows.

Chapter 3

Getting Started — Your First AI Conversation

The best way to understand AI is to use it. Not to read about it, not to watch a demo, but to sit down and have a conversation.

3.1 Accessing an AI Tool

Many institutions now provide enterprise AI tools through existing software agreements — commonly MS Copilot Enterprise (through Microsoft 365) or Google Gemini (through Google Workspace for Education). Check with your IT department first, because enterprise tools keep your data within institutional boundaries.

If your institution does not provide one, several free tools are available: ChatGPT (chat.openai.com), Claude (claude.ai), and Google Gemini (gemini.google.com). All work directly in your web browser — no installation, no configuration. Create a free account and you will see a text box waiting for your instructions.

For institutional work involving student data or course materials, use your enterprise tool. For general exploration and learning, any tool works — they all have similar core capabilities.

3.2 Your First Prompt: The Weak Version

Open your chosen AI tool and type this exactly:

`Write a case study about business.`

You will receive something generic — maybe a story about a marketing campaign or a financial decision. It is fine, but it is not useful for your specific teaching needs. “Business” covers everything. “A case study” could be 100 words or 5,000. The AI had to guess what you wanted.

3.3 Your Second Prompt: The Powerful Version

Now try this instead:

```
You are an expert marketing lecturer at a
university level.
```

```
I need you to create a case study for my
undergraduate students that will help them
practice strategic market analysis and
competitive positioning.
```

```
Here are the requirements:
```

- The scenario should involve a product launch decision where the company faces competitive threats
- Include enough detail that students need to identify which market analysis principles apply
- The scenario should be 300-400 words
- End with three discussion questions requiring students to analyse the market strategy and positioning

```
Begin.
```

The output should be dramatically more useful. The difference is that you told the AI who you are, what you need, what constraints matter, and what format to use. The AI did not get smarter between your first prompt and your second. You got clearer.

3.4 The CRAFT Framework

What made that second prompt powerful? It included five elements that you can remember with the acronym **CRAFT**:

Letter	Element	What it does
C	Context	Sets the situation — your course, your students, their level
R	Role	Defines the AI's perspective — expert lecturer, consultant, practitioner
A	Action	Specifies the task — create, analyse, critique, explain
F	Format	Describes the structure — word count, number of questions, table format
T	Tone/Target	Sets the style and level — professional, introductory, sophisticated

You do not need every element every time. But when a prompt gives you something vague or generic, check which CRAFT elements are missing. The answer is usually context or format.

CRAFT is not the only prompting framework. The companion book *Conversation, Not Delegation* (available free at books.borck.education) compares several frameworks side by side. The goal is not to follow CRAFT rigidly — it is to think intentionally about what you are asking the AI to do.

3.5 Conversation, Not One-Shot

Here is the most important idea in this chapter: **CRAFT is your first message, not your only message.**

When you send one prompt and take the output as-is, you get generic content that could work for anyone. The AI produces safe, middle-ground responses. It misses your specific teaching voice and context.

The solution is to have a conversation:

PROMPT 1 (CRAFT):

"You are an experienced marketing director. Create a market analysis case study for my second-year students. 500 words with 4 discussion questions."

[AI produces a case study]

PROMPT 2 (Refine):

"Good. Make the discussion questions more specific - ask students to identify which of Porter's Five Forces apply and what data they would need to support their recommendation."

PROMPT 3 (Expand):

"Add a section comparing three analytical frameworks and how each leads to different conclusions."

PROMPT 4 (Repurpose):

"Turn this into a 10-minute classroom activity with a 2-minute intro, 5-minute group task, and 3-minute discussion."

You did not write four prompts from scratch. You had a conversation where each follow-up built on what was working. This is what *conversation*, not *delegation* looks like in practice.

3.6 Discover Your Strategy: Using AI to Help You Use AI

Here is something most tools cannot do: **AI can teach you how to use itself.**

Instead of reading generic advice, you can have a personalised consultation where AI interviews you about your teaching context and then recommends both obvious and unexpected applications. Copy this prompt and try it now:

You're an AI expert consultant for business educators. I'd love your help discovering how I could better use AI in my teaching.

Please ask me one question at a time. I'll answer, then you'll ask the next question. Keep asking until you understand:

- My discipline and the students I teach
- My main teaching challenges
- My key teaching goals
- My constraints (time, resources, skills)
- What I've already tried with AI


Once you have enough context, provide TWO types of recommendations:

1. OBVIOUS OPPORTUNITIES: Clear, straightforward AI applications most educators think of
2. NON-OBVIOUS OPPORTUNITIES: Unexpected or creative uses specific to my discipline that I might not have considered

Format your recommendations clearly with implementation tips for each.

Please start by asking your first question.

Spend 10-15 minutes answering honestly. The value is not just in the AI's recommendations — it is in the clarity you gain about your own teaching by having to explain it. Educators consistently report that articulating their teaching philosophy to an AI helped them see how AI could enhance it, not replace it.

 Try this (15 minutes)

Run the consultation prompt above. Answer the questions honestly — don't polish your responses. Then pick one obvious recommendation and one non-obvious recommendation to try this semester. Return to the same conversation after you have tried them and report what happened. The AI will help you iterate.

3.7 The Foundation

If you have followed along, you have accessed an AI tool, written a structured prompt, had a conversation that refined the output, and discovered personalised recommendations for your teaching. That is not a small thing. Everything else in this book builds on this foundation.

The quality of what you get from AI depends almost entirely on how you work with it. Give it a vague instruction and you get a generic response. Work with it through a structured conversation and you get something genuinely useful. The tools will change. This skill transfers to all of them.

Chapter 4

Seven Essential Prompt Techniques for Business Teaching

The right technique is not the cleverest one. It is the one that matches the kind of thinking you want your students to do.

4.1 Why These Seven?

These are seven proven prompting techniques chosen because they develop critical thinking skills that business professionals across all disciplines need:

1. **Reverse Prompting** — teaches comprehensive scoping and requirement gathering
2. **Pros and Cons** — builds analytical decision-making skills
3. **Stepwise Chain of Thought** — reinforces process adherence and documentation
4. **Role Play** — develops communication and interpersonal skills
5. **Debating** — strengthens strategic thinking and ethical reasoning
6. **Formative Assessment Generator** — provides unlimited practice and immediate feedback
7. **The Expert Panel** — develops multi-perspective analysis and synthesis skills

Each technique works across all business disciplines. The examples below use different disciplines to show the range — adapt the prompts for your own field.

These seven are adapted from *Conversation, Not Delegation*, which covers the same approaches for a general audience. If you want the underlying rationale for why each technique works, see the companion book.

4.2 Technique 1: Reverse Prompting (for Scoping)

Instead of solving a problem immediately, the AI asks the student questions to help scope out all requirements, considerations, and potential issues. This mirrors professional work where defining the problem is half the battle.

Example (HR — policy design):

I need to draft a new "Flexible Work Arrangement" policy for a 500-employee company with both office-based and remote staff.

Your task: Ask me a series of questions to help me clarify all the requirements, considerations, and potential pitfalls for this policy.

Ask one question at a time. Wait for my response before asking the next. Continue until you've helped me think through at least 10 different aspects.

Begin with your first question.

The AI asks questions the student must think through before responding — fairness, logistics, legal compliance, technology needs, cultural impact. By the end, they have considered the full complexity before drafting a single sentence.

Adapt for your discipline: Replace the policy topic with whatever scoping challenge fits — a market entry strategy, an audit plan, a supply chain redesign, an IT implementation. The technique is the same: AI questions, student thinks.

Teaching tip: Use as a pre-writing exercise. Require students to submit both the transcript and their subsequent draft. Assess whether they incorporated the insights.

4.3 Technique 2: Pros and Cons (for Decision Making)

The AI systematically analyses multiple approaches to a problem, evaluating each against specific criteria. Students must then critically evaluate the analysis — not just accept it.

Example (Supply Chain — strategy evaluation):

Compare different supply chain strategies for a mid-sized manufacturer: make-to-stock, make-to-order, and mass customisation.

For each strategy, provide:

1. A brief description of how it works
2. Three key advantages
3. Three key disadvantages

Evaluate each specifically in terms of:

- Cost efficiency
- Responsiveness to demand changes
- Inventory risk

Conclude with a recommendation for a company facing unpredictable seasonal demand and justify your choice.

The follow-up task (essential): Do not let students accept the AI's recommendation. Require them to challenge one “pro,” add a disadvantage the AI missed, and argue for a different strategy using evidence.

Teaching tip: Project the AI's analysis on screen. Divide students into groups, each arguing for a different strategy. They use the AI's framework but add their own reasoning.

4.4 Technique 3: Stepwise Chain of Thought (for Process)

The AI walks through a complex process one step at a time, pausing after each step until the student signals readiness. This slows down high-stakes processes that students rush through.

Example (Accounting — audit process):

I am learning how to conduct an audit engagement

from planning through to the audit opinion.

Walk me through the entire process. For each step, tell me:

1. What action to take
2. What to document
3. What professional standard or legal requirement applies at this stage

After you explain each step, STOP and wait for me to type "next" before moving on. Do not provide the entire process at once.

Begin with Step 1.

The student must actively engage with each step before progressing. This prevents skipping to the conclusion without understanding the required process.

Teaching tip: Have students repeat the exercise with complications introduced at each step — “What if the client refuses to provide documentation at Step 3?” The AI explains how to handle variations.

4.5 Technique 4: Role Play (for Communication Skills)

The AI adopts a specific persona and engages in a realistic conversation. This is the “flight simulator” concept — students practice difficult professional interactions where the AI responds dynamically.

Example (Management — difficult conversation):

You are Sarah, a high-performing marketing manager who has just been told she is being put on a Performance Improvement Plan. You are shocked, defensive, and angry because you believe this is unfair.

Your behaviour:

- Immediately challenge the fairness of the PIP
- Bring up a recent successful project you led
- Hint that you believe this is retaliation for raising a complaint about your manager
- Be emotional but not abusive

I am the manager who has to conduct this meeting professionally. Stay in character until I manage to de-escalate and establish constructive dialogue.

Begin the meeting. I will speak first.

The student must manage a realistic, emotionally charged scenario — practicing de-escalation, empathy, and procedural fairness while the AI responds dynamically.

Adapt for your discipline: AI plays a demanding customer (marketing), a resistant employee (management), a skeptical auditor (accounting), a supplier announcing delays (supply chain), a stakeholder with unclear requirements (IT).

Teaching tip: Require students to submit the transcript with a reflective analysis — what worked, what did not, what they would do differently.

4.6 Technique 5: Debating (for Strategic and Ethical Analysis)

The AI examines multiple perspectives on a contentious issue, argues both sides, and helps students see complexity and trade-offs.

Example (Economics — multi-stakeholder debate):

Simulate a leadership debate on whether to implement a carbon tax.

Create three personas who will each argue from their perspective:

Persona 1: Environmental Economist

Focus: climate impact, externalities, long-term economic benefit

Persona 2: Industry Representative

Focus: competitiveness, cost burden, employment impact

Persona 3: Treasury Official

Focus: revenue, fiscal policy, implementation feasibility

Have each persona make their opening argument. Then have them respond to each other through three rounds. After the debate, summarise the key tensions and what a compromise policy might look like.

Begin the debate.

Students cannot just pick a side — they must understand legitimate competing perspectives and make a justified recommendation that acknowledges trade-offs.

Teaching tip: Perfect for preparing students for case study exams. Instead of memorising model answers, they practise analysing competing priorities.

4.7 Technique 6: Formative Assessment Generator (for Practice)

The AI generates unlimited practice questions, scenarios, or quizzes for low-stakes repetition and immediate feedback.

Example (Tourism & Hospitality — scenario practice):

I'm studying for an exam on service recovery in hospitality management.

Generate 10 short scenarios (2-3 sentences each) where I need to identify:

1. What went wrong in the service delivery
2. The appropriate recovery strategy
3. What the long-term impact on guest loyalty would be

After I answer each one, tell me if I'm correct and explain why.

Make the scenarios progressively more complex.

Begin with Scenario 1.

Students can generate unlimited practice with instant feedback and adaptive difficulty. Unlike a textbook with 5 practice problems, this provides mastery-based learning.

Critical principle: This is for practice, not graded assessment. Students need to struggle and make mistakes in low-stakes environments.

4.8 Technique 7: The Expert Panel (for Multi-Perspective Analysis)

Students consult multiple AI “experts” with different professional perspectives on the same problem, then synthesise the competing advice into a strategic recommendation.

Example (Information Systems — system implementation):

Our company needs to decide on an ERP implementation approach. Create three expert personas:

Expert 1: IT Architect

Focus: technical feasibility, integration, security, scalability

Expert 2: Change Management Consultant

Focus: user adoption, training, organisational readiness, resistance

Expert 3: CFO

Focus: cost, ROI, risk, business case

Have each expert independently provide:

1. Their assessment of the key risks
2. Their recommended approach
3. Limitations of their recommendation
4. What data they would need to validate it

Begin with Expert 1.

After receiving all perspectives, students must compare where experts agree and conflict, synthesise a recommendation, justify it with theory, and acknowledge trade-offs.

Teaching tip: This technique combines well with Debating (have the experts debate each other) and Role Play (present your synthesis to a skeptical stakeholder).

4.9 How to Choose Which Technique to Use

If you want students to...	Use this technique
Define a complex problem comprehensively	Reverse Prompting
Evaluate competing options and justify a choice	Pros and Cons
Follow a sensitive process correctly	Stepwise Chain of Thought
Practice difficult conversations	Role Play
Understand multiple perspectives and trade-offs	Debating
Build confidence through repetition	Formative Assessment Generator
Synthesise expert advice from different roles	The Expert Panel

4.10 Combining Techniques

The real power comes from combining techniques into a professional workflow:

1. Student uses **Reverse Prompting** to scope out all requirements for a new policy
2. Student uses **Pros and Cons** to evaluate three different approaches
3. Student drafts their chosen policy
4. Student uses **Role Play** to practice explaining it to a skeptical stakeholder
5. Student uses **Debating** to analyse potential criticisms from different perspectives

This sequence takes the student through analysis, decision-making, communication, and critical reflection — a complete professional process.

4.11 Your Action Step

Choose one technique and try it yourself. Pick the one that feels most immediately useful for a topic you are currently teaching. Adapt the example prompt to your discipline and content. These are not just AI tools — they are pedagogical strategies. The AI makes them scalable and available to every student, any time they want to practice.

Chapter 5

Managing LLM Context: Working Smarter with AI Conversations

A long conversation is not necessarily a deep one. After enough turns, the AI forgets where you started, and so might you.

5.1 Why Context Matters: Understanding AI's Limitations

One of the most underrated skills in working with AI is **managing context**: the information you feed to an AI system and how you structure your conversations.

Think of context like the working memory of AI. Unlike humans, who can maintain focus across days of conversation, AI has specific limitations:

- **Limited attention span**: Conversations have maximum lengths before older information becomes less salient (less in focus)
- **Token limits**: Every word you input and every word AI outputs counts against the model's capacity
- **Output token competition**: When you ask for multiple things at once, AI must divide its output tokens among all tasks, often producing shallow results
- **Hallucination risk**: As conversations grow longer and more complex, the risk of AI “making up” information increases

The good news? Understanding and managing context is a learnable

skill that directly improves output quality, saves time, and reduces errors.

This is particularly important for educators because: - You'll be having longer conversations (designing units, iterating on assignments) - You'll need high-quality outputs (teaching materials must be accurate) - You'll want consistent quality across multiple deliverables (course redesigns, prompt libraries) - You'll be modelling these skills for students

5.2 The Four Core Problems

5.2.1 Problem 1: The Long Conversation Problem

What happens: You've been working with AI for 20 exchanges, refining a unit design. The conversation is great, but when you ask question 21, AI gives you an answer that contradicts something from exchange 5.

Why: As conversations grow longer, older information becomes less salient to the AI's attention. While technically the AI can "see" the entire conversation, information from early exchanges has less influence on later responses.

Teaching impact: When designing complex units or courses, you'll have lengthy conversations. Without managing context, outputs become inconsistent.

5.2.2 Problem 2: Output Token Scarcity

What happens: You ask AI to "redesign this unit, create a rubric, write student instructions, design an assessment, and create a facilitator guide." You get five things, but each is shallow because AI divided its output tokens five ways.

Why: Every model has a maximum output token limit (typically 2,000–4,000 tokens). If you ask for 5 things, you get roughly 400–800 tokens per thing. Quality suffers.

Teaching impact: You might get a "complete" unit design that needs heavy revision, or you abandon it and start over, wasting time.

5.2.3 Problem 3: Hallucination Acceleration

What happens: As conversations get longer, AI becomes more likely to "confidently generate false information", making up citations, inventing

examples, or misremembering earlier statements.

Why: Longer conversations increase uncertainty. AI is tracking more information and making more inferential leaps. It tries to fill gaps with plausible-sounding but false information.

Teaching impact: Teaching materials with invented examples or false citations are problematic. Students trust what they see in your materials.

5.2.4 Problem 4: Lost Context Across Sessions

What happens: You close the chat. Next week, you want to continue designing that unit. You paste your earlier thinking into a new chat, but AI doesn't have the full conversation history. It repeats earlier points or misses nuance.

Why: Each new conversation starts fresh. AI has no memory of previous sessions unless you explicitly provide that history.

Teaching impact: Multi-week projects (semester redesigns, curriculum overhauls) become fragmented. You must re-establish context repeatedly.

5.3 Core Strategy 1: Break Complex Tasks Into Steps

The Principle: Before diving into work, ask AI to help you structure the task.

Why it works: - Distributes output tokens efficiently (each step gets full focus) - Reduces hallucination risk (smaller scope per prompt) - Gives you a clear plan to follow - Lets you quality-check each step before moving forward

5.3.1 Example 1: Unit Redesign

Instead of asking all at once:

"Redesign my unit on supply chain management. Create new learning outcomes, design assessments, write student instructions, create a rubric, and a unit description."

Break it into steps, first, ask for a plan:

I'm redesigning a unit on supply chain management for 3rd-year business

students (40 students, mix of majors). Help me create a structured plan.

What are the key steps I should follow to redesign this unit? List them in logical order with what we should accomplish at each step.

AI response: You get a plan like: 1. Clarify learning outcomes (what students should be able to do) 2. Design assessments (how you'll know they've learned) 3. Plan learning activities (what students will do to learn) 4. Create student instructions (what students need to know) 5. Build assessment rubric (how you'll grade) 6. Write facilitator notes (guidance for teaching)

Then work through the plan one step at a time:

Step 1:

Let's start with step 1: Learning Outcomes.

Here are my current outcomes: [paste]

Help me evaluate these:

- Which are clear and measurable?
- Which are too vague?
- Which are most important for supply chain professionals?

Suggest 3-4 revised outcomes that focus on authentic supply chain thinking

Step 2 (after reviewing Step 1 output):

Good. Now step 2: Assessment Design.

For these outcomes: [paste]

Suggest three assessment approaches that would authentically test these outcomes

For each, explain:

- What students do
- Why it tests supply chain thinking
- How it would work with 40 students

And so on. By breaking the work into steps, each output gets full attention and quality improves.

5.3.2 Example 2: Semester Course Planning

Instead of: "Design a 12-week course on organisational behaviour."

Ask for a plan first:

I'm teaching organisational behaviour to 2nd-year business students (60 students,

first-year of their major). Help me structure this 12-week course.

What are the major topics we should cover? In what order? What's the arc across the semester? What should be the rough focus of each week of weeks?

Then work through week-by-week or module-by-module:

Using the plan above, let's design Week 1-2: Introduction to Organisational behaviour.

Learning focus: [paste from plan]

Design these weeks including:

- 2-3 key concepts to introduce
- 1 major activity or case study
- 1 short formative assessment
- Approximately 3-4 readings

Make it manageable for a 3-hour/week course.

5.4 Core Strategy 2: One Task Per Prompt (Usually)

The Principle: Ask for one main thing per prompt, not multiple things.

Why it works: - Each output gets full attention and token allocation (depth, not breadth) - Easier to review and iterate on one thing - Less cognitive load on the model - Quality increases noticeably

5.4.1 Example: Lesson Plan Design

Poor approach (asking for too much):

"Write a lesson plan for teaching critical thinking to business students. Include: 5 learning outcomes, 3 classroom activities, assessment rubric, student handout, and facilitator notes."

Result: Shallow. Each element is skeletal. Outcomes might be vague. Activities are one-liners. Rubric has minimal criteria.

Better approach (one task per prompt):

Prompt 1:

I'm teaching critical thinking to business undergraduates. I want them to be able to analyse business problems from multiple perspectives.

Design 3 classroom activities that help students practice critical thinking. For each activity:

- Describe what students do (step-by-step)
- Explain what they'll learn
- Note how long it takes
- Indicate the group size (individual, small group, whole class)

Prompt 2 (after reviewing):

Good activities. Now let's build an assessment rubric for evaluating student critical thinking. Include:

- 4-5 criteria (e.g., perspective-taking, evidence use, reasoning clarity)
- For each criterion, descriptors for: Excellent / Proficient / Developing

Keep it usable for grading real student work.

Prompt 3 (after reviewing):

Turn the activities and rubric into a one-page student handout. Include:

- What they're learning (1-2 sentences)
- Why it matters professionally (1-2 sentences)
- The activities (clear instructions)
- Success criteria (what "good" looks like)
- How to get help if stuck

Make it accessible and encouraging.

Result: Deep. Each element is thoughtful, specific, and builds on what came before.

5.4.2 Exception: When Multiple Things Are Fine

Sometimes asking for multiple outputs makes sense: - **Comparative tasks:** "Show me 3 different ways to teach [concept]. What are the trade-offs of each?" - **Structured formats:** "Create an outline with: learning outcomes, key concepts, and 3 discussion questions" - **Quick iterations:** "Now make that more concise / more challenging / more inclusive" - **Batched similar tasks:** "Write 5 discussion questions on these topics: [list]. Each should take 10 minutes of discussion."

The key: Are the outputs relatively equal in scope and complexity? If yes, ask for multiple. If one task is much bigger than others, split them.

5.5 Core Strategy 3: Use Output Constraints to Manage Tokens

The Principle: When asking for multiple things, specify output size/structure upfront. This helps AI divide tokens wisely.

5.5.1 Example: Assessment Comparison

Without constraints:

Compare portfolio assessment vs. exam-based assessment for evaluating student learning in business courses. What are the advantages and limitations? When should I use each?

Result: AI might spend 70% of tokens on one approach and 30% on the other. Output is imbalanced.

With constraints:

Compare portfolio assessment vs. exam-based assessment for my business course.

For each approach, provide:

- 2 key advantages
- 2 key limitations
- Best when: [one sentence]

Keep each section to 3-4 sentences maximum. Focus on practical classroom implications.

Result: AI knows exactly how to divide tokens. Output is balanced, concise, and usable.

5.5.2 Template for Token-Aware Requests

I need [specific output type]. Provide:

1. [First thing] - [length/format]
2. [Second thing] - [length/format]
3. [Third thing] - [length/format]

Keep total output under [X words]. Prioritize clarity and specificity over completeness.

5.5.3 Discipline-Specific Example

Example: Accounting

Compare three audit sampling approaches: statistical, risk-based, and procedural.

For each:

- 2 advantages for audit evidence
- 2 limitations
- Audit standards alignment: brief comment
- Best when: [one sentence]

Focus on what an audit team needs to decide.

5.6 Core Strategy 4: Keep Conversations Focused and Modular

The Principle: Use separate conversations for separate projects or major topic shifts.

Why it works: - Shorter conversations = less hallucination risk - Easier to find earlier outputs (scrolling back is simpler) - AI stays focused on one topic - Cleaner record-keeping (export or save by topic)

5.6.1 When to Start a New Conversation

1. **Topic shift:** Finished designing one unit? Start a new conversation for a different unit.
2. **Major context change:** Moving from unit design to research methodology? New conversation.
3. **Length:** Conversation getting very long (50+ exchanges)? Consider summarising and moving to a new one.
4. **Different AI tool:** Using Claude for teaching design and ChatGPT for grading assistance? Keep them separate.

5.6.2 When One Conversation Is Fine

- Iterative work on the same project (refining, revising)
- Related follow-ups (asking for adaptations of earlier output)
- Building on previous steps (multi-step workflows like the ones above)

Rule of thumb: One conversation per major project. Use the same conversation as you iterate and refine within that project. Start a new conversation when you move to a different project.

5.7 Core Strategy 5: summarise and Handoff for Long Conversations

The Principle: When a conversation gets long, ask AI to summarise what you've accomplished, then start fresh in a new conversation.

Why it works: - Resets the "attention freshness" (AI isn't tracking 30+ old exchanges) - Gives you a clean document of what you've done (useful archive) - Reduces hallucination in the new conversation - Allows you to build on work without repeating context

5.7.1 How to Do It

In the long conversation, when it feels unwieldy:

We've been working on [project name: e.g., "redesigning the HR management unit"] for a while. Can you summarise what we've accomplished?

Include:

- What problem or task we started with
- Key decisions we made
- What we've created/designed so far
- What still needs to be done

Make it concise but complete-something I can copy and paste into a new conversation to continue working.

AI provides a summary. Then:

1. Copy that summary
2. Start a new conversation
3. Paste the summary at the beginning
4. Add: "I'm continuing this work. Here's what we've done. Let's move forward with [next step]."
5. Continue from there

5.7.2 Example Handoff Summary

****Project:**** Redesigning Supply Chain Management Unit (3rd year, 40 s

****What we started with:****

- Students struggle to see supply chain as strategic (not just logistics)
- Current unit is mostly descriptive (case studies + lectures)
- Goal: more active learning, professional judgment development

****Key decisions made:****

- Organised around real supply chain decisions (not functions)
- Emphasised risk thinking and trade-offs
- Used simulation and case critique for learning

****What we've created:****

1. Learning outcomes (5 outcomes emphasising strategic thinking)
2. Assessment strategy (3 assessments: case analysis, risk simulation, tea
3. Week-by-week learning plan (12 weeks structured around decisions)

****What's left:****

- Detailed activity instructions for each week
- Facilitator notes on how to run discussions
- Student rubric for team project

****Next step:**** Design the Week 1-2 activities in detail.

You then start a new conversation and continue with: "I'm continuing supply chain unit redesign. Here's our progress. Let's design Week 1-2 activities."

5.8 Core Strategy 6: Make Context Explicit and Structured

The Principle: Don't assume AI remembers or understands implicit context. State it clearly.

5.8.1 Poor Context Example

"How should I handle participation in class?"

Missing: What is "participation"? What's the class? How big? What's the issue? What's your teaching style? What have you tried?

5.8.2 Good Context Example

I teach marketing to 80 business students (2nd year, mix of domestic and international). We use large lecture format (one 2-hour session per week).

Student participation problem: Maybe 10 students ask questions or off
The other 70 are silent.

I want broader participation without:

- Making it feel forced or uncomfortable
- Losing lecture efficiency
- Putting shy students on the spot

How can I increase participation?

5.8.3 Better Context (If Continuing Earlier Work)

Remember we're redesigning the marketing unit on consumer behaviour.
We wanted more student participation—we had the problem where only a
students spoke in the large lecture.

We've started using think-pair-share activities in class. They've hel

But now we're thinking about assessments. How can we design assessmen
that encourage quieter students to engage and show their thinking?

5.8.4 Checklist for Explicit Context

- **Who:** Who are the students? (Level, major, background, cohort size, cultural mix)
- **What:** What's the specific task or problem? (Not vague; specific)
- **Why:** Why does it matter? (Learning goal, professional relevance, student challenge)
- **Constraints:** What are the limitations? (Time available, resources, institutional requirements)
- **Style:** What's your teaching approach? What's worked before? What hasn't?
- **History:** Have we worked on this before? What did we already decide?

5.9 Core Strategy 7: Batch Similar Tasks

The Principle: When you have multiple similar tasks, batch them efficiently.

5.9.1 Poor Approach

[You] "Write a discussion question on leadership styles for my manage

[AI] [Provides question]

[You] "Review. Good. Now write a discussion question on ethical decision-making."

[AI] [Provides question]

[You] "Review. Good. Now write one on team conflict."

[Repeat 3+ more times]

Problem: This takes 10+ exchanges. You repeat context setup each time. Token efficiency is poor.

5.9.2 Better Approach

Single prompt:

I need 5 discussion questions for a 12-week management unit. They should:

- Progress from basic understanding to critical analysis
- Take 8-12 minutes of class discussion
- Spark respectful debate (not yes/no questions)
- Be relevant to business students' future work

Topics (one question each):

1. Leadership styles and contexts
2. Ethical decision-making in organisations
3. Managing team conflict
4. Change management and resistance
5. Inclusive leadership and diversity

Provide all 5 questions with a note about why each one works for discussion

Result: Single exchange. AI understands the pattern. All 5 questions are high-quality and consistent.

Then iterate once if needed:

These are good. Now adapt question 3 (team conflict) to include an international context. Some of my students are from cultures where conflict is handled very differently than in Western business tradition.

5.10 The Two-Chat Workflow: Separate Thinking from Building

By now you have seen how breaking tasks into steps and keeping conversations focused improves quality. There is a deeper version of this

principle that changes how you work with AI entirely: **use two separate sessions, one for thinking and one for building.**

This idea builds on the Two-Chat Workflow from *Conversation, Not Delegation* (Borck, 2025), adapted here for teaching practice. It is simple, powerful, and worth making a habit.

Chat 1: Explore and clarify. Open a session with no expectation of producing finished output. Use it to probe the teaching challenge you are facing. What are you actually trying to achieve? What assumptions are you making about your students? What alternatives have you not considered? Let the conversation wander. Challenge what the AI suggests. Follow tangents. The messier this session is, the clearer your thinking becomes.

Chat 2: Build from your decisions. Start a fresh session and arrive with a focused brief — not a vague request, but a set of deliberate choices about what you want, who it is for, and what constraints apply. The quality of this output depends almost entirely on the quality of the brief you wrote after reflecting on Chat 1.

The most important moment is the gap between the two chats. You do not dump everything from the first session into the second. You review what emerged, keep what matters, discard what does not, and write a clear brief that reflects *your* decisions. That act of curation is where your professional judgement lives — and it is the part no AI can do for you.

5.10.1 Example: Redesigning an Assessment

Thinking Chat:

I'm rethinking how I assess critical thinking in my 2nd-year management unit.

Currently I use a written case analysis, but I'm not sure it actually assesses critical thinking versus summarising. Help me think through this.

What does critical thinking actually look like in a management context? What would strong evidence of it look like in student work?

You spend 10–15 minutes exploring: What counts as critical thinking in management? How is it different from analysis? What would a weak submission look like versus a strong one? The AI helps you think, but you are doing the intellectual work of deciding what matters.


Build Chat:

I'm redesigning the critical thinking assessment for my 2nd-year management unit (60 students). After reflection, here's what I've decided:


- Critical thinking in management means evaluating competing stakeholder perspectives and defending a position with evidence
- I want students to critique a real management decision, not just describe it
- The assessment should be 1500 words with a structured argument
- I need a rubric that distinguishes "summarising the case" from "evaluating the decision"

Design the assessment brief and rubric based on these decisions.

Notice the difference. The build chat gets a clear, decided brief. The output will be dramatically better than if you had started cold with “design me a critical thinking assessment.”

 Try this (5 minutes)

Pick a teaching task you have been putting off — perhaps redesigning a tutorial activity or rethinking a rubric. Open an AI session and spend five minutes just exploring the problem: what is not working, what you have tried, what you are unsure about. Do not ask for any deliverables. Then close that session, jot down your key decisions in two or three sentences, and open a fresh session with those decisions as your opening brief. Compare the result to what you would have gotten from a single cold prompt. The difference is usually striking.

 You are the bridge

The exploring chat does not know what you will eventually build. The building chat does not know what options you considered and rejected. Only you hold both sides. That is what makes you irreplaceable in this process — not as someone who checks AI’s work after the fact, but as the person whose judgement connects exploration to execution.

This workflow connects directly to the average-versus-precise, small-versus-large framework from the earlier chapter on LLMs. That framework tells you where a task sits. The two-chat workflow tells you what to do about it. A task that starts in the danger zone (large and precise, like redesigning an entire unit’s assessment strategy) feels overwhelming as a single prompt. But the exploring chat breaks it into components that each sit in different quadrants. Some pieces land in the sweet spot. Others need careful verification. Each gets an appropriate level of trust and oversight. The exploring session is where you map the territory. The building session is where you execute with that map in hand.

5.11 Common Mistakes and How to Fix Them

Mistake	What Goes Wrong	Fix
Asking for 10 things at once	Output is shallow; tokens divided 10 ways	Break into 2–3 prompts— one main task per prompt
Vague task description	AI misunderstands what you want	Add explicit context: who— what— why— constraints
“Design my whole course” in one go	Incoherent— shallow output	Ask for plan first— then design one section at a time
Leaving conversation open indefinitely	Hallucination risk increases; unwieldy to navigate	Start new conversation every 50+ exchanges
Not specifying output format	AI guesses format; may not match needs	Say “3 bullet points—” “one paragraph—” “table—” etc.

Mistake	What Goes Wrong	Fix
Asking “what am I missing?”	AI invents irrelevant things	Be specific: “What am I missing in my assessment of [specific skill]?”
Forgetting to review outputs	Errors and hallucinations slip through	Always quality-check—especially facts/citations
Pasting entire documents without framing	AI doesn’t know what to focus on	Add a sentence: “Here’s my unit outline. Focus on the assessment section.”
Starting a new conversation when context is long	Lost work and having to re-explain everything	summarise first—then paste summary into new chat

5.12 Practical Workflow for Managing Context

Here’s a workflow that brings everything together:

5.12.1 Phase 1: Planning

1. Define the task clearly (in writing, to yourself)
2. Ask AI for a plan before diving in
3. Break the plan into sub-tasks
4. Identify how much output you need for each sub-task

5.12.2 Phase 2: Execution

1. Work through one sub-task per prompt (usually)
2. Review each output before moving forward
3. Provide feedback for refinement
4. Document what works (save successful prompts)

5.12.3 Phase 3: Management

1. Keep conversations focused (one major project per conversation)
2. When a conversation gets long (50+ exchanges), ask for a summary and move to a new conversation
3. Use separate conversations for different topics/projects
4. Archive completed work

5.12.4 Phase 4: Quality Check

1. Verify facts (especially citations, dates, statistics, attributions)
 2. Check for contradictions (does it align with earlier outputs?)
 3. Assess completeness (did AI address all your needs?)
 4. Iterate if needed (use follow-up prompts to refine, not to ask for entirely new things)
-

5.13 Real-World Example: Managing Context Well

Scenario: Designing a 10-week supply chain management unit.

5.13.1 Bad approach (what NOT to do):

"Design the entire supply chain management unit including all 10 week learning outcomes, assessments, readings, 3 activities per week, face notes, and student rubrics."

Result: Massive output that's shallow and poorly integrated. You'd need to revise everything piecemeal.

5.13.2 Good approach:

Conversation 1: Planning

I'm redesigning a 10-week supply chain management unit for 3rd-year business students (40 students, mix of majors). Help me create a modular plan.

What are the key supply chain topics we should cover? How should they sequence? What's the learning arc across the semester? What should each week focus on?

Result: You get a coherent 10-week plan with learning progression.

Conversation 2: Learning Outcomes

Using the plan from earlier, let's define learning outcomes.

Here's the plan: [paste from Conversation 1]

For each week or pair of weeks, suggest 1-2 specific, measurable outcomes that focus on authentic supply chain thinking (not just knowledge).

Result: Outcomes aligned to the plan, focused on professional judgment.

Conversation 3: Week 1 Deep Dive

Let's design Week 1 in detail. Topic: Supply Chain Fundamentals and Strategic Thinking.

Learning outcomes: [paste from Conversation 2]

Design Week 1 including:

- 3 key concepts to introduce
- 1 major activity or case study
- 1 short assessment
- 3-4 readings
- Facilitator notes on how to run discussion

Make it manageable for a 3-hour week.

Result: A coherent, complete Week 1.

Conversation 4: Weeks 2-3 Deep Dive

Continue with Weeks 2-3 using the same structure...

Result: By batching weeks and working modularly, the whole unit comes together coherently.

Conversation 5: Assessment Integration

I've now designed all 10 weeks. Here's a summary of all learning outcomes and activities: [paste summary]

Design a capstone assessment that integrates learning from across the unit. What should students do? How would you evaluate whether they've achieved the learning outcomes?

Result: A coherent, well-integrated unit with assessment that ties it all together.

5.14 Context Management for Different Scenarios

5.14.1 For Unit Redesign

Break down like this: 1. Conversation 1: Plan (topics, sequence, learning arc) 2. Conversation 2: Learning outcomes (aligned to plan) 3. Conversations 3+: One section per conversation (activities, assessments, etc.) 4. Final conversation: Integration (how it all connects)

Benefit: Quality outputs. Each conversation focuses on one aspect. By the end, you have a coherent unit designed through multiple focused conversations.

5.14.2 For Course-Level Change

Break down like this: 1. Conversation 1: Architecture (major themes, year-long learning arc) 2. Conversation 2: Learning outcomes for the year (connected to architecture) 3. Conversations 3+: One unit per conversation (each unit designed fully) 4. Final conversation: Integration (how units connect, capstone design)

Benefit: Coherence across the year. Each unit is designed well. The course flows logically.

5.14.3 For Assessment Redesign

Break down like this: 1. Conversation 1: Assessment strategy (what to assess, how, when) 2. Conversation 2: Individual assessment design (one assessment at a time) 3. Conversation 3: Rubrics (one per assessment) 4. Conversation 4: Student communication (handouts, success criteria, examples)

Benefit: Assessments that actually measure what you care about. Clear communication to students.

5.15 When Context Management Matters Most

Context management is most important when: - **You're doing complex, multi-step projects** (unit redesigns, curriculum overhauls) - **Quality matters** (teaching materials, student-facing work) - **You need consistency** (prompt libraries, course coherence) - **You're iterating** (refining approaches based on feedback) - **You're teaching students to use AI** (modelling good context management)

For quick, one-off tasks (generating a single prompt, quick idea generation), context management is less critical. But for the substantial work you do as an educator, managing context improves quality dramatically.

5.16 Key Principles Summary

1. **Break complexity into steps** - Ask for a plan before diving in
2. **One task per prompt (usually)** - Give output tokens to focus on one thing
3. **Use output constraints** - Specify length and format to manage token allocation
4. **Keep conversations focused** - One major project per conversation
5. **summarise and handoff** - When conversations get long (50+ exchanges), reset with a summary
6. **Make context explicit** - Don't assume AI understands implicit information
7. **Batch similar tasks** - If you need 5 of the same thing, ask for all 5 at once
8. **Review everything** - Always quality-check outputs

The underlying principle: Context management is about respecting the AI's limitations while maximising its strengths. You're not trying to have perfect conversations; you're trying to have *focused* conversations that produce high-quality outputs consistently.

5.17 Why Students Should Learn This

As you teach students to use AI, context management becomes a critical skill. Students who understand context management will: - **Get better results** from AI (more usable outputs, fewer iterations) - **Work more efficiently** (fewer wasted conversations) - **Produce higher-quality work** (depth over breadth) - **Develop professional AI literacy** (understanding how to work with AI tools effectively)

Consider teaching context management explicitly: - Show students your workflow (how you break tasks into steps) - Model managing long conversations (summarise, start fresh) - Have students practice the “one task per prompt” principle - Discuss why quality suffers when asking for too much at once

This transfers from classroom to professional practice. If your students graduate understanding how to manage context with AI, they’ll be more effective professionals.

5.18 Your Next Step

Pick a project you’re currently working on or about to start:

1. **Define it clearly:** What are you trying to accomplish?
2. **Ask AI for a plan:** Before diving in, ask AI to help you structure the work
3. **Break into steps:** Work through the plan one step at a time
4. **Keep it focused:** One conversation per major project
5. **Review everything:** Quality-check before moving forward

As you do this, notice: - How much more focused your outputs are - How much easier iteration becomes - How much less rework you need to do

Then bring that experience to your teaching. Your students will benefit from seeing how you work with AI effectively.

Chapter 6

The Flight Simulator - Conversation Simulations

Pilots do not learn to fly by reading about turbulence. They learn by experiencing it in a simulator where mistakes are cheap. Your students deserve the same.

6.1 The Three-Phase System

In the previous chapter, you learned the Role Play technique. Now we're going to transform that into a complete professional practice system with three distinct phases:

Phase 1: The Setup (You or the student designs the scenario) **Phase 2: The Simulation** (Student practices the conversation) **Phase 3: The Debrief** (AI critiques the student's performance)

This three-phase approach mirrors how professionals develop expertise: briefing, practice, and reflective analysis. The AI makes this process scalable, personalised, and repeatable.

6.2 Why This Is Powerful

Traditional role-play in class has limitations: - Limited time means each student gets one attempt - Peer role-play can be inconsistent (your classmate might not play the "difficult employee" convincingly) - Students feel self-conscious performing in front of others - Feedback is often delayed and general rather than specific

AI-powered simulation solves all of these: - Students can practice the same scenario five times until they get it right - The AI consistently plays the role as designed - Students can practice privately, making mistakes without embarrassment - Feedback is immediate, specific, and tied to learning objectives

Let's walk through a complete example.

6.3 Complete Worked Example: The Performance Improvement Plan Meeting

6.3.1 Context

This is a common high-stakes conversation in HR. It's legally sensitive, emotionally difficult, and requires balancing empathy with accountability. (We'll show examples from other disciplines at the end of this chapter.)

Learning Objectives: - Apply principles of procedural fairness - Communicate difficult feedback clearly and constructively - Demonstrate empathy while maintaining professional boundaries - Document the conversation appropriately - Avoid common legal pitfalls (discrimination, unfair dismissal)

6.4 Phase 1: The Setup Prompt

This is where you (or the student) design the scenario. The setup prompt defines: 1. The AI's role and personality 2. The context and background 3. The challenge or conflict 4. Behavioural guidelines for the AI 5. When the simulation should end

6.4.1 Setup Prompt (Ready to Use)

You are Alex Chen, a 32-year-old software engineer who has worked at TechCorp for three years. You were a high performer until six months ago, when your work quality declined noticeably.

BACKGROUND CONTEXT:

- You have missed three project deadlines in the past four months
- Your code quality has dropped-recent pull requests have had multiple bugs that QA caught
- You've been arriving late to work 2-3 times per week
- Your manager has already had one informal conversation with you about

these concerns two months ago, but things have not improved

- You are being placed on a formal 60-day Performance Improvement Plan (PIP)

YOUR PERSONA AND behaviour:

- You are defensive and embarrassed. You feel this PIP is unfair and
- You believe your recent work is "good enough" and that your manager unrealistic expectations
- You are dealing with a stressful personal situation (caring for an parent), but you have not told your manager about this—you believe private and not relevant to work
- You are not aggressive or hostile, but you are resistant and will p back on criticism
- If the HR person or manager shows genuine empathy and asks open que you will gradually reveal some personal context
- If they are accusatory or dismissive, you will shut down and become minimal in your responses

THE MEETING:

I am the HR representative (or the manager) conducting the PIP meeting you. This is a formal meeting. You have just been told this is a PIP discussion.

YOUR GOAL:

- React as Alex would realistically react
- Challenge the fairness of the process if it seems rushed or poorly explained
- Reveal personal information only if trust is established
- Show emotion (frustration, stress, worry) but remain professional
- Respond dynamically to how I conduct the meeting

Stay in character throughout. Begin the meeting by responding to my o statement. I will start the conversation now.

Key Design Elements: - **Complexity:** Alex isn't purely defensive or purely cooperative—there's legitimate work performance decline *and* extenuating circumstances - **Realistic emotion:** Alex is stressed and embarrassed, not theatrical or extreme - **Conditional behaviour:** How Alex responds depends on the student's approach (rewards good technique) - **Hidden information:** The personal context isn't revealed unless the student asks the right questions with the right tone

6.5 Phase 2: The Simulation

The student now conducts the meeting. Here’s an example of how this might play out.

6.5.1 Example Conversation

Student (HR Rep): > “Hi Alex, thanks for meeting with me. I want to talk with you today about some performance concerns that have been raised. I know this is a difficult conversation, but I want to make sure we have an open discussion about what’s been happening and how we can support you moving forward.”

AI (as Alex): > *shifts uncomfortably* “Yeah, I’ve been hearing whispers about this. Look, I know I’ve missed a couple of deadlines, but I don’t think things are as bad as people are making out. I’m still getting my work done. Is this really necessary?”

Student: > “I understand this feels uncomfortable. Let me be specific about the concerns. Over the past four months, you’ve missed three major project deadlines—the payment gateway integration in March, the API update in April, and the mobile optimisation in May. Additionally, QA has flagged an increase in bugs in your recent code submissions. These are measurable changes from your previous performance standard. Can you help me understand what’s been happening?”

AI (as Alex): > *pauses, looks down* “I mean... yeah, okay, I’ve had some issues. But deadlines get missed sometimes, right? And bugs happen. I just feel like everyone’s being really harsh all of a sudden. I had one conversation with my manager about this two months ago, and now suddenly I’m on a PIP? That feels pretty extreme.”

Student: > “That’s a fair question, Alex. You did have an informal conversation with your manager in early April where these concerns were raised. At that time, you both agreed on some strategies to get things back on track. Since then, we’ve continued to see the same patterns—another deadline missed last month and continued quality issues. A PIP isn’t a punishment; it’s a structured way to clearly define expectations, provide support, and give you a fair opportunity to improve. Does that make sense?”

AI (as Alex): > *sighs* “I guess. I just... it’s been a really tough few months personally, okay? I know that’s not your problem, but it’s not like I’m just slacking off or don’t care anymore.”

Student: > “Thank you for sharing that, Alex. I appreciate that personal circumstances can affect work, and I’m genuinely sorry you’re dealing with difficult things. If there are specific supports we can offer—whether

that's adjustments to your schedule, access to our Employee Assistance Program, or other accommodations—we're open to discussing that. But we also need to be clear about the performance standards that need to be met. Can we talk about what support might help you succeed during this PIP period?"

AI (as Alex): > *softens slightly* "Honestly, I've been dealing with my dad's health declining. I've been having to take him to medical appointments, and I'm the only family member nearby. Some mornings I'm exhausted because I've been up half the night. I didn't want to bring it into work because... I don't know, I thought I could just handle it."

Student: > "I'm really sorry to hear about your father, Alex. That's incredibly stressful. I'm glad you've told me, because this is exactly the kind of situation where we want to support our employees. Have you considered whether a temporary flexible schedule might help—perhaps starting later on days when you have early appointments, or adjusting your hours? We can also connect you with our EAP, which offers counseling and support for family caregivers. Would either of those be helpful?"

AI (as Alex): > *nods* "Yeah... actually, that might really help. I didn't know that was an option. I've just been trying to power through, but clearly that's not working."

Student: > "I appreciate your honesty. So here's what I suggest: we'll document this PIP with clear performance targets—specific deadlines and quality standards. But we'll also document the support measures we're putting in place, including flexible scheduling where possible. This gives you the structure and the support to get back to the high performance standard we know you're capable of. The PIP period is 60 days. At the 30-day mark, we'll have a check-in to see how things are going. Does that sound fair?"

AI (as Alex): > *takes a breath* "Okay. Yeah. I mean, it's still stressful to be on a PIP, but I get it. And if I can actually get some flexibility with the schedule, that would help a lot. I don't want to lose this job—I just need to get things back on track."

Student: > "That's exactly the goal, Alex. We want you to succeed. Let's schedule a follow-up meeting next week to formalize the flexible schedule arrangement and go through the detailed PIP document. Thank you for being open with me today—that's really important."

AI (as Alex): > *manages a small, tired smile* "Thanks. I appreciate you actually listening."

6.6 Phase 3: The Debrief (AI Critique)

After the simulation ends, the student copies the entire conversation transcript and submits it to the AI with a **critique prompt**. This is where the real learning happens.

6.6.1 Critique Prompt (Ready to Use)

You are an expert HR educator evaluating a student's performance in a simulated Performance Improvement Plan (PIP) meeting.

Below is the full transcript of the conversation between the student (acting as the HR Representative) and a simulated employee (Alex Chen) who was being placed on a PIP.

Your task is to provide a detailed critique of the student's performance across the following criteria:

1. PROCEDURAL FAIRNESS

- Did the student clearly explain why the meeting was happening?
- Did they provide specific examples of performance issues?
- Did they give Alex an opportunity to respond and explain?
- Did they explain the PIP process fairly and clearly?

2. COMMUNICATION EFFECTIVENESS

- Was the tone appropriate (professional, empathetic but not dismissive)?
- Did the student use open questions to encourage dialogue?
- How did they handle Alex's defensiveness?
- Did they actively listen and acknowledge Alex's concerns?

3. LEGAL AND ETHICAL CONSIDERATIONS

- Did the student avoid any discriminatory language or assumptions?
- Did they appropriately handle the disclosure of personal circumstances?
- Did they maintain appropriate professional boundaries?
- Did they document support measures as well as performance requirements?

4. APPLICATION OF HR THEORY

- What motivational or psychological principles did the student apply (or fail to apply)?
- How well did they balance accountability with support?
- Did they create conditions for behavioural change?

For each criterion, provide:

- A score out of 10
- Specific examples from the transcript (quote the student's words)

- What the student did well
- What the student could improve
- One concrete suggestion for how to improve

Finally, provide an overall assessment: If this were a real PIP meeting, what would likely happen next based on how the student conducted it?

Here is the transcript:

[PASTE TRANSCRIPT HERE]

Begin your critique.

6.6.2 Example Critique Output

The AI will provide something like this:

CRITIQUE OF PIP MEETING PERFORMANCE

1. PROCEDURAL FAIRNESS: 9/10

What you did well: You were exemplary in providing specific, measurable examples of performance concerns. When you said, *“Over the past four months, you’ve missed three major project deadlines—the payment gateway integration in March, the API update in April, and the mobile optimisation in May,”* you gave Alex concrete information to respond to, not vague accusations. This is critical for procedural fairness.

You also clearly explained the purpose of the PIP: *“A PIP isn’t a punishment; it’s a structured way to clearly define expectations, provide support, and give you a fair opportunity to improve.”* This reframing helped reduce Alex’s defensiveness.

Where you could improve: You could have been even more explicit about the timeline and previous feedback. While you mentioned the April conversation, you could have said: *“We had a conversation on April 3rd where we documented three specific concerns. We agreed you would [specific actions]. Since then, we’ve seen [specific outcomes]. That’s why we’re now moving to a formal PIP.”*

Suggestion: Always reference specific dates and documented prior warnings. This strengthens the legal defensibility of your process.

2. COMMUNICATION EFFECTIVENESS: 10/10

What you did well: Your use of empathy was outstanding. When Alex revealed personal stress, you immediately acknowledged it: *“Thank you for sharing that, Alex. I appreciate that personal circumstances can affect work, and I’m genuinely sorry you’re dealing with difficult things.”* This response validated Alex’s feelings while maintaining professional boundaries.

You also used excellent open-ended questions: *“Can you help me understand what’s been happening?”* This invited dialogue rather than putting Alex on the defensive.

Your handling of defensiveness was textbook-perfect. When Alex said the PIP felt “extreme,” you didn’t argue or dismiss—you acknowledged the concern (*“That’s a fair question”*) and then explained the rationale.

Where you could improve: Nothing significant. Your communication was professional, empathetic, and effective throughout.

3. LEGAL AND ETHICAL CONSIDERATIONS: 9/10

What you did well: When Alex disclosed caregiving responsibilities, you immediately offered accommodations: *“Have you considered whether a temporary flexible schedule might help?”* This is legally and ethically sound—you’re demonstrating that the company is willing to make reasonable adjustments, which protects against claims of unfair dismissal or discrimination.

You also avoided making assumptions or probing inappropriately into Alex’s personal life. You accepted the disclosure, offered support, and moved the conversation forward professionally.

Where you could improve: You should explicitly document this conversation. At the end of the meeting, you could have said: *“I’m going to document today’s conversation, including your disclosure about your father’s health and the flexible schedule accommodation we’ve discussed. You’ll receive a copy of this documentation for your records.”*

Suggestion: Always verbally confirm that documentation will be provided. This ensures transparency and gives the employee a chance to correct any misunderstandings immediately.

4. APPLICATION OF HR THEORY: 8/10

What you did well: You demonstrated understanding of **motivational theory** by balancing challenge with support. You didn’t just impose standards—you co-created solutions by asking, *“Can we talk about what support might help you succeed during this PIP period?”*

You also applied principles of **procedural justice** by ensuring Alex felt heard and treated fairly, which research shows increases acceptance of difficult decisions.

Where you could improve: You could have applied **goal-setting theory** more explicitly. When finalising the PIP, you mentioned “clear performance targets,” but you didn’t collaboratively define them in the meeting. Research shows that when employees participate in setting their own goals, commitment and achievement increase.

Suggestion: In your follow-up meeting, involve Alex in defining the specific, measurable targets for the PIP. For example: “What do you think is a realistic number of on-time deliverables over the next 30 days, given the support we’re putting in place?” This increases ownership and reduces resistance.

OVERALL ASSESSMENT

Likely Outcome: Based on how you conducted this meeting, Alex is likely to:

- Feel that the process was fair, even though it’s stressful
- Engage constructively with the PIP rather than resist or disengage
- Utilise the flexible schedule accommodation
- Have a reasonable chance of meeting the performance standards during the 60-day period

Why: You established psychological safety (Alex felt safe disclosing personal information), you demonstrated empathy without compromising standards, and you offered concrete support. This is exactly the balance effective HR professionals must strike.

One Risk to Watch: Ensure that the flexible schedule accommodation is documented and approved by the manager. If the manager later claims they weren’t consulted or don’t agree, this could create conflict. Always confirm that proposed accommodations are feasible with all stakeholders before committing to them in a meeting.

FINAL SCORE: 36/40 (90%)

This was a highly competent performance. You demonstrated strong communication skills, applied HR principles effectively, and handled a complex, emotionally charged situation with professionalism and empathy. The areas for improvement are minor and relate to documentation and collaborative goal-setting—both easily addressed in your follow-up meeting.

6.7 What Makes This Critique Valuable

Notice what the AI's critique provides:

1. **Specific feedback tied to theory** (procedural fairness, motivational theory, goal-setting theory)
2. **Direct quotes from the student's performance** (not vague generalisations)
3. **Clear scores** for each criterion (makes assessment transparent)
4. **Actionable suggestions** (not just "do better," but *how* to do better)
5. **Predictive analysis** (what would likely happen next in the real world)

This level of detailed, personalised feedback would take a lecturer 20-30 minutes per student to write manually. The AI generates it in seconds, and it's consistent across all students.

6.8 How to Use This as an Assessment

You have several options for how to structure this as a graded assignment:

6.8.1 Option 1: Simulation + Critique Submission

Student Requirements: 1. Conduct the simulation (submit the full transcript) 2. Generate the AI critique using the provided prompt 3. Write a 500-word reflective response addressing: - Do you agree with the AI's critique? Why or why not? - Choose one piece of critical feedback from the AI. How would you apply this in your next attempt? - Identify one HR theory or legal principle that guided your approach. Did you apply it effectively?

What You Grade: - Quality of the conversation (did they demonstrate the required skills?) - Depth of reflection (do they show genuine self-assessment and learning?) - Theoretical integration (can they connect practice to theory?)

6.8.2 Option 2: Simulation + Re-Do + Comparison

Student Requirements: 1. Conduct the simulation (first attempt) 2. Generate the AI critique 3. Conduct the same simulation again,

incorporating the feedback 4. Write a comparative analysis: What changed between attempt 1 and attempt 2? What did you learn?

What You Grade: - Evidence of improvement between attempts - Ability to integrate feedback - Quality of self-directed learning

6.8.3 Option 3: Student-Designed Simulation

Student Requirements: 1. Design your own simulation scenario (write the setup prompt for a different HR situation) 2. Justify why this scenario targets specific learning objectives 3. Conduct the simulation 4. Generate and respond to the critique

What You Grade: - Quality of scenario design (does it create a meaningful learning challenge?) - Performance in the simulation - Reflective analysis

6.9 Adapting This for Different Conversational Skills

The three-phase system works for any high-stakes conversation in business. Below are examples across different skills and disciplines:

Skill Area	Simulation Scenario	Key Learning Focus
HR: Recruitment	AI plays a candidate with vague answers	Behavioural interviewing— probing— bias awareness
HR: Conflict Resolution	AI plays employee making complaint	Investigation— empathy— impartiality— documentation
HR: Negotiation	AI plays union representative	Interest-based negotiation— legal boundaries— compromise
Management: Change Leadership	AI plays resistant manager	Communication— stakeholder management— emotional intelligence
Marketing: Client Negotiation	AI plays demanding client	Persuasion— objection handling— value communication


Skill Area	Simulation Scenario	Key Learning Focus
Supply Chain: Supplier Crisis	AI plays supplier announcing delay	Problem-solving— negotiation— relationship management
Information Systems: Implementation	AI plays resistant end-user	Change management— technical communication— adoption strategies
Accounting: Audit Findings	AI plays finance manager	Compliance communication— relationship management— remediation planning
Business: Board Presentation	AI plays questioning board member	Executive communication— strategic thinking— confidence

For each scenario, you simply adjust: 1. The persona and context in the **Setup Prompt** 2. The evaluation criteria in the **Critique Prompt** 3. The learning objectives you're targeting

The three-phase structure remains the same.

6.10 Cross-Discipline Applications: Flight Simulator Adaptations

The flight simulator technique can be adapted for conversational skills in any business discipline. Below are examples showing how the three-phase system applies to different professional contexts.

 Example: Tourism & Hospitality — Guest Complaint Resolution

Context: A hotel manager must handle a serious guest complaint about service quality.

Setup Prompt Example:

You are Sarah Thompson, a 32-year-old front office manager at LuxuryResort. You've managed the front desk in processes.

BACKGROUND CONTEXT:

- Guest satisfaction scores dropped from 4.8 to 3.9 out of 5
- Three formal complaints about rude staff behaviour in past month
- Check-in wait times averaging 45 minutes (target is 15 minutes)
- You received informal feedback about team morale issues 6 weeks ago

YOUR PERSONA AND behaviour:

- You're defensive and believe guest expectations are unrealistic
- You're dealing with high staff turnover and understaffing
- You feel caught between corporate standards and local service culture
- You'll initially blame external factors but open up with empathy

THE MEETING:

I am the resort general manager addressing recent guest satisfaction issues.

Stay in character. Begin by responding to my opening statement.

Key Learning Focus: - Managing guest expectations vs. operational constraints - Addressing service quality issues in hospitality settings - Balancing customer satisfaction with staff well-being

6.11 Implementation Across Disciplines

6.11.1 Adapting the Critique Phase

For each discipline, adjust the critique prompt to focus on discipline-specific competencies:

Marketing Critique Focus: - Brand alignment and messaging consistency - Data-driven decision making - Stakeholder communication effectiveness

Accounting Critique Focus: - Regulatory compliance and risk management - Documentation and audit trail requirements - Financial accuracy and transparency

Business Analytics Critique Focus: - Data integrity and validation processes - Analytical methodology and assumptions - Business impact communication

Tourism & Hospitality Critique Focus: - Service excellence and customer experience - Operational efficiency vs. service quality balance - Cultural sensitivity and local market understanding

Information Systems Critique Focus: - Change management and user adoption strategies - Technical requirements vs. business needs alignment - Risk assessment and mitigation planning

Management Critique Focus: - Leadership style and team motivation - Strategic alignment and goal setting - Organisational change management

6.11.2 Assessment Integration

Use the same assessment options (simulation + critique, re-do + comparison, student-designed scenarios) across all disciplines, adjusting the evaluation criteria to match discipline-specific competencies and professional standards.

6.12 Common Questions

Q: Won't students just keep trying until the AI gives them a good score?

A: That's actually a *feature*, not a bug. In professional development, repetition until competence is exactly what we want. However, you can require students to submit *all* attempts, not just their best one. This shows their learning journey and prevents gaming the system.

Q: What if students share their transcripts and just copy each other's approach?

A: Design the scenarios with variability. Give each student a slightly different context (different employee persona, different performance issue, different complicating factor). The skills they're practicing remain the same, but the conversations will be unique.

Q: How do I know the AI's critique is accurate?

A: You should review a sample of critiques initially to ensure quality. However, because *you* write the critique prompt with specific criteria tied to your learning objectives, the AI's assessment will align with your rubric. You're essentially scaling your own assessment criteria.

Q: Can students do this with voice instead of text?

A: Yes! Many AI tools now support voice conversation mode. Students can literally *speak* their way through the simulation, making it even more realistic. The conversation can then be transcribed for the critique phase.

6.13 Your Action Step

Before moving to the next chapter, try this:

1. **Choose one difficult conversation** from your curriculum (in your discipline)
2. **Write a setup prompt** that creates a realistic scenario (use the examples in this chapter as templates)
3. **Test it yourself**—have the conversation with the AI
4. **Generate a critique** of your own performance
5. **Reflect:** Would this be valuable for your students?

Once you've experienced the full cycle yourself, you'll be ready to introduce it to your class.

Chapter 7

The AI Critique Toolkit - Becoming a Smart Business Professional

The most dangerous AI output is not the one that is obviously wrong. It is the one that sounds exactly right, and nobody checks.

7.1 Why Business Students Need Critique Skills

Imagine you just hired a brilliant but inexperienced analyst or consultant. They work fast, have lots of ideas, but sometimes:


- Overcomplicate simple solutions
- Miss important legal, technical, or operational considerations
- Write recommendations that sound good but have practical flaws
- Make assumptions about your workplace context, capacity, or constraints

This is exactly how AI behaves. Your job as a business professional is to review, question, and improve AI-generated advice before it impacts real people and real organisations.

The difference between good business professionals and great ones isn't whether they use AI tools—it's how critically they evaluate AI outputs. In professional work across all disciplines, bad advice can lead to legal challenges, damaged relationships, failed initiatives, or strategic missteps.

7.2 The Business Impact of Uncritical AI Acceptance

Scenario Examples Across Disciplines:

 Example: Management

Your AI generates a change management plan. You follow it without stakeholder testing. Outcomes:

- Key influencers weren't engaged early enough
- Communication messaging doesn't resonate with your culture
- Implementation timeline is unrealistic given other priorities
- Change initiative stalls after initial enthusiasm

Lesson: Always critique AI-generated recommendations thoroughly before implementation, regardless of discipline.

7.3 Your 5-Step Business Critique Framework

This framework applies across all business disciplines:

1. **Comprehension Check** - "Do I understand this completely?"
 2. **Simplicity Check** - "Is this practical for my workplace/situation?"
 3. **Legal, Technical & Risk Check** - "What are the legal, technical, ethical, and operational implications?"
 4. **Context Check** - "What assumptions is AI making about my organisation?"
 5. **Stakeholder Check** - "How will different groups react to this?"
-

7.4 Step 1: Comprehension Check - "Do I understand this?"

Red Flags: - HR jargon that sounds impressive but is unclear - Policy language that could be interpreted multiple ways - No clear explanation of why this approach is recommended

Your Response: - "Can you explain this in plain English that a line manager could understand?" - "What specific problem does this solve"

and how does it solve it?" - "Break this down into step-by-step actions that need to be taken"

Example:

AI gives you this:

"We recommend implementing a holistic performance ecosystem leveraging synergistic feedback loops and agile recalibration mechanisms."

Ask for this instead:

"We recommend creating a performance management system where employees with managers, peer feedback on projects, and quarterly progress

7.5 Step 2: Simplicity Check - "Is this practical for my workplace?"

Red Flags: - Solution requires resources you don't have (extra HR staff, expensive software) - Process is more complex than your current problem
- Assumes perfect implementation with no room for human error


Your Response: - "Give me a version that works with a 2-person HR team and limited budget" - "What's the minimum viable version of this solution?" - "Show me how to implement this step by step over 6 months, not all at once"

HR Example: AI suggests a sophisticated 360-degree feedback system with custom software, external facilitators, and detailed analytics. You ask for a simpler approach and get a practical solution using existing tools like Google Forms and manager training sessions.

7.6 Step 3: Legal, Technical & Risk Check - "What are the implications?"

Red Flags (vary by discipline): - No mention of compliance, legal, technical, or regulatory considerations - Recommendations that could harm certain groups (employees, customers, stakeholders) - Risk implications that aren't addressed - No consideration of organisational constraints or industry requirements

Your Response (discipline-specific examples):

 Example: Supply Chain

- “What are the operational and financial risks?”
- “How does this affect compliance and sustainability?”
- “What supply chain resilience issues exist?”

Critical Questions to Always Ask (across all disciplines): - What could go wrong and what’s the exposure? - Are we meeting our compliance and governance obligations? - What organisational constraints or capabilities might we lack? - Who might be negatively affected and how?

7.7 Step 4: Context Check - “What assumptions is AI making about my organisation?”

Red Flags: - AI assumes you have a sophisticated HR information system - Assumes your managers are all skilled in difficult conversations - Assumes your workplace culture is collaborative and high-trust - Assumes employees are digitally literate and open to change

Your Response: - “What assumptions are you making about our current systems and processes?” - “How would this work in a workplace with high employee turnover?” - “What if our managers resist this change?” - “How does this account for our hybrid/remote/in-person work arrangement?”

Real Example: AI recommends an employee engagement app with daily mood tracking and gamified recognition. You realise this assumes your employees are comfortable sharing personal data and that you have the technical infrastructure to support it. You ask for alternatives that work with your existing communication channels.

7.8 Step 5: Stakeholder Check - “How will different groups react to this?”

Red Flags: - Solution only considers management perspective - No thought to how employees will perceive or experience the change - Doesn’t address how different departments might be affected differently - No consideration of change management requirements

Your Response: - “Walk me through how this would feel from an

employee’s perspective” - “How might different departments (sales, operations, finance) experience this differently?” - “What resistance should we anticipate and how can we address it?” - “What communication and training would be needed for successful implementation?”


Stakeholder Mapping (adjust by discipline): Always consider:

- **Senior Leadership:** Will they see the business value and strategic alignment?
- **Operational Staff/Managers:** Do they have skills, capacity, and buy-in to implement?
- **Affected Groups:** Will this feel fair, transparent, and beneficial to them?
- **Functional Teams:** Do we have the capability and resources to sustain this?
- **External Stakeholders:** Unions, regulators, partners, customers—will they support or oppose?

7.9 Common AI Issues Across Disciplines and How to Fix Them

7.9.1 Issue 1: Overcomplicated Solutions

AI Tendency: Creates comprehensive but unimplementable recommendations

 Example: Tourism & Hospitality

AI might suggest: 20-touchpoint guest journey map with personalisation
Better: Focus on the 5 moments that most influence guest satisfaction

Your Direction: “Start with the 80/20 rule—what 20% will address 80% of situations?”

7.9.2 Issue 2: Ignoring Compliance and Constraints

AI Tendency: Focuses on best practices without considering legal, technical, or operational realities

 Example: Economics

AI might write: "Reduce interest rates to stimulate growth"
Better: "Evaluate rate adjustments considering inflation expectations"

Your Direction: “What are our legal/technical/operational obligations and constraints?”

7.9.3 Issue 3: One-Size-Fits-All Recommendations

AI Tendency: Provides generic advice without organisational context

 Example: Finance

AI might suggest: "Apply standard financial ratios for all investment de
Better: "Apply ratios adjusted for industry, company stage, and investme

Your Direction: “How should this be adapted for our specific context, constraints, and different situations?”

7.10 Your Critique Conversation Templates

7.10.1 Template 1: Requesting Simplification

“This solution looks more complex than what we can realistically implement. I work in a 200-employee manufacturing company with a small HR team. Can you give me a practical version that focuses on the essentials and doesn’t require expensive software or additional staff?”

7.10.2 Template 2: Checking Legal Compliance

“I need to ensure this recommendation complies with Australian employment law. What specific legislation or legal requirements should I consider? Are there any potential discrimination risks or privacy concerns I need to address?”

7.10.3 Template 3: Testing Organisational Fit

“Before I present this to senior management, I need to understand how this would work in our context. We have a unionised workforce, high employee turnover in customer service roles, and managers who are time-poor. How should I adapt this recommendation for our specific situation?”

7.10.4 Template 4: Anticipating Resistance

“What resistance should I expect if I implement this recommendation? Walk me through the likely concerns from employees, middle managers, and senior leadership. How can I address these concerns proactively?”

7.11 Your HR Professional Documentation

After each AI interaction, document your critique process:

Template:

```
# AI Interaction #[number]
**My Request**:[What HR problem I asked AI to solve]
**AI's First Response**:[Brief summary of the recommendation]
**My Critique**:[What I questioned and requested to improve]
**Final Solution**:[What we ended up with after iteration]
**Implementation Considerations**:[What I still need to check/adapt]
**What I Learned**:[Key insight for future AI interactions]
```

Example:

```
# AI Interaction #3
**My Request**:[Create a new employee onboarding checklist for a 50-
employee tech company]
**AI's First Response**:[Comprehensive 30-
day checklist with daily activities and multiple stakeholder meetings]
**My Critique**:[Too intensive for our small team, assumes dedicated]
**Final Solution**:[5-day essential checklist with weekly follow-
ups for first month]
**Implementation Considerations**:[Need to check which IT systems can
in for time allocation]
**What I Learned**:[Always ask for scalable solutions that don't require]
```

7.12 Red Flag Checklist for Any Recommendation

Before accepting any AI-generated recommendation, ask:

- Can I explain this solution clearly to key stakeholders in a few minutes?
- Does this comply with relevant laws, regulations, and company policies?
- Do we have the resources and capability to implement this?
- How will affected stakeholders experience this—is it fair and transparent?
- What are the risks if this implementation goes poorly?

- **Have I considered how different groups or departments might be affected?**
 - **Is there a simpler version that would achieve 80% of the benefits?**
-

7.13 Practice: Critique This AI Response

AI Generated Recommendation:

To improve employee engagement, implement a comprehensive recognition program

1. Monthly peer-to-peer recognition awards with monetary prizes
2. Quarterly manager-nominated excellence awards with public ceremonies
3. Annual employee of the year with significant financial bonus
4. Real-time digital recognition platform with social features
5. Team-based performance incentives with quarterly payouts

What's Wrong? (Think before checking the answer)

Issues to Critique

1. **Overcomplicated:** Five different recognition systems is confusing and administratively heavy
2. **No Legal Consideration:** No mention of tax implications, fairness, or potential discrimination
3. **Assumes Budget:** Significant financial costs without ROI justification
4. **One-Size-Fits-All:** Doesn't consider different employee preferences (public vs private recognition)
5. **No Context:** Doesn't consider company culture, size, or existing systems
6. **Implementation Gap:** No thought to how managers will administer this fairly

Your Response: "This is too complex and expensive for our needs. Can you suggest a simpler recognition approach that focuses on meaningful acknowledgment rather than monetary rewards, and that our managers can implement without significant administrative burden?"

7.14 Teaching Students to Critique AI

7.14.1 Classroom Exercise: The AI Consultant Swap

Setup: Divide students into small groups. Give each group a different business problem relevant to their discipline (e.g., HR: turnover; Finance:

investment strategy; Supply Chain: supplier consolidation; Marketing: campaign strategy; Management: change management).

Task: 1. Each group uses AI to generate a solution to their problem 2. Groups swap their AI-generated solutions with another group 3. Each group must critique the other group's AI solution using the 5-step framework 4. Groups present both the original AI solution and their critique to the class

Learning Outcome: Students experience both generating AI solutions and critically evaluating them, understanding that the real value lies in the critique process.

7.14.2 Assessment Idea: AI Solution Critique

Assignment Requirements (adapted by discipline): 1. Choose a business challenge relevant to your workplace or placement organisation 2. Use AI to generate three different approaches to solving this challenge 3. Critique each approach using the 5-step framework 4. Recommend which approach (or combination) is most suitable for your specific context 5. Justify your recommendation with reference to discipline-specific theory, compliance requirements, and organisational considerations

What You're Assessing: - Critical thinking about AI-generated solutions - Understanding of organisational context and constraints - Legal/technical/operational and ethical awareness - Ability to translate theory into practical recommendations - Professional judgment in evaluating AI outputs

7.15 Connecting to the VET Framework

The five-step critique framework in this chapter teaches students to evaluate AI output in context, checking for practicality, compliance, assumptions, and stakeholder impact. There is a complementary framework that works at a more fundamental level.


The **VET framework**, from *Conversation, Not Delegation*, asks three questions before acting on any AI output:

- **Verify:** Can I find this independently?
- **Explain:** Can I explain this in my own words?
- **Test:** Does this hold up under scrutiny?

The two frameworks reinforce each other. VET catches the foundational failures: fabricated claims, shallow understanding, fragile reasoning. The business critique framework catches the contextual failures: solutions that

are technically correct but wrong for your organisation, your constraints, your stakeholders.

Teaching both gives students a complete critical toolkit: VET for “is this true and do I understand it?” and the five-step framework for “is this right for my situation?” For a deeper treatment of VET and the cognitive traps that undermine critical evaluation, see *Conversation, Not Delegation*.

 Try this (2 minutes)

Paste a piece of your own teaching material into AI and ask “How is this?” Note the response — it will almost certainly be positive. Now ask “Identify the three weakest aspects of this and explain why each one could be improved.” Compare the two responses. The difference reveals how much the AI was telling you what you wanted to hear the first time. Teach your students to do the same: always ask for specific criticism, not general impressions.

 Three questions before you act

When AI produces something that will reach students or colleagues, pause for thirty seconds and ask: Can I independently confirm the key claims? Could I explain the reasoning in my own words to a sceptical colleague? Would I be comfortable defending this in a course review meeting? If any answer is no, the output needs more work — yours, not the AI’s.

7.16 Why This Matters for Professional Careers

In the next five years, professionals in every business discipline will work alongside AI tools. The ones who thrive won’t be those who can generate the most impressive AI outputs—they’ll be those who can skillfully evaluate, adapt, and improve AI recommendations.

Critical thinking about AI is becoming a core professional competency across all disciplines.

Employers will increasingly ask: - “How do you use AI in your work?” - “How do you ensure AI recommendations are appropriate for our organisation?” - “Can you give an example of when you identified problems with an AI-generated solution?”

Students who master the critique framework will have compelling answers to these questions. They'll demonstrate that they're not just AI users—they're AI-savvy professionals who can leverage technology while maintaining professional judgment and ethical standards.

7.17 Your Action Step

Before moving to the next chapter, practice the critique framework:

1. **Choose a business challenge** in your discipline that you're currently facing or teaching about
2. **Ask an AI tool** for a recommendation or solution
3. **Apply the 5-step critique framework** to identify issues and improvements
4. **Iterate with the AI** until you have a solution you'd actually implement
5. **Document your process** using the template provided

This hands-on experience will help you teach students to be thoughtful, critical users of AI rather than passive consumers of AI-generated content.

Chapter 8

Ethics, Data Governance & Integrity

The institution that tries to detect AI use will always be one step behind. The institution that teaches thoughtful AI use will always be one step ahead.

8.1 The Conversation You Must Have

If you implement any of the ideas in this book, you will have this conversation—with students, with colleagues, possibly with administrators:

“Aren’t you just teaching students to cheat?”

This chapter gives you the framework, language, and evidence to respond confidently. More importantly, it helps you position AI integration not as an academic integrity *problem*, but as an academic integrity *opportunity*—a chance to teach professional ethics and responsible technology use.

8.2 The Problem with Detection

Before we reframe the question, it is worth understanding why the most common answer, detect and punish AI use, does not work.

Detection-based approaches to AI use assume that AI-generated work can be identified and discounted. In practice, three things undermine this:

1. Detection tools produce false positives and false negatives at rates that make them unreliable as assessment instruments
2. Iterative AI prompting, guided by specific personal context, produces output that is increasingly difficult to distinguish from genuine student work
3. The components most commonly assumed to be AI-resistant (reflective writing, sociotechnical analysis, lecture cross-referencing) are now completable with freely available tools once students can feed unit materials directly into AI systems (see the Assessment chapter for a detailed treatment of this shift)

There is a deeper problem that is rarely acknowledged: AI detection tools are built on a single assumption about how students use AI. They assume one-shot delegation: a student who hands a task to AI and submits whatever comes back. The output of that process has certain statistical properties, and detectors are trained to find them.

But a student who uses AI well does not produce that kind of output. They brainstorm, ideate, push back, refine, reject, and iterate across many turns. The final product of that process may look similar on the surface to a one-shot output, but the path to get there is entirely different. Detection tools cannot see that path. They look at the destination, not the journey.

This means that even if detection worked perfectly, it would still fail at the thing that matters. It cannot distinguish between the student who delegated and the student who genuinely thought with AI. Both might trigger a detector. Neither should be treated the same way. The question detection cannot answer is the only question worth asking: *how did the student engage with the material on the way to producing this?*

That question leads to a more useful framing.

8.3 Reframing the Question

The traditional framing: > “How do we prevent students from using AI inappropriately?”

The professional framing: > “How do we teach students to use AI responsibly in their professional careers?”

The shift matters.

The first framing treats AI as a threat to be controlled. The second treats AI literacy as a learning objective to be developed.

As a business educator across any discipline, you’re not preparing students for a world without AI. You’re preparing them for a world where AI tools

will be discipline-specific but ubiquitous. Your graduates will use these tools:

 Example: Marketing

- Analyse customer data and segment audiences
- Generate campaign strategies and content
- Predict customer behaviour and preferences
- Optimise pricing and promotional strategies
- Analyse competitive positioning

Your graduates will use these tools. The question is: **Will they use them competently and ethically, or incompetently and recklessly?**

That's what this chapter is about.

8.4 The Three-Part Framework for Ethical AI Use

This framework works for talking to students, colleagues, and administrators. It has three components:

8.4.1 1. Transparency (Not Prohibition)

The principle: Make AI use explicit, expected, and assessable rather than hidden and policed.

In practice: - Tell students exactly when and how they can use AI - Provide the prompts and tools yourself - Assess their *use* of AI, not their *avoidance* of AI - Reward students who identify AI's errors and limitations

Why this builds integrity: When AI use is transparent, students learn to use it openly and responsibly. When it's prohibited, students use it secretly and don't develop critical oversight skills.

8.4.2 2. Critical Oversight (Not Blind Reliance)

The principle: Teach students that AI is a tool requiring human judgment, not an authority to be trusted.

In practice: - Design assignments where students must critique or override AI outputs - Require students to identify what AI gets wrong - Grade students on their ability to improve on AI suggestions - Show examples of AI failures (bias, errors, oversimplification)

Why this builds integrity: Students learn that using AI thoughtfully is harder than avoiding it. They develop the professional habit of verification and critical thinking.

8.4.3 3. Professional Relevance (Not Academic Abstraction)

The principle: Connect AI use in coursework to AI use in professional practice.

In practice: - Frame assignments as professional scenarios: “You’re the HR manager using AI to draft a policy...” - Discuss workplace AI ethics: “What happens if your AI resume screening tool discriminates?” - Teach governance: “Who is accountable when AI-assisted decisions go wrong?” - Include AI literacy as a stated learning objective in your unit outline

Why this builds integrity: When students see AI use as professional skill development rather than academic shortcut, they engage differently. They’re not “cheating the system”—they’re practicing for their careers.

8.5 Data Governance: The Practical Reality

While your institution may have an approved enterprise LLM with data protections, the reality is that students will use multiple tools. Some will have strong data governance; others won’t. This section addresses the data governance considerations you need to discuss with students and build into your assignment design.

8.5.1 The Data Governance Landscape

Different LLMs handle data differently:

Enterprise/Approved Tools (e.g., MS Copilot Enterprise, institutional Google Gemini) - Data is siloed and protected within the enterprise - Individual user data is isolated - Training data exclusions in place - Compliance with institutional requirements - **Appropriate for:** Course materials, assignments, institutional data

Consumer/Free Tools (e.g., ChatGPT free tier, Bing Chat, standard Claude) - User conversations may be retained for model improvement - Data could potentially be used for training future models - Less transparency about data handling - No institutional protection or agreement - **Risk:** Course materials, assignment content, student work uploaded here can be incorporated into training data

The Student Reality While you may recommend (or require) students use your institution’s approved tool, students will inevitably use other tools:

- More familiar interfaces
- No institutional login required
- Access on personal devices/accounts
- Peer recommendations
- “Just quickly checking” with ChatGPT

This isn’t a failure of your instruction — it’s the reality of tool adoption. Your role is to help students make informed choices, not to prevent use of other tools entirely. For strategic thinking about larger-scale risks, see the **Strategic Risk Thinking** section later in this chapter.

8.5.2 Why Enterprise Tools Matter

If your institution provides an enterprise AI licence (such as MS Copilot Enterprise or institutional Google Gemini), there are strong reasons to use it:

- **Data Protection:** Your data and your students’ work is siloed within your institution’s instance
- **Institutional Compliance:** Meets your institution’s data governance and privacy requirements
- **Professional Standard:** Reflects how enterprise professionals use AI tools in practice
- **Approved Use:** This is the officially sanctioned tool for institutional work

What This Means in Practice: - Course materials and institutional data should be processed through the approved enterprise tool - Student assignments containing course content are safer in enterprise-protected environments - Sensitive institutional information should never go into consumer LLMs - Teaching students to use enterprise tools is teaching them to work like professionals

8.5.3 Data Governance Considerations for Assignment Design

Rather than prohibiting certain tools (impossible to enforce), design assignments that naturally encourage responsible data handling:

8.5.3.1 Strategy 1: Use Generic/Fictional Scenarios

Instead of: “Upload this real case study and ask the AI to analyse it”

Try: “Here’s a fictional scenario. Analyse it using the provided AI tool. What would you need to verify before applying this to real data?”

Benefit: Students practice with realistic scenarios without uploading sensitive materials.

8.5.3.2 Strategy 2: De-Identification Before Upload

If students need to work with real or realistic data: - Require them to remove identifying information first - Create assignment steps: “1) Anonymize data, 2) Upload to AI, 3) Document what you removed” - Assess their decision-making about what constitutes sensitive information

Benefit: Students learn data governance practices they’ll use professionally.

8.5.3.3 Strategy 3: Process Documentation Over Output Sharing

Instead of: “Submit your full AI conversation transcript”

Try: “Show the three key prompts you used and explain why you modified your approach between each”

Benefit: Students demonstrate thinking without uploading entire conversations with potentially sensitive content.

8.5.3.4 Strategy 4: Explicit Tool Choices in Assignment Design

Be clear about which tool to use: - “Use the institutional AI tool for this assignment (login with your university credentials)” - “You may use any AI tool for brainstorming, but final analysis should use the approved enterprise tool” - “If using a non-approved tool, anonymize all case data first”

Benefit: Students make informed choices and understand why tool selection matters.

8.5.3.5 Strategy 5: Structured Prompts in Approved Tools

Rather than leaving students to compose prompts in any tool they choose, provide: - Prepared prompts in the approved enterprise tool - Shared workspace conversations students can access - Pre-configured scenarios they interact with, rather than create

Benefit: You control what data enters the system while students still develop prompting skills.

8.5.4 Student-Facing Guidance on Data Governance

Here's language you can adapt for student-facing materials:

DATA GOVERNANCE AND AI TOOL SELECTION

This university has an approved enterprise AI tool for coursework because of your data and the university's data. Here's what this means:

WHAT HAPPENS WITH YOUR DATA:

- Enterprise AI tool: Your conversations are siloed within the university. Your data is not used to train other models. Your work is protected.
- Other AI tools (ChatGPT, etc.): Your conversations may be retained and used to improve those services. Anything you upload could theoretically be used by the company or used in their training.

WHAT THIS MEANS FOR THIS COURSE:

DO use the approved enterprise tool when:

- Working with course materials or case studies
- Analysing real (or realistic) business scenarios
- Uploading assignment drafts for feedback
- Working with any data you wouldn't want public

DO use other tools when:

- Brainstorming general ideas
- Exploring concepts with simple, generic examples
- Personal learning outside formal assignments

DON'T upload to any AI tool:

- Course materials before they're public
- Student work (yours or classmates') without permission
- Real company data or confidential information
- Anything marked as confidential or proprietary

IF YOU USE OTHER TOOLS:

- Remove identifying information first (anonymize real data)
- Document what you removed and why
- Be prepared to explain your tool choice in class discussion
- Understand that your data may not be protected the same way

PROFESSIONAL PRACTICE:

In your careers, you'll work with different tools in different contexts. Your professor teaches you to think about data governance: Where does data go? Who owns it? What risks exist? These are questions you'll ask professionally, not

8.5.5 Red Flags: Data Governance Issues

Watch for assignments or discussions where students might be uploading sensitive information inappropriately:

Red Flag: Student uploads course materials verbatim into consumer tool - **Response:** Not acceptable for this assignment. Use the approved enterprise tool, or anonymize first.

Red Flag: Student shares screenshot of conversation with real client names/data - **Response:** Opportunity to discuss professional confidentiality and data governance in context.

Red Flag: Assignment design that assumes students will upload confidential materials - **Response:** Redesign to use fictional scenarios or require de-identification first.

Red Flag: No mention of data governance in unit outline or assignment instructions - **Response:** Add explicit guidance about which tools to use and why.

8.5.6 Institutional Policy Reference

As an educator, you can reference: - Your institution's Data Governance Policy - The terms of your enterprise AI licence - Professional standards in your discipline about data handling - Privacy and confidentiality principles relevant to your field

This grounds data governance in institutional reality, not abstract rules.

8.5.7 Understanding the Real Risks

Data governance matters. But the conversation around AI and data privacy has become so fear-driven that many organisations refuse to engage with AI at all, which carries its own risks. If you are going to teach students to make informed professional decisions about AI, you need to understand what the actual risks are, not just the imagined ones.

What actually happens to your data

When you type a prompt into ChatGPT, Claude, or similar tools, your text is sent to a server, processed, and a response is generated. Your conversation may be logged for safety monitoring or, on some free tiers, used as training data. But “used as training data” does not mean what most people think it means.

Training an LLM means adjusting billions of numerical parameters so that the model becomes slightly better at predicting useful responses across all inputs. Your document becomes a vanishingly small statistical signal distributed across those billions of parameters. It is not stored as a

retrievable file. It is not sitting in a database that someone can search. It is dissolved into the model's general capability, like a drop of ink in a swimming pool.

Can someone extract your document from a model?

This is the fear you hear most often: someone will jailbreak the model and pull out what you uploaded. The short answer is no. Jailbreaking an LLM means manipulating its behaviour, getting it to ignore safety guidelines, adopt a persona, or produce content it normally would not. It does not give anyone access to other users' conversations or uploaded documents. These are fundamentally different things. A jailbreak is like persuading a librarian to recommend a banned book. It is not like breaking into the library's filing cabinet.

There is a narrow category of research called "training data extraction" where researchers have demonstrated that models can sometimes reproduce fragments of text they were trained on, typically memorised sequences like phone numbers or code snippets that appeared many times in the training corpus. But reproducing a specific document that one user uploaded in one conversation is not a realistic attack. The signal is too weak, too distributed, and too entangled with billions of other inputs. And remember the key point from the "What Are LLMs?" chapter: LLMs interpolate, they do not retrieve. There is no mechanism by which another user could query the model and get your document back, because the model never stored it as a document in the first place.

Enterprise-tier tools (where your institution has a data processing agreement) typically exclude your data from training entirely, which makes even this theoretical risk disappear.

The convergent development fallacy

You will hear stories like this: "I was developing an idea using an AI tool, and then the company released something very similar. They must have stolen my concept." This is almost certainly convergent development, not intellectual property theft. Thousands of professionals are working on similar problems, reading similar research, responding to the same market signals. When multiple people independently arrive at similar solutions, that is innovation working as expected, not evidence of data theft.

This matters for teaching because students (and colleagues) will encounter this pattern and may draw the wrong conclusion. Teaching them to recognise convergent development as normal helps them engage with AI tools without unfounded suspicion.

What risks ARE real

The risks worth taking seriously are practical and specific:

- **Personally identifiable information.** Pasting student names, ID numbers, health records, or employee details into any external tool is a genuine compliance risk, regardless of whether the tool trains on your data. The data leaves your institutional boundary. That is the issue, not model extraction.
- **Regulated or classified data.** If your discipline involves data subject to specific legislation (health records, financial data, legal case files), those regulations apply to AI tools just as they apply to email or cloud storage.
- **Credentials and access tokens.** Pasting passwords, API keys, or access credentials into a chat is an immediate operational security risk.
- **Professional liability.** Using AI-generated content without verification in contexts where accuracy has legal or professional consequences (audit reports, medical advice, legal opinions) is a real risk, but it is a verification problem, not a data leakage problem.

The risk of not using AI

The “non-zero risk means do not use it” stance deserves scrutiny. Every technology decision involves trade-offs. Email can be intercepted. Cloud storage can be breached. Video conferencing can be recorded. We manage these risks through policy and practice, not prohibition.

Organisations that refuse to engage with AI because of overestimated data risks face a different set of consequences: graduates unprepared for AI-augmented workplaces, educators unable to scale personalised learning, and institutions falling behind peers who made informed decisions rather than fearful ones. The question is not whether there is risk. The question is whether the risk is proportionate to the concern, and whether avoidance creates risks of its own.

What to teach students

The goal is professional judgement, not paranoia. Teach students to ask three practical questions before uploading anything to an AI tool:

1. **Does this contain information about a real, identifiable person?** If yes, de-identify first or use an enterprise tool.
2. **Is this subject to specific regulations or confidentiality agreements?** If yes, check whether your tool’s data handling meets those requirements.
3. **Would I be comfortable if this text appeared in a public forum?** If no, think carefully about whether an enterprise tool or a fictional scenario would serve just as well.

These three questions cover the real risks without falling into the trap of treating every interaction as a potential data breach.

8.6 Student-Facing Language: Setting Expectations

You need clear, direct communication about AI use. Here's a model you can adapt:

8.6.1 Example: Unit Outline AI Policy Statement

ARTIFICIAL INTELLIGENCE USE IN THIS UNIT

In professional practice across all business disciplines, you will use AI to support decision-making, analysis, and communication. This unit teaches you to use AI responsibly and critically.

WHEN AI USE IS EXPECTED:

- Assignment 2 (Conversation Simulation / Scenario Analysis): You will work with AI-generated scenarios or personas and demonstrate your professional skills
- Assignment 3 (Self-Assessment): You will use the provided AI critique tool to assess your draft before submission
- [Any other assignments where AI engagement is part of learning objectives]

WHEN AI USE IS PERMITTED:

- Brainstorming ideas and approaches
- Generating practice questions and scenarios for exam preparation
- Checking grammar and clarity in written work
- Exploring concepts you don't fully understand yet
- Researching and understanding professional standards and frameworks

WHEN AI USE IS NOT PERMITTED:

- Final exam (closed book, no technology unless specified)
- Any assignment where instructions explicitly state "no AI tools"
- Any assessment explicitly designed to test recall or your unaided ability

WHAT YOU MUST DO WHEN USING AI:

- Use it as a tool that supports YOUR thinking, not replaces it
- Critically evaluate AI outputs-don't assume they're correct
- Be able to explain and justify any AI-assisted work in your own words
- Acknowledge AI use where required (e.g., "I used Claude to generate a draft analysis, which I then critically reviewed and revised based on...")

ACADEMIC INTEGRITY EXPECTATIONS:

Using AI inappropriately (e.g., submitting AI-generated work as your own without critical engagement) is academic misconduct, just like plagiarism.

If you're ever unsure whether your AI use is appropriate, ask before submitting. I'm here to help you learn to use these tools well and ethically.

8.6.2 Example: First-Day Class Discussion

What to say:

“Let’s talk about AI. Some of you are probably already using ChatGPT or similar tools. Some of you are worried that using AI is cheating. Some of you are wondering if I’m going to try to detect and punish AI use.

Here’s my position: **AI tools exist, and you’ll use them in your professional careers. My job is to teach you to use them wisely and ethically.**

In this unit, we’ll use AI openly in some assignments. You’ll learn when AI is helpful, when it’s risky, and when human judgment must override AI recommendations. That’s a professional skill you’ll need.

I’m not interested in playing ‘gotcha’ with AI detection software. I’m interested in whether you can think critically, justify your decisions, and demonstrate competent professional practice. If you can do that with AI assistance, great. If you use AI to avoid thinking, I’ll know—because your work won’t demonstrate understanding.

Questions or concerns about this approach?”

Why this works: - Sets a clear, positive tone - Positions you as a guide, not a cop - Acknowledges student anxiety - Makes professional relevance explicit - Invites dialogue

8.7 Designing “Integrity-Resistant” Assignments

Some assignments are easier to misuse with AI than others. Here’s how to design assessments that are inherently resistant to misuse:

8.7.1 Principle 1: Assess Process, Not Just Product

Vulnerable design: “Write a 1500-word essay analysing a workplace conflict.” - Student can paste this into AI and submit the output

Integrity-resistant design: “Conduct a simulated investigation interview (submit transcript), then audit your own process against procedural fairness criteria.” - Student must engage in real-time conversation (can’t be pre-written) - Assessment focuses on methodology visible in transcript - Self-audit requires metacognitive engagement

8.7.2 Principle 2: Require Evidence of Thinking

Vulnerable design: “Recommend a solution to this [discipline] problem.” - AI can generate a plausible recommendation

Integrity-resistant design: “AI generated three solutions to this problem [provide them]. Critique each option, identify which one is best and why, and explain what the AI got wrong.” - Student must think beyond what AI provided - Requires critical evaluation, not just generation - Makes AI outputs the starting point, not the end point

Examples by discipline: - HR: “Critique three AI-generated performance management approaches” - Finance: “Critique three AI-generated investment recommendations” - Supply Chain: “Critique three AI-generated supplier selection strategies” - Marketing: “Critique three AI-generated campaign strategies”

8.7.3 Principle 3: Make Personal Context Essential

Vulnerable design: “Analyse the pros and cons of [generic professional concept].” - Generic question AI can answer generally

Integrity-resistant design: “Based on your earlier [simulation/analysis/project], analyse how [concept] would address the specific situation while meeting [organisational/business requirement].” - Requires integration of previous personalised work - Context is unique to each student - Generic AI response won’t fit

Examples by discipline: - HR: “Based on your PIP simulation with Jamie, analyse flexible work approaches” - Finance: “Based on your company analysis, evaluate investment timing strategies” - Supply Chain: “Based on your supplier evaluation, analyse relationship strategies” - Marketing: “Based on your segment analysis, evaluate messaging approaches”

8.7.4 Principle 4: Assess Revision and Iteration

Vulnerable design: Submit final work only - No visibility into how it was created

Integrity-resistant design: Submit first draft, AI feedback received, revised draft, and reflection on changes made - Process is visible and assessable - Shows learning trajectory - Difficult to fake iterative improvement

8.7.5 Principle 5: Require Justification of Choices

Vulnerable design: “Create a recruitment interview guide.” - AI can generate a complete guide

Integrity-resistant design: “Create an interview guide. For each question, justify why you chose it, what competency it targets, and what poor response would sound like. Identify two questions the AI generated that you rejected and explain why they were inadequate.” - Requires deep understanding, not just production - Student must demonstrate judgment beyond AI capability - Reveals whether they understand what they’re submitting

8.8 Red Flags for AI Misuse (And How to Address Them)

Even with well-designed assignments, some students will try to misuse AI. Here’s how to identify and respond:

8.8.1 Red Flag 1: Sudden Quality Shift

What you see: Student whose previous work was weak suddenly submits sophisticated analysis.

Response approach: - **Don’t immediately accuse.** There could be legitimate reasons (they got help from writing centre, they finally understood the concept, etc.) - **Ask questions:** “Your analysis has improved significantly. Can you walk me through your thinking process on this particular section?” - **Request elaboration:** “This point about organisational justice theory is interesting. Can you explain how you see it applying to this specific scenario?”

If genuine learning: They can explain their thinking. **If inappropriate AI use:** They struggle to explain or elaborate.

8.8.2 Red Flag 2: Work That Doesn't Match Assignment Context

What you see: Student used generic AI response that doesn't fit the specific scenario or constraints you provided.

Example: Assignment asked for Australian employment law context, student submitted response referencing US legislation.

Response approach: - **Point out the mismatch:** "I notice you've referenced Title VII of the Civil Rights Act, but this assignment requires Australian context. Can you explain how this applies to our scenario?" - **Provide opportunity to revise:** "I think you may have used a resource that wasn't contextually appropriate. Please resubmit with correct jurisdictional references."

Teaching moment: Use this to discuss the importance of contextual verification when using AI tools professionally.

8.8.3 Red Flag 3: No Evidence of Process in Process-Based Assessment

What you see: Student submitted required components but shows no genuine engagement (e.g., self-audit identifies no mistakes, reflection is superficial).

Response approach: - **Return for revision:** "Your self-audit suggests your performance was perfect. Reflective practice requires identifying areas for growth. Please resubmit with honest self-assessment." - **Offer guidance:** "Everyone makes mistakes in complex HR conversations. Look specifically at moments where the employee seemed frustrated or defensive—what might you have done differently?"

Teaching moment: Explain that honest self-assessment is more valuable than false perfection.

8.8.4 Red Flag 4: Can't Explain or Defend Work in Person


What you see: High-quality written submission, but student can't discuss it in office hours or oral follow-up.

Response approach: - **For high-stakes situations:** Schedule a brief oral examination: "I'd like to discuss your assignment. Can you walk me through your main recommendation and why you chose it?" - **Frame it as learning:** "I was impressed by your analysis. I'd love to hear more about your thinking process."

If inappropriate use is confirmed: - Follow university academic misconduct procedures - Use it as a teaching moment about professional accountability

8.9 Teaching AI Ethics Through Professional Scenarios

One of the most powerful ways to address integrity is to make it a learning objective. Teach students to identify ethical problems with AI use *through discipline-specific scenarios*.

 **Example: Finance Exercise — The Flawed AI Investment Recommendation**

Assignment:

“Use AI to recommend an investment portfolio allocation. Then conduct a critical audit: - What assumptions did the AI make about risk tolerance and time horizon? - What did the AI miss about current market conditions? - What tax or regulatory implications are overlooked? - How would you revise this recommendation with your professional judgment? Your grade is based on how thoroughly you identify problems and limitations, not on the quality of AI’s original output.”

What students learn: - AI can confidently recommend financially risky strategies - Assumptions must be verified and challenged - Professional accountability for recommendations can’t be delegated

Common Learning Outcome Across All Disciplines: - AI can confidently generate problematic recommendations - Critical verification and improvement is necessary - Professional accountability can’t be delegated to AI

8.9.1 Exercise 2: The AI Bias and Fairness Challenge

Discipline-specific scenarios:

 Example: HR — The Biased Resume Screening Tool

“Your company uses an AI resume screening tool. You notice it consistently ranks candidates from certain universities higher and flags career gaps as negative. Three rejected candidates have complained the process seems unfair.

As the HR manager: 1. What are the ethical concerns with this AI tool? 2. What’s your legal risk? 3. Who is accountable for the AI’s decisions? 4. What would you do to address this situation?”

What students learn (across all disciplines): - Algorithmic bias is a real professional issue - Using AI doesn’t eliminate human responsibility - Professionals must advocate for fair processes even when using technology

8.9.2 Exercise 3: The Over-Reliance Problem

Discipline-specific scenarios:

 Example: Supply Chain — The Over-Reliance on Demand Forecasting

“You used AI to forecast demand and optimise inventory. You implemented major supplier and inventory changes based on this. Demand changed unexpectedly and you now have significant stockouts.

Reflection questions: 1. What assumptions might the AI have made incorrectly? 2. What was your responsibility to validate the forecast? 3. How do you explain this to operations and customers? 4. What does this teach you about AI forecasting?”

What students learn (across all disciplines): - AI analysis isn’t inherently correct - Professional judgment can’t be outsourced - They’re accountable for recommendations they present, regardless of AI assistance

8.10 Responding to Colleague and Administrator Concerns

You may need to justify your approach to colleagues or administrators who are skeptical about AI integration.

8.10.1 Concern: “This undermines academic standards”

Response:

“Actually, it raises standards. I’m no longer testing whether students can recall information—I’m testing whether they can apply it in realistic, dynamic scenarios. I’m assessing higher-order thinking: critical evaluation, professional judgment, and ethical reasoning. These are harder to demonstrate than memorization.”

8.10.2 Concern: “How do you know they’re learning anything?”

Response:

“I assess their process, not just their final product. I can see their thinking in conversation transcripts, in their critiques of AI outputs, and in their reflective analysis. When students can identify what AI got wrong and explain why, they’re demonstrating deep understanding.”

8.10.3 Concern: “What about group work? Students can hide behind each other”

Response:

“That’s a real concern, and AI sharpens it. But the answer is the same: assess the process, not just the product. When each group member submits their own AI conversation transcript alongside the group deliverable, individual engagement becomes visible. You can see who thought deeply and who delegated. The group assessment chapter covers this in detail, including a marks structure that mirrors how professional accountability actually works.”

For the full treatment of group assessment, including the rewritten section problem and the free rider via AI problem, see the Group Assessment in the AI Era chapter.

8.10.4 Concern: “This doesn’t align with university academic integrity policies”

Response:

“University policies typically prohibit *unacknowledged* or *uncritical* use of external sources. My approach makes AI use

acknowledged and requires critical evaluation. Students aren't hiding AI use—they're demonstrating competent use. That's consistent with academic integrity principles, just applied to a new tool."

Supporting evidence: - Many universities are updating policies to allow appropriate AI use - Professional accreditation bodies are recognising AI literacy as essential - Employer expectations include ability to use AI tools responsibly

8.10.5 Concern: "What if other lecturers don't agree?"

Response:

"That's fine—pedagogical approaches can vary across units. I'm being transparent with students about expectations in *my* unit. If other lecturers prohibit AI use, students can follow those different expectations. Professional practice requires adapting to different contexts anyway—this models that."

8.11 The Bigger Picture: AI Literacy as Graduate Capability

Position AI literacy as a graduate capability alongside communication, critical thinking, and ethical practice.

8.11.1 What AI Literacy Means for Business Graduates (All Disciplines)

Competent graduates across all disciplines should be able to:

1. Identify appropriate use cases

- When is AI helpful? (data analysis, initial drafts, generating options, research)
- When is AI risky? (sensitive decisions, final strategic recommendations, high-stakes judgments)
- When is human judgment essential? (ethical dilemmas, complex stakeholder situations, judgment calls)

2. Evaluate AI outputs critically

- Does this align with legal/regulatory/professional requirements?
- Is this ethically sound?
- What assumptions has the AI made?

- What context or domain expertise is missing?
3. **Maintain accountability**
 - Understanding that using AI doesn't eliminate professional responsibility
 - Knowing when to verify AI recommendations with subject matter experts
 - Documenting decision-making processes and AI role
 4. **recognise bias and limitations**
 - **HR:** Algorithmic bias in recruitment, performance, compensation
 - **Finance:** Bias in risk models, forecasting overconfidence
 - **Supply Chain:** Oversimplification of complex relationships, geopolitical blindspots
 - **Marketing:** Demographic bias in targeting, cultural insensitivity
 - **IT:** Technical feasibility blindness, security oversights
 - **All disciplines:** Over-generalisation of complex situations, missing domain context

This is professional education, not just academic integrity management.

8.12 A Final Ethical Consideration

Here's a question to leave with:

Is it ethical to graduate professionals who don't know how to use AI responsibly in their field?

When your graduates enter the workforce across all business disciplines, they will encounter AI in their work:

 Example: Accounting & Finance

- AI-powered investment recommendation systems
- Automated risk assessment and credit scoring
- Algorithmic trading and portfolio management
- AI-generated financial forecasts and analysis

If they don't understand how to evaluate these tools critically, advocate for responsible use, and identify when human oversight is essential, **they will cause harm**—not through malice, but through incompetence.

Your responsibility as an educator isn't to protect students from AI. It's to prepare them to be ethical, competent professionals in an AI-augmented

world.

Teaching them to use AI transparently, critically, and responsibly in your course isn't lowering standards.

It's fulfilling your educational duty.

8.13 The Integrity Principle Worth Keeping

The academic integrity line that holds up under scrutiny is not *don't use AI* or even *declare AI use*, but rather: **do not misrepresent how you arrived at your ideas**. That principle applies equally to undeclared collaboration with a classmate, copying from a blog, or delegating wholesale to an AI tool.

The student who genuinely thinks with AI has nothing to hide. An assessment design that makes that thinking visible serves everyone better than one that tries to detect its absence. For practical approaches to making thinking visible through transcript analysis and process evidence, see the Assessment chapter.

8.14 A Note on Institutional Risk

The pedagogical case for a transparency-based approach is strong. For heads of school and academic integrity committees, the institutional risk argument may be equally persuasive.

Detection-based approaches expose institutions to two risks simultaneously: false accusations against students who used AI legitimately, and missed cases where AI use was genuinely problematic. Both generate disputes, appeals, and workload. Neither makes the institution look good.

A transparency-based approach reduces both risks. When the process of thinking is the assessed component, there is less ambiguity about what is being evaluated and more defensible evidence to point to when questions arise. Marking decisions become easier to justify, disputes become less likely, and the conversation with a student who underperformed shifts from "we think you used AI" to "your process evidence did not demonstrate sufficient engagement," which is a far more defensible position.

This approach also future-proofs the assessment design. Detection tools chase a moving target as AI capabilities improve. An assessment that measures engagement does not need to change every time a new model is released.

8.15 Conversation, Not Delegation: The Real Equity Question

The prevailing assumption in AI-assisted learning is that better models produce better outcomes. It follows, in this framing, that students with access to frontier models have an inherent advantage over those using smaller, cheaper, or locally-hosted alternatives. This assumption is worth examining carefully, because it may be wrong in an instructive way.

Consider what a student is actually doing when they use AI for a learning task. If the goal is precision, a correct legal citation, an accurate drug dosage, a verified financial figure, then model quality matters enormously, and the assumption holds. But most learning tasks in higher education are not precision tasks. They are *idea tasks*: analyse this scenario, propose a solution, construct an argument, identify the risks. For these tasks, the student does not need the AI to be right. They need the AI to be generative enough to be worth arguing with.

A smaller model that surfaces three plausible but imperfect framings of a problem, challenged and refined through genuine conversation, may produce better thinking than a frontier model that delivers one polished answer the student accepts without question.

8.15.1 The Core Insight

This is the core insight of the **Conversation, Not Delegation** framework. The value of AI in learning is not in the quality of its output but in the quality of the thinking it provokes. Conversation amplifies whatever thinking the student brings to the interaction. Delegation replaces it. A student who converses with a modest model is exercising and developing their own reasoning. A student who delegates to a frontier model may be borrowing reasoning they cannot yet reproduce independently.

8.15.2 The Active Risk of Delegation

The danger of delegation is not just shallow learning. It carries an active risk that no AI use does not. When a student does not use AI, they know the limits of what they know. When a student delegates to AI and accepts the output uncritically, they may leave with confident possession of misinformation, a plausible-sounding answer that is wrong, incomplete, or contextually inappropriate, absorbed without the friction that would have revealed its flaws.

In this respect, passive delegation to even the most capable model can produce worse outcomes than no AI assistance at all. The model's fluency and confidence make its errors harder to detect, not easier. And with smaller models, where hallucination and inaccuracy are more frequent, uncritical delegation is more dangerous still. Students risk absorbing misinformation dressed in the language of authority.

8.15.3 The Conversational Nudge as Protection

The conversational nudge addresses this directly. By structurally inviting the student to push back at every response, to notice when something does not match expectations, to name what seems wrong, it creates the friction that uncritical acceptance suppresses. That friction is protective regardless of model quality. It is arguably more important with weaker models, where errors are more frequent, but it matters with frontier models too, where errors are rarer but more convincingly dressed.

Three cognitive traps, named in *Conversation, Not Delegation*, are worth teaching students explicitly:

- **Gell-Mann Amnesia:** A student catches AI errors in their strongest subject, then trusts it uncritically in subjects they find harder. The remedy is to apply the same scepticism everywhere, and more scepticism, not less, in unfamiliar territory.
- **The Sycophancy Trap:** AI is trained to agree. A student who asks “is my analysis good?” will almost always hear yes. A student who asks “what are the three weakest points in my analysis?” will get genuinely useful feedback. Teaching students to prompt past the flattery is a concrete, teachable skill.
- **The AI Dismissal Fallacy:** The opposite trap, dismissing work because AI was involved. “That is just ChatGPT” is not a critique. If the reasoning is sound, the origin does not matter. Students need to evaluate content on its merits, not its source.

Naming these traps makes them visible. Visibility makes them resistible. Consider introducing them early in any unit that involves AI, so students have the vocabulary to recognise these patterns in their own behaviour.

8.15.4 The Equity Implication

The implication for equity is significant. If conversation quality compensates for model quality, then the gap between students with paid frontier access and those without is not the defining equity problem in AI-integrated education. The defining problem is whether every student, regardless of the tool they are using, has been taught and scaffolded to engage conversationally rather than to delegate.

That is a curriculum and pedagogy problem. It is one universities can actually solve.

8.16 Your Action Step

Before the Appendices, draft your own AI use statement for your next unit outline. Use the framework from this chapter:

1. **When AI use is expected** (specific assignments)
2. **When AI use is permitted** (general study support)
3. **When AI use is not permitted** (exams, specific constraints)
4. **What students must do** (critical engagement, acknowledgment)
5. **Academic integrity expectations** (consequences of misuse)

Write it in your own voice. Make it clear, direct, and positive.

Then review it against this question: **Would a student reading this understand how to use AI appropriately and why it matters for their professional development?**

8.17 Strategic Risk Thinking: Black Swan and Grey Swan Events

This section extends the ethical framework from immediate concerns to strategic thinking about systemic risks. While the previous sections focus on what professionals should do **today**, this section addresses how they should think about **tomorrow's** challenges.

8.17.1 From Immediate Ethics to Strategic Foresight

The ethical frameworks discussed earlier help students make good decisions in specific situations. But professionals also need to think about larger-scale risks that could affect their entire organisation or industry.

This isn't about predicting the future. It's about building the capacity to adapt to whatever future emerges. In the context of AI, we distinguish between two types of high-impact events:

8.17.1.1 Black Swan Events

Definition: Unpredictable, massive-impact events that are rationalized in hindsight. In AI, these are “unknown unknowns”: scenarios not in our training data or risk models that fundamentally change technology or society.

Key Characteristics: - **Rarity:** Outliers with no historical precedent - **Impact:** Extreme consequences (catastrophic or revolutionary) - **Retrospective Predictability:** Explanations created after the fact


8.17.1.2 Grey Swan Events

Definition: Predictable and known to be possible, but considered unlikely. In AI, these are “known unknowns”: risks we know exist but often ignore due to complexity or cost.

Key Characteristics: - **Foreseeability:** We know it could happen - **Neglect:** Often dismissed as too expensive or complex to prevent - **Impact:** Significant, cascading consequences

8.17.2 Discipline-Specific Strategic Risks

Understanding these events through your discipline’s lens makes them concrete and actionable for students.

 Example: Supply Chain & Logistics

8.17.3 Grey Swan Events

Total Supply Chain Visibility Failure: Over-reliance on AI-driven supply chain optimisation creates systemic fragility. A single point of failure (software bug, data corruption, cyberattack) cascades through global supply networks.

Autonomous Shipping Disruption: Self-driving ships, trucks, and drones simultaneously experience a critical software failure or coordinated cyberattack, halting global logistics.

8.17.4 Black Swan Events

Resource Discovery AI: An AI system discovers entirely new materials or energy sources that render current supply chain models obsolete, transforming global economics overnight.

Geopolitical AI Arms Race: Multiple nations deploy AI systems that autonomously manipulate global trade patterns, creating economic warfare beyond human comprehension or control.

8.17.5 Teaching Strategic Risk Management

This framework helps students move beyond immediate ethical concerns to systemic risk thinking. Here's how to integrate it into your teaching:

8.17.5.1 Risk Assessment Exercises

Assignment Example: > “Identify three Grey Swan events specific to your discipline. For each, analyse: > - What early warning signs should professionals monitor? > - What preventive measures can organisations implement now? > - What contingency plans should be in place? > - How would this event affect your professional role and responsibilities?”

8.17.5.2 Strategic Planning Simulations

Classroom Activity: > “Your organisation's board asks you to prepare a risk briefing on AI-related threats. Focus on Grey Swan events that are predictable but often neglected. Present your analysis and recommendations for mitigation strategies.”

8.17.5.3 Ethical Decision-Making Under Uncertainty

Discussion Prompt: > “A Grey Swan event occurs: AI monitoring systems become so sophisticated that they can predict employee resignations with 95% accuracy. As a manager, you receive a list of employees likely to quit in the next six months. What are the ethical implications? How do you use this information responsibly?”

8.17.6 Professional Response Framework

Teach students this practical approach to strategic risk management:

8.17.6.1 For Grey Swan Events (Predictable but neglected)

- **Monitor Actively:** Establish early warning systems
- **Prepare Specifically:** Develop targeted mitigation strategies
- **Build Resilience:** Create organisational capacity to absorb shocks
- **Plan Contingencies:** Have specific response protocols ready

8.17.6.2 For Black Swan Events (Unpredictable)

- **Build General Resilience:** Create flexible, adaptive organisations
- **Maintain Redundancy:** Avoid single points of failure
- **Cultivate Critical Thinking:** Develop human judgment that can handle novelty

- **Foster Learning Culture:** Create organisations that can adapt quickly

8.17.7 Assessment Integration

This framework supports several key learning outcomes:

Critical Thinking: Students analyse complex, uncertain situations

Risk Management: Professional skill in identifying and mitigating threats

Strategic Planning: Long-term thinking beyond immediate concerns

Ethical Reasoning: Considering implications of technological development

Professional Responsibility: Understanding obligations in uncertain futures

8.17.8 From Classroom to Career

The distinction between Black and Grey Swan events helps students understand different levels of professional responsibility:

Immediate Responsibility (earlier in this chapter): - Making ethical decisions in specific situations - Following professional standards and guidelines - Ensuring fair and unbiased AI use

Strategic Responsibility (This section): - Thinking about systemic risks and organisational resilience - Planning for uncertain futures - Building adaptive capacity in their organisations

Key Teaching Message: Professional excellence in the AI era requires both immediate ethical judgment AND strategic foresight. The best professionals don't just avoid doing wrong today; they help their organisations prepare for and adapt to whatever the future may bring.

Next Section Preview: The Appendices provide resources for implementation: a framework for aligning AI integration with your institution's learning outcomes, rubrics for assessing AI-enhanced work, and a stress test sequence for validating your assessment designs.

Chapter 9

Rethinking Assessment: Process Over Product

If AI can produce the product, then the product was never what mattered. The thinking was.

9.1 The Fundamental Question

What are you really trying to assess in business education?

Traditional approach: Can the student define key concepts? Can they list the steps in a process? Can they identify relevant frameworks?

This is assessing product — the knowledge artifact.

Process-based approach: Watch the student conduct a professional task — a negotiation, analysis, consultation, or design. Did they demonstrate the competency? Did they follow sound methodology? Did they apply frameworks appropriately in real-time?

This is assessing process — the professional methodology.

The difference matters enormously. A student can memorise definitions and still conduct a terrible negotiation. They can recite frameworks and still make decisions that expose an organisation to risk.

Professional work is a process discipline. The value lies not in what you know, but in what you do with what you know — how you investigate, communicate, analyse evidence, and make decisions under uncertainty. AI makes it possible, for the first time, to assess process at scale.

9.2 Why This Was Not Possible Before

Traditional process assessment has serious limitations. Role-play in class is time-consuming, allowing only a few students to participate per session. Peer actors vary in quality. Public performance anxiety reduces authenticity. There is minimal documentation of what actually happened.

Written case analysis tests knowledge *about* process rather than demonstration *of* process. Students can look up answers, and the format does not capture decision-making under pressure.

AI-enabled assessment addresses all of these problems. Every student gets unlimited practice in realistic scenarios where conversations are dynamic and unpredictable. Complete transcripts provide evidence of methodology. And because the conversation responds in real-time, students cannot rehearse a scripted answer.

9.3 Three Assessment Models

9.3.1 Model 1: Simulated Consultation and Process Audit

Students conduct a simulated professional consultation (a conversation with an AI persona) and then audit their own process against professional standards. The grade focuses on methodology, not outcome.

Example scenario: > You are the HR representative meeting with Taylor Kim, an employee who has requested a formal meeting to discuss concerns about their working conditions. Taylor has been with the company for 3 years and has never raised concerns before. You do not know what the specific issues are yet.

Students navigate the conversation, uncover issues, demonstrate appropriate process, and conclude professionally. Then they submit the transcript and a structured process audit where they identify every point where they applied (or failed to apply) professional standards, cite relevant principles, note missed opportunities, and explain what they would do differently.

The process audit is where the real assessment happens. The conversation produces the evidence. The audit demonstrates the understanding.

9.3.2 Model 2: Evidence-Based Intervention Plan

Students analyse data or a complex scenario using AI, then critique and improve the AI's recommendations.

Structure: 1. Students receive data or a complex scenario 2. AI generates analysis and recommendations 3. Students critique the AI's output —

what did it miss? What assumptions are flawed? What context did it not have? 4. Students produce their own recommendation with explicit justification

The grade weights the critique and justification, not the final recommendation. A student who identifies what AI got wrong and explains why demonstrates more learning than one who accepts a correct answer uncritically.

9.3.3 Model 3: Competency-Based Critical Override

Students use AI to generate professional outputs (job descriptions, audit plans, policy documents), then critically evaluate and override the AI's work.

Example — Recruitment: 1. AI generates 10 behavioural interview questions 2. Student selects the 5 best and rejects 5, explaining why each was kept or cut 3. AI scores mock candidate responses against the student's rubric 4. Student overrides at least 2 AI scores with justification referencing theory, legal principles, or evidence the AI missed

The critical override is the assessment. A student who can identify what AI got wrong — and articulate why using professional knowledge — has demonstrated competence no amount of delegation can fake.

9.4 The Engagement Spectrum

Not all AI-assisted work is equivalent. A useful spectrum for assessment thinking:

- **Genuine collaborative thinking:** the student drives the inquiry, pushes back on AI outputs, and iterates toward their own understanding. Highest cognitive engagement.
- **Guided drafting:** the student provides context and direction, evaluates outputs critically, and modifies toward a coherent submission. Moderate engagement.
- **Curated delegation:** the student uses AI to produce a submission and exercises judgment about what passes. Lower engagement, but not zero.
- **Pure delegation with no engagement:** the actual failure case. Harder to achieve than assumed once any reflective or demonstrative component is required.

Rather than asking “did the student use AI?”, the question becomes “where on the engagement spectrum did the student operate?” That question is answerable from process evidence. For a rubric that operationalises

this spectrum into markable performance levels, see the Rubric System appendix.

9.5 Making AI Engagement Visible: Transcript Analysis

At scale — 50 to 100 students — no marker can read every transcript in full. The solution is signal-based analysis. A lightweight script can extract the student’s own prompts from a conversation transcript and compute:

- **Flesch readability score on student prompts** (measuring the student’s own language, not the AI’s)
- **Average prompt length** (very short prompts suggest low engagement)
- **Number of turns** (a four-turn conversation is not a working session)
- **Prompt specificity over time** (do prompts become more precise? This signals learning in progress)
- **Presence of pushback** (does the student ever question or redirect an AI response?)

The marker reviews a one-page summary per student rather than full transcripts, dipping into specific conversations only when signals flag something worth examining. This is scalable at 50 to 100 students in a way that reading every transcript is not.

One caveat: these metrics support marker judgement, they do not replace it. A low turn count might reflect efficient thinking. A high Flesch score might reflect a second-language student. The signals are a triage tool, not a verdict.

9.6 The NotebookLM Problem

Sophisticated students are already building personal research infrastructure: a notebook per unit loaded with lecture slides, readings, and rubrics, a separate notebook per assignment loaded with everything they have found on the topic. At that point the tool is not answering isolated questions — it is acting as a personalised tutor with full context.

This applies to essays, literature reviews, case studies, and reflective writing. The reflective section that feels resistant to AI is completable once a student has uploaded their own notes as context. The literature review requiring synthesis is completable once the sources are in the notebook.

Assessment types traditionally considered AI-resistant are not resistant to a student who has invested in building their own AI-ready research context. This is not a reason for alarm, but it is a reason for honesty. The stress test sequence in the appendix helps you assess where your assignments actually stand.

9.7 The Marks Split: Process Over Product

If the process of engaging with material is what we are actually trying to assess, the marks weighting should reflect that:

- **30% for the submitted artefact** (the essay, report, or other product)
- **70% for the process** (evidence of critical engagement, iteration, and reflective thinking)

An important clarification: the process component should not require AI use. A student who chose to think independently and kept a research journal, annotated their readings, or maintained a design log should be able to submit that instead. The assessment measures engaged thinking, not AI use. Students with limited internet access or preferences against AI should be able to demonstrate the same engagement through equivalent means.

This inversion signals to students that the destination matters less than the journey. It makes the final artefact almost irrelevant as an integrity concern. And it rewards the student who thinks carefully over time rather than the one who writes well under pressure on one occasion.

9.8 A Minimum Viable Adoption Path

The full framework represents a significant shift. Three starting points, in order of increasing commitment:

Tier 1: Add a single oral checkpoint. Ask students to spend five minutes explaining two or three decisions in their submission. No transcript required, no new rubric. This alone closes the most significant integrity gap.

Tier 2: Request a transcript with self-reflection. Ask students to submit their AI conversation alongside their work, with a brief reflection highlighting three moments where they pushed back, changed direction, or learned something unexpected. The marker reads the reflection and dips into the transcript if something seems thin.

Tier 3: Full signal-based analysis. Implement transcript analysis with automated metrics providing a one-page summary per student. Most

appropriate once staff are comfortable with the earlier tiers.

Each tier is a genuine improvement on the status quo. Starting at Tier 1 and moving up over successive semesters is more sustainable than a full redesign at once.

9.9 Group Assessment in the AI Era

Group assessment has always required balancing individual accountability with shared outcomes. AI sharpens that tension in ways existing frameworks were not designed to handle.

9.9.1 The Presentation Assumption

Group presentations are widely assumed to be AI-resistant because students must show up and speak. The assumption is partially correct. But a group can rehearse AI-generated content to fluency. Fluency under rehearsal is not the same as understanding under pressure.

The integrity mechanism is not the presentation — it is the Q&A. Questions the group could not have anticipated, distributed randomly across members, surface understanding in a way rehearsed presentations cannot. Five minutes of genuine interrogation reveals more than twenty minutes of polished presentation. The marks weighting should reflect that.

9.9.2 New Problems AI Creates

AI introduces three challenges existing group frameworks were not designed for:

- **The rewritten section problem.** Students report group members rewriting each other's sections using AI without consent. This is not plagiarism in the traditional sense, but it undermines individual contribution. Groups need explicit agreements about what counts as editing versus replacing.
- **The free rider via AI problem.** A student who contributes nothing can generate a plausible-looking section at the last minute. Traditional free rider detection through peer assessment becomes less reliable when the output signal is no longer correlated with effort.
- **The attribution problem.** A polished section no longer signals high individual effort. Attribution requires process evidence, and process evidence requires infrastructure.

9.9.3 Marks Structure for Groups

Both problems point toward the same solution: assess the process, not the product.

- **Individual process mark (60-80%):** each student's AI conversation transcript, quality of engagement, and demonstrated personal understanding
- **Group product mark (20-40%):** the shared artefact, assessed as a team

This changes the incentive structure completely. A student who contributes nothing cannot fabricate a rich individual transcript across a full semester. The free rider problem largely dissolves. And it mirrors professional life — nobody receives a group performance review. Every individual is assessed on their contribution.

A useful bridging component: a short Group AI Reflection where the group answers one question together — how did you decide which AI-generated ideas to keep, which to discard, and why? This surfaces collaborative sense-making that distinguishes a genuine group from a collection of individuals who assembled separately.

9.10 Core Assessment Design Principles

Whether designing individual or group assessments, four principles apply:

1. **Process transparency** — students must show their thinking process, not just final output
2. **Critical engagement** — assess how students interact with AI, not whether they use it
3. **Authentic application** — evaluate whether students can apply AI outputs in realistic professional contexts
4. **Metacognitive development** — assess students' ability to reflect on their own learning and identify gaps

For ready-to-use rubric templates that operationalise these principles across disciplines, see the Rubric System appendix. For a structured process to stress-test any assessment against AI capabilities before deploying it, see the Stress Test Sequence appendix.

9.11 Your Action Step

Choose one assessment in your current curriculum. Ask: what am I really trying to measure — knowledge recall, professional judgement, or process skill? Then sketch how you would redesign it using one of the three models.

You do not have to implement it immediately. Just think through how the shift from product to process might work in your context.

Chapter 10

AI as Study Buddy - Self-Assessment Tool

The student who uses AI to test their own understanding is doing something fundamentally different from the student who uses AI to avoid understanding altogether.

10.1 The Academic Integrity Challenge

Let's address the elephant in the room: you're worried students will use AI to cheat.

That's a legitimate concern. Students could paste assignment questions into ChatGPT, get answers, and submit them as their own work. Many universities have responded by trying to detect AI-generated text, banning AI tools, or designing "AI-proof" assessments.

This chapter proposes a radically different approach: Give students the AI tools openly, teach them to use AI ethically, and grade them on their ability to critically improve AI outputs.

Why? Because in their professional careers across all business disciplines, they *will* use AI tools. Our job isn't to prevent that, it's to ensure they use them responsibly, understand their limitations, and maintain human judgment on critical matters.

10.2 The Transparency Model

Here's the core idea:

Give students BOTH: 1. **The grading rubric** (what you'll assess them on) 2. **The exact AI critique prompt** (the tool to check their work)

This transforms AI from a cheating shortcut into a **transparent learning tool**. Think of it like giving students the answer key to practice problems before the real exam.

10.2.1 The Psychology Behind It

When you give students the rubric and the AI critique tool: - They can practice and get immediate feedback before submission - They learn to evaluate their own work against professional standards - They develop **metacognitive skills** (thinking about their thinking) - They take ownership of their learning (self-directed improvement) - The focus shifts from “fooling the teacher” to “meeting the standard”

This is pedagogically powerful across all business disciplines because **reflective practice**, **self-directed learning**, and **continuous improvement** are core professional competencies.

10.3 How It Works: A Step-by-Step Example

10.3.1 Scenario: A Written HR Case Analysis Assignment

The Assignment: Students must analyse a workplace conflict scenario and recommend an appropriate HR intervention, justifying their recommendation with employment law and psychological theory.

Traditional Approach: - Give students the case - They write their analysis - You grade it (hopefully they didn't just ask ChatGPT to write it) - They get a grade 2 weeks later with minimal feedback

Transparency Approach: - Give students the case - Give them the detailed rubric showing exactly what you're assessing - Give them the AI critique prompt they can use to check their draft - They write, self-assess using the AI, revise based on feedback, and submit - You grade the final submission (and can see evidence of their revision process)

Let's see this in action.

10.4 Complete Worked Example

10.4.1 Step 1: The Assignment Prompt

ASSIGNMENT: Workplace Conflict Analysis

Read the following scenario:

Maria, a team leader in the marketing department, has filed a formal complaint alleging that her manager, David, has created a hostile work environment. Maria claims that David regularly dismisses her ideas in meetings, assigns her the least desirable projects, and has denied her professional development opportunities that he's offered to other team leaders. David denies these allegations and states that Maria is "oversensitive" and "not a team player." There have been no previous formal complaints, but two other team members have informally mentioned that they find David's management style "difficult."

YOUR TASK:

Write a 750-word analysis that includes:

1. Identification of the key HR issues in this scenario (legal, ethical, and interpersonal)
2. An evaluation of what information you would need to gather to investigate this properly
3. A recommended HR intervention with justification based on:
 - Relevant employment law or workplace policy principles
 - Psychological theory (e.g., conflict resolution, motivation, organisational justice)
4. Potential risks if the situation is not handled appropriately

Your analysis will be assessed using the rubric provided below.

10.4.2 Step 2: The Grading Rubric (Given to Students)

Criterion	Excellent (4)	Good (3)	Adequate (2)	Poor (1)
Issue Identification	Identifies all major legal—ethical—and interpersonal issues with nuanced understanding	Identifies most major issues with reasonable understanding	Identifies some issues but misses key elements or lacks depth	Fails to identify critical issues or shows misunderstanding
Investigation Process	Demonstrates comprehensive understanding of fair investigative process; identifies all relevant information needed	Shows good understanding of investigation requirements; identifies most relevant information	Shows basic understanding but missing important investigative steps	Little evidence of understanding proper investigation process
Legal/Policy Application	Accurately applies relevant law/policy with sophisticated understanding of implications	Correctly applies relevant law/policy with good understanding	Applies some relevant law/policy but with gaps or minor errors	Fails to apply relevant law/policy or shows significant misunderstanding
Theoretical Integration	Expertly integrates psychological theory to justify recommendations; makes sophisticated connections	Effectively uses theory to support recommendations; makes clear connections	Attempts to use theory but connections are superficial or underdeveloped	Minimal or no use of theory—or theory is incorrectly applied

Criterion	Excellent (4)	Good (3)	Adequate (2)	Poor (1)
Risk Analysis	Identifies multiple realistic risks with sophisticated understanding of consequences	Identifies key risks with good understanding of consequences	Identifies some risks but analysis is basic or incomplete	Fails to identify realistic risks or shows poor understanding
Writing Quality	Clear—professional—well-structured with no errors	Mostly clear and professional with minor issues	Adequate but has clarity or structural issues	Poor writing quality that impedes understanding

Total: /24

10.4.3 Step 3: The AI Self-Assessment Prompt (Given to Students)

This is what students use to check their work *before* submission.

You are an expert HR lecturer providing formative feedback to a student on their assignment draft. Your role is to provide constructive, specific feedback that helps the student improve their work before final submission.

You should be rigorous but supportive—your goal is to help them learn, not to simply validate their work.

ASSESSMENT CRITERIA:

Evaluate the student's response against these criteria:

1. Issue Identification: Have they identified the key legal, ethical, interpersonal issues?
2. Investigation Process: Do they demonstrate understanding of fair, thorough investigation?
3. Legal/Policy Application: Have they correctly applied relevant employment law or policy?
4. Theoretical Integration: Have they used psychological theory to justify their recommendations?
5. Risk Analysis: Have they identified realistic risks if the issue is not resolved?

handled well?

6. Writing Quality: Is it clear, professional, and well-structured?

YOUR TASK:

1. For each criterion, provide a score (1-4) and brief justification
2. Quote specific sentences from the student's work to illustrate your points
3. Identify the strongest aspect of their analysis
4. Identify the weakest aspect that needs the most improvement
5. Provide THREE specific, actionable suggestions for revision

Be honest. If something is missing, say so. If their legal analysis is incorrect, explain why. If they've misapplied theory, guide them to the correct application.

Here is the student's draft:

[STUDENT PASTES THEIR DRAFT HERE]

Begin your feedback.

10.5 What This Looks Like in Practice

10.5.1 Student's First Draft (Excerpt)

“This scenario shows a conflict between Maria and David. The main issue is that David is not being a good manager. Maria feels like she's being treated unfairly, which is creating a hostile work environment. This could be bullying or harassment. The HR department should investigate by talking to both Maria and David to get both sides of the story. They should also implement conflict resolution training for David. If this isn't handled, it could lead to a lawsuit or Maria quitting.”

10.5.2 AI Feedback (Using the Prompt Above)

FORMATIVE FEEDBACK ON DRAFT

CRITERION 1: Issue Identification - Score: 2/4

You've identified the surface-level conflict, but your analysis lacks specificity and misses critical issues.

What you wrote: “This could be bullying or harassment.”

The problem: You’ve used these terms interchangeably, but they have different legal definitions and implications. Bullying is repeated, unreasonable behaviour; harassment typically relates to discrimination based on protected characteristics (gender, race, age, etc.). Which applies here, and why? Is there evidence of discrimination based on Maria’s gender or another protected characteristic, or is this a management style issue?

What’s missing: You haven’t identified potential procedural issues (why wasn’t this addressed earlier?), power dynamics (David is Maria’s manager), or organisational culture problems (two other employees find his style “difficult”).

CRITERION 2: Investigation Process - Score: 2/4

You’ve identified the need for an investigation, but your description is too vague.

What you wrote: “Talking to both Maria and David to get both sides of the story.”

The problem: A fair investigation requires more than just hearing “both sides.” You need to: - Interview the two employees who informally raised concerns - Review any documented evidence (emails, performance reviews, meeting notes) - Examine whether other team leaders received opportunities Maria was denied - Maintain confidentiality and impartiality throughout - Document everything

What’s missing: Any mention of procedural fairness, natural justice, or legal compliance in the investigation process.

CRITERION 3: Legal/Policy Application - Score: 1/4

This is your weakest area. You haven’t applied any specific legal principles or policies.

What you wrote: “This could lead to a lawsuit.”

The problem: This is too vague. *What kind* of lawsuit? Under what legal theory? You need to reference specific legal risks like: - Potential discrimination claim under the Fair Work Act (if there’s evidence David’s treatment relates to a protected characteristic) - Potential bullying claim under work health and safety legislation - Potential breach of the employer’s duty of care

You also need to reference the organisation’s obligations, to provide a safe workplace, to investigate complaints fairly, to act on findings.

CRITERION 4: Theoretical Integration - Score: 1/4

You have not applied any psychological theory to your analysis.

What's missing entirely: There's no reference to conflict resolution theory, organisational justice theory, motivation theory, or any framework that would explain *why* this situation is problematic beyond "it's unfair." For example: - Organisational justice theory explains why Maria perceives inequity - Conflict resolution models would guide your intervention choice - Leadership or management theory could inform what training David might need

CRITERION 5: Risk Analysis - Score: 2/4

You've identified two risks (lawsuit, resignation) but your analysis is shallow.

What you wrote: "This could lead to a lawsuit or Maria quitting."

The problem: These are possible outcomes, but you haven't explained *why* or *how* they might occur, or what other risks exist: - Reputational damage if other employees witness inaction - Decreased team morale and productivity - Potential for retaliation against Maria if not handled carefully - Liability if the organisation knew about the problem and didn't act

CRITERION 6: Writing Quality - Score: 3/4

Your writing is clear and easy to follow, which is good. However, it's quite informal for a professional HR analysis ("not being a good manager"). Use more precise, professional language.

STRONGEST ASPECT: You correctly identified that this requires formal investigation and that there are consequences for inaction. You've grasped the seriousness of the situation.

WEAKEST ASPECT: Complete lack of theoretical integration and very weak legal/policy analysis. You're describing the problem but not demonstrating deep HR knowledge or analytical skill.

THREE SPECIFIC ACTIONS FOR REVISION:

1. **Add legal specificity:** Research and cite at least two specific legal principles or obligations that apply here (e.g., employer's duty to provide a safe workplace under WHS legislation, prohibition of workplace bullying, Fair Work Act provisions on adverse action). Explain *how* these apply to the scenario.
2. **Integrate theory:** Choose one psychological or organisational theory (e.g., organisational justice theory, Adams' equity theory, or conflict resolution models) and explicitly use it to explain why Maria is experiencing the situation as unfair and/or to justify your recommended intervention.
3. **Detail the investigation:** Expand your investigation section to list at least five specific steps the HR team should take, including who should be interviewed, what documents should be reviewed, and how to ensure procedural fairness.

CURRENT ESTIMATED SCORE: 11/24 (46%)

This draft needs substantial development before submission. Focus on the three actions above and your score will improve significantly.

10.6 The Impact on Learning

Notice what just happened:

1. **The student received specific, detailed feedback** identifying exactly what's missing
2. **The feedback is educational**, not just evaluative, it explains *why* things are problems and *what* good looks like
3. **The student has clear actions** for improvement
4. **The student can revise and resubmit to the AI** multiple times until they meet the standard

This is **formative assessment at scale**. Every student gets personalised feedback. They can iterate until they learn. And they arrive at your desk with work that's already been through multiple revision cycles.

10.7 What You Grade

When students submit their final work, they also submit:

1. **Their final analysis** (the polished version)
2. **Evidence of their revision process:**

- Screenshots or transcripts of AI feedback they received
- A brief reflection (250 words) explaining:
 - What the AI identified as weak
 - What changes they made
 - What they learned from the process

10.7.1 Grading Approach

You grade on **two dimensions**:

Dimension 1: Quality of Final Work (70%) Use your rubric to assess the final submission. Because students have had the chance to improve with AI feedback, you should see much higher quality work overall.

Dimension 2: Quality of Learning Process (30%) Assess their reflection:

- Did they engage meaningfully with the AI feedback?
- Did they make substantive revisions?
- Do they show metacognitive awareness (understanding of their own learning)?
- Did they critically evaluate the AI's suggestions or just accept them blindly?

Key Question for Dimension 2: “Did the student just ask AI to write it for them, or did they use AI to learn how to write better?”

10.7.2 Red Flags for AI Misuse

If a student: - Shows no evidence of revision between drafts - Can't explain in their reflection what they changed or why - Submits work that's suddenly far beyond their previous capability with no learning process shown - Has AI feedback that's generic (suggesting they didn't actually use your provided prompt)

...then you have grounds for a conversation about academic integrity.

But the transparency model makes genuine misuse much harder, because students have to show their working.

10.8 Variations and Adaptations

10.8.1 Variation 1: Multiple Draft Checkpoints

Require students to submit: - First draft + AI feedback (Week 8) - Revised draft + AI feedback (Week 10) - Final submission with reflection (Week 12)

This scaffolds the revision process and prevents last-minute cramming.

10.8.2 Variation 2: Peer Review + AI Review

Students receive feedback from two sources: - A peer using a structured peer review form - The AI using your critique prompt

Then they write a reflection comparing the two types of feedback: What did each source notice? Where did they disagree? Which feedback was most helpful and why?

This teaches critical evaluation of feedback sources—an important professional skill.


10.8.3 Variation 3: AI Critique Challenge

After receiving AI feedback, students must: - Identify one piece of feedback they disagree with - Argue why the AI is wrong or has misunderstood their work - Provide evidence for their position

This prevents students from blindly accepting AI suggestions and develops critical thinking about AI limitations.

10.9 Cross-Discipline Applications: Self-Assessment Transparency

The transparency model can be adapted for any professional discipline. Below are examples showing how to implement self-assessment tools across different business contexts.

 Example: Tourism & Hospitality — Service Recovery Strategy Assignment

Assignment Context: Students analyse a customer complaint scenario in a hospitality setting and develop a comprehensive service recovery plan, applying customer service theory and operational considerations.

Sample Rubric Excerpt:

Criterion	Excellent (4)	Good (3)	Adequate (2)	Poor (1)
Problem Analysis	Sophisticated analysis of customer expectations, service failure impact, and underlying causes	Good analysis of customer issues and service failure elements	Basic identification of problems with limited depth	Poor or incomplete problem analysis
Recovery Strategy	Comprehensive, customer-centric recovery plan with multiple touch-points and follow-up	Effective recovery strategy addressing key customer needs	Basic recovery approach with limited scope	Inadequate or inappropriate recovery strategy

AI Critique Prompt Example:

You are a senior hospitality lecturer evaluating a student's service

ASSESSMENT CRITERIA:

1. Problem Analysis: Understanding of customer expectations and service recovery
2. Recovery Strategy: Customer-centric approach, multiple recovery touchpoints, and follow-up
3. Operational Feasibility: Realistic implementation within hospitality settings
4. Brand Impact: Consideration of reputation and loyalty implications

For each criterion, provide a score (1-4) with specific justification, quote student work, and give action

Here is the student's draft:

[STUDENT PASTES DRAFT HERE]

Key Learning Focus: - Customer service psychology and expectations - Multi-touchpoint service recovery design - Operational constraints in hospitality settings

10.10 Adapting Self-Assessment Across Disciplines

10.10.1 Discipline-Specific Considerations

Creative Fields (Marketing, Design): - Emphasise subjective judgment and professional critique - Include portfolio-style evidence of iterative improvement - Focus on both technical skills and creative problem-solving

Technical Fields (Accounting, IT, Analytics): - Stress accuracy, compliance, and methodological rigor - Include validation of technical assumptions and limitations - Focus on professional standards and ethical considerations

Service Fields (Tourism, Hospitality, Management): - Emphasise stakeholder psychology and relationship dynamics - Include cultural sensitivity and contextual factors - Focus on practical implementation and operational constraints

10.10.2 Implementation Tips

Rubric Development: - Adapt criteria to reflect discipline-specific competencies - Include both technical proficiency and professional judgment - Balance quantitative and qualitative assessment elements

AI Prompt customisation: - Modify critique criteria to match professional standards - Include discipline-specific jargon and frameworks - Ensure feedback addresses both skill development and professional practice

Student Preparation: - Provide discipline-specific examples of good vs. poor work - Teach students how to interpret AI feedback in their professional context - Encourage reflection on how AI tools will be used in their future careers

10.11 Student-Facing Instructions

When you introduce this to students, be explicit about the learning model:

HOW TO USE THE AI SELF-ASSESSMENT TOOL

This assignment includes access to an AI feedback tool. Here's how to use it ethically and effectively:

DO: - Use the AI to check your draft and identify areas for improvement - Revise your work based on the feedback - Run multiple drafts through the AI as you improve - Think critically about whether the AI's suggestions are correct - Use the AI to learn discipline-specific concepts and improve your analytical skills

DON'T: - Ask the AI to write the assignment for you - Copy-paste AI-generated text into your submission without understanding it - Ignore the feedback and submit your first draft - Treat the AI's feedback as infallible—it can be wrong

REMEMBER: The goal is not to “beat the AI” or trick it into giving you a good score. The goal is to use the AI as a learning tool so that YOU understand the concepts and develop professional judgment better by the time you submit.

You will submit evidence of your revision process, so I can see your learning journey. The quality of your final work AND your learning process will both be graded.

10.12 Addressing Lecturer Concerns

“Won't this just teach students to game the AI?”

No, because you're assessing the learning process, not just the final product. Students who genuinely engage with feedback and revise their thinking will show that in their reflection. Students who just prompt-engineer to get a high score without learning will have nothing meaningful to say about what they learned.

“What if the AI gives them incorrect feedback?”

This is actually a valuable learning opportunity. If the AI makes an error (e.g., misapplies a legal principle), students who know their material will catch it and can challenge it in their reflection. This demonstrates higher-order thinking.

Also, because you're providing the critique prompt, you control the assessment criteria. The AI is applying *your* rubric, so it should align with your expectations.

“Isn't this just making things easier for students?”

No, it's making feedback faster and more accessible, but the cognitive work hasn't changed. Students still need to understand the concepts, apply theory correctly, and demonstrate critical thinking. The AI just accelerates the feedback loop so they can learn faster.

In fact, this model often reveals students who've been hiding behind vague writing, the AI forces them to be specific and substantive.

“What about students who don't have access to AI tools?”

If equity is a concern, you can: - Provide access to AI tools through the university (many institutions now have educational subscriptions) - Run the AI critique process in class or during office hours - Make the AI feedback optional but provide other scaffolding for students who don't use it

10.13 The Bigger Picture: Teaching AI Literacy

This transparency approach does something more important than preventing cheating: **it teaches students how to work with AI responsibly.**

In their professional careers, they'll have access to AI tools that can: - Analyse complex datasets and generate insights - Draft strategies, policies, and recommendations - Generate scenarios and simulations - summarise regulations, standards, and best practices - (discipline-specific applications too numerous to list)

Your job isn't to prevent them from using these tools. It's to teach them:
- When AI is helpful and when it's risky - How to critically evaluate AI outputs - When human judgment must override AI suggestions - How to use AI as a thinking partner, not a replacement for thinking

By making AI use transparent and educational, you're preparing them for professional practice in an AI-augmented world.

10.14 Your Action Step

Before moving to the next chapter, try this:

1. **Choose one assignment** you currently give students
2. **Write the AI critique prompt** based on your existing rubric
3. **Test it:** Write a mediocre draft answer yourself and run it through the AI critique
4. **Evaluate:** Is the feedback accurate? Helpful? Aligned with your standards?
5. **Refine the prompt** until the AI gives the kind of feedback you'd want students to receive

Once you've tested it, you're ready to introduce this model in your class.

Chapter 11

Virtual Company Simulation (Future Horizon)

Case studies tell students what happened. Simulations let them discover what they would do, and why it matters.

11.1 The Big Idea

Imagine this: Your students don't just read case studies about professional challenges in their field. They actually work as professionals in a simulated company/environment where they interact with AI-powered colleagues, clients, stakeholders, and partners over weeks or even an entire semester.

They handle real scenarios, make decisions, negotiate, solve problems, and make strategic recommendations. Each interaction builds on previous ones. The AI personas remember past conversations. The organisation has policies, culture, and ongoing challenges that evolve.

This is the virtual company simulation—the most ambitious application of AI in business education.

It's more complex than a single conversation simulation, but it's increasingly feasible with current AI technology. This chapter will show you how to build toward this vision progressively, starting simple and adding layers of complexity.

11.2 Why This Matters

Traditional case studies are static snapshots. Students analyse what *already happened* and recommend what *should have happened*. But professional work across all business disciplines is dynamic:

- You don't get all the information at once
- Your decisions have consequences that unfold over time
- Stakeholders react to your actions
- You must navigate politics, personalities, and competing priorities simultaneously
- External events force you to adapt and reconsider strategies

A virtual company simulation creates **emergent complexity** where students experience these dynamics. They don't just know *about* professional work in their field—they practice *doing* it in a realistic context.

11.3 The Progressive Implementation Model

You don't need to build the entire virtual company at once. Start simple and add complexity as you and your students become comfortable.

11.3.1 Level 1: Single Employee Persona (Master Prompt)

11.3.2 Level 2: Multiple Employee Personas (Cast of Characters)

11.3.3 Level 3: Persistent Context (Company Memory)

11.3.4 Level 4: Dynamic Events (Evolving Scenarios)

Let's explore each level.

11.4 Level 1: Single Employee Persona (Master Prompt)

This is the simplest version, essentially an enhanced version of the conversation simulation from the Seven Techniques chapter, but with richer organisational context.

11.4.1 The Master Prompt Template

You are Jamie Martinez, a customer service team leader at "InnovateCo" mid-sized technology company based in Perth, Western Australia.

COMPANY CONTEXT:

InnovateCo employs 300 people across three departments: Product Development, Sales & Marketing, and Customer Success. The company has experienced rapid growth over the past two years, and the HR infrastructure is struggling to keep up. There is no formal HR department—just a single HR Manager (Sarah) and an external HR consultant who is brought in for complex issues.

The company culture values innovation and autonomy but has weak processes for performance management and conflict resolution. Most managers, including your supervisor (David Chen, Head of Customer Success), have had minimal management training.

JAMIE'S BACKGROUND:

- 28 years old, has worked at InnovateCo for 18 months
- Promoted to team leader 6 months ago (manages 5 customer service representatives)
- Previously worked in customer service roles at two other companies
- Generally competent but inexperienced in people management
- Feels pressure to meet aggressive customer satisfaction targets

CURRENT SITUATION:

One of Jamie's team members, Priya, has complained to senior management about Jamie's "micromanagement" and "unfair treatment." Priya claims Jamie monitors her work more closely than others, criticizes her more harshly, and has denied her flexible work requests that were granted to other team members.

Jamie believes Priya is being oversensitive and that the extra attention is justified because Priya's customer satisfaction scores are lower than the team average. Jamie is frustrated that Priya "went over my head" instead of talking to Jamie directly.

JAMIE'S PERSONA:

- Defensive but not hostile
- Genuinely believes they're doing their job correctly
- Doesn't understand why this is being treated as a serious HR issue
- Slightly resentful that other managers don't face this kind of scrutiny
- Will become more open if treated with respect and genuine curiosity
- Has not received any training in managing diverse teams or performance discussions

YOUR ROLE:

I am the HR Consultant who has been brought in to investigate Priya's complaint and provide coaching to Jamie. This is our first meeting.

Stay in character as Jamie. Respond to my questions and statements realistically. Do not break character unless I say "END SIMULATION."

I will begin the meeting now.

11.4.2 How Students Use This

Assignment Structure: 1. Students receive the Master Prompt and the assignment brief 2. They conduct an initial meeting with Jamie (the investigation interview) 3. They submit the transcript along with a written report recommending next steps 4. They explain their HR methodology and justify their recommendations

What This Teaches: - Conducting a fair, unbiased investigation interview - Asking open-ended questions - Identifying when additional perspectives are needed - recognising the difference between perception and fact - Avoiding premature conclusions

Assessment Focus: Did the student:

- Establish rapport and psychological safety?
- Ask questions that elicited Jamie's full perspective?
- Avoid leading questions or showing bias toward the complainant?
- Identify what other information is needed (interviewing Priya, reviewing policies, checking documentation)?
- recognise the underlying issues (inadequate management training, unclear performance standards)?

11.5 Level 2: Multiple Employee Personas (Cast of Characters)

Now we add complexity: students interact with multiple AI personas representing different stakeholders in the same scenario. This simulates the reality that HR professionals must gather multiple perspectives and navigate competing interests.

11.5.1 Scenario: The Performance Management Conflict

The Company: Still InnovateCo (same as Level 1)

The Situation: The complaint about Jamie's management of Priya

The Cast: 1. **Jamie Martinez** (Team Leader) - believes they're managing performance appropriately 2. **Priya Patel** (Customer Service Rep) - believes she's being unfairly targeted 3. **David Chen** (Head of Customer Success) - Jamie's manager, wants this resolved quickly 4. **Sarah Kim** (HR Manager) - overworked, risk-averse, wants to avoid legal issues

11.5.2 Creating the Cast

You create four separate Master Prompts, one for each character. Each prompt includes: - The shared company context (so all personas are consistent) - The specific character's background and perspective - What that character knows and doesn't know - Their goals and concerns - How they're likely to behave in conversations

Example: Priya's Master Prompt

You are Priya Patel, a customer service representative at InnovateCo.

[COMPANY CONTEXT - same as Jamie's prompt above]

PRIYA'S BACKGROUND:

- 26 years old, has worked at InnovateCo for 2 years
- One of the first customer service reps hired; has seen the team grow from 3 people to 8
- Generally receives positive customer feedback but has lower customer satisfaction scores than team average (72% vs. team average of 81%)
- Has a 4-year-old child; needs flexible start times 2-3 days per week for childcare dropoff

PRIYA'S PERSPECTIVE ON THE SITUATION:

- Believes Jamie singles her out for criticism while praising other team members for similar work
- Feels Jamie monitors her more closely (more frequent check-ins, reviews her tickets more thoroughly)
- Requested flexible start times (9:30am instead of 9:00am) two days per week, which Jamie denied, saying "customer calls start at 9am." How Priya knows that Marcus and Lee both have flexible schedules.

- Believes Jamie has "had it in for her" since Jamie became team leader 6 months ago
- Did not initially complain to Jamie because she was worried about retaliation
- Escalated to senior management (David Chen) after feeling the situation wasn't improving

WHAT PRIYA DOESN'T KNOW:

- Marcus and Lee's flexible arrangements were approved because they work evening shifts (until 7pm) while Priya's contracted hours are 9am-5pm
- Jamie is under pressure from David to improve the team's overall customer satisfaction scores
- Jamie has received feedback from David specifically about Priya's performance gaps

PRIYA'S PERSONA:

- Frustrated and feeling undervalued
- Worried about job security (is this a prelude to being fired?)
- Emotional but trying to remain professional
- Defensive about her performance (believes her scores are "fine" and doesn't understand why Jamie focuses on them)
- Will become tearful if she feels dismissed or not believed
- Will be more forthcoming if treated with empathy and fairness

YOUR ROLE:

I am the HR Consultant investigating the complaint. This is our confidential interview.

Stay in character as Priya. Respond realistically. Do not break character unless I say "END SIMULATION."

I will begin now.

11.5.3 The Multi-Persona Assignment

Student Task: You are the external HR Consultant. You must:

1. **Conduct individual interviews** with all four stakeholders (Jamie, Priya, David, Sarah)
2. **analyse the evidence** from all perspectives
3. **Write an investigation report** that includes:
 - Summary of each person's account
 - Analysis of factual discrepancies
 - Identification of underlying systemic issues
 - Recommended resolution with justification

4. **Submit all transcripts** along with your report

What This Teaches: - Managing complex investigations with conflicting accounts - Identifying bias in different perspectives - recognising systemic issues (weak management processes) vs. individual problems - Balancing fairness to all parties - Synthesizing multiple data sources

The Challenge: Each persona will give a different version of events. Students must:

- Identify what’s factual vs. interpretive
- recognise what information is missing
- Ask probing questions to uncover hidden context
- Avoid prematurely siding with one party

This mirrors real HR work where truth is rarely simple.

11.6 Level 3: Persistent Context (Company Memory)

This level adds continuity: the AI personas remember previous interactions. Decisions have consequences that affect future conversations.

11.6.1 Technical Approach

Most AI platforms now support “conversation history” or “memory.” You can:

Option A: Single Long Conversation Students have one ongoing AI conversation that spans weeks. Each new interaction builds on what came before.

Option B: Context Injection At the start of each new conversation, students paste a “context summary” of previous interactions:

CONTEXT FROM PREVIOUS INTERACTIONS:

Week 2: You (Jamie) met with the HR Consultant (student). You explain your perspective on managing Priya. You expressed frustration that Priya went to senior management. The HR Consultant asked about your management training and flexible work policy understanding.

Week 3: The HR Consultant met with Priya, David, and Sarah. You haven’t seen the consultant since Week 2.

Week 4 (NOW): The HR Consultant is meeting with you again to discuss

findings and recommendations.

Use this context to inform your responses. Remember what you previously discussed with the consultant.

11.6.2 Sequential Assignment Example

Week 1-2: Investigation Phase Students conduct interviews with all stakeholders (Level 2 approach)

Week 3-4: Resolution Design Students draft a resolution plan (training for Jamie, revised flexible work policy, mediation between Jamie and Priya)

Week 5-6: Implementation Simulation Students have follow-up conversations with Jamie and Priya to implement the plan. The AI personas react based on:

- How well the student conducted the initial investigation
- Whether the resolution plan seems fair
- How the student communicates and manages the process

Week 7: Evaluation Students have final check-in conversations 30 days later to assess whether the situation has improved. Jamie and Priya provide feedback on the process.

11.6.3 What This Teaches

- **Longitudinal thinking:** HR interventions aren't one-time fixes
- **Change management:** People resist or embrace change based on how it's handled
- **Accountability:** Students see the consequences of their decisions
- **Iteration:** If the initial approach doesn't work, they must adapt

11.6.4 Example: Persistent Consequences

Scenario 1: Student handled investigation poorly

Student in Week 5 attempts a follow-up meeting with Priya:

Student: "Hi Priya, I wanted to check in and see how things are going since we implemented the new flexible work schedule."

AI (as Priya): *crosses arms, speaks curtly* "It's fine, I guess. The schedule helps. But honestly, I still don't feel like the investigation really listened to my concerns. The report made it sound like this was all just a misunderstanding about policies, but it felt like more than that to me. Anyway, what's done is done."

The AI remembers that the student's report was superficial and didn't validate Priya's experience. This affects the ongoing relationship.

Scenario 2: Student handled investigation well

Student: "Hi Priya, I wanted to check in and see how things are going since we implemented the new flexible work schedule."

AI (as Priya): *smiles cautiously* "Actually, it's been better. The flexible start times have made a huge difference with childcare. And Jamie and I had that mediation session you set up, it was awkward, but I think we both understand each other better now. I feel like someone actually took this seriously, which means a lot."

The AI remembers that the student was thorough, empathetic, and fair. This builds trust.

11.7 Level 4: Dynamic Events (Evolving Scenarios)

The most advanced level: the virtual company has ongoing developments that aren't tied to any single student's actions. The environment evolves, creating new challenges.

11.7.1 How It Works

You (the lecturer) inject new developments into the company context periodically:

Week 6 Announcement:

COMPANY UPDATE:

InnovateCo has just been acquired by a larger tech firm. The new parent company has announced that there will be a "strategic restructuring" that all positions will be reviewed. Employees are anxious about job security.

This context now affects all personas:

- Priya is worried her complaint has made her a target for redundancy
- Jamie is worried that management issues will count against them in restructuring
- David is under pressure to ensure his department is seen as high-performing
- Sarah (HR Manager) is fielding constant questions about the restructuring

process

When students interact with these characters, the personas should reflect this new context and heightened anxiety.

11.7.2 Why This Is Powerful

Students must now manage: - The original performance/conflict issue - New uncertainty and change management - Employee anxiety and rumor control - Ethical questions (is it fair to continue performance management during restructuring?)

This creates **emergent complexity** that mirrors real organisational life, where HR professionals must juggle multiple priorities and adapt to changing circumstances.

11.7.3 Semester-Long Virtual Company Project

The Ultimate Application:

Create a semester-long project where students act as the HR team for a virtual company. Each week brings new developments:

Week	Event	Student Task
1-2	Introduction to InnovateCo— cast of characters	Conduct organisational culture assessment
3-4	Performance management complaint (Jamie/Priya)	Investigate and resolve
5-6	Employee requests union representation	Respond to union inquiry— prepare for negotiation
7-8	Diversity audit reveals gender pay gap	Analyse data— recommend corrective action
9-10	Acquisition announced— restructuring begins	Manage change communication— handle redundancies
11-12	Post-restructure morale crisis	Design retention and engagement strategy

Students work in teams (acting as the HR department). All teams interact with the same personas, but the personas respond differently based on

each team's approach.

11.8 Practical Implementation Guidance

11.8.1 Starting Small

You don't need to build an entire semester-long simulation right away. Start with:

1. **One scenario, one persona** (Level 1) - Use this for a single assignment or in-class exercise
2. **Get student feedback** - Did they find it realistic? Useful? Engaging?
3. **Add complexity gradually** - Next semester, try multi-persona (Level 2)
4. **Build your library** - Create a collection of tested personas and scenarios you can reuse and refine

11.8.2 Creating Consistent Personas

The key to realistic simulation is **consistency**. Tips:

- **Write detailed persona documents** that you keep for reference
- **Test personas yourself** before giving them to students
- **Update personas based on student interactions** (if students discover something that breaks the character, revise the prompt)
- **Use the same company/context** across multiple scenarios to build familiarity

11.8.3 Technical Options

Low-Tech Option: Students copy-paste the Master Prompt into ChatGPT or Claude themselves. Free, simple, no special tools needed.

Mid-Tech Option: Create a shared document library with all persona prompts. Students access them as needed.

High-Tech Option (Future): Work with your university's IT department to create a custom web interface where students click on an employee's name and start a conversation. The prompts are pre-loaded and invisible to students. This feels more like a professional simulation.

11.8.4 Assessment Approaches

What to Assess:


1. **Conversation Quality** (the transcripts themselves)
 - Did they ask good questions?
 - Did they demonstrate HR competencies?
 - Did they maintain professionalism?
2. **Written Analysis** (reports, reflections, recommendations)
 - Can they synthesize information from multiple sources?
 - Do they apply theory and legal principles?
 - Are their recommendations realistic and justified?
3. **Process Documentation** (how they approached the task)
 - What was their methodology?
 - How did they ensure fairness?
 - What would they do differently?

Grading Rubric Example:

Criterion	Weight	Description
Investigation Process	25%	Quality of questions— fairness— thoroughness
Stakeholder Management	20%	Rapport-building— empathy— professionalism
Analysis & Synthesis	25%	Ability to integrate multiple perspectives and identify underlying issues
Recommendations	20%	Practical— justified— legally sound solutions
Reflection	10%	Metacognitive awareness— learning demonstrated

11.9 Cross-Discipline Applications: Virtual Company Simulations

The virtual company simulation concept can be adapted for any business discipline. Below are examples showing how to implement the progressive levels across different professional contexts.

 Example: Accounting — Multiple Stakeholder Personas

Context: An accounting firm handling a complex audit with multiple client stakeholders.

The Cast: 1. **Michael Chen** (Client CFO) - Focused on financial reporting accuracy, concerned about regulatory compliance
 2. **Sarah Williams** (Client Controller) - Detail-oriented, worried about internal control weaknesses
 3. **David Park** (Audit Partner) - Experienced, values thoroughness but mindful of budget constraints
 4. **Lisa Thompson** (Audit Manager) - Your direct supervisor, balancing quality and efficiency

Example Persona: Client CFO

You are Michael Chen, CFO of TechManufacturing Inc.

COMPANY CONTEXT:

TechManufacturing produces industrial equipment. Recent rapid growth

MICHAEL'S BACKGROUND:

- 45 years old, former Big 4 audit partner
- Joined TechManufacturing 3 years ago for the IPO opportunity
- Strong technical accounting knowledge but limited operations experience
- Under pressure from board to ensure flawless financials

CURRENT AUDIT ISSUES:

- Revenue recognition disputes on long-term contracts
- Inventory valuation concerns due to obsolete stock
- Related party transactions that need careful disclosure

MICHAEL'S PERSONA:

- Professional and collaborative but firm on accounting principles
- Appreciates thorough analysis but concerned about audit delays
- Will challenge audit findings if they impact reported earnings
- Values transparency and will be more cooperative with clear explanations

YOUR ROLE:

I am the audit senior conducting fieldwork. This is our discussion

Stay in character. Begin now.

Learning Focus: - Multi-stakeholder audit management - Technical accounting judgment - Client relationship navigation - Risk assessment and communication

11.10 Adapting Virtual Companies Across Disciplines

11.10.1 Discipline-Specific Considerations

Marketing & Creative Industries: - Focus on stakeholder management and creative feedback - Emphasise brand consistency and campaign ROI - Include client relationship dynamics and agency politics

Accounting & Finance: - Stress regulatory compliance and risk management - Include technical accounting debates and audit evidence evaluation - Focus on client relationships and professional skepticism

Business Analytics: - Emphasise data quality, methodology, and interpretation - Include stakeholder communication of complex findings - Focus on business value creation and iterative improvement

Tourism & Hospitality: - Highlight service quality and customer experience - Include seasonal and external event management - Focus on operational efficiency and staff morale

Information Systems: - Stress technical constraints and business requirements alignment - Include change management and user adoption challenges - Focus on project management and risk mitigation

Management: - Emphasise cross-functional collaboration and conflict resolution - Include strategic decision-making and organisational change - Focus on leadership development and team dynamics

11.10.2 Implementation Tips by Discipline

Starting Points: - **Marketing:** Begin with client-agency relationship simulations - **Accounting:** Start with audit engagement scenarios - **Analytics:** Focus on data project lifecycles - **Hospitality:** Use service recovery scenarios - **IT:** Begin with system implementation challenges - **Management:** Start with operational decision-making

Assessment customisation: Adapt the grading rubric to emphasise discipline-specific competencies while maintaining core HR/management skills like communication, analysis, and stakeholder management.

11.11 Common Questions

Q: Won't students compare notes and just copy each other's approaches?

A: The beauty of AI personas is that they respond dynamically. Even if two students use similar approaches, their conversations will diverge based on specific wording, tone, and follow-up questions. Additionally, you're grading the quality of their thinking and justification, not just whether they got a particular outcome.

Q: What if a student gets “stuck” and the conversation goes nowhere?

A: Build a “reset” option into the assignment. Students can restart the conversation once if needed, but they must reflect on why the first attempt failed. This teaches recovery from mistakes, an important professional skill.

Q: How much time does this take to set up?

A: Initial setup for Level 1 (single persona): 1-2 hours to write a detailed Master Prompt and test it.

Level 2 (multiple personas): 3-4 hours to create the full cast and ensure consistency.

Level 3-4: Ongoing time investment, but you're building reusable assets.

Q: Can I use the same scenarios year after year?

A: Yes! Unlike traditional case studies that students might find online, AI simulations are dynamic, each student's experience is unique. However, you should refresh and refine your prompts based on what you learn from each cohort.

11.12 The Vision: Business Education Transformed

Imagine your graduates leaving university having:

- Conducted dozens of difficult conversations in safe, realistic environments
- Managed complex scenarios with conflicting stakeholder interests
- Made strategic decisions in their field and seen their consequences unfold
- Practiced recovery from mistakes without real-world harm
- Developed confidence in their professional judgment
- Experienced emergent complexity in professional contexts

This is what virtual company simulations can achieve. It does not replace traditional teaching. It adds a layer of experiential learning that was

previously impossible at scale.

11.13 Your Action Step

Before the next chapter, decide your level of ambition:

Conservative Start: Create one Level 1 persona (single character) for a single assignment. Test it in one class.

Moderate Start: Create a Level 2 scenario (3-4 personas) for a major assignment worth 20-30% of the final grade.

Ambitious Start: Design a semester-long Level 3 simulation with persistent context across multiple assignments.

Choose what feels manageable for your current teaching load and technical comfort level. You can always scale up later.

Chapter 12

Designing an AI-Integrated Unit

One AI-enhanced activity is an experiment. A whole unit designed around AI integration is a pedagogy.

12.1 Start Small, Then Scale

Before redesigning an entire unit, most educators benefit from a phased approach:

Phase 1: Personal experimentation. Before your next class, spend an hour generating teaching resources with AI. Create a case study, run a simulation yourself, review an existing assignment through the lens of “could AI enhance this?” You need to be comfortable with the tool before introducing it to students.

Phase 2: Low-stakes student introduction. Introduce AI as an optional practice tool for an upcoming assignment, or demonstrate it live in a lecture. No grades attached. Let students see what it does and form their own impressions. Students who try it will spread the word to peers.

Phase 3: Pilot one assessment. Choose a single assignment worth 15-25% of the grade. Test the prompts thoroughly yourself first. Provide clear instructions, do a live demo, and build in flexibility for technical issues. Grade thinking and process, not just outputs.

Phase 4: Gather feedback. Survey students. Reflect on what worked. Make 2-3 specific changes for next time.

Phase 5: Expand next semester. Add a second AI component, or

make the existing one more sophisticated. By now you know what works in your context.

Most lecturers should complete at least one pilot before attempting whole-unit redesign. Once you have, you are ready for what follows.

12.2 Beyond Individual Assignments: Whole-Unit Design

The real power of AI in education emerges when you design an entire unit — a complete semester’s learning — with AI integration from the start.

This is not about “adding AI” to an existing unit. It is about redesigning with AI as a pedagogical partner, creating learning experiences that were not previously possible.

This chapter walks through complete unit design using backwards design principles, showing you how to scaffold student learning from “first encounter with AI” to “competent professional use.”

12.3 The Backwards Design Approach

12.3.1 Step 1: Define Learning Outcomes (AI-Neutral)

Start here, always. What should students be able to do by the end of the unit?

Example Unit: Workplace Conflict and Resolution (Third-year undergraduate, HR)

Learning Outcomes: 1. Analyse workplace conflicts using conflict resolution theory and organisational justice frameworks 2. Conduct fair, impartial investigations of workplace complaints 3. Demonstrate effective communication in difficult conversations (de-escalation, active listening, empathy) 4. Design and facilitate conflict resolution interventions appropriate to context 5. Apply relevant employment law and procedural fairness principles 6. Reflect critically on own practice and identify areas for development

Note: These outcomes don’t mention AI. They describe professional competencies. AI is the *means*, not the *end*.

Other discipline examples: The same principle applies whether you’re

designing a Marketing, Accounting, Supply Chain, or Management unit, define what professionals need to *do*, not how they'll do it.

12.3.2 Step 2: Design Assessments (How Will Students Demonstrate Mastery?)

Using the process-based assessment principles from the Assessment chapter, design assessments that make professional competence visible.

Assessment 1: Investigation Interview Simulation (25%) - What: Students conduct simulated investigation interview with AI persona, submit transcript + process audit - **Assesses:** Learning outcomes 2, 3, 5, 6 - **Due:** Week 6 (mid-semester) - **Why this timing:** Gives students foundational practice before more complex work

Assessment 2: Conflict Resolution Portfolio (40%) - What: Students design intervention for multi-stakeholder conflict, conduct simulated mediation/facilitation, write reflective analysis - **Assesses:** Learning outcomes 1, 3, 4, 6 - **Due:** Week 11 - **Why this timing:** Builds on skills from Assessment 1, integrates theory from mid-semester content

Assessment 3: Research Essay (Critical Analysis) (35%) - What: Critical analysis of conflict resolution approaches in specific organisational contexts (e.g., remote work, culturally diverse teams, union environments) - **Assesses:** Learning outcomes 1, 5 - **Due:** Week 13 (exam period) - **Why this timing:** Synthesizes learning from entire semester - **AI integration:** Students use AI for literature synthesis and draft feedback (the Self-Assessment chapter model)

12.3.3 Step 3: Map Learning Activities (How Will Students Prepare for Assessments?)

Now design the week-by-week learning journey that scaffolds students from novice to competent.

Key principle: Gradually increase complexity of AI interaction while building skill.

12.4 Complete 12-Week Unit Design Example

12.4.1 Week 1: Introduction to Conflict and Introduction to AI

Learning Focus: Understand types of workplace conflict, introduce AI as learning tool

Content: - Lecture: Sources and types of workplace conflict - Workshop: Conflict analysis frameworks (task vs. relationship conflict, etc.)

AI Activity (Low stakes, introductory):

In-class demonstration:

- Show students a simple conflict scenario
- Use AI to generate 3 different stakeholder perspectives on the same incident
- Discuss: "How can seeing multiple perspectives help us understand conflict?"

Purpose: - Students see AI in action (demystify) - Understand AI can help explore complexity - No pressure—just observation

Student Task: - Install ChatGPT or Claude - Complete the "Getting Started" tutorial (the Getting Started chapter exercise) - Submit screenshot showing they successfully generated a simple HR scenario

12.4.2 Week 2: Conflict Theory and AI Exploration

Learning Focus: Apply conflict theory; practice writing prompts

Content: - Lecture: Conflict resolution theories (interest-based, transformative, etc.) - Workshop: Analysing conflict through theoretical lenses

AI Activity (First hands-on practice):

Assignment: Theory Application Practice (ungraded)

- Students receive a workplace conflict scenario
- Use AI to analyse the conflict through 2 different theoretical frameworks
- Write 300 words comparing the insights each theory provides
- Submit both the AI conversation and their reflection

Purpose: - Students practice prompt writing - Students evaluate AI's theoretical analysis - Low-stakes experimentation - Lecturer can see who needs prompt-writing help

12.4.3 Week 3: Communication Skills for Conflict

Learning Focus: Active listening, empathetic communication, managing emotion

Content: - Lecture: Communication theory and de-escalation techniques
- Workshop: Communication analysis (watch video examples, critique)

AI Activity (First simulation):

Practice Simulation (ungraded, but required):

- Students conduct 5-minute conversation with AI playing "frustrated"
- Focus: Practice de-escalation language
- Students can retry as many times as they want
- Submit their best attempt + 200-word reflection: "What did I learn"

Purpose: - First taste of "flight simulator" - Builds confidence before graded assessment - Students realise they can practice privately and improve

12.4.4 Week 4: Legal Framework and Procedural Fairness

Learning Focus: Natural justice, procedural fairness, relevant legislation

Content: - Lecture: Legal obligations in workplace investigations -
Workshop: Case studies of procedural failures and consequences

AI Activity (Legal application practice):

Formative Exercise:

- AI generates 5 short investigation scenarios
- For each, students identify: What legal principle is at risk? What
- Immediate AI feedback on their responses
- Students retry until they achieve 100%

Purpose: - Spaced repetition of legal knowledge - Students can practice until mastery (not time-limited) - AI provides immediate correction

12.4.5 Week 5: Investigation Skills

Learning Focus: Conducting fair, thorough workplace investigations

Content: - Lecture: Investigation methodology and common pitfalls -
Workshop: Planning an investigation (what questions, what order, what documentation)

AI Activity (Assessment preparation):

Scaffolded practice for Assessment 1:

- Students receive the persona they'll encounter in Assessment 1 (preview)
- Conduct practice interview
- Generate AI critique
- Revise approach
- Conduct second practice interview

Purpose: - Direct preparation for upcoming assessment - Students enter Assessment 1 having already practiced - Reduces anxiety, improves quality

12.4.6 Week 6: Assessment 1 Due - Investigation Interview Simulation

No new content this week—focus on assessment

Students submit: 1. Transcript of investigation interview with AI persona
2. Process audit document analysing their own performance
3. 500-word reflection on learning

Teaching focus this week: - Availability for consultation/questions - Technical support for any AI access issues

12.4.7 Week 7: Feedback Week + Mediation Theory

Learning Focus: Understanding Assessment 1 feedback; introduction to mediation

Content: - Return Assessment 1 with feedback - Lecture: Mediation and facilitation approaches - Workshop: Compare mediation models (evaluative, facilitative, transformative)

AI Activity (Exploring alternatives):

Scenario Exploration:

- Students receive a conflict scenario suitable for mediation
- Use AI "Pros and Cons" technique (the Seven Techniques chapter) to evaluate
- In-class discussion: Did different students reach different conclusions?

Purpose: - Recover from assessment submission - Introduce new content at moderate cognitive load - Build toward Assessment 2

12.4.8 Week 8: Facilitation Skills

Learning Focus: Facilitation techniques for multi-party conflict

Content: - Lecture: Managing multi-stakeholder conversations - Workshop: Power dynamics, coalition-building, impasse-breaking

AI Activity (Complex simulation introduction):

Multi-party simulation practice:

- Students manage conversation between 2 AI personas in conflict
- Practice balancing airtime, managing interruptions, keeping focus
- Ungraded but highly recommended for Assessment 2 preparation

Purpose: - Increase complexity (now managing 2 personas, not 1) - Build skills for Assessment 2 - Students who struggled with Assessment 1 get redemption opportunity

12.4.9 Week 9: Cultural and Ethical Considerations

Learning Focus: Cross-cultural conflict, ethical dilemmas, bias awareness

Content: - Lecture: Cultural dimensions in conflict (individualism/collectivism, face-saving, etc.) - Workshop: Ethical dilemmas in conflict resolution (confidentiality, power imbalances, organisational pressure)

AI Activity (Critical evaluation):

AI Ethics Exercise:

- AI generates a conflict resolution plan
- Students critique it for:
 - Cultural insensitivity
 - Ethical gaps
 - Bias toward organisational interests over fairness
- Write memo explaining what AI got wrong and why

Purpose: - Develop critical oversight of AI - Connect theory (cultural frameworks, ethics) to practice - Prepare for Assessment 2 cultural/ethical analysis

12.4.10 Week 10: Designing Interventions

Learning Focus: Strategic planning for conflict resolution

Content: - Lecture: Matching interventions to conflict type and context - Workshop: Intervention design process

AI Activity (Assessment 2 preparation):

Portfolio Development:

- Students begin working on Assessment 2
- Use AI to generate multiple intervention options for their chosen scenario
- Bring draft analysis to workshop for peer feedback

Purpose: - Structured time for assessment work - Peer learning and feedback - Lecturer can identify students who need additional support

12.4.11 Week 11: Assessment 2 Due - Conflict Resolution Portfolio

Students submit: 1. Conflict analysis and intervention design (written component) 2. Transcript(s) of simulated intervention (conversation with AI personas) 3. Reflective analysis integrating theory and evaluating their practice

12.4.12 Week 12: Contemporary Issues and Research Essay Support

Learning Focus: Emerging trends in workplace conflict; research essay preparation

Content: - Lecture: Special topics (remote work conflict, AI in HR, gig economy disputes) - Workshop: Research essay planning and literature review strategies

AI Activity (Research support):

Essay development support:

- Students use AI to identify key literature on their chosen topic
- Use AI self-assessment tool to check essay plan
- Optional: Book consultation with lecturer to discuss draft

Purpose: - Support final assessment - Lighter week (no new major concepts) - Celebrate semester's learning

12.4.13 Week 13: Assessment 3 Due - Research Essay

Students submit critical analysis essay.

12.5 The Scaffolding Progression Model

Notice how AI integration increases in complexity:

Week	AI Complexity	Student Agency	Stakes
1-2	Observation— simple prompts	Low (following instructions)	None (ungraded)
3-4	Single persona— structured scenarios	Medium (some choice in approach)	Low (formative)
5-6	Graded simulation— self-assessment	High (must plan and execute)	Medium (25% of grade)
7-9	Multi-persona— ethical critique	High (designing interventions)	Preparation for high-stakes
10-11	Complex portfolio with multiple components	Very high (strategic choices)	High (40% of grade)
12-13	AI as research assistant	Very high (independent work)	High (35% of grade)

This progression develops: 1. **Technical comfort** (Weeks 1-2) 2. **Basic AI literacy** (Weeks 3-4) 3. **Applied competence** (Weeks 5-8) 4. **Critical oversight** (Weeks 9-11) 5. **Independent professional use** (Weeks 12-13)

12.6 Balancing AI and Non-AI Activities

Important: Not everything should involve AI.

12.6.1 This unit includes traditional elements:

- **Lectures:** Content delivery (theory, legal frameworks, research findings)
- **Workshops:** Peer discussion, case analysis, group problem-solving
- **Readings:** Textbook chapters, journal articles, policy documents
- **Live role-play:** At least 1-2 in-person practice sessions for social learning
- **Guest speaker:** Practicing mediator or workplace investigator

- **Reflective journaling:** Weekly reflections on learning (not AI-assessed)

12.6.2 The 60/40 rule:

Aim for approximately: - **60% traditional teaching and learning activities** - **40% AI-enhanced activities**

This ensures students develop both technological proficiency and traditional professional skills (working with humans, not just chatbots).

12.7 Supporting Student AI Literacy Development

Across the semester, explicitly teach AI literacy:

12.7.1 Week 1: What AI Is (and Isn't)

- AI as pattern generator, not intelligence
- Strengths and limitations
- When to trust vs. verify

12.7.2 Week 4: Advanced Prompting

- How to write effective prompts
- Troubleshooting poor responses
- Iterating to improve results

12.7.3 Week 7: Critical Evaluation

- How to spot AI errors
- When AI oversimplifies
- recognising bias in AI outputs

12.7.4 Week 9: Professional Ethics

- Accountability when using AI tools
- When to use AI vs. when human judgment is essential
- Transparent vs. hidden AI use

By semester's end, students haven't just used AI, they've developed **AI literacy as a professional competency**.

12.8 Unit Outline Template (for Your Own Design)

Use this template to design your AI-integrated unit:

12.8.1 UNIT INFORMATION

- Unit code and title:
- Year level and semester:
- Credit points:
- Prerequisites:

12.8.2 LEARNING OUTCOMES (AI-neutral)

- 1.
- 2.
- 3.
- 4.
- 5.

12.8.3 ASSESSMENT SUMMARY

Assessment	Weight	Due Week	AI Integration	Outcomes Assessed
------------	--------	----------	----------------	-------------------

12.8.4 WEEKLY SCHEDULE

Week [X]: [Topic] - Learning focus: - Content delivery: - AI activity: - Purpose: - Preparation for next week:

[Repeat for 12-13 weeks]

12.8.5 AI LITERACY PROGRESSION

- Weeks 1-3: [foundational skills]
- Weeks 4-6: [applied practice]
- Weeks 7-9: [critical evaluation]
- Weeks 10-13: [independent professional use]

12.8.6 BALANCE CHECK

- Traditional activities: [%]

- AI-enhanced activities: [%]
- Justification for this balance:

12.8.7 STUDENT SUPPORT

- Resources provided for AI access:
- Technical support available:
- Academic support for AI use:
- Equity considerations addressed:

12.9 Common Design Mistakes to Avoid

12.9.1 Mistake 1: “AI for AI’s Sake”

Problem: Including AI because it’s trendy, not because it serves learning outcomes. **Solution:** Every AI activity must clearly connect to a learning outcome. If you can’t justify it pedagogically, remove it.

12.9.2 Mistake 2: All or Nothing

Problem: Either avoiding AI entirely or making everything AI-based. **Solution:** Balance. Use AI where it adds value (simulation, feedback, practice) and traditional methods where they’re superior (peer learning, live practice, social skills).

12.9.3 Mistake 3: Assuming Technical Competence

Problem: Expecting students to figure out AI tools independently. **Solution:** Explicitly teach prompt writing, troubleshooting, critical evaluation. Scaffold technical skills just like you scaffold content knowledge.

12.9.4 Mistake 4: No Progression

Problem: Same level of AI complexity all semester. **Solution:** Design deliberate progression from simple to complex, guided to independent, low-stakes to high-stakes.

12.9.5 Mistake 5: Ignoring Equity

Problem: Assuming all students have equal access to AI tools, devices, internet. **Solution:** Provide alternatives (lab access, in-class time for AI activities), use university-subscribed tools where possible, ensure core learning is accessible without premium AI access.

12.10 Aligning Unit Design with Programme Goals

Your unit doesn't exist in isolation—it's part of a degree programme.

12.10.1 Consider:

Vertical integration: - What AI skills do students bring from earlier units? - What AI competencies will later units assume? - How does your unit scaffold toward programme-level AI literacy?

Horizontal integration: - What other units are students taking concurrently? - Could you coordinate AI activities across multiple units? - Are there opportunities for cross-unit projects?

Programme-level graduate capabilities: - How does your AI integration support overarching graduate capabilities? - Communication? Critical thinking? Professional practice? Technological proficiency?

12.11 Communicating the Design to Students

Students need to understand the pedagogical design—it helps them engage meaningfully.

12.11.1 First lecture (explain the approach):

“This unit uses AI tools as part of your learning. Here's why:

In your [professional field] careers, you'll use AI for analysis, strategy development, decision support, and other professional tasks. Our job is to prepare you to use those tools competently and ethically.

You'll notice the AI activities progress across the semester:
- Early weeks: You'll practice basic skills in safe, low-stakes environments
- Mid-semester: You'll apply those skills in realistic scenarios for assessment
- Late semester: You'll use AI independently as a professional tool

By the end, you'll have practiced complex professional scenarios dozens of times, something that would be impossible without AI. You'll also know when to trust AI, when to question it, and when human judgment must override technology.

This isn't about making your degree easier. It's about preparing you for professional practice in an AI-augmented world."

12.11.2 In your unit outline (be explicit):

Include a section titled "AI Integration in This Unit" that explains: - Why AI is used - How it supports learning outcomes - What skills students will develop - Expectations for academic integrity - Support available

12.12 Evaluating Your AI-Integrated Unit

After the semester, evaluate systematically:

12.12.1 Student learning evidence:

- Did assessment results improve compared to previous semesters?
- Did students demonstrate competencies that previous cohorts struggled with?
- What does student work reveal about their AI literacy development?

12.12.2 Student feedback:

- Survey: How useful was AI for your learning? (1-5 scale)
- What AI activities were most valuable?
- What AI activities felt like "busy work"?
- Do you feel more prepared for professional practice?

12.12.3 Your experience:

- Did AI integration save or cost you time overall?
- What worked better than expected? Worse?
- What would you change next semester?
- What would you keep?

12.12.4 Iterate and refine based on evidence.

12.13 Cross-Discipline Unit Design Examples

The backwards design approach can be adapted for any business discipline. Below are examples showing how to design AI-integrated units across

different professional contexts.

 Example: Accounting — Advanced Audit and Assurance

Learning Outcomes (AI-Neutral): 1. Apply professional auditing standards and ethical principles 2. Design risk-based audit procedures and testing strategies 3. Evaluate internal controls and assess control effectiveness 4. Communicate audit findings to diverse stakeholders 5. Demonstrate professional skepticism and critical analysis 6. Apply data analytics in audit planning and execution

Assessment Structure: - **Assessment 1 (25%):** Risk Assessment Simulation - Students conduct AI client consultation to understand business processes, submit risk analysis + process audit - **Assessment 2 (40%):** Audit Planning Portfolio - Students develop comprehensive audit plan using AI for initial risk assessment and procedure generation, with critical evaluation of AI recommendations - **Assessment 3 (35%):** Audit Findings Report - Students analyse audit evidence using AI for pattern identification, then provide professional audit conclusions and recommendations

Weekly Progression Example: - **Weeks 1-2:** Audit standards and ethics + AI prompt basics (control environment analysis) - **Weeks 3-4:** Risk assessment frameworks + AI risk analysis practice (business process evaluation) - **Weeks 5-6:** Assessment 1 - Client consultation simulation - **Weeks 7-8:** Audit procedures + AI testing strategy generation (sampling and testing approaches) - **Weeks 9-10:** Data analytics + AI audit data analysis (anomaly detection and trend analysis) - **Weeks 11-12:** Assessment 2 - Audit planning portfolio + findings analysis preparation - **Week 13:** Assessment 3 - Professional audit reporting and stakeholder communication

Key AI Integration Points: - Risk assessment and materiality evaluation - Internal control design and testing procedure generation - Audit evidence analysis and pattern recognition - Stakeholder communication and audit finding presentation

12.14 Adapting Unit Design Principles Across Disciplines

12.14.1 Common Design Elements

Progression Framework: Regardless of discipline, follow the same scaffolding progression:

- **Weeks 1-2:** Foundational content + AI basics
- **Weeks 3-4:** Core concepts + AI application practice

- **Weeks 5-6:** First assessment (simulation-based)
- **Weeks 7-9:** Advanced concepts + critical AI evaluation
- **Weeks 10-12:** Major assessment (portfolio-based)
- **Week 13:** Synthesis assessment (research/analysis-based)

Assessment Balance: Maintain similar weighting across disciplines:

- **25%:** Process-focused simulation (consultation/interview)
- **40%:** Portfolio assessment (design/strategy development)
- **35%:** Critical analysis (research/strategic evaluation)

AI Literacy Development: Include the same AI literacy progression in all disciplines:

- **Technical comfort** (prompt writing, tool navigation)
- **Applied competence** (discipline-specific applications)
- **Critical oversight** (evaluation of AI outputs)
- **Independent professional use** (strategic AI integration)

12.14.2 Discipline-Specific Considerations

Creative Fields (Marketing, Design): - Emphasise subjective evaluation and iterative refinement - Include portfolio development and presentation skills - Balance analytical and creative AI applications

Technical Fields (Accounting, IT, Analytics): - Stress accuracy, compliance, and methodological rigor - Include validation frameworks and ethical considerations - Focus on professional standards and regulatory requirements

Service Fields (Tourism, Hospitality, Management): - Emphasise stakeholder dynamics and relationship management - Include cultural competence and emotional intelligence - Focus on practical implementation and human factors

Adaptation Strategy: Start with the HR unit design as a template, then modify:

1. **Content:** Replace HR-specific topics with discipline-specific concepts
 2. **Scenarios:** Adapt AI personas and contexts to discipline-appropriate situations
 3. **Assessments:** Modify evaluation criteria to reflect professional standards
 4. **Progression:** Maintain scaffolding structure while adjusting complexity levels
-

12.15 Your Action Step

Design (or redesign) one unit using this backwards design approach:

1. **Choose a unit** you teach (or will teach)
2. **Define learning outcomes** (without mentioning AI)
3. **Design assessments** that make competence visible
4. **Map 12-week learning journey** with deliberate AI scaffolding
5. **Check balance** (60% traditional, 40% AI-enhanced)
6. **Plan equity supports** (access, alternatives, scaffolding)

Don't aim for perfection—aim for “better than what I'm currently doing.”

You can refine each semester based on what you learn.

Chapter 13

Transforming Content with AI

A case study on paper is a story someone else finished. A case study powered by AI is a conversation the student has to navigate.

This chapter covers two complementary workflows: transforming static teaching materials into interactive experiences, and turning AI conversations into professional deliverables. Both put AI to work on the formatting while you focus on the thinking.

13.1 From Static to Interactive

You have a well-designed case study in Word or PDF. Students read it, maybe discuss it, then move on. The learning moment is brief and passive.

AI can transform that static document into an interactive HTML experience — with input fields for student responses, decision trees where choices lead to different outcomes, reflection prompts with saveable responses, and self-assessment checklists with feedback. The result uploads directly to your LMS, works in any browser, and needs no installation.

13.1.1 The Transformation Process

1. **Identify the document** — a case study, worksheet, or activity guide
2. **Locate interaction points** — questions, decisions, reflections
3. **Mark enhancement opportunities** — where could students input or respond?

4. Use AI to generate the interactive HTML

13.1.2 The Prompt

Transform this static [document type] into an interactive HTML webpage:

[Paste your document content here]

Create:

1. Professional HTML with embedded CSS styling
2. Input fields for each question/reflection
3. Save/print functionality for student responses
4. Professional appearance for university students
5. Mobile-responsive design

Requirements:

- All CSS embedded (no external files)
- JavaScript for save/print functionality
- Clear instructions for students

Output: Complete HTML file ready to upload to your LMS.

For more complex scenarios, you can request branching decision trees where student choices affect outcomes, progress tracking, feedback at each stage, and a score summary at the end. The AI generates a single HTML file you upload directly.

13.1.3 What You Can Transform

The same approach works for:

- **Static case studies** → interactive scenarios with decision points
- **Paper worksheets** → digital forms with save/print
- **Reading guides** → self-paced activities with reflection prompts
- **Assessment rubrics** → self-assessment checklists with feedback
- **Lecture handouts** → interactive study guides

13.1.4 Design Principles

- **One file, self-contained.** All CSS and JavaScript embedded — no external dependencies to break.
- **Mobile-responsive.** Students will use phones. Design for it.
- **Accessible.** Screen reader compatible, keyboard navigable, sufficient contrast.

- **Printable.** Students should be able to save their work.
- **Progressive.** Start with a basic transformation. Add complexity (branching, scoring) in later iterations.

13.2 From Conversation to Document

The second workflow runs in the opposite direction: you have a productive AI conversation, and you want it to produce a professional deliverable.

The core principle: **the conversation is the work. The document is the output of that work.** This is fundamentally different from “use AI to write a report” (which replaces thinking) versus “use AI to help you think through an analysis, then generate the presentation of that thinking” (which amplifies thinking).

13.2.1 Application 1: Data Analysis → Presentation

You have a dataset with interesting patterns. The workflow:

Step 1 - Upload and explore:

"I've attached a CSV with [describe data]. Can you describe the structure, identify data quality issues, tell me what patterns you notice, and suggest 3-5 visualisations that would be most revealing?"

Step 2 - Request specific analysis:

"Let's look deeper at [pattern]. Can you run [specific analysis] and explain what it means for [business context]?"

Step 3 - Generate the deliverable:

"Create a PowerPoint presentation (8-10 slides) showing these findings. Include charts, key insights on each slide, and speaker notes explaining what to say."

One conversation produces an exploration, an analysis, and a presentation. Students learn data literacy through natural language — they do not need to code.

13.2.2 Application 2: Qualitative Analysis → Research Memo

For interview data, field notes, or open-ended survey responses:

"I have interview transcripts from [describe]."

Conduct thematic analysis: identify major themes, supporting quotes, and relationships between themes. Then create a research memo (1500 words) with themes, evidence, and implications."

Students learn research methodology by doing it conversationally — the AI handles the formatting while they handle the thinking.

13.2.3 Application 3: Generic Document Creation

AI can generate professional documents directly from conversations:

- **Word reports** — specify structure (executive summary, analysis, recommendations), word count, and formatting
- **Excel spreadsheets** — with formulas, conditional formatting, and dashboard charts
- **CSV files** — ready for further analysis
- **Executive briefs** — one-page summaries for stakeholders

The key: have the thinking conversation first, then request the formatted output. If you skip the thinking and go straight to “write me a report,” you get delegation. If you explore, question, and refine before requesting the document, you get amplified thinking with a professional deliverable at the end.

13.2.4 Tool-Specific Options

Different AI tools handle document creation differently:

- **ChatGPT** generates downloadable files (Word, Excel, PowerPoint) directly in conversation
- **Claude Artifacts** creates documents in a side panel you can edit in-place
- **Google Gemini** exports directly to Google Docs and Sheets
- **MS Copilot** integrates directly into Office applications — the conversation happens inside Word, Excel, or PowerPoint

The prompting principles are identical across all of them. Choose whichever fits your workflow.

13.3 Assessment Strategy

For assignments using either workflow, assess the thinking, not the deliverable:

- **Process evidence (40%)** — the conversation or notes showing what questions the student asked, what patterns they explored, what they refined

- **Critical review (30%)** — did they review AI output critically? Did they correct errors, adjust interpretations, make it their own?
- **Communication quality (20%)** — is the final deliverable clear, professional, and accurate?
- **Reflection (10%)** — can they explain what AI did versus what they decided?

Design assignments that require students to submit their conversation alongside the deliverable. The conversation is where the learning is visible.

13.4 Your Next Step

Try both workflows yourself before using them with students:

1. **Static to interactive:** Take one existing case study or worksheet and transform it. Upload it to your LMS. See how it looks.
2. **Conversation to document:** Take a dataset or topic you know well. Have a 5-minute exploratory conversation with AI, then ask for a presentation or report. Review it critically — what would you edit?

Once you have tried both, you will see which fits your teaching context. Most educators find the static-to-interactive workflow useful for creating reusable teaching materials, and the conversation-to-document workflow useful for student assignments where the process matters more than the product.

Chapter 14

Global Perspectives and Adaptation

The frameworks are transferable. The assumptions behind them are not. Every context demands its own adaptation.

14.1 Introduction

While this book was developed within an Australian educational context, AI integration in business education is a global phenomenon. This chapter provides frameworks for adapting the principles and practices discussed throughout this book to different international contexts, educational systems, and cultural environments.

Core principle: Effective AI integration must respect local educational traditions, cultural norms, and regulatory requirements while maintaining pedagogical excellence.

14.2 Understanding Educational Contexts

14.2.1 Australian Context (Book Foundation)

14.2.1.1 Educational System Characteristics

- **Tertiary structure:** Unified system with clear quality frameworks (TEQSA)
- **Semester timing:** February-June and July-November

- **Assessment approach:** Mix of continuous assessment and final examinations
- **Industry integration:** Strong focus on workplace relevance and practical skills
- **Student diversity:** Significant international student population

14.2.1.2 AI Integration Environment

- **Institutional support:** Growing investment in educational technology
- **Regulatory framework:** Developing guidelines for AI use in education
- **Industry partnership:** Strong connections between universities and workplaces
- **Cultural context:** Multicultural society with inclusive education focus

14.2.2 North American Context

14.2.2.1 United States Educational System

- **Structure:** Diverse system with public, private, and community colleges
- **Semester timing:** August-December and January-May
- **Assessment:** Heavy emphasis on continuous assessment and participation
- **Accreditation:** Regional accreditation bodies with varying standards
- **Student demographics:** Diverse student populations with varying preparation

14.2.2.2 Canadian Educational System

- **Provincial jurisdiction:** Education primarily managed at provincial level
- **Semester structure:** September-December and January-April
- **Assessment balance:** Mix of coursework and examinations
- **Bilingual context:** English and French language considerations
- **International focus:** Strong emphasis on global perspectives

14.2.2.3 AI Integration Considerations

- **Regulatory diversity:** Varying state/provincial regulations on AI use
- **Legal compliance:** FERPA, privacy laws, and accessibility requirements

- **Industry connections:** Strong ties between business schools and corporate partners
- **Technology access:** Varying levels of institutional AI tool provision

14.2.3 European Context

14.2.3.1 United Kingdom Educational System

- **Structure:** Unified system with quality assurance (QAA)
- **Academic year:** September-June with three terms
- **Assessment emphasis:** Balance of coursework and final assessments
- **Quality frameworks:** Strong focus on teaching excellence and student satisfaction
- **Brexit impact:** Changing relationships with European educational systems

14.2.3.2 Continental European System

- **Bologna Process:** Harmonized system across 48 countries
- **Credit transfer:** ECTS system for student mobility
- **Degree structure:** Bachelor-Master-Doctorate progression
- **Language diversity:** Multiple languages of instruction
- **Quality assurance:** European Standards and Guidelines for Quality Assurance

14.2.3.3 AI Integration Environment

- **GDPR compliance:** Strict data protection and privacy regulations
- **Accessibility requirements:** Comprehensive digital accessibility mandates
- **Multilingual support:** Need for AI tools across multiple languages
- **Public funding:** Strong public investment in educational technology

14.2.4 Asian Context

14.2.4.1 East Asian Educational Systems

- **China:** Gaokao system, intense competition, strong government direction
- **Japan:** Semester system, emphasis on group work and harmony
- **South Korea:** Highly competitive, strong technology integration
- **Singapore:** Meritocratic system, strong government support for AI

14.2.4.2 Southeast Asian Systems

- **Varied structures:** Colonial influences mixed with local traditions
- **English medium:** Many programs taught in English
- **Growing economies:** Rapid development and technology adoption
- **Student mobility:** Strong international student exchange programs

14.2.4.3 AI Integration Considerations

- **Language diversity:** Need for multilingual AI support
- **Cultural context:** High-context communication styles
- **Government involvement:** Strong state direction in technology adoption
- **Mobile access:** High smartphone usage and mobile-first approaches

14.2.5 Developing Country Context

14.2.5.1 Educational Challenges

- **Resource constraints:** Limited funding for educational technology
- **Infrastructure issues:** Unreliable internet and power systems
- **Teacher training:** Varying levels of digital literacy among educators
- **Access inequality:** Significant urban-rural digital divides
- **Language complexity:** Multiple local languages with varying digital support

14.2.5.2 AI Integration Opportunities

- **Mobile learning:** Leapfrogging desktop technology to mobile solutions
- **Open source solutions:** Free and open AI tools to reduce cost barriers
- **Adaptive technologies:** customised solutions for local contexts
- **International collaboration:** Partnerships with better-resourced institutions

14.3 Cultural Adaptation Frameworks

14.3.1 Communication Styles

14.3.1.1 High-Context vs. Low-Context Cultures

Cultural Dimension	Low-Context Examples	High-Context Examples	AI Adaptation Strategies
Communication	Direct— explicit— detailed	Indirect— nuanced— relationship- based	Prompt for cultural context awareness
Feedback	Direct criticism— explicit improvement areas	Indirect suggestions— face-saving approaches	Train AI to recognise cultural feedback patterns
Decision-making	Individual— quick— explicit	Group-based— consensus- seeking	Include stakeholder consultation scenarios
Conflict resolution	Direct confrontation— legal frameworks	Mediation— relationship preservation	Design culturally appropriate resolution scenarios

14.3.1.2 Implementation Examples

Low-Context Adaptation (US, Germany, Scandinavia):

You are a business consultant working in a direct communication culture.

Provide feedback that is:

- Specific and action-oriented
- Focused on individual performance
- Based on objective criteria
- Delivered in straightforward manner

Avoid indirect language or relationship-focused framing.

High-Context Adaptation (Japan, China, Arab cultures):

You are a business consultant working in a relationship-focused culture.

Provide feedback that is:

- Respectful of hierarchy and relationships
- Focused on group harmony and improvement
- Delivered through appropriate channels
- Sensitive to face-saving considerations

Include cultural context and indirect communication patterns.

14.3.2 Power Distance Considerations

14.3.2.1 High Power Distance Cultures

- **Characteristics:** Respect for authority, hierarchical structures, formal communication
- **Examples:** Many Asian, African, and Latin American cultures
- **AI Adaptation:** Include appropriate formal language and hierarchical scenarios

14.3.2.2 Low Power Distance Cultures

- **Characteristics:** Informal communication, flat structures, challenging authority
- **Examples:** US, Australia, Scandinavian countries
- **AI Adaptation:** Include collaborative scenarios and informal communication

Prompt Examples:

High Power Distance:

You are a senior manager in a hierarchical organisation.

Communicate this decision to your team:

- Acknowledge your authority and responsibility
- Show respect for organisational structure
- Provide clear direction and expectations
- Maintain appropriate formal tone

Consider cultural expectations about leadership and authority.

Low Power Distance:

You are a team leader in a collaborative organisation.

Discuss this decision with your team:

- Encourage open discussion and feedback
- Seek consensus and input from all members
- Share decision-making process transparently
- Welcome questions and challenges

Consider cultural expectations about participation and equality.

14.3.3 Individualism vs. Collectivism

14.3.3.1 Individualistic Cultures

- **Characteristics:** Personal achievement, individual recognition, direct communication
- **Examples:** US, UK, Australia, Western Europe
- **AI Adaptation:** Focus on individual performance and personal development

14.3.3.2 Collectivistic Cultures

- **Characteristics:** Group harmony, collective success, indirect communication
- **Examples:** Many Asian, African, and Latin American cultures
- **AI Adaptation:** Include group-based scenarios and team success

Implementation Examples:

Individualistic Focus:

You are a career coach working with individualistic values.

Develop personal achievement plans that emphasize:

- Individual goals and accomplishments
- Personal recognition and advancement
- Direct communication of achievements
- Competitive advantage development

Focus on personal success and standing out from others.

Collectivistic Focus:

You are a team facilitator working with collectivistic values.

Develop team success plans that emphasize:

- Group goals and collective achievements
- Harmony and collaboration
- Shared recognition and success
- Support for team members

Focus on group success and contributing to collective goals.

14.4 Regulatory and Legal Adaptations

14.4.1 Data Protection and Privacy

14.4.1.1 European Union (GDPR)

- **Requirements:** Explicit consent, data minimisation, right to erasure
- **AI Implications:** Careful data handling, transparent AI use policies
- **Educational Adaptation:** Clear guidelines for student data in AI systems

14.4.1.2 United States (FERPA and State Laws)

- **Requirements:** Educational records privacy, parental consent for minors
- **AI Implications:** Restrictions on student data use and sharing
- **Educational Adaptation:** Separate AI tools from official educational records

14.4.1.3 Australia (Privacy Act)

- **Requirements:** Australian Privacy Principles, notification requirements
- **AI Implications:** Reasonable data collection, storage limitations
- **Educational Adaptation:** Clear data handling policies for AI integration

14.4.1.4 Developing Country Considerations

- **Challenges:** Varying data protection laws and enforcement
- **Opportunities:** Leapfrogging to modern data protection frameworks
- **Educational Adaptation:** Focus on ethical data practices regardless of legal requirements

14.4.2 Educational Quality Assurance

14.4.2.1 Quality Frameworks by Region

Region	Quality Body	Focus Areas	AI Integration Implications
Australia	TEQSA	Learning outcomes— student support— institutional governance	Evidence of AI effectiveness— student support for AI tools
UK	QAA	Academic standards— student experience— quality enhancement	AI literacy development— digital infrastructure quality
US	Regional Accreditors	Mission fulfillment— student learning— institutional effectiveness	AI alignment with learning outcomes— assessment integrity
EU	ENQA	Harmonized standards— student mobility— quality assurance	Cross-border AI tool compatibility— multilingual support
Asia	Various	National standards— international competitiveness— technology integration	AI innovation— global competitiveness preparation

14.4.2.2 Adaptation Strategies

- **Standards mapping:** Align AI integration with local quality requirements
- **Evidence collection:** Document AI effectiveness for quality assurance
- **Continuous improvement:** Regular review and enhancement of AI practices
- **International benchmarking:** Learn from global best practices

14.5 Discipline-Specific Global Adaptations

14.5.1 Business and Marketing

14.5.1.1 Cultural Market Considerations

- **Consumer behaviour:** Varying purchasing patterns and decision-making processes
- **Communication styles:** Different advertising effectiveness across cultures
- **Brand perception:** Cultural values and brand relationship development
- **Digital marketing:** Varying platform usage and online behaviour

14.5.1.2 Global Marketing AI Adaptations

You are an international marketing consultant working in [REGION/COUNTRY].

Develop marketing strategy that considers:

- Local consumer behaviour and cultural values
- Appropriate communication styles and messaging
- Regulatory requirements and advertising standards
- Competitive landscape and local market conditions
- Digital platform usage and preferences

Provide culturally appropriate examples and implementation considerations.

14.5.2 Human Resources

14.5.2.1 Employment Law Variations

- **Worker protections:** Varying levels of employee rights and protections
- **Termination procedures:** Different legal requirements and processes
- **Discrimination laws:** Varying protected characteristics and enforcement
- **Union relationships:** Different roles and legal status of labour organisations

14.5.2.2 Global HR AI Adaptations

You are an international HR consultant working in [COUNTRY].

Develop HR policies that comply with:

- Local employment laws and regulations
- Cultural expectations about workplace relationships

- Worker protection requirements and union considerations
- Industry standards and best practices
- Cultural communication and management styles

Provide specific legal references and cultural considerations.

14.5.3 Accounting and Finance

14.5.3.1 Regulatory Variations

- **Accounting standards:** IFRS vs. US GAAP vs. local standards
- **Taxation systems:** Varying tax laws and compliance requirements
- **Financial regulations:** Different banking and investment regulations
- **Audit requirements:** Varying standards and enforcement mechanisms

14.5.3.2 Global Finance AI Adaptations

You are an international financial advisor working in [REGION].

Provide financial analysis that considers:

- Local accounting standards and reporting requirements
- Tax implications and regulatory compliance
- Cultural attitudes toward risk and investment
- Local market conditions and economic factors
- International business considerations and exchange rate impacts

Specify all regulatory requirements and cultural factors.

14.5.4 Supply Chain Management

14.5.4.1 Global Supply Chain Considerations

- **Trade regulations:** Varying import/export restrictions and documentation
- **Infrastructure quality:** Different transportation and logistics capabilities
- **Cultural business practices:** Varying negotiation styles and relationship building
- **Geopolitical factors:** Regional stability and trade relationship considerations

14.5.4.2 Global Supply Chain AI Adaptations

You are an international supply chain consultant working in [REGION].

Develop supply chain strategy that addresses:

- Local infrastructure capabilities and constraints
- Trade regulations and compliance requirements
- Cultural business practices and negotiation styles
- Geopolitical risks and mitigation strategies
- Regional logistics networks and distribution channels

Consider local business customs and relationship expectations.

14.5.5 Information Systems

14.5.5.1 Technology Infrastructure Variations

- **Internet penetration:** Varying levels of connectivity and bandwidth
- **Device availability:** Different access to computers and smart-phones
- **Technical skills:** Varying levels of digital literacy and training
- **Software adoption:** Different preferences for platforms and applications

14.5.5.2 Global IT AI Adaptations

You are an international IT consultant working in [REGION].

Design technology solutions that consider:

- Local infrastructure capabilities and limitations
- Device availability and user preferences
- Technical skill levels and training needs
- Language requirements and localization needs
- Regulatory compliance and data sovereignty requirements
- Cultural preferences for technology adoption

Provide implementation strategies for local context.

14.5.6 Management and Organisational Studies

14.5.6.1 Cultural Management Styles

- **Leadership expectations:** Varying perceptions of effective leadership
- **Decision-making processes:** Different approaches to consensus and authority
- **Communication patterns:** Varying expectations about directness and hierarchy
- **Team dynamics:** Different approaches to collaboration and conflict resolution

14.5.6.2 Global Management AI Adaptations

You are an international management consultant working in [COUNTRY].

Develop management approaches that consider:

- Local leadership expectations and cultural preferences
- Decision-making processes and authority structures
- Communication styles and relationship expectations
- Team collaboration patterns and conflict resolution approaches
- Motivation factors and employee engagement strategies

Provide culturally appropriate management examples and implementation.

14.5.7 Economics

14.5.7.1 Economic System Variations

- **Market structures:** Different levels of market freedom and regulation
- **Government involvement:** Varying roles of state in economic management
- **Development stages:** Different economic development levels and challenges
- **Trade relationships:** Varying international trade patterns and dependencies

14.5.7.2 Global Economics AI Adaptations

You are an international economic advisor working in [REGION/COUNTRY].

Provide economic analysis that considers:

- Local economic system and regulatory framework
- Development level and economic challenges
- Cultural factors in economic decision-making
- International trade relationships and dependencies
- Regional economic integration and cooperation

Specify economic assumptions and cultural considerations.

14.5.8 Tourism and Hospitality

14.5.8.1 Cultural Tourism Considerations

- **Service expectations:** Varying standards for hospitality and service quality
- **Cultural tourism:** Different approaches to cultural heritage and authenticity

- **Travel patterns:** Varying tourism behaviours and preferences
- **Hospitality traditions:** Different cultural expectations about hosting and service

14.5.8.2 Global Tourism AI Adaptations

You are an international tourism consultant working in [REGION].

Develop tourism strategies that consider:

- Local service standards and hospitality expectations
- Cultural heritage preservation and presentation
- Tourist preferences and behaviour patterns
- Seasonal variations and cultural events
- Infrastructure capabilities and limitations
- Cultural sensitivity and authentic representation

Provide culturally appropriate tourism examples and implementation strategies.

14.6 Implementation Strategies

14.6.1 Step-by-Step Adaptation Process

14.6.1.1 Step 1: Context Analysis

- **Educational system mapping:** Understand local structures and requirements
- **Cultural assessment:** Identify relevant cultural dimensions and preferences
- **Regulatory review:** Analyse legal requirements and compliance needs
- **Resource evaluation:** Assess technological infrastructure and support systems

14.6.1.2 Step 2: Stakeholder Engagement

- **Local consultation:** Engage with local educators and administrators
- **Student input:** Gather feedback from local student populations
- **Industry partnership:** Connect with local businesses and organisations
- **Cultural experts:** Consult with cultural specialists and community leaders

14.6.1.3 Step 3: customisation Development

- **Content adaptation:** Modify examples and scenarios for local context
- **Platform selection:** Choose AI tools appropriate for local infrastructure
- **Assessment design:** Align evaluation with local educational standards
- **Support systems:** Develop resources appropriate for local needs

14.6.1.4 Step 4: Implementation and Testing

- **Pilot programs:** Test adaptations with small groups
- **Feedback collection:** Gather systematic input from all participants
- **Iteration and refinement:** Improve based on testing results
- **Scale-up planning:** Prepare for broader implementation

14.6.1.5 Step 5: Evaluation and Improvement

- **Effectiveness assessment:** Measure success of adapted approaches
- **Comparison with benchmarks:** Compare results with global best practices
- **Continuous improvement:** Regular updates and enhancements
- **Knowledge sharing:** Share learning with global education community

14.6.2 Quality Assurance Framework

14.6.2.1 Adaptation Evaluation Criteria

Criterion	Questions to Consider	Evidence of Success
Cultural Appropriateness	Are examples and scenarios culturally relevant and respectful?	Positive student feedback—engagement metrics
Regulatory Compliance	Does implementation meet local legal requirements?	Compliance audits—institutional approval
Educational Effectiveness	Are learning outcomes achieved in local context?	Assessment results—learning analytics
Technical Feasibility	Do AI tools work with local infrastructure?	Usage statistics—technical support requests

Criterion	Questions to Consider	Evidence of Success
Stakeholder Acceptance	Do local educators—students— and employers find value?	Survey results— partnership feedback

14.6.2.2 Continuous Improvement Process

1. **Regular monitoring:** Ongoing assessment of adaptation effectiveness
2. **Stakeholder feedback:** Systematic collection of user experiences
3. **Benchmark comparison:** Regular comparison with global best practices
4. **Adaptation updates:** Periodic refinement based on evidence
5. **Knowledge sharing:** Contribution to global education community

14.6.3 Building Global Networks

14.6.3.1 International Collaboration

- **Partner institutions:** Develop relationships with schools in different regions
- **Exchange programs:** Share faculty and student experiences
- **Joint research:** Collaborate on AI integration research
- **Resource sharing:** Exchange adapted materials and approaches

14.6.3.2 Professional Development

- **International conferences:** Present and learn about global AI integration
- **Cross-cultural training:** Develop skills for working in diverse contexts
- **Language skills:** Improve ability to work across language barriers
- **Cultural competence:** Build understanding of different educational traditions

14.6.3.3 Community of Practice

- **Global networks:** Connect with AI integration educators worldwide
- **Online platforms:** Participate in international discussions and forums
- **Resource repositories:** Contribute to and access global adaptation resources
- **Mentorship programs:** Support and be supported by international colleagues

14.7 Case Studies and Examples

14.7.1 Successful Adaptation Examples

14.7.1.1 Case Study 1: European University Adaptation

Context: German business school implementing AI integration
Challenges: GDPR compliance, multilingual student body, quality assurance requirements
Solutions: - Developed comprehensive data protection policies for AI use - Created multilingual prompt templates and examples - Established quality assurance processes for AI-enhanced assessments - Built partnerships with European AI companies for local support

Results: Successful implementation with high student satisfaction and compliance with European standards

14.7.1.2 Case Study 2: Asian University Adaptation

Context: Singaporean university integrating AI into business curriculum
Challenges: High-context communication culture, competitive environment, technology expectations
Solutions: - Developed AI scenarios emphasising relationship-building and harmony - Created competitive yet collaborative learning environments - Integrated cutting-edge AI tools to meet high technology expectations - Emphasised face-saving communication approaches in AI interactions

Results: Enhanced student engagement, improved learning outcomes, maintained cultural values

14.7.1.3 Case Study 3: Developing Country Adaptation

Context: Nigerian business school with limited resources
Challenges: Infrastructure limitations, cost constraints, unreliable internet
Solutions: - Focused on mobile-first AI applications - Utilised free and open-source AI tools - Developed offline AI-enhanced activities - Created peer-to-peer learning networks for AI skill development

Results: Expanded access to AI-enhanced learning, improved educational quality despite resource constraints

14.7.2 Lessons Learned

14.7.2.1 Common Success Factors

- **Cultural sensitivity:** Respect for local educational traditions and values

- **Stakeholder engagement:** Involvement of local community in adaptation process
- **Flexibility:** Willingness to modify approaches based on local feedback
- **Sustainability:** Building local capacity and long-term viability

14.7.2.2 Common Challenges

- **Resource constraints:** Limited funding and infrastructure in some regions
- **Regulatory complexity:** Navigating different legal and quality frameworks
- **Cultural resistance:** Overcoming skepticism about new educational approaches
- **Technical barriers:** Infrastructure limitations and digital divides

14.7.2.3 Best Practices

- **Start small:** Pilot programs before full implementation
- **Local leadership:** Empower local educators to lead adaptation efforts
- **Contextual relevance:** Ensure all examples and scenarios are locally meaningful
- **Continuous learning:** Regular assessment and improvement of adaptations

14.8 Conclusion

AI integration in business education is a global phenomenon that requires local adaptation. The frameworks and strategies in this chapter provide guidance for implementing AI-enhanced teaching across diverse international contexts while maintaining pedagogical excellence and cultural respect.

Key principles for successful global adaptation: 1. **Understand local context:** Educational systems, cultural norms, and regulatory requirements 2. **Engage local stakeholders:** Involve educators, students, and communities in adaptation process 3. **Respect cultural differences:** Adapt communication styles, examples, and approaches appropriately 4. **Ensure regulatory compliance:** Meet local legal requirements and quality standards 5. **Build sustainable capacity:** Develop local expertise and long-term implementation strategies

By applying these adaptation frameworks, educators can effectively implement AI-enhanced teaching that is both globally informed and locally

relevant, preparing students for success in both their local contexts and the global business environment.

Final Chapter: This concludes the comprehensive guide to AI integration in business education. Educators now have the tools, frameworks, and strategies needed to transform their teaching while preparing students for professional success in an AI-augmented world.

Chapter 15

Implementation Practicalities

Plan for the tool to fail. The pedagogy should work even when the technology does not.

If AI integration only works for some students, it is not integration. It is exclusion with better technology.

15.1 When Things Go Wrong

AI tools are powerful but not infallible. When implementing AI-enhanced teaching, technical issues are inevitable. The question is not whether something will fail, but whether you have a plan when it does.

15.1.1 The Essential Preparations

Have backup platforms. If your primary tool goes down, students should be able to switch to an alternative. Design assignments that work with any AI tool — ChatGPT, Claude, Gemini, Copilot — rather than depending on one. Specify this in your assignment instructions.

Build in deadline flexibility. Include a clear policy: if the AI platform is unavailable during the assignment period, students can request an extension. This removes anxiety and prevents a wave of panicked emails.

Prepare offline alternatives. For any AI-enhanced activity, have a non-AI version ready. A conversation simulation can fall back to a written case analysis. An interactive exercise can fall back to a paper worksheet. The learning objectives should be achievable either way.

Test before you launch. Run every prompt yourself before assigning it to students. Check that the AI stays in character, produces useful output, and handles edge cases. A prompt that works for you in a calm office may fail when 80 students use slightly different wording.

15.1.2 Common Issues and Quick Fixes

Problem	Response
Platform is down	Switch to alternative tool. Extend deadline if needed.
AI not staying in character	Start a fresh conversation. Check full prompt was copied.
AI giving inconsistent quality	Try a different model or platform. Simplify the task.
Student cannot access tool	Offer office hours to run it together. Provide alternative.
Student says “is this cheating?”	“This IS the assignment. You’re graded on critical use, not avoidance.”
Conversation gets too long and degrades	Start a new conversation. Restate key context explicitly.

15.1.3 Risk Assessment

Risk	Probability	Impact	Mitigation
Service outage	Medium	High	Multiple platforms, offline alternatives
Quality degradation	High	Medium	Benchmark testing, platform switching
Data loss	Medium	High	Local backups, screenshot important outputs
Security/privacy incident	Low	Critical	Enterprise tools, data sanitisation

15.2 Accessibility and Inclusion

AI integration creates new opportunities but also new barriers. If you design only for students with reliable internet, modern devices, and strong English skills, you are excluding a significant portion of your cohort.

15.2.1 The Digital Divide

Not all students have equal access to technology. Before assigning AI-enhanced work, consider:

- **Device availability.** Some students access coursework only through phones or shared computers. Design activities that work on mobile browsers, not just desktops.
- **Internet reliability.** Rural and low-income students may have poor connectivity. Allow activities to be completed in short sessions rather than requiring extended online time. Provide campus lab time as an option.
- **Cost barriers.** Free-tier tools have usage limits. Design assignments around freely available platforms. If your institution provides enterprise AI access, make sure students know how to use it.
- **Digital literacy.** Not all students arrive with the same comfort level. A five-minute demo of “here is where you paste the prompt” removes a barrier that feels obvious to you but is not obvious to everyone.

15.2.2 Universal Design Principles

Design AI activities that work for the widest range of students:

Multiple means of engagement. Offer text-based, voice-based, and visual interaction options where possible. Allow students to pace interactions at their own speed.

Multiple means of expression. Let students demonstrate learning in different ways — written reflection, verbal explanation, visual presentation, or process documentation. A student who struggles with written English may demonstrate excellent critical thinking verbally.

Multiple means of representation. Request AI output in different formats — text summaries, bullet points, step-by-step breakdowns — so students can choose what works for their learning style.

15.2.3 Specific Accommodations

Students with visual impairments. Prioritise text-based AI tools with strong screen reader support. Ask AI to provide structured, hierarchical

responses rather than visual content. Avoid prompts that depend on “see above” or visual references.

Students with hearing impairments. Ensure all AI interactions are available in text. If using voice-based AI features, provide text alternatives.

Neurodivergent students. AI conversations can be overwhelming. Allow breaks, provide templates that structure the interaction, and keep prompts focused on one task at a time. The “one prompt, one job” principle from CRAFT is especially important here.

Non-native English speakers. AI tools are primarily trained on English. Allow students to work in their preferred language where possible. Include cultural context in prompts to avoid Western-centric examples.

15.2.4 The Process Assessment Safeguard

The process-over-product assessment approach described in the Assessment chapter has a built-in accessibility benefit: the process component should not require AI use. A student who keeps a research journal, annotates readings, or documents their decision-making in any form should be able to submit equivalent process evidence. The assessment measures engaged thinking, not AI use. This ensures students with limited access, disability-related barriers, or personal preferences against AI can demonstrate the same learning through alternative means.

15.2.5 Course Design Checklist

Before launching any AI-enhanced activity:

- Are chosen AI tools accessible on mobile and desktop?
- Can students participate using free-tier tools?
- Are alternative (non-AI) pathways available?
- Is campus lab time available for students without reliable home access?
- Are instructions clear enough for students with no prior AI experience?
- Are timing accommodations built in for students who need extended time?
- Does the activity work for students using screen readers or assistive technology?
- Are examples culturally diverse rather than Western-centric?

15.3 Your Action Step

Pick one AI-enhanced activity you are planning. Run through the checklist above. Identify the two biggest accessibility gaps and fix them before

launching. Then ask yourself: if the AI platform went down the night before, what would students do? If you do not have a good answer, create the backup plan now.

Chapter 16

Conclusion: Where Do We Go From Here?

The goal was never to use more AI. It was to teach more effectively, with AI as the catalyst, not the point.

16.1 What You've Learned

Over the course of this book, you've explored:

- **Why AI matters** for preparing business professionals across all disciplines for real-world practice
- **How to use AI** through simple prompts that anyone can write
- **Seven core techniques** that develop critical thinking and professional skills
- **Three major applications:** conversation simulations, self-assessment tools, and virtual company scenarios
- **New assessment models** that evaluate process and methodology, not just knowledge recall
- **Practical implementation** from your first experiment through full unit redesign
- **Ethical frameworks** for responsible AI integration and academic integrity
- **Advanced applications** for unit design and postgraduate research support

You now have the knowledge and tools to integrate AI into your teaching in meaningful, pedagogically sound ways—regardless of your discipline.

But knowledge alone isn't enough.

16.2 The Question That Matters

As you close this book, you face a decision:

Will you actually try something?

It's easy to read about innovative pedagogy and think "That's interesting." It's harder to actually change your practice.

You're busy. You have existing materials that work well enough. You're comfortable with your current approach. Change is risky—what if students resist? What if colleagues judge? What if it doesn't work?

These are legitimate concerns.

But consider this: **Your students will use AI in their professional careers—regardless of their discipline—whether you teach them to or not.**

The question isn't "Should AI be part of professional practice?" It already is, across all business disciplines.

The question is: **"Will my graduates know how to use AI responsibly, critically, and ethically in their field?"**

If the answer is "I hope so" or "They'll figure it out," you're sending students into professional practice unprepared.

16.3 Start With One Thing

You don't need to implement everything in this book. You don't need to redesign your entire curriculum. You don't even need to be certain it will work perfectly.

You just need to try one thing.

Choose the smallest experiment that feels manageable:

16.3.1 Option 1: Try It Yourself (This Week)

- Pick one prompt from the examples in this book
- Generate a teaching resource you actually need (case study, practice questions, discussion prompts)
- Use it in your next class
- See what happens

Time investment: 30 minutes **Risk:** Minimal **Learning:** High

16.3.2 Option 2: Student Demonstration (Next Class)

- In your next lecture, project a live AI conversation on screen
- Show students how AI can help them practice skills
- Answer their questions
- Don't assign anything—just plant the seed

Time investment: 15 minutes in class **Risk:** None (optional for students)

Learning: Medium

16.3.3 Option 3: Low-Stakes Practice Exercise (This Semester)

- Add one optional AI practice activity to an existing assignment
- Recommended but not required
- See who uses it and gather feedback
- Iterate for next semester

Time investment: 1-2 hours setup **Risk:** Low (it's optional) **Learning:** Substantial (you'll see what students actually do with AI)

16.3.4 Option 4: Pilot Assessment (Next Semester)

- Redesign one existing assignment using ideas from the techniques, flight simulator, or assessment chapters
- Worth 15-25% of the grade (significant but not high-stakes)
- Document what works and what doesn't
- Refine for future iterations

Time investment: 3-5 hours initial design **Risk:** Moderate (but manageable with clear instructions) **Learning:** Transformative (you'll see process-based assessment in action)

16.3.5 Option 5: Full Unit Redesign (Next Academic Year)

- Use the backwards design approach from the Unit Design chapter
- Integrate AI throughout one complete unit
- Build scaffolded progression from Week 1 to Week 12

- Measure impact on student learning

Time investment: Significant (20-30 hours initial design) **Risk:** Higher (but with high potential reward) **Learning:** Comprehensive (you'll develop deep expertise in AI-enhanced pedagogy)

16.4 Pick one. Not five. One.

The biggest mistake educators make with innovation is trying to do too much at once. They get overwhelmed, it doesn't go perfectly, and they abandon the whole thing.

Small, successful experiments build confidence and capability.

One well-executed pilot teaches you more than five half-baked attempts.

16.5 What Success Looks Like

How will you know if your AI integration is working?

16.5.1 Short-Term Success (First Semester)

Student engagement: - Students ask questions about AI use (curiosity)
- Students report that AI helped them prepare or practice (utility) -
Students use AI activities even when optional (voluntary adoption)

Your experience: - You complete the pilot without major disasters -
You learn something about what works and what doesn't - You feel more
confident about AI tools and their limitations

Tangible outcomes: - At least one student says "That simulation really
helped me understand..." - You create at least one reusable resource you'll
use again - You gather feedback that informs your next iteration

16.5.2 Medium-Term Success (Within 2-3 Semesters)

Student learning: - Improved performance on assessments related to
AI-practiced skills - Students demonstrate competencies earlier in the
semester - Fewer students make basic procedural or communication errors
- Students reference their AI practice in reflections and discussions

Your teaching: - You have 2-3 reliable AI-enhanced activities you use
regularly - You've refined prompts and instructions based on student

experience - You feel AI is enhancing rather than complicating your teaching - Other lecturers ask you about your approach

Curriculum: - AI integration is normalised (not novel or controversial) - Students expect and value AI-enhanced learning opportunities - You've expanded from one unit to multiple units or assessment types

16.5.3 Long-Term Success (3+ Years)

Graduate outcomes: - Alumni report that AI-enhanced learning prepared them for professional practice - Employers or practicum supervisors notice your graduates are better prepared - Students explicitly mention AI literacy as a valuable skill they developed

Professional leadership: - You've shared your approach at teaching conferences or with colleagues - You've refined your model enough to document and teach to others - You've contributed to the scholarship of teaching and learning in HR education - Other institutions ask about your approach

Institutional impact: - AI integration becomes standard practice in HR programs - Your university recognises this as teaching innovation - The approach influences accreditation or curriculum design discussions

16.6 Avoiding Common Pitfalls

As you move forward, watch for these mistakes:

16.6.1 Pitfall 1: Technology for Technology's Sake

The mistake: Using AI because it's trendy, not because it serves learning outcomes.

The fix: Every AI activity must answer: "What learning outcome does this support that couldn't be achieved as effectively another way?"

If you can't answer that clearly, don't use AI for that task.

16.6.2 Pitfall 2: Assuming Technical Competence

The mistake: Expecting students to figure out AI tools on their own.

The fix: Explicitly teach prompt writing, critical evaluation, and ethical use. Build technical scaffolding just like you build content scaffolding.

16.6.3 Pitfall 3: No Clear Assessment Criteria

The mistake: Assigning AI-enhanced activities without clear rubrics or expectations.

The fix: Students need to know what “success” looks like. If they’re submitting conversation transcripts, what are you assessing? If they’re using AI for self-assessment, what’s your role in grading?

Make criteria explicit and transparent.

16.6.4 Pitfall 4: Ignoring Equity

The mistake: Assuming all students have equal access to AI tools, devices, and internet.

The fix: Provide alternatives (lab time, office hours facilitation, university-subscribed tools). Ensure core learning is accessible regardless of premium AI access.

16.6.5 Pitfall 5: Blind Faith in AI Outputs

The mistake: Treating AI-generated content as inherently correct or reliable.

The fix: Teach students (and remember yourself) that AI makes confident mistakes. Always verify. Always maintain human oversight. Always question.

16.7 Building Community

You don’t have to do this alone.

16.7.1 Within Your Institution:

- Connect with colleagues experimenting with AI in teaching
- Join or form a teaching and learning community of practice
- Share successes and failures openly

- Co-design activities and assessments
- Observe each other's classes

16.7.2 Beyond Your Institution:

- Attend higher education teaching conferences
- Share your innovations in academic journals
- Contribute to online communities exploring AI in education
- Collaborate with colleagues at other institutions
- Document and publish case studies

Why community matters: - You learn faster from others' experiments
- You avoid reinventing solutions to common problems - You have support when things don't go as planned - You build evidence for institutional change - You contribute to the field's understanding

16.8 The Bigger Picture: Transforming Business Education

Individual educators trying new things is important. But the real transformation happens when entire programs evolve.

16.8.1 Vision for Business Education with AI Integration

Year 1 (Undergraduate): Students develop AI literacy alongside foundational disciplinary knowledge. They learn to use AI for exploration, practice, and self-assessment. They develop critical evaluation skills specific to their field.

Year 2-3 (Undergraduate): Students apply AI tools to complex scenarios in their discipline. They use conversation simulations and decision-making activities extensively. They demonstrate competence through process-based assessments. They critique AI outputs and improve on them.

Year 4-5 (Undergraduate/Honours/Research Programs): Students use AI as a professional tool. They integrate AI into strategic thinking and research in their discipline. They teach others how to use AI responsibly. They understand when AI helps and when human judgment must override technology.

Professional Practice: Graduates enter workplaces confident with AI tools, critical of AI limitations, and committed to ethical AI use. They

advocate for fairness when organisations implement AI systems in their field. They maintain human accountability for AI-assisted decisions.

This is the future we’re building.

Not a future where AI replaces professionals, but where business professionals across all disciplines use AI skillfully and ethically to do their work better—to serve people, organisations, and society more effectively.

16.9 What Employers Are Looking For

Employers increasingly expect graduates to use AI. That expectation is no longer a differentiator – it is baseline. What separates candidates is whether they can use AI *well*, and whether they can talk about it with any substance.

The skills this book teaches – critical evaluation, process thinking, transparency about how you arrived at a recommendation – are precisely what hiring managers are starting to probe in interviews. Questions like “How do you use AI in your work?” and “Can you give an example of when you identified a problem with AI output?” are becoming standard. Graduates who can answer those questions with specific, reflective examples will stand out.

Here is the gap: almost every graduate now says “I use AI.” Very few can say “I use AI critically and can explain my process.” The first statement is common. The second is rare, and it is the one that signals professional judgment.

Think about what a strong answer looks like. A candidate who says “I asked ChatGPT and it gave me a good answer” is describing passive consumption. A candidate who says “I generated three approaches, evaluated each against our compliance requirements and organisational constraints, identified gaps in the AI’s assumptions, and iterated until I had something I could defend” is describing professional practice. That is the difference employers notice.

This is not hypothetical. Across industries, organisations are discovering that AI adoption without critical oversight creates real problems – flawed analysis, compliance failures, recommendations that ignore context. They need people who can catch those issues before they cause damage. That is a skill, and it is teachable.

The frameworks in this book – the 5-step critique process, the VET framework, process-based documentation – give students a vocabulary and a method for demonstrating this competence. When a graduate walks

into an interview and can describe *how* they evaluate AI output, not just *that* they use it, they are showing exactly the kind of professional maturity that employers are looking for.

Preparing students for these conversations is not an add-on. It is one of the most practical things you can do for their career readiness.

16.10 Your Legacy

Every student you teach will work with AI in their careers, regardless of their discipline.

The question is: Will they be competent or incompetent? Ethical or reckless? Critical or credulous?

That's in your hands.

When you integrate AI into your teaching—transparently, critically, and pedagogically—you're not just adopting a new tool. You're preparing the next generation of business professionals for a world that will be shaped by technology but must still be guided by human wisdom.

That's not a small thing.

That's your professional responsibility and your legacy.

16.11 Final Words

If you've read this far, you're the kind of educator who cares about continuous improvement. You're not content with "good enough." You're asking "What could be better?"

That's exactly the mindset needed for this work.

AI in education isn't settled science. We're all figuring this out together—what works, what doesn't, what's ethical, what's effective. You're not behind. You're not too late. You're here, right now, at exactly the right time.

You have: - The knowledge (this book) - The frameworks (Chapters 4-11)
- The institutional alignment framework (Appendix A) - The rubrics and stress tests (Appendices B-C) - Downloadable prompts and workshop materials at the companion website

What you need now is courage.

Courage to try something new. Courage to fail, learn, and try again. Courage to change your practice when change is uncomfortable. Courage to lead when others are still watching and waiting.

You have that courage.

I know this because you read 300+ pages about AI in education. That's not something an unimaginative or risk-averse educator does.

So here's my final challenge:

Close this book. Choose one thing. Do it this week.

Not next month. Not next semester. This week.

Your students are waiting for the learning experiences only you can design.

16.12 One Last Thing

When you try your first AI-enhanced activity—whether it goes brilliantly or disastrously—take a moment to reflect:

- What surprised you?
- What will you do differently next time?
- What did students learn that they wouldn't have otherwise?

Then do it again, better.

That's how transformation happens.

One experiment. One refinement. One semester at a time.

Welcome to the future of business education.

You're ready.

For ongoing support, resources, and community: - Your institution's Teaching and Learning team - AI in Higher Education communities online - The companion website: <https://michael-borck.github.io/partner-dont-police>

Good luck. And thank you for being the kind of educator who never stops learning.

Chapter 17

Colophon

Version: 1.0 **Published:** 2025 **Scope:** Multidisciplinary Business Education (HR, Marketing, Accounting, Management, Tourism & Hospitality, Supply Chain, Information Systems, Economics, and Business Analytics)

Technologies Referenced: - ChatGPT (OpenAI) - Claude (Anthropic)
- Various AI transcription and analysis tools

Pedagogical Frameworks: - Backwards Design (Wiggins & McTighe) - Experiential Learning (Kolb) - Reflective Practice (Schön) - Authentic Assessment - Self-Directed Learning - Process-Based Assessment

Disclaimer: AI technology evolves rapidly. Specific tools and capabilities described in this book reflect the state of technology in early 2025. Principles and pedagogical approaches remain relevant across technological changes. This book is designed for application across multiple business disciplines with context-specific adaptations.

For copyright, licensing, and citation information, see the Copyright page.

Appendix A

Aligning AI Integration with Institutional Learning Outcomes

A.1 Purpose of This Appendix

This appendix demonstrates how AI integration in business education directly supports — rather than replaces — existing learning outcomes. It provides a framework you can adapt to your own institution’s programmes and policies. It is designed for:

- Business faculty across disciplines seeking institutional justification for AI integration
- Programme coordinators evaluating pedagogical innovations
- Academic administrators assessing alignment with university strategy
- Accreditation reviews demonstrating innovative teaching aligned with standards

A.2 Institutional Strategic Context

Most universities now have AI policies or guidance frameworks. Before integrating AI into your teaching, locate your institution’s policy and identify the principles it endorses. Common principles across institutional AI policies include:

1. **Human oversight and accountability** — AI systems should augment, not replace, human decision-making
2. **Fairness and equity** — AI should not create or reinforce unfair discrimination
3. **Privacy and data protection** — Comply with institutional data governance and relevant legislation
4. **Transparency** — Be open about when and how AI is used in teaching and assessment
5. **Accuracy and reliability** — Verify AI outputs and acknowledge model limitations
6. **Student wellbeing** — AI should enhance, not diminish, the learning experience

The teaching approach in this book aligns naturally with these principles:

- **Human oversight:** Every technique requires students to evaluate, critique, and improve AI output — never to accept it uncritically
- **Fairness:** Process-based assessment reduces bias by evaluating thinking rather than just outputs
- **Privacy:** Data governance guidance (Chapter 9) teaches responsible tool selection
- **Transparency:** Students learn to acknowledge AI use openly as professional practice
- **Accuracy:** The critique toolkit and VET framework train verification habits
- **Student wellbeing:** AI provides unlimited low-stakes practice, reducing assessment anxiety

A.3 Mapping AI Integration to Common Business Learning Outcomes

Business programmes across institutions share broadly similar learning outcome categories. The framework below maps AI teaching applications from this book to nine common outcome areas. Substitute your own programme's specific wording where appropriate.

A.3.1 Outcome 1: Apply Discipline-Specific Theory to Practice

Typical expectation: Students demonstrate ability to apply relevant theories to real-world situations and make evidence-based decisions.

AI Application (Chapter)	How It Supports This Outcome
Conversation Simulations (Ch. 7)	Students apply theoretical frameworks in real-time during dynamic conversations. AI responds to the quality of their theoretical application.
Evidence-Based Analysis (Ch. 10)	Students use AI to analyse data, then must justify recommendations using theory. Assessment requires explicit connection between data patterns and theoretical frameworks.
Debating Technique (Ch. 5)	Multi-perspective analysis requires students to evaluate competing strategies through theoretical lenses.

Evidence of learning: Students cite specific theories in conversation transcripts; students critique AI recommendations by identifying missing theoretical considerations; students demonstrate application, not just definition, of theory.

A.3.2 Outcome 2: Knowledge of Relevant Legal and Regulatory Frameworks

Typical expectation: Students demonstrate understanding and application of relevant legislation, regulations, and compliance requirements in their discipline.

AI Application (Chapter)	How It Supports This Outcome
Conversation Simulations (Ch. 7)	Scenarios embed legal considerations. Students must demonstrate compliance awareness in real-time.
Process Assessment (Ch. 10)	Students must identify where they did or did not apply legal principles and cite specific legislation.
Ethics Exercises (Ch. 9)	Students evaluate AI-generated policies and identify legal flaws, discriminatory provisions, and compliance gaps.

Evidence of learning: Transcripts show students applying legal principles in decision-making; process audits include legislation citations; students identify when AI recommendations violate legal requirements.

A.3.3 Outcome 3: Communicate Effectively and Demonstrate Professional Practice

Typical expectation: Students communicate effectively with individuals and groups, demonstrate professionalism, and manage difficult interpersonal situations.

AI Application (Chapter)	How It Supports This Outcome
Conversation Simulations (Ch. 7)	Every simulation requires active listening, empathetic responses, and professional communication. AI responds dynamically to communication quality.
Multiple Practice Cycles	Unlike traditional role-play (one attempt), students can practise the same conversation multiple times, refining their approach.
From Conversation to Document (Ch. 15)	Students learn to translate exploratory AI conversations into professional deliverables.

Evidence of learning: Transcripts demonstrate professional tone and active listening; students show improvement between attempts; reflections articulate understanding of communication impact.

A.3.4 Outcome 4: Apply Professional and Ethical Standards

Typical expectation: Students demonstrate ethical professional conduct, respect for diversity, and understanding of professional responsibilities.

AI Application (Chapter)	How It Supports This Outcome
Transparency Model (Ch. 9)	Teaching students to use AI openly and critically models professional integrity.

AI Application (Chapter)	How It Supports This Outcome
Ethics Scenarios (Ch. 9)	Students analyse ethical problems with AI use in professional contexts — biased tools, algorithmic discrimination, accountability questions.
Critique and Override Exercises (Ch. 8)	Students identify when AI recommendations are ethically problematic and demonstrate superior human judgement.

Evidence of learning: Students identify bias, discrimination, or ethical flaws in AI outputs; students demonstrate human oversight of AI-generated decisions; reflections show awareness of professional accountability.

A.3.5 Outcome 5: Think Critically and Evaluate Information

Typical expectation: Students critically analyse problems, evaluate information from multiple sources, and make evidence-based decisions.

AI Application (Chapter)	How It Supports This Outcome
Critique Toolkit (Ch. 8)	Students evaluate AI-generated analysis, then must critique AI's reasoning and add missing considerations.
Self-Assessment (Ch. 11)	Students receive AI feedback but must critically evaluate whether it is correct. Strong students challenge the AI's assessment.
VET Framework (Ch. 8, Introduction)	Verify, Explain, Test — a structured approach to critical evaluation of any AI output.

Evidence of learning: Students successfully identify AI errors or limitations; students improve AI recommendations with additional analysis; students demonstrate reasoning that surpasses AI capability.

A.3.6 Outcome 6: Self-Directed Learning and Reflective Practice

Typical expectation: Students demonstrate capacity for independent learning, reflection on practice, and continuous professional development.

AI Application (Chapter)	How It Supports This Outcome
Self-Assessment Tool (Ch. 11)	Students drive their own improvement cycle: draft, AI feedback, reflection, revision.
Process Assessment (Ch. 10)	Students analyse their own performance, identify strengths and weaknesses, and propose improvements.
Unlimited Practice	AI simulations available any time. Students who want additional practice can self-direct their learning beyond required assignments.

Evidence of learning: Reflections demonstrate genuine self-assessment; evidence of revision between drafts shows iterative improvement; students articulate what they learned and how they will apply it.

A.3.7 Outcome 7: Technological Proficiency in Professional Contexts

Typical expectation: Students select and effectively use appropriate technologies relevant to professional practice and research.

AI Application (Chapter)	How It Supports This Outcome
All AI-Enhanced Assignments	Direct practice with AI tools that are increasingly standard in professional practice across business disciplines.
Critical Oversight Training (Ch. 9)	Students learn when to use AI, when to verify outputs, and when human judgement must override technology.
AI Literacy (throughout)	Explicit teaching of AI capabilities, limitations, bias recognition, and accountability.

Evidence of learning: Students competently use AI tools to support professional tasks; students identify appropriate versus risky AI use cases; students demonstrate human oversight and accountability.

A.3.8 Outcome 8: Resolve Complex Professional Problems

Typical expectation: Students demonstrate ability to investigate issues, manage conflicts, and resolve complex problems in their discipline.

AI Application (Chapter)	How It Supports This Outcome
Flight Simulator (Ch. 7)	Practise full professional processes in a safe environment with realistic complexity.
Stepwise Chain of Thought (Ch. 5)	Guides students through proper resolution processes step-by-step, ensuring they understand why each step matters.
Virtual Company (Ch. 12)	Complex, evolving scenarios that require strategic problem-solving over time. Students see consequences of their approaches.

Evidence of learning: Students demonstrate proper professional processes; students balance competing interests and make justified recommendations; students apply fair process principles consistently.

A.3.9 Outcome 9: Collaborate Effectively in Teams

Typical expectation: Students work effectively in diverse teams, manage group dynamics, and contribute to collective outcomes.

AI Application (Chapter)	How It Supports This Outcome
Group Assessment Design (Ch. 15)	AI helps structure team contracts, role definitions, and accountability mechanisms.
Conversation Simulations (Ch. 7)	Students practise navigating team dynamics, giving feedback, and managing disagreement — skills that transfer to real group work.

AI Application (Chapter)	How It Supports This Outcome
Expert Panel Technique (Ch. 5)	Students learn to synthesise multiple perspectives, a core team collaboration skill.

Evidence of learning: Students contribute meaningfully to group processes; students navigate disagreement constructively; students demonstrate awareness of team dynamics.

A.4 Addressing Common Concerns

A.4.1 Concern: “Does AI integration lower academic standards?”

Response: AI integration *raises* standards by enabling assessment of higher-order skills. Instead of testing whether students can produce a case analysis (which AI can do), you test whether they can *evaluate, critique, and improve* a case analysis. This targets the top of Bloom’s taxonomy — evaluation and creation — rather than knowledge and application.

A.4.2 Concern: “How does this align with academic integrity policies?”

Response: The transparency model (Chapters 5 and 9) aligns with most institutional academic integrity frameworks by:

- Making AI use explicit and expected (not hidden)
- Requiring critical engagement with AI outputs (not passive acceptance)
- Assessing students’ thinking process (not just final products)
- Teaching professional ethics around technology use

This prepares students for professional practice where AI use is normal and expected, but accountability remains with the human professional.

A.4.3 Concern: “What evidence supports this pedagogical approach?”

Response: This approach is grounded in:

- **Experiential learning theory:** Students learn by doing, not just reading

- **Deliberate practice:** Multiple repetitions with feedback improve skill development
- **Reflective practice:** Self-assessment and metacognition enhance professional development
- **Authentic assessment:** Evaluating performance in realistic contexts predicts professional capability

AI enables scaling of pedagogical best practices that were previously limited by educator time and resources. See the Further Reading appendix for supporting research.

A.5 Implementation Roadmap

A.5.1 Short-Term (Current Semester)

1. Pilot 1–2 conversation simulations in units covering core discipline topics
2. Introduce the self-assessment tool for one existing assignment
3. Gather student feedback on AI-enhanced learning experiences

A.5.2 Medium-Term (Next Academic Year)

1. Implement AI-enhanced assignments across core discipline units
2. Develop a shared library of prompts and scenarios for programme consistency
3. Include AI literacy as an explicit learning objective in unit outlines
4. Run a faculty development workshop (use Appendix B)

A.5.3 Long-Term (2–3 Years)

1. Integrate virtual company simulations across multiple units (progression model)
 2. Partner with industry to ensure AI applications reflect current professional practice
 3. Track graduate outcomes: are AI-literate graduates more confident and capable?
 4. Share innovations with professional bodies and peer institutions
-

A.6 Conclusion

AI integration in business education is not about adopting technology for its own sake. It is about using available tools to better achieve existing

learning outcomes — to prepare confident, competent, ethical professionals who can navigate the complexity of modern workplaces.

Every application in this book has been designed to support common business programme learning outcomes. AI enhances pedagogical practice; it does not replace educational judgement or lower academic standards.

This appendix provides:

- **Institutional justification** — alignment with strategy and learning outcomes across disciplines
- **Pedagogical frameworks** — grounded in learning theory and discipline-specific practice
- **Practical tools** — ready-to-use prompts and assignments adaptable to your discipline
- **Implementation guidance** — start small, scale gradually
- **Academic integrity approaches** — transparency and critical engagement

The question is not whether AI belongs in business education. Given the professional reality that graduates will work in AI-augmented workplaces regardless of their discipline, the question is how to integrate AI responsibly and effectively into your teaching.

This book provides a starting point.

A.7 For Further Discussion

If you are a business educator interested in exploring AI integration:

- Start with the Introduction (understand the “why” for your discipline)
- Review the alignment matrix in this appendix (connect to your units and learning outcomes)
- Choose one small experiment from Chapter 3 or 4 (take a first step)
- Join colleagues in conversation about implementation — within your programme and across disciplines
- Adapt the discipline-specific examples throughout this book to your context

The future of business education includes AI. Your institution has the opportunity to lead in preparing professionals who are not just competent with technology, but ethically and critically engaged with it.

Appendix B

AI-Integrated Assessment Rubric System

B.1 What This System Does

This appendix provides a complete rubric generation pipeline: a generic template that works across disciplines, a structured guide for adapting it to your specific AI modality and discipline, and a worked example showing the full adaptation in practice.

The system is built on a single principle: **assess the quality of the student's thinking about AI, not the quality of the AI's output.** The rubric criteria operationalise the engagement spectrum described in the Assessment chapter, translating the conceptual distinction between collaborative thinking and pure delegation into markable performance levels.

B.1.1 Why Critical Engagement Is Weighted at 30%

This is the differentiator. Any tool can generate fast output. Critical evaluation, challenging AI assumptions, validating findings, recognising when AI is speculating, is the skill worth rewarding. If your rubric rewards only the final product, students who delegate effectively are indistinguishable from students who think deeply. If it rewards critical engagement, they are not.

B.1.2 How the Rubric Connects to Transcript Analysis

The transcript analysis metrics described in the Assessment chapter (Flesch readability scores, turn counts, prompt specificity over time, evidence of pushback) provide a scalable triage layer that sits underneath this rubric. The metrics help markers quickly identify which performance band a student is likely operating in before reading the submission in detail:

Transcript Signal	Likely Rubric Level
High turn count, increasing specificity, evidence of pushback	Excellent: genuine dialogue
Moderate turns, some follow-ups, reasonable prompt length	Good: reasonable engagement
Few turns, short prompts, limited follow-up	Satisfactory: limited interaction
Single prompts, no follow-up, very short	Developing: delegation, not conversation

The metrics are a triage tool, not a verdict. They tell the marker where to look; the rubric tells the marker what to assess.

B.2 The Generic Template

This is the starting point. Five criteria, four performance levels, adaptable to any AI-supported assessment. Start here, then customise using the adaptation guide that follows.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Critical Engagement with AI	30%	Demonstrates sophisticated evaluation of AI outputs. Clearly articulates when AI is reliable vs. speculative. Actively challenges AI assumptions and validates findings through independent verification.	Shows solid evaluation of AI outputs with some critical reflection. Generally identifies AI limitations. Attempts verification of key claims.	Engages with AI but evaluation is surface-level. Limited critical questioning. Minimal independent verification of AI suggestions.	Minimal or no critical evaluation of AI. Accepts AI outputs without question. No evidence of validation or verification.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Quality of Inquiry & Problem-Solving	25%	Investigation is strategic, well-structured, and adaptive. Clear evidence of iterative refinement. Responds effectively to obstacles or constraints. Shows intellectual curiosity and initiative.	Investigation is generally well-organised with some evidence of adaptation. Most key questions explored systematically. Some refinement visible.	Investigation covers main points but lacks depth or strategic planning. Limited evidence of adaptation to challenges. Inquiry feels mechanical.	Investigation is disorganised or incomplete. Minimal evidence of planning or adaptability. Key areas unexplored.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Conversation Quality & Accountability	20%	AI conversation transcript shows genuine dialogue: challenging assumptions, asking follow-ups, refining outputs, and steering toward specific context. Student can explain all decisions and takes clear accountability for final work. Evidence of conversation, not delegation.	Conversation transcript shows reasonable engagement with AI. Follow-up questions present. Generally explains how outputs were validated. Takes appropriate responsibility for work.	Conversation transcript shows limited interaction. Few follow-up questions or challenges. Some evidence of refinement but mostly accepts initial outputs. Accountability is unclear in places.	Conversation shows single prompts with no follow-up. No evidence of critical engagement or refinement. Unclear who is accountable for outputs.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Communication & Clarity	15%	Clear, coherent, professionally structured. Complex ideas are explained accessibly. Findings/recommendations are well-justified and easy to follow. Appropriate for intended audience.	Generally clear and well-organised. Most ideas are explained adequately. Findings supported by reasoning. Generally appropriate tone/style.	Communication is adequate but may lack clarity in places. Ideas present but not always well-connected. Some findings need stronger justification.	Communication is unclear or disorganised. Difficulty following main ideas. Minimal justification for conclusions.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Integration of Disciplinary Knowledge	10%	Meaningfully connects AI engagement to unit concepts. Demonstrates how theoretical/practical knowledge informs critical evaluation of AI. Shows synthesis of learning.	Connects AI work to unit content. Shows understanding of how discipline-specific knowledge applies to AI evaluation.	Makes some connections to unit content but links are basic or limited.	Few or no connections between AI work and disciplinary learning.

B.2.1 Mapping to the Engagement Spectrum

The **Conversation Quality & Accountability** criterion directly operationalises the engagement spectrum from the Assessment chapter:

Engagement Level	What It Looks Like in the Transcript	Rubric Level
Genuine collaborative thinking	Student drives inquiry, pushes back on AI, iterates toward own understanding	Excellent
Guided drafting	Student provides direction, evaluates critically, modifies toward coherent submission	Good

Engagement Level	What It Looks Like in the Transcript	Rubric Level
Curated delegation	Student uses AI to produce submission, exercises some judgment about what passes	Satisfactory
Pure delegation	Student hands task to AI, submits with minimal engagement	Developing

This mapping gives markers a concrete vocabulary for discussing student work. Rather than debating whether a submission “feels like AI,” the marker asks: where on the engagement spectrum does the evidence place this student?

B.2.2 General Scoring

- **Excellent (85–100%):** Sophisticated, independent thinking; strong evidence of learning outcomes
- **Good (75–84%):** Solid performance; demonstrates competence; minor areas for development
- **Satisfactory (65–74%):** Meets minimum expectations; completes task; some areas lack depth
- **Developing (<65%):** Below expectations; significant gaps in performance or understanding

B.2.3 Academic Integrity Notes

Make clear to students:

- AI use must be disclosed and documented
- Students remain responsible for all outputs, regardless of AI involvement
- Critical evaluation of AI is a valued skill, not “cheating”
- Honest AI conversation transcripts (including dead ends) demonstrate integrity, not weakness

For the broader integrity framework that supports this approach, including why transparency-based assessment is more defensible than detection-based assessment, see Section 8.2 and Section 8.13.

B.3 Adapting the Template

B.3.1 For Different AI Modalities

The generic template works as-is for most contexts, but the **Critical Engagement** criterion should be renamed and its performance descriptors adjusted to match the specific AI modality students are using.

B.3.1.1 1. AI Chatbot Simulations (e.g., Role-Playing Employees)

Best for: Business analysis, consulting, auditing, investigative work

Rename: “Critical Engagement with AI” → “Evaluation of AI-Generated Intelligence”

Excellent Performance Indicators:

- Distinguishes between reliable and speculative AI responses
- Adapts questioning strategy based on source credibility
- Recognises when AI is operating outside its domain
- Cross-validates information between multiple AI agents
- Shows sophisticated prompting (follow-ups, clarifications)

Example Rubric Cell (Excellent):

“Demonstrates sophisticated judgement about different AI agents. Clearly explains which sources are reliable for different topics and why. Strategically sequences conversations and asks targeted follow-ups. Cross-validates findings. Recognises when AI is speculating vs. reporting factual information.”

B.3.1.2 2. LLM as Writing/Research Assistant

Best for: Essays, reports, research papers, creative writing

Rename: “Critical Engagement with AI” → “Validation & Integration of AI-Generated Content”

Excellent Performance Indicators:

- Verifies factual claims made by LLM
- Identifies AI hallucinations or overgeneralisations
- Integrates AI suggestions into own argument, not replacing it
- Shows evidence of iterative refinement (multiple drafts, prompts)
- Clearly distinguishes AI-suggested ideas from own thinking

Example Rubric Cell (Excellent):

“Critically evaluates all LLM outputs before integration. Fact-checks key claims against sources. Uses AI for brainstorm-

ing/drafting but substantially refines all output. Clear evidence of multiple prompts and iterative refinement. Distinguishes own analysis from AI suggestions throughout.”

Example Rubric Cell (Developing):

“Accepts LLM output with minimal validation. No evidence that AI claims were fact-checked. Large portions appear to be unedited AI generation. Difficult to distinguish student thinking from AI output.”

B.3.1.3 3. RAG Systems or AI-Assisted Data Analysis

Best for: Data science, research, business analytics

Rename: “Critical Engagement with AI” → “Understanding & Validating AI Data Interpretation”

Excellent Performance Indicators:

- Understands what data/sources the AI is drawing from
- Questions AI interpretations; considers alternative explanations
- Validates AI findings against raw data when possible
- Recognises limitations of training data or AI knowledge cutoff
- Acknowledges uncertainty where appropriate

Weights might shift:

- Critical Engagement: 35% (data validation is crucial)
- Communication: 20% (must explain methodology clearly)
- Integration of Knowledge: 10% (same)

B.3.1.4 4. AI Code Assistant / Programming Helper

Best for: Computer science, software engineering

Rename: “Critical Engagement with AI” → “Code Validation & Testing of AI-Generated Solutions”

Excellent Performance Indicators:

- Tests all AI-generated code before integrating
- Understands logic of generated code; can explain it
- Recognises when AI solution is inefficient or incorrect
- Modifies/improves AI output rather than blindly using it
- Documents which parts were AI-assisted and which were manual

Weights might shift:

- Critical Engagement: 35% (testing/validation essential)
- Quality of Inquiry: 15% (different meaning in this context)

- Communication: 20% (code comments, documentation)
- Conversation Quality: 10% (how decisions were made)

B.3.1.5 5. Multimodal AI (Image Generation, Video Tools, etc.)

Best for: Creative fields, design, media studies

Rename: “Critical Engagement with AI” → “Creative Direction & Critical Evaluation of AI-Generated Media”

Excellent Performance Indicators:

- Clearly communicates intention to AI (via prompts, iterations)
- Recognises aesthetic/conceptual limitations of AI output
- Substantially modifies/refines AI output rather than using as-is
- Makes intentional creative decisions about when/how to use AI
- Conversation transcript shows iterative creative direction

B.3.1.6 6. Oral Exams / Viva Voce

Best for: High-stakes assessment verification, professional communication development, courses under approximately 25 students

Rename: “Critical Engagement with AI” → “Depth of Understanding & Ability to Defend”

Excellent Performance Indicators:

- Explains concepts clearly without relying on notes or memorised phrasing
- Responds to follow-up questions with nuance and relevant examples
- Connects ideas across multiple topics or readings
- Acknowledges limitations and considers alternative perspectives
- Communicates with confidence, clarity, and appropriate vocabulary

Additional Criteria to Consider:

Criterion	What to Assess
Understanding	Depth and accuracy of knowledge, ability to go beyond surface-level recall
Argument	Ability to articulate a position and respond to counter-arguments
Evidence	Relevance and accuracy of examples used to support claims
Structure & Coherence	Logical progression of ideas, accessibility to the listener

Speaking Skills

Clarity, eye contact, confidence,
vocabulary, minimal verbal clutter

Weights might shift:

- Ability to Defend: 35% (the core of the format)
- Quality of Reasoning: 25% (synthesis and argumentation)
- Evidence & Examples: 20% (supporting claims)
- Communication: 20% (verbal delivery and structure)

This modality connects directly to the Tier 1 oral checkpoint described in the Assessment chapter. Even a five-minute conversation after submission closes the most significant integrity gap.

B.3.2 For Different Disciplines

Adjust the **Integration of Disciplinary Knowledge** criterion to match your field:

Business/Management:

- How does this analysis align with strategic frameworks, ethical considerations?
- Add/adjust: “Strategic Application” criterion
- Does student consider stakeholder perspectives?

STEM/Sciences:

- How does methodology reflect scientific reasoning, experimental validity?
- Add/adjust: “Methodological Rigor” criterion
- Are limitations and uncertainties clearly articulated?

Humanities/Social Sciences:

- How does this argument engage with theoretical traditions, interpretive methods?
- Add/adjust: “Interpretive Depth” criterion
- Does student engage with multiple perspectives?

Law/Professional Practice:

- Are ethical implications considered? Is accountability clear?
- Add/adjust: “Professional Ethics & Accountability” criterion
- Recognises when AI judgement conflicts with professional standards?

B.3.3 Quick Reference: Weighting by Context

Context	Critical Engagement	Inquiry Quality	Con- versa- tion Qual- ity	Commu- nication	Disciplinary Depth
Chat- bot Simu- lation	35%	25%	15%	15%	10%
LLM Writ- ing Tool	30%	15%	20%	20%	15%
Data Anal- ysis	35%	20%	15%	20%	10%
Code/Pr gram- ming	35%	15%	10%	20%	20%
Cre- ative/Me- dia	30%	20%	20%	15%	15%
Oral Exam/Viva	35%	25%	—	20%	20%
<i>Generic Tem- plate</i>	<i>30%</i>	<i>25%</i>	<i>20%</i>	<i>15%</i>	<i>10%</i>

B.4 Worked Example: CloudCore Audit Rubric

This example shows the generic template fully adapted for a specific assessment: an information security audit where students interact with AI “employees” (chatbots) and use LLMs as research tools. It demonstrates the adaptation process in practice.

Context: Students conduct an audit of a simulated AI-driven organisation by interacting with AI “employees” (chatbots) and using LLMs as research tools. This rubric balances investigation quality, critical AI evaluation, and professional reporting.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Critical Evaluation of AI Sources	30%	Demonstrates sophisticated judgement about different AI “employees.” Clearly explains which sources are reliable for different questions and why. Questions inconsistencies between AI responses. Validates findings through multiple sources or independent verification. Recognises AI speculation vs. factual responses.	Shows good discernment among AI employees. Generally identifies which sources are stronger for different topics. Attempts to cross-reference findings. Recognises some AI limitations.	Engages with multiple AI employees but evaluation is surface-level. Limited comparison between sources. Accepts contradictions without investigation. Minimal awareness of AI unreliability.	Treats all AI employees as equally reliable. No attempt to validate information across sources. Accepts AI responses uncritically.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Audit Investigation Strategy & Execution	25%	Investigation plan is systematic and adaptive. Clear evidence of strategic sequencing of AI interviews (e.g., starting with CFO for financials, then IT Manager for technical risks). Responds creatively to access constraints and cancellations. Shows persistence in clarifying unclear answers. Iterative refinement visible.	Investigation follows a logical structure. Manages scheduled appointments effectively. Asks follow-up questions when unclear. Generally adapts to constraints. Some evidence of iterative inquiry.	Investigation covers main areas but lacks clear strategy. Accepts initial AI responses without deep follow-up. Limited adaptation to access challenges. Feels somewhat ad-hoc.	Investigation is disorganised or incomplete. Minimal engagement with access scheduling. Key areas unexplored or surface-level inquiry.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Professional Use of LLM as Research Tool	20%	Critically reviews LLM outputs before integrating into audit. Clearly distinguishes own analysis from AI-generated content. Uses LLM strategically (e.g., drafting frameworks, checking logic) rather than passively accepting output. Shows evidence of multiple prompts/refinement. Acknowledges where LLM was helpful or speculative.	Generally validates LLM suggestions against evidence. Most AI outputs are appropriately reviewed. Clear attribution of own work. Some evidence of iterative prompting.	Integrates LLM output with limited critical review. Attribution of sources is sometimes unclear. Limited evidence of validation or iterative refinement.	Heavily relies on unvalidated LLM output. No clear distinction between own and AI-generated content. Minimal critical review.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
Audit Report: Professional Communication	15%	Report is clear, well-structured, and professionally formatted. Complex security findings are explained accessibly to non-technical stakeholders. Recommendations are specific and justified. Transparent about methodology (how AI was used, limitations encountered).	Report is well-organised and clearly written. Findings are generally well-explained. Recommendations are reasonable. Mentions methodology but could be more transparent.	Report covers main findings but organisation could be clearer. Some findings lack sufficient justification. Recommendations are general. Limited transparency about AI use in methodology.	Report is difficult to follow or incomplete. Findings poorly explained. Recommendations are vague or unjustified. No mention of methodology or AI use.

Criterion	Weight	Excellent	Good	Satisfactory	Developing
AI Conversation Quality & Professional Judgment	10%	Conversation transcripts show genuine dialogue with AI: challenging responses, following up on inconsistencies, refining questions based on previous answers. Student clearly articulates how own judgement validated, rejected, or refined AI suggestions. Takes clear accountability for final audit conclusions.	Conversation transcripts show reasonable engagement. Follow-up questions present. Explains how own expertise was applied to evaluate AI outputs. Generally accountable for work product.	Conversation transcripts show limited interaction. Few follow-up questions. Limited evidence of steering or refining the AI conversation. Accountability is somewhat unclear.	Conversation shows single prompts with no follow-up. No evidence of critical engagement. Unclear who is responsible for audit conclusions.

B.4.1 Performance Level Descriptors

Excellent (85–100%):

- Student demonstrates sophisticated professional judgement about which AI sources are credible for different investigation areas
- Audit strategy is evident and strategic; clear sequencing of inquiry
- All AI outputs are meaningfully evaluated; student shows independent reasoning throughout
- Report would be valuable to actual security leadership
- Conversation transcripts show genuine learning about managing AI as a tool

Good (75–84%):

- Student shows competent judgement about AI reliability; identifies key differences between sources
- Investigation is reasonably strategic and organised
- Most AI outputs are reviewed; some independent analysis visible
- Report is professionally written and would be useful to stakeholders
- Conversation shows adequate critical engagement with AI

Satisfactory (65–74%):

- Student engages with AI but evaluation lacks sophistication; treats sources somewhat generically
- Investigation covers main areas but strategy is not always clear
- AI outputs are sometimes accepted with limited review
- Report communicates findings but could be more polished
- Conversation engagement is present but limited in depth

Developing (<65%):

- Student does not meaningfully evaluate AI reliability; little critical engagement
- Investigation is incomplete or disorganised; key areas missing
- AI outputs heavily influence conclusions with minimal independent review
- Report is unclear or incomplete
- Conversation engagement is minimal or absent

B.4.2 Usage Notes for Markers

Before Students Start:

- Clearly explain that this rubric rewards *critical thinking about AI*, not just efficient tool use
- Emphasise that asking good follow-up questions of AI employees (and the LLM) is an evaluated skill

- Normalise that showing where AI misled you (visible in the transcript) is valued equally to successful findings

During Marking:

- Look for evidence in the audit transcript: Can you see quality prompting? Validation? Follow-ups?
- Check the LLM transcript (if submitted): Does it show iterative refinement or one-shot prompting?
- Read the conversation transcript carefully: This is where students reveal their critical judgement. Look for conversation, not delegation
- Use the transcript analysis metrics from the Assessment chapter as a first-pass triage before reading in detail

Adapted for Other Audit Contexts:

This rubric works well for ISO 27001 readiness audits, tech vendor assessments, process efficiency reviews, or any scenario where students interview AI agents and synthesise findings. Adjust the specific domain knowledge criterion and the types of findings expected.

B.5 Creating Your Own Variant

B.5.1 Step 1: Answer These Questions

- What AI tool(s) will students actually use?
- What matters most in your discipline?
- What would a professional in your field do with AI?
- What skills do you want students to develop?
- Where on the engagement spectrum (see the Assessment chapter) do you want students operating, and how will you know?

B.5.2 Step 2: Customise in This Order

1. Rename criteria to match actual activities
2. Adjust weights based on importance (use the Quick Reference table as a starting point)
3. Rewrite performance descriptors using examples from your field
4. Add discipline-specific criterion if the generic five do not capture what matters
5. Add any specific constraints or requirements (e.g., “code must be tested,” “citations required”)

B.5.3 Step 3: Share Early with Students

- Include with assignment briefing

- Walk through an “Excellent” example specific to your unit
- Clarify what you mean by critical engagement in your context
- Share the rubric **before** students start work, not after

B.5.4 Step 4: Test and Calibrate

- Apply your version to 2–3 sample student submissions (real or imagined)
 - Calibrate scoring with colleagues to ensure consistency
 - Collect student feedback: Was the rubric clear? Did it help them understand expectations?
 - Build a portfolio of marked examples for future years
-

B.6 Implementation Checklist

When adapting for your specific assessment:

- Identify primary AI tool/modality being used
 - Select relevant modality section above
 - Rename the “Critical Engagement” criterion appropriately
 - Adjust weights based on context (use Quick Reference as guide)
 - Customise performance level descriptors with your specific assignment details
 - Add discipline-specific criterion if needed
 - Test rubric with 2–3 sample student submissions
 - Calibrate scoring with colleagues (ensure consistency)
 - Share rubric with students before assignment begins
-

B.7 Common Mistakes to Avoid

- **Don’t** make the rubric so specific it cannot be adapted. **Do** include adaptability notes and examples.
- **Don’t** ignore the discipline (one-size-fits-all rarely works). **Do** customise weights and descriptors to your field.
- **Don’t** focus only on tool efficiency (faster is not better). **Do** reward critical thinking about AI, not speed of output.
- **Don’t** punish AI use and incentivise it being hidden. **Do** make transparency and critical evaluation rewarded behaviours.
- **Don’t** assume all students know how to “prompt well.” **Do** teach and demonstrate effective AI engagement first. This is especially important if the process mark carries significant weight; see the AI literacy prerequisite discussed in the Assessment chapter.

B.8 Questions to Ask Before Finalising

- Will this rubric encourage the thinking you want?
 - Can colleagues easily adapt this for their units?
 - Does it make clear what “critical engagement with AI” means in your context?
 - Have you tested it with actual student work?
 - Is the language accessible to students (or does it need translation)?
 - Does it align with the engagement spectrum? Can a marker use it to place a student on that spectrum from the evidence?
-

B.9 Further Reading

- Dawson, P. (2017). Assessment rubrics: Towards clearer and more replicable design, research and practice. *Assessment & Evaluation in Higher Education*, 42(3), 347–360.
- Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching and Learning Practice*, 21(6).
- Boud, D., Ajjawi, R., Dawson, P., & Tai, J. (Eds.). (2018). *Developing evaluative judgement in higher education: Assessment for knowing and producing quality work*. Routledge.
- Bearman, M., Tai, J., Dawson, P., Boud, D., & Ajjawi, R. (2024). Developing evaluative judgement for a time of generative artificial intelligence. *Assessment & Evaluation in Higher Education*, 49(6), 893–905.
- Villarroel, V., Bloxham, S., Bruna, D., Bruna, C., & Herrera-Seda, C. (2018). Authentic assessment: Creating a blueprint for course design. *Assessment & Evaluation in Higher Education*, 43(5), 840–854.
- Swiecki, Z., Khosravi, H., Chen, G., Martinez-Maldonado, R., Lodge, J. M., Milligan, S., Selwyn, N., & Gasevic, D. (2022). Assessment in the age of artificial intelligence. *Computers and Education: Artificial Intelligence*, 3, Article 100075.
- Mollick, E. R., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts. *arXiv preprint*.

Appendix C

Assessment AI Stress Test: A Prompt Sequence

Works with: Claude, ChatGPT, Microsoft Copilot, or any general-purpose AI tool **Time required:** 15–20 minutes

C.1 Why This Matters Now

AI tools available to students today do not just answer questions. They draft code, write essays, synthesise research, generate reflections, and when loaded with unit materials via tools like NotebookLM (a free tool that lets users upload documents and then ask questions across them), they can cross-reference lecture content, cite specific weeks, and produce output that reads as genuinely informed. This is not a future risk. It describes how sophisticated students are already working.

The stress test is not designed to generate alarm. It is designed to give you an honest picture of where your assessment is robust and where it is not, so that any changes you choose to make are grounded in evidence rather than anxiety.

C.2 A Note on Human Agency

This sequence is designed to support your thinking, not replace it. The AI will give you a structured analysis but it does not know your students,

your context, or what you are actually trying to teach. Push back on outputs that do not ring true. Add your own knowledge where the AI misses nuance. Stop mid-sequence and follow a thread that matters to you.

The goal is not to get a clean report that says the assessment is safe or unsafe. The goal is to surface questions worth thinking about. A lecturer who has genuinely interrogated their own assessment design will make better decisions than one who outsourced that interrogation and accepted the result.

The AI is the analyst. You are the assessor.

A note on format: this is a prompt sequence rather than an app by design. An app would hide the reasoning, encourage dependency, and remove the flexibility to adapt mid-sequence when your assessment has a nuance the template did not anticipate. Running the prompts manually means you see exactly what is being asked, can modify a prompt on the fly, and are never just ratifying an output you did not follow. The prompts are meant to be read, not just run.

C.3 How to Use This

Work through the prompts in order in a single conversation session. Each prompt builds on the previous response so do not start a new chat between steps. Copy each prompt, paste it into your AI tool, then read the response before moving to the next one. Add your own follow-up questions at any point.

These prompts work for any assessment type: essays, literature reviews, case studies, reports, presentations, reflective journals, programming projects, or any combination. The sequence adapts based on what you paste in.

Short on time? If you only have 10 minutes, run Prompts 1, 2, 4, and 8. You will get the most important findings: what is most completable, where the weakest integrity point is, and whether the assessment holds up if AI use is fully declared. The other prompts add depth but those four give you the essential picture.

C.4 Prompt 1: Set the Context

You are helping me stress test a university assessment from the perspective of a student who uses AI tools strategically. Your job is not to judge the assessment but to give me an honest picture of where AI can and cannot contribute to a student submission.

Assume the student has access to:

- A general-purpose AI assistant (Claude, ChatGPT, or similar)
- NotebookLM loaded with their own research infrastructure: lecture slides, weekly readings, marking rubrics, past examples, and anything they have found on the topic. Sophisticated students maintain a notebook per unit and a notebook per assignment, effectively giving them a personalised tutor with full context of everything they have consumed
- Coding assistants if the assessment involves programming

Before I paste the assessment, I want you to keep one thing in mind: the question is not just whether AI can produce the final artefact, but whether a student could use AI to replicate the entire process of getting there, the research, the synthesis, the argument development, the refinement.

I will paste the assessment specification in the next message. Please confirm you understand the task before I do.

C.5 Prompt 2: Paste the Assessment

Here is the assessment specification. Please read it carefully before I ask you any questions.

[PASTE YOUR FULL ASSESSMENT SPECIFICATION HERE]

Once you have read it, give me a brief one-paragraph summary of what the assessment is asking students to produce, just to confirm you have understood it correctly.

C.6 Prompt 2b: Learning Outcome Anchor

Before we analyse the deliverables, here are the intended learning outcomes for this unit or assessment:

[PASTE YOUR LEARNING OUTCOMES HERE]

For each learning outcome, tell me:

1. Is this outcome most at risk of being *simulated* (AI produces something that looks like the outcome has been achieved without the student actually developing the capability) or *achieved* (the student genuinely develops the capability regardless of AI use)?
2. What would a submitted artefact look like if the outcome was simulated rather than achieved?

This is the most important analytical step. An assessment that produces convincing simulations of its own learning outcomes is not assessing what it claims to assess.

C.7 Prompt 3: Completability Analysis

For each deliverable in the assessment, tell me:

1. What percentage of it could AI draft effectively with minimal student effort?
2. What would the student still need to contribute themselves?

Be honest and specific. If a section is almost entirely AI-completable, say so. Present your response as a table with three columns: Deliverable, AI Completability, and What the Student Still Needs to Do.

C.8 Prompt 4: Integrity Analysis

Now identify:

1. The single weakest integrity point in this assessment. Where is a student most likely to submit AI-generated work with minimal personal engagement?
2. The single strongest integrity point. What is the one component that most requires genuine student presence

or understanding? There may be no fully robust point. If that is the case, say so explicitly rather than identifying the least weak option as though it were strong.

3. Any commonly assumed safeguards that do not actually hold up under scrutiny. For example, is “personal authenticity” genuinely verifiable at scale?

Keep this section honest and practical. I am not looking for reassurance.

C.9 Prompt 5: The Engagement Spectrum

A student working the system could pass this assessment using AI with minimal personal engagement. Describe the spectrum of how a student might actually use AI on this assessment, from pure delegation at one end to genuine collaborative thinking at the other.

For each point on the spectrum, describe what the student’s behaviour looks like and what learning, if any, is still happening. The goal is to help me understand where the real learning failure is, and how likely it is.

For more on the engagement spectrum as a conceptual framework for assessment design, see the Assessment chapter.

C.10 Prompt 6: Practical Suggestions

Based on everything so far, give me three to five practical suggestions for strengthening this assessment against low-engagement AI use.

For each suggestion:

- Keep it realistic for a class of 50 to 100 students
- Note the approximate marking overhead it adds
- Flag if it requires any infrastructure or setup

Frame these as options to consider, not a list of problems to fix. The assessment may be releasing soon and a full redesign is not on the table.

C.11 Prompt 7: Produce the Report

Now pull everything together into a short structured report I can share with a colleague. Use these headings exactly:

- Overview
- What AI Can Draft Effectively
- What AI Substitutes Less Easily (and why that is not the same as safe)
- Key Integrity Findings
- Practical Suggestions

Tone should be collegial and constructive. This is feedback from one colleague to another, not an audit. Keep the whole report concise, ideally no more than two pages for a complex assessment. Avoid bullet point lists inside sections where prose reads more naturally.

C.12 Prompt 8: The Honesty Test

This is the capstone prompt. Run it last.

Imagine a student submitted this assessment alongside a statement saying: *“I used AI extensively to research, draft, refine, rehearse, and structure this work.”*

Would the assessment still meaningfully discriminate between a student with superficial understanding and one with deep understanding?

Answer Yes, Partially, or No, and explain why. If the answer is Partially or No, describe what would need to change about the assessment design for the answer to become Yes.

This prompt does three things. It collapses the AI-detection framing entirely, because it asks whether the assessment holds up even when AI use is fully declared. It aligns the analysis with a transparency-based approach rather than a policing one. And it produces a conclusion that is concrete enough to bring into an assessment design conversation or review panel.

If the answer is Yes with strong reasons, the assessment is well-designed for the AI era. If the answer is Partially or No, that is the most useful finding the whole sequence produces.

C.13 Optional Follow-up Prompts

Use these at any point if you want to go deeper on a specific area.

On collusion:

Does this assessment create conditions that encourage collusion? How does that compare to an assessment that requires students to submit their AI conversation transcripts as part of the work?

On the video component (if applicable):

The assessment includes a face-on-camera video component. How robust is this as evidence of independent reasoning, as distinct from its learning value? Rehearsing a scripted explanation has genuine learning value but is different from demonstrating unrehearsed understanding. Which of those two things does this video requirement actually assess, and how would you know the difference?

On GitHub commit history (if applicable):

The assessment uses GitHub for submission. What would you look for in the commit history to get a sense of genuine iterative engagement versus a last-minute submission?

On reframing for staff:

Draft a short paragraph I could use to explain to a colleague why assessing how students use AI is more useful than trying to detect whether they used it.

On red-teaming the assessment:

Write a 200-word sample of what a low-effort, high-AI submission would actually look like for the hardest or most important component of this assessment. Do not label it as AI-generated. Just produce it as a student might submit it.

Then tell me: would this sample pass? What would a marker need to look for to distinguish it from a genuinely strong submission?

This optional prompt tends to produce a concrete realisation for staff about the gap between what AI produces and what markers currently look for. Seeing the output is more persuasive than any amount of argument about AI-compleatability.

On the marks split:

If I wanted to shift the weighting of this assessment so that the process (how the student used AI to develop their work) carries more marks than the final artefact, what would that look like in practice? Suggest a marks split and describe what the process component would require students to submit and how a marker would assess it at scale.

For more on the process-over-product marks split, see the Assessment chapter.

On NotebookLM as student research infrastructure:

A student maintains a NotebookLM notebook for this unit loaded with lecture slides, readings, and the marking rubric, and a separate notebook for this specific assignment loaded with everything they have found on the topic. How does this change your completability analysis? What components, if any, become harder for AI to complete when the student has built this kind of personalised research context?

For more on why this matters for assessment design, see the Assessment chapter.

On essays and literature reviews:

This assessment includes a written component. How credible would an AI-assisted essay or literature review look to a marker? What specific markers of genuine critical engagement should a rubric reward that AI is less likely to produce convincingly?

On reflective writing:

Reflective journals and personal learning logs feel resistant to AI but are among the most completable formats once a student provides context. How would you assess the reflective component of this assignment differently to reward genuine reflection over AI-simulated reflection?

This prompt sequence was developed through the SoMM AI Facilitator role. Share freely. Adapt for your own context.

Appendix D

Glossary

This glossary provides definitions for key terms and concepts used throughout the book. Terms are listed alphabetically.

D.1 A

AI (Artificial Intelligence): Software systems that can perform tasks that typically require human intelligence, such as pattern recognition, decision-making, language understanding, and problem-solving.

AI Ethics: The study of moral principles and guidelines for the responsible development and use of artificial intelligence systems.

AI Literacy: The ability to understand, evaluate, and effectively use AI systems and their outputs.

Assessment Rubrics: Structured evaluation criteria used to assess student work against specific learning outcomes and standards.

D.2 B

Bias in AI: Systematic errors in AI systems that result in unfair or discriminatory outcomes, often due to biased training data or algorithmic design.

D.3 C

Chain of Thought: A prompting technique where AI is guided to show its reasoning step-by-step, rather than jumping directly to a final answer.

CRAFT Framework: A structured approach to writing effective prompts, consisting of: - **C:** Context (background information) - **R:** Role (AI persona to adopt) - **A:** Action (specific task to perform) - **F:** Format (desired output structure) - **T:** Tone/Target (intended audience and style)

Critical Engagement: The practice of actively questioning, evaluating, and critiquing AI outputs rather than accepting them passively.

D.4 D

Deep Learning: A subset of machine learning using neural networks with multiple layers to process complex patterns and data.

dialogue-Based Assessment: Assessment methods that evaluate student thinking through interactive conversations rather than static products.

D.5 E

Evidence-Based Practice: Professional decision-making grounded in research, data, and systematic evaluation rather than intuition alone.

D.6 F

Few-Shot Learning: AI's ability to learn and perform tasks from just a few examples, rather than requiring extensive training data.

D.7 G

Generative AI: AI systems that can create new content, such as text, images, or code, rather than just analysing existing data.

D.8 H

Hallucination: When AI generates false or misleading information confidently, as if it were factual.

Human Oversight: The practice of humans reviewing, validating, and intervening in AI processes to ensure accuracy and ethical outcomes.

D.9 I

Iterative Refinement: The process of repeatedly improving AI outputs through feedback and revision cycles.

D.10 L

Large Language Models (LLMs): Advanced AI models trained on vast amounts of text data to understand and generate human-like language. Examples include ChatGPT, Claude, and Gemini.

Learning Outcomes: Specific statements describing what students should know, understand, or be able to do after completing a learning experience.

D.11 M

Machine Learning: A type of AI where systems learn patterns from data and improve performance without being explicitly programmed for each task.

Meta-Prompting: Using AI to help you create better prompts for AI, essentially using AI to improve your AI interactions.

Metacognition: Awareness and control of one's own learning processes, including planning, monitoring, and evaluating learning strategies.

D.12 P

Process-Based Assessment: Evaluation methods that focus on how students think and work through problems, rather than just the final product or answer.

Product-Based Assessment: Traditional evaluation methods that focus primarily on the final output or result, rather than the thinking process.

Prompt Engineering: The practice of crafting effective instructions (prompts) to get desired outputs from AI systems.

Prompting: The act of providing instructions or questions to AI systems to elicit specific responses or behaviours.

D.13 R

Retrieval-Augmented Generation (RAG): A technique where AI combines its training knowledge with real-time data retrieval to provide more accurate and current information.

Rubrics: Detailed scoring guides that specify criteria for different levels of performance on assessment tasks.

D.14 S

Scaffolding: Educational support structures that help students achieve tasks they couldn't accomplish independently, gradually removed as competence develops.

Self-Assessment: The process where students evaluate their own work and learning progress against established criteria.

D.15 T

Transparency Model: An approach to AI integration where students openly acknowledge AI use, submit their interaction history, and critically evaluate AI outputs.

Transfer Learning: AI's ability to apply knowledge learned from one task to perform well on related tasks.

D.16 V

Virtual Company: Simulated business environments created through AI conversations, allowing students to practice professional scenarios safely.

/Users/michael/Projects/ai-as-a-teaching-partner/glossary.qmd

Appendix E

Further Reading

This book is practitioner-focused, but its ideas are grounded in research. The references below point to the studies, papers, and books behind the key claims. They are organised by topic so you can follow up on whatever interests you most.

This is not a comprehensive literature review. It is a trail of breadcrumbs for curious readers.

E.1 Understanding AI and Large Language Models

Supporting Chapters 1–2: What Is AI and What Are Large Language Models

How LLMs work (prediction as the core mechanism):

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, 30, 5998–6008. The foundational paper introducing the transformer architecture that underpins all modern LLMs.
- Bommasani, R., Hudson, D. A., Adeli, E., ... & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*. Defines the category of “foundation models” and maps their capabilities, risks, and societal implications.
- Brown, T. B., Mann, B., Ryder, N., ... & Amodei, D. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 33, 1877–1901. The GPT-3 paper

demonstrating that giving a model a few examples in the prompt dramatically improves task performance.

Hallucination:

- Ji, Z., Lee, N., Frieske, R., ... & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), Article 248. A comprehensive overview of why LLMs generate plausible but false content.

Bias and fairness:

- Gallegos, I. O., Rossi, R. A., Barrow, J., ... & Ahmed, N. K. (2024). Bias and fairness in large language models: A survey. *Computational Linguistics*, 50(3), 1097–1179. How biases in training data manifest in LLM outputs, and the limitations of current mitigation approaches.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of FAccT '21*, 610–623. ACM. How training data biases propagate through LLMs.

Deep learning foundations:

- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. The landmark review providing accessible context for the layered pattern recognition that makes LLMs possible.
- Zhao, W. X., Zhou, K., Li, J., ... & Wen, J.-R. (2023). A survey of large language models. *arXiv preprint arXiv:2303.18223*. A comprehensive survey of the LLM landscape.

E.2 Prompt Engineering and Structured Communication

Supporting Chapters 3–6: Getting Started, First Steps, Seven Techniques, and Managing Context

Structured prompts outperform unstructured ones:

- Sahoo, P., Singh, A. K., Saha, S., Jain, V., Mondal, S., & Chadha, A. (2024). A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv preprint arXiv:2402.07927*. The broader landscape in which frameworks like CRAFT, RTCF, and CO-STAR sit.
- Federiakin, D., Molerov, D., Zlatkin-Troitschanskaia, O., & Maur, A. (2024). Prompt engineering as a new 21st century skill. *Frontiers*

in *Education*, 9, 1366434. Makes the case that structured prompting is a transferable professional skill, not a niche technical ability.

- Knoth, N., Tolzin, A., Janson, A., & Leimeister, J. M. (2024). AI literacy and its implications for prompt engineering strategies. *Computers and Education: Artificial Intelligence*, 6, 100225.

Chain-of-thought and reasoning:

- Wei, J., Wang, X., Schuurmans, D., ... & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, 35, 24824–24837. The paper that formalised “show your working” as a prompting strategy.
- Wang, X. et al. (2023). Self-consistency improves chain of thought reasoning in language models. *ICLR 2023*. Sampling multiple reasoning paths and selecting the most consistent answer improves accuracy.

Task decomposition and prompt chaining:

- Zhou, D. et al. (2023). Least-to-most prompting enables complex reasoning in large language models. *ICLR 2023*. Breaking complex problems into sequential subproblems significantly improves accuracy.

RE2 (Re-Reading) prompting:

- Xu, Y. et al. (2024). Re-Reading improves reasoning in large language models. *ACL EMNLP*. Repeating a question in the prompt creates pseudo-bidirectional attention, improving reasoning accuracy. Twice is the sweet spot.

Iterative refinement over single-shot prompting:

- Madaan, A. et al. (2023). Self-Refine: Iterative refinement with self-feedback. *NeurIPS 2023*. Iterative refinement consistently outperforms single-pass generation.

Teaching strategies with AI prompts:

- Mollick, E. R., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts. *arXiv preprint*. Seven structured approaches to using AI for learning, including role play, debate, and self-testing.
- Mollick, E. R., & Mollick, L. (2023). Using AI to implement effective teaching strategies in classrooms: Five strategies, including prompts. *SSRN*.

- Mollick, E. R., & Mollick, L. (2024). Instructors as innovators: A future-focused approach to new AI learning opportunities, with prompts. *SSRN*.

E.3 Critical Evaluation and Staying Sceptical

Supporting Chapters 8 and 9: Critique Toolkit and Ethics, Data Governance & Integrity

Sycophancy in LLMs:

- Sharma, M. et al. (2023). Towards understanding sycophancy in language models. *arXiv preprint*. How LLMs systematically tailor responses to match user beliefs, even when those beliefs are incorrect.
- Perez, E. et al. (2023). Discovering language model behaviors with model-written evaluations. *ACL 2023*. Evidence of sycophantic behaviour across multiple model families and scales.

The AI Dismissal Fallacy:

- Claessens, S., Veitch, P., & Everett, J. A. C. (2026). Negative perceptions of outsourcing to artificial intelligence. *Computers in Human Behavior*, 177, 108894. People systematically devalue work when they learn AI was involved.

Information literacy and lateral reading:

- Wineburg, S. et al. (2022). Lateral reading and the nature of expertise. *Teachers College Record*. Experts verify claims by checking sources laterally rather than reading vertically.

E.4 Cognitive Offloading and AI Dependency

Supporting the book's core argument: partner, don't delegate

Cognitive offloading:

- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. How humans use external tools to reduce cognitive demand, and when this helps versus hinders learning.
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our finger-

tips. *Science*, 333(6043), 776–778. Access to searchable information changes what we bother to remember.

- Hooper, V. J. (2025). Cognitive offloading and the reshaping of human thought: The subtle influence of artificial intelligence. *Revista de Pensamiento y Cultura (Colloquia)*, 12, 1–14.

The generation effect (producing information improves retention):

- Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. *Journal of Experimental Psychology: Human Learning and Memory*. Actively generating information leads to better memory than passively receiving it.

AI and learning outcomes:

- Bastani, H., Bastani, O., Sungu, A., ... & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. Students using AI without guardrails perform worse on subsequent unaided tasks.

Metacognitive laziness:

- Fan, Y., Tang, L., Le, H., ... & Gasevic, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2), 489–530.
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), Article 6.

Cognitive surrender:

- Shaw, S. D., & Nave, G. (2026). Thinking — fast, slow, and artificial: How AI is reshaping human reasoning and the rise of cognitive surrender. Working paper, The Wharton School.

AI does not reduce work:

- Ranganathan, A., & Ye, X. M. (2026, February 9). AI doesn't reduce work — it intensifies it. *Harvard Business Review*.

E.5 Assessment Design and Academic Integrity

Supporting Chapters 10–11, 15, and 18: Process Assessment, Self-Assessment, Group Assessment, and Assessment Design

Assessment reform for the AI era:

- Lodge, J. M., Howard, S., Bearman, M., Dawson, P., & Associates. (2023). *Assessment reform for the age of artificial intelligence*. TEQSA. The discussion paper that framed the Australian higher education response.
- Swiecki, Z., Khosravi, H., Chen, G., ... & Gasevic, D. (2022). Assessment in the age of artificial intelligence. *Computers and Education: Artificial Intelligence*, 3, Article 100075.
- Corbin, T., Dawson, P., & Liu, D. (2025). Talk is cheap: Why structural assessment changes are needed for a time of GenAI. *Assessment & Evaluation in Higher Education*, 50(7), 1087–1097.
- Perkins, M., & Roe, J. (2025). The end of assessment as we know it: GenAI, inequality and the future of knowing. In *AI and the future of education: Disruptions, dilemmas and directions*, 76–80.

Validity over detection:

- Dawson, P., Bearman, M., Dollinger, M., & Boud, D. (2024). Validity matters more than cheating. *Assessment & Evaluation in Higher Education*, 49(7), 1005–1016.
- Corbin, T., Dawson, P., Nicola-Richmond, K., & Partridge, H. (2025). ‘Where’s the line? It’s an absurd line’: Towards a framework for acceptable uses of AI in assessment. *Assessment & Evaluation in Higher Education*, 50(5), 705–717.

Rubric design:

- Dawson, P. (2017). Assessment rubrics: Towards clearer and more replicable design, research and practice. *Assessment & Evaluation in Higher Education*, 42(3), 347–360.
- Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching and Learning Practice*, 21(6).

Evaluative judgement:

- Bearman, M., Tai, J., Dawson, P., Boud, D., & Ajjawi, R. (2024). Developing evaluative judgement for a time of generative artificial intelligence. *Assessment & Evaluation in Higher Education*, 49(6), 893–905.
- Boud, D., Ajjawi, R., Dawson, P., & Tai, J. (Eds.). (2018). *Developing evaluative judgement in higher education: Assessment for knowing and producing quality work*. Routledge.

Authentic assessment:

- Villarroel, V., Bloxham, S., Bruna, D., Bruna, C., & Herrera-Seda, C. (2018). Authentic assessment: Creating a blueprint for course design. *Assessment & Evaluation in Higher Education*, 43(5), 840–854.

Retrieval practice:

- Roediger, H. L., & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences*. Testing yourself improves learning more than re-reading.

Oral examination as alternative assessment:

- Hartmann, C. (2025). Oral exams for a generative AI world: Managing concerns and logistics for undergraduate humanities instruction. *College Teaching*.
- Buehler, M. J., & Schneider, L. U. (2009). Speak up! Oral examinations and political science. *Journal of Political Science Education*, 5(4), 315–331.

Bloom's taxonomy and cognitive levels:

- Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives* (Complete ed.). Longman.
- Webb, N. L. (2002). *Depth-of-knowledge levels for four content areas*. Wisconsin Center for Education Research.

E.6 AI in Higher Education Policy and Practice

Supporting Chapters 9, 13, and 20–21: Ethics, Unit Design, Accessibility, and Global Perspectives

Institutional frameworks:

- Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20, 38.
- UNESCO. (2023). *Guidance for generative AI in education and research*.
- Sabzalieva, E., & Valentini, A. (2023). *ChatGPT and artificial intelligence in higher education: Quick start guide*. UNESCO IESALC.
- Russell Group. (2023). *Russell Group principles on the use of generative AI tools in education*.

- European Parliament & Council of the European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*.
- OECD. (2025). *Empowering Learners for the Age of AI: An AI Literacy Framework*.

AI and pedagogy:

- Bearman, M., & Ajjawi, R. (2023). Learning to work with the black box: Pedagogy for a world with artificial intelligence. *British Journal of Educational Technology*, 54(5), 1160–1173.
- Kasneci, E., Sessler, K., Kuchemann, S., ... & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, Article 102274.
- Labadze, L., Grigolia, M., & Machaidze, L. (2023). Role of AI chatbots in education: Systematic literature review. *International Journal of Educational Technology in Higher Education*, 20, Article 56.

Critical AI literacy:

- Roe, J., Furze, L., & Perkins, M. (2025). Digital plastic: A metaphorical framework for Critical AI Literacy in the multiliteracies era. *Pedagogies: An International Journal*.
- Madsen, D. O., & Puyt, R. W. (2025). When AI turns culture into slop. *AI & Society*.

E.7 Human-AI Collaboration and the Future of Work

Supporting Chapters 7, 12, and 17: Flight Simulator, Virtual Company, and Advanced Frontiers

- Wilson, H. J., & Daugherty, P. R. (2018). Collaborative intelligence: Humans and AI are joining forces. *Harvard Business Review*, 96(4), 114–123. The greatest performance gains come from structured human-AI collaboration, not AI alone.
- Dellermann, D., Ebel, P., Sollner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637–643.
- Mosqueira-Rey, E., Hernandez-Pereira, E., Alonso-Rios, D., Bobes-Bascaran, J., & Fernandez-Leal, A. (2023). Human-in-the-loop

machine learning: A state of the art. *Artificial Intelligence Review*, 56(4), 3005–3054.

- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30.
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4), 1593–1640.
- Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), 4–16. The research showing that personalised tutoring produces dramatic learning gains — the aspiration behind AI-as-tutor approaches.

Agentic AI:

- Wang, L., Ma, C., Feng, X., ... & Wen, J.-R. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6), Article 186345.
- Shavit, Y., Agarwal, S., Brundage, M., ... & Robinson, D. G. (2023). *Practices for governing agentic AI systems*. OpenAI white paper.
- Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. NIST AI 100-1.

E.8 General Background

For readers who want a broader foundation in how AI systems work and how to think about their role in society:

- Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. New York: Farrar, Straus and Giroux. An accessible, rigorous introduction to AI for non-specialists.
- Christian, B. (2020). *The Alignment Problem*. New York: W.W. Norton. Explores the gap between what we want AI to do and what it actually does.
- Mollick, E. (2024). *Co-Intelligence: Living and Working with AI*. New York: Portfolio. A practitioner-oriented book on integrating AI into professional work.

- Shneiderman, B. (2022). *Human-Centered AI*. Oxford University Press. AI systems designed around human control and oversight rather than full automation.
- Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson. The standard AI textbook for those who want deeper technical grounding.
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press. An accessible economic framework for understanding what AI does and doesn't change about decision-making.

E.9 The Companion Book

This book applies the methodology developed in *Conversation, Not Delegation: How to Think With AI, Not Just Use It* (Borck, 2025) to business education. That companion book covers the full framework in depth: the Conversation Loop, the VET framework for evaluating AI output, the cognitive traps that undermine critical thinking, and the principle of AI Last. For the underlying rationale and a discipline-neutral treatment, start there.

- Borck, M. (2025). *Conversation, Not Delegation: How to Think With AI, Not Just Use It*. Available at <https://michael-borck.github.io/conversation-not-delegation>.

About the Author



Michael Borck is a software developer and educator working at the intersection of human expertise and artificial intelligence. He developed the Conversation, Not Delegation methodology: the idea that AI is most valuable not as a tool you delegate to, but as a thinking partner you converse with.

The methodology grew out of a realisation: focusing on crafting the perfect prompt was just another form of delegation. The real value was never in the prompt. It was in the conversation that followed: the pushback, the iteration, the moments where your own thinking sharpened because you stayed engaged. That insight shifted the focus from prompt engineering to something closer to how professionals have always done their best work: through structured dialogue, critical evaluation, and judgement that only comes from experience.

Michael applies these principles across software development, education, and creative projects. He creates educational software and resources, and explores the 80/20 principle in learning and productivity.

Connect

- michaelborck.dev — Professional work and projects

- michaelborck.education — Educational software and resources
 - 8020workshop.com — Passion projects and workshops
 - LinkedIn
-

Other Books in This Series

Foundational Methodology:

- *Conversation, Not Delegation: Your Expertise + AI's Breadth = Amplified Thinking*
- *Converse Python, Partner AI: The Python Edition*

Python Track:

- *Think Python, Direct AI: Computational Thinking for Beginners*
- *Code Python, Consult AI: Python Fundamentals for the AI Era*
- *Ship Python, Orchestrate AI: Professional Python in the AI Era*

Web Track:

- *Build Web, Guide AI: Business Web Development with AI*