

The Judgment Classroom

How Educators Use Generative AI
to Build Thinking, Discussion, and Signature Work



Sonic Wu 吳相勳

元智大學管理學院

The Judgment Classroom

How Educators Use Generative AI to Build Thinking, Discussion, and Signature Work

Sonic Wu 吳相勳

Published by 元智大學管理學院

Responsible Editor: Codex

Manuscript Edition / 2026-04-30

The Judgment Classroom

How Educators Use Generative AI to Build Thinking, Discussion,
and Signature Work

Contents

Part I Reframe The Problem

- Chapter 1 The Wrong First Question About AI** 5
- 1. in Higher Education**
The central design problem is not whether students can use AI, but where AI belongs in the work of learning
- 2. Chapter 2 AI × Discipline** 10
Generic prompt fluency matters far less than the ability to translate disciplinary judgment into tasks, evidence, and standards
- 3. Chapter 3 Put AI in the Right Place** 15
When the tool sits before thinking, it becomes a shortcut; when it sits inside thinking, it can become leverage

Part II Build The Method

- 4. Chapter 4 Thinking Before Output** 20
The first move is not to ask for an answer but to force frameworks, comparisons, and criteria into the open
- 5. Chapter 5 The Real Threat Is False Mastery** 25
Smooth prose and quick completion often hide cognitive drift, not learning
- 6. Chapter 6 Prompting Is Work Design** 30
A good prompt does not decorate a question; it allocates roles, evidence, sequence, and critique

Part III Redesign The Course

- 7. Chapter 7 Start the Semester with a Contract** .. 38
If the rules, tool access, and evidence expectations are vague on day one, the rest of the semester becomes performative policing

8. **Chapter 8 The Teacher Cannot Be Weaker Than the Tool** 45
Faculty do not need to worship new software, but they do need enough command to judge, redirect, and demonstrate
9. **Chapter 9 Move from Lecture to Probe** 51
AI changes the classroom rhythm from one-way transmission to iterative questioning, testing, and revision
10. **Chapter 10 The Discussion Classroom** 57
Generative AI is most useful when it widens the field of perspectives and least useful when it closes the question too early

Part IV Make Better Student Work

11. **Chapter 11 Build Signature Work** 62
The goal is not cleaner output but visible authorship, stronger judgment, and work a student can defend in public
12. **Chapter 12 Grade the Process, Not Just the Product** 68
Logs, checkpoints, and oral defense reveal thinking; polished text by itself reveals very little
13. **Chapter 13 A Course Needs an Operating System** 74
Projects, instructions, archives, routines, and review loops matter more than any single model release

Part V Keep Judgment Human

14. **Chapter 14 Cases, Simulations, and Deliberate Friction** 80
The point of a good exercise is not speed but productive struggle, adversarial testing, and revision under pressure
15. **Chapter 15 The Last Right That Cannot Be Outsourced** 85
A university can delegate drafting, sorting, and searching, but it cannot delegate the judgment of what is worth trusting and acting on

Appendices

A.	Appendix A Course Contract and AI Use Policy	90
	Pack	
	Clarity about tools, boundaries, process evidence, and privacy should be built into the course before the first assignment begins	
B.	Appendix B Prompt Design Toolkit	98
	Prompting becomes educationally serious when it is treated as work design rather than verbal cleverness	
C.	Appendix C Case Library for AI-Rich Teaching	107
	Real classroom cases become more valuable when the teacher can specify the placement of AI and the friction that preserves learning	
D.	Appendix D Assessment Toolkit for AI-Rich	116
	Courses	
	Better grading in an AI-rich classroom requires stronger process evidence, clearer rubrics, and visible ownership	
E.	Appendix E Building the Course Operating	123
	System	
	Source spaces, instruction layers, prompt books, archives, and review loops should be designed as reusable infrastructure	
F.	Appendix F A Sixteen-Week Teaching Map	130
	A semester becomes coherent when AI use is staged across orientation, workflow, discussion, authorship, and judgment	
G.	Appendix G Reusable Prompt Books	137
	Prompt books are valuable when they encode repeatable teaching moves rather than one-off clever wording	
H.	Appendix H System Instruction Library	146
	System instructions convert course philosophy into tool behavior and preserve role boundaries over time	
I.	Appendix I Student Playbook for Working	153
	with Generative AI	
	Students need practical protocols for how to work with AI without becoming passive transporters of model output	
J.	Appendix J Faculty Demo and Workshop	160
	Scripts	
	Departments become more competent when they can demonstrate better moves in public rather than debate AI in the abstract	

K.	Appendix K International Trade and Compliance Master Prompt Book	166
	Trade and compliance teaching becomes stronger when concepts, calculations, thresholds, and debriefs are staged rather than collapsed	
L.	Appendix L NotebookLM Workshop and Assistant Modes	175
	Source-bounded assistants are most valuable when they guide reading and reasoning without becoming answer keys	
M.	Appendix M Digital Transformation, Data, and Simulation Prompt Book	183
	Advanced classroom prompt books can encode repeatable moves for data work, qualitative analysis, transformation conflict, and role-based simulation	

This manuscript is organized as a practical higher-education book rather than a prompt catalog. It is meant for teachers who want to redesign classrooms, assignments, and discussion routines without handing judgment to the machine.

Preface

Most books about generative AI in education arrive with one of two moods. The first is evangelical. The machine has arrived, history has turned, and the sensible adult must now clap on command. The second is defensive. Students are cheating, standards are collapsing, and the only serious response is prohibition dressed up as integrity. Both moods are understandable. Neither is good enough.

The real problem is less theatrical and more managerial. Teachers now work in an environment where searching, drafting, summarizing, role-playing, pattern finding, and stylistic imitation have become cheap. What remains scarce is judgment: the ability to define the problem correctly, choose the right framework, test evidence, resist the first smooth answer, and decide what deserves to count as learning. That is the work this book is about.

The argument is simple. We should not teach students to “use AI” in the abstract. We should teach them to use it inside a discipline, against a question, under constraints, with evidence, and with an explicit standard of judgment. Put differently, the right unit of design is not the prompt. It is the learning task. The prompt matters only because it expresses how that task has been thought through.

This book grew out of live teaching rather than conference vocabulary. The materials behind it include prompt books, workshop scripts, classroom simulations, role-play exercises, case discussions, and course operating rules actually used with students. Some were built for international business, some for digital transformation, some for marketing and media cases. They all point to the same conclusion. The machine is rarely the hard part. The hard part is deciding where the machine belongs, what it may touch, what it may not touch, and what the teacher will still demand from the student.

The voice of this book is deliberately practical. I am not trying to help faculty sound modern. I am trying to help them run a better classroom. The reader I have in mind is a university teacher who does not want to become a full-time prompt engineer, but also has no interest in surrendering the course to black-box convenience. That teacher needs a method.

Preface

This is that method.

Author: Sonic Wu 吳相勳

Note to Readers

This is a practical book for educators, especially in higher education. It is written in the style of a management book because the problem is managerial. A course is an operating system. It has inputs, routines, constraints, bottlenecks, and failure modes. If generative AI enters that system, the sensible response is not panic and not worship. It is redesign.

The book is also intentionally tool-light. Model names will change. Interfaces will change. Pricing tiers will change. Corporate promises will change, then change again. What does not change as quickly is the underlying logic of better teaching: clear goals, staged reasoning, evidence requirements, visible judgment, and assignments that students can defend as their own (Ambrose et al., 2010; Bain, 2004).

Read the book in sequence if you are redesigning a course from scratch. Read it selectively if you are trying to solve a particular problem such as shallow discussion, AI-written assignments, uneven tool access, or weak assessment design.

Note on Cases

The classroom materials behind this manuscript came from real workshops, courses, and case discussions. Where necessary, company names, file names, course identifiers, and operational details have been generalized or anonymized. The goal is to preserve the teaching move without exposing internal materials or commercially sensitive details.

The cases in this book should therefore be read as teaching cases, not as investigative reporting. Their purpose is to show how an educator can turn generative AI into a disciplined environment for reasoning, comparison, and argument rather than an answer vending machine.

Chapter 1 The Wrong First Question About AI in Higher Education

The first question most universities ask about generative AI is some version of this: can students use it? It is an understandable question. It is also the wrong one.

I have now heard the question in enough rooms to know its usual choreography. Someone from administration asks whether a university-wide rule should permit, limit, or prohibit AI use. One faculty member worries about cheating. Another worries about being left behind. Someone else raises equity: not every student has the same tools, the same paid tier, the same device, or the same tolerance for technical friction. The conversation feels serious because everyone is anxious. Then the room does what institutions often do under pressure. It compresses a design problem into a permissions problem.

That is the beginning of the mistake.

The Familiar Meeting

Imagine a typical faculty meeting in the second or third month after a new AI release gains traction. The room is full of sensible people. A dean wants clarity. A committee chair wants consistency. A younger faculty member wants the institution to be less performative and more realistic. A senior faculty member wants to know whether students are still writing anything on their own.

The draft policy on the screen says something like this: “Students may use generative AI only with instructor permission and proper acknowledgment.” Everyone nods because the sentence sounds balanced. It is balanced in the way a chair with three legs is balanced.

The sentence does not answer the questions that actually govern learning. May students use AI before they have read the assigned case? May they use it

to summarize the case they were supposed to read? May they use it to generate counterarguments after they have produced a first position of their own? May they use it to clean up grammar after the reasoning work has already been done? May they use it to build a role-play scenario that will later be debated in class? Each of these uses sits at a different point in the learning process. Each creates a different risk. Each may be educationally sound or unsound depending on the task.

Yet the meeting wants one answer. That desire is bureaucratically understandable and pedagogically disastrous.

Why the Permission Question Fails

The permission question fails for three reasons.

First, it treats AI as one thing. It is not one thing. It is a bundle of capabilities: summarizing, drafting, questioning, role-playing, classifying, translating, searching, comparing, visualizing, coding, and stylistically imitating. A student who asks an AI system to produce a polished final memo before thinking has not done the same thing as a student who asks the same system to simulate a hostile regulator after the student has already written a recommendation. Both are “using AI.” Only a lazy institution would stop there.

Second, the permission question treats the assignment as fixed. It assumes the task is already well designed and that the only matter left is whether a tool should be allowed near it. But AI has changed the meaning of weak tasks. A vague essay prompt once produced mediocre student prose. Now it produces mediocre machine prose with better punctuation. The tool did not create the weakness. It exposed it.

Third, the permission question encourages institutions to speak morally where they should speak operationally. They tell students to “use AI responsibly” and faculty to “maintain academic integrity.” Fine. Responsible how? Integrity with respect to which steps of the process? What evidence should remain visible? What part of the work is still expected to remain stubbornly human? A moral posture without workflow design quickly degenerates into ritual scolding.

Good teaching has never been merely a matter of whether a student arrives at a product. It has always depended on sequence, struggle, feedback, revision, and visible improvement (Ambrose et al., 2010; Bain, 2004). AI changes none of that. It simply makes bad sequencing faster, prettier, and easier to hide.

The Better Question

The better question is this: where does AI belong in the work of learning?

That is a different question because it asks about placement rather than permission. Once placement becomes the issue, the teacher can begin designing. What must happen before AI enters the task? What may AI do once it is allowed in? What evidence of process must survive? What should the student still have to defend without assistance?

That shift immediately produces better course questions:

Weak question	Better question
Can students use AI?	At which stages of this assignment may AI be used?
Should faculty allow ChatGPT?	What part of the learning sequence benefits from AI, and what part is damaged by it?
Do students have to disclose use?	What record of reasoning, revision, and evidence do I need in order to assess this task honestly?
Is AI cheating?	Which uses replace judgment, and which uses support it?

The teacher who asks the better question is already doing something important. The teacher is defining the task as a sequence of intellectual acts rather than as a finished object. That matters because AI is strongest when the task is treated as a finished object. It is much weaker when the task is treated as a chain of reasoned moves.

A Tale of Two Students

The difference becomes clearer if we compare two students in the same class.

Both students are assigned a strategic case about entering a foreign market. Both are told to submit a recommendation memo. Student A opens a model immediately and types: “What is the best market entry mode for this company?” The model returns a crisp answer: perhaps a joint venture because the market is risky but promising, local knowledge matters, and full acquisition may be too costly. The answer is readable. It may even be reasonable.

But the model has already made the first conceptual cut. It has decided what kind of problem this is. Student A experiences relief and mistakes that relief for progress.

Student B is told to do something different first. Before using AI, the student has to identify the decision criteria: control, capital burden, reversibility, speed, institutional risk, partner risk, and learning value. Only then is AI allowed in, and only for a specific purpose: generate two competing frameworks for choosing an entry mode, compare what each framework highlights and ignores, and produce one strong counterargument to the student's initial leaning. Student B still uses AI. But now the tool enters after the problem has been framed, not before.

In administrative language, both students "used AI." In educational reality, they did not do the same work at all.

The Cost of Vagueness

This is why generic campus guidance so often feels both noisy and thin. It is noisy because it speaks in the register of disruption, crisis, and innovation. It is thin because it does not specify learning design. A student may use AI to widen a comparison, surface a blind spot, generate a rival stakeholder position, or prepare a structured rebuttal. Used in those places, AI may improve learning. The same system may be used to produce a passable first draft before the student has read closely, chosen a framework, or identified the real tension in the case. Used that way, AI may hollow the task out from the start.

Vague policy cannot distinguish these uses well. Good course design can.

Vagueness has other costs too. It pushes faculty into covert inconsistency. One professor quietly tolerates AI summaries because the reading load is heavy. Another bans AI but never redesigns the assignments, so students simply use it more discreetly. A third wants to experiment but has no clear language for process documentation, source control, or oral defense. The institution then wonders why "policy compliance" looks uneven. The answer is simple. The policy is pretending to do the work of pedagogy.

What the Better Question Reveals

Once the question changes, the design space opens. The teacher must now decide at least five things.

First, what intellectual work must be done before the machine enters. In some courses that means reading the primary text. In others it means framing the problem, defining criteria, or generating one provisional claim without assistance.

Second, what kind of help the machine is allowed to provide. Is it being used to widen the set of plausible perspectives? To organize raw evidence? To simulate opposition? To support revision? These are different uses and should not be collapsed.

Third, what evidence must remain visible. If the final product is now easy to polish, the teacher has to ask what traces of process are necessary: logs, revision notes, oral explanation, annotated drafts, or structured checkpoints.

Fourth, what counts as authorship. If the model produces fluent prose, what makes the resulting work the student's rather than simply the student's submission? The answer will not come from intuition. It has to be designed.

Fifth, what part of the work still belongs to judgment. This is the deepest point. If the machine may summarize, compare, and draft, what exactly remains the teacher's and student's irreducible responsibility? The answer, I will argue throughout this book, is judgment: the right to define the problem, choose the frame, weigh the evidence, reject the convenient answer, and take responsibility for the conclusion.

What This Book Will Not Do

This book will not offer therapeutic slogans about the future of education. It will not pretend that every faculty member needs to become an expert in every model release. It will not sell prompt tricks as a substitute for course design. And it will not tell teachers to return heroically to a pre-AI classroom that was often less rigorous than memory suggests.

What it will do is more demanding and, I hope, more useful. It will treat AI as a design pressure on higher education. It will ask where the machine belongs, where it does not belong, how assignments should change, how discussion should change, how assessment should change, and how students can still produce work that shows judgment rather than merely polished completion.

That is a less glamorous conversation than most technology launches prefer. It is also the only one worth having.

Chapter 2 AI × Discipline

The phrase “AI literacy” is useful up to a point. Beyond that point it begins to mislead. It suggests that there is a portable skill called using AI and that once students acquire it they can carry it unchanged from marketing to engineering, from teacher education to finance, from international business to public policy. That is convenient language for workshops. It is weak language for curriculum.

The more accurate view is that serious AI use is always nested inside a discipline. It depends on what counts as a worthwhile question, what counts as acceptable evidence, what counts as a sound explanation, and what kind of mistake is unforgivable in that field. A student in corporate strategy does not ask the same kind of question as a student learning qualitative methods. Nor should they.

Three Rooms, One Model

Consider three classrooms using the same underlying generative model.

In the first room, students are studying international trade and the rules of origin that determine whether a product qualifies for tariff treatment under an agreement. The teacher asks the model to become a novice trade coach and explain the problem through a campus shopping district analogy. Later the same exercise becomes more technical. Students feed in the cost structure of an e-bike. The question is no longer “What is trade?” It becomes: can the model help calculate regional value content without skipping the logic or inventing arithmetic?

In the second room, students are studying the collision between Hollywood, streaming platforms, and creator ecosystems. The model is asked to find five facts that violate common intuition about the Oscars, Netflix, and YouTube. Then students are told to verify the sources, sort official data from journalistic claims, and generate an opposing explanation for each supposedly counterintuitive insight. The educational problem here is not explanation in the abstract. It is how to use the machine to investigate platform change without becoming the victim of its first neat narrative.

In the third room, students are doing qualitative case work. They are not asking for a finished answer at all. They are trying to design system instructions for a research assistant that can move from interview transcript to keywords, from keywords to hypotheses, and from hypotheses to evidence without inventing facts. Later, the same students compare a live case with IDEO method cards to ask what research methods were used, what was missing, and what another methodology would surface.

All three rooms are “using AI.” That description is almost useless. The real educational work is different in each case. Trade reasoning demands precision about rules, thresholds, and calculations. Media and platform reasoning demand source scrutiny, multiple interpretations, and a sensitivity to changing business models. Qualitative work demands disciplined interpretation, evidence handling, and protection against over-reading. One model, three disciplines, three different standards of seriousness.

Generic Fluency Is Not Enough

A student can be fluent in prompt syntax and still be intellectually empty. The student may know how to ask for bullet points, role-play, tone shifts, charts, or polished summaries. Fine. If the student cannot distinguish a mechanism from a symptom, or a framework from a slogan, or evidence from commentary, prompt fluency simply accelerates confusion.

This is the central reason I insist on AI × discipline rather than AI alone. In one classroom, students used AI to compare market-entry frameworks. In another, they used it to pressure-test industrial policy and de-risking. In another, they used it to contrast research methods in a digital-transformation case. The prompts looked different on the surface. The real difference sat underneath: each task required disciplinary standards, not general cleverness.

That is why the best prompt in the room is often written by the teacher who knows the field well, not by the person most impressed by the tool. The teacher knows where the naive answer will fail. The teacher knows which variable is decisive and which merely sounds technical. The teacher knows what a field treats as evidence, and what it treats as noise.

Case: Trade, Rules, and the Cost of Vagueness

Take the trade-compliance exercises built around free trade agreements and rules of origin. A superficial AI interaction might ask: “Can this product qualify for tariff exemption?” The model will likely answer quickly and confidently. That confidence is educationally dangerous because the real work lies in specifying the product architecture, the origin of components, the labor allocation, the threshold rule, and the logic of the agreement itself.

In one classroom sequence, students were first asked to use a campus alliance analogy so they could understand why origin rules exist at all. That was the accessible layer. Then the exercise became concrete. One student group worked on Bluetooth earphones. Another worked on e-bikes. They had to explain what happens when batteries come from one country, frames from another, and assembly work from a third. They discovered that the apparent simplicity of the rule hid a demanding evidentiary problem: in real supply chains, upstream cost data are often incomplete, strategic, or hard to verify.

This is what AI × discipline means in practice. The machine can help stage the explanation, visualize the comparison, or even draft a computational plan. But the disciplinary demand remains. The student has to know that a calculation without a clearly specified threshold is worthless, that a tariff outcome without verifiable cost components is suspect, and that legal compliance cannot be reduced to rhetorical clarity.

Case: Hollywood, Streaming, and Counterintuitive Inquiry

Now compare that with a classroom exploring Hollywood, streaming, and creator media. Here the educational risk is different. The model’s default behavior is to return median narratives: streaming changed audience behavior, awards matter less, creators are rising, studios face pressure. None of that is necessarily false. All of it may be boring.

So the teacher redesigns the task. Students must ask for counterintuitive facts. They must force web search. They must run the same task through multiple settings: conventional search, deeper research, and tail-end distribution prompts that try to pull unusual but still defensible hypotheses from the model. Then they must verify source quality and produce opposite explanations.

That workflow teaches a disciplinary habit. In business and media analysis, the first plausible explanation is often not the one worth trusting. What matters is not whether the model can speak fluently about Netflix or the Oscars. What matters is whether students can investigate power, attention, and strategy without confusing a tidy explanation for a tested one.

In the classroom record, students surfaced insights they did not begin with: low-budget films outperforming assumptions about prestige production, ad-supported streaming tiers succeeding more than ideology suggested, mobile gaming functioning as an attention-retention strategy rather than an identity shift, legacy intellectual property re-entering strategic importance after original-content enthusiasm cooled. Those are not just “AI insights.” They are examples of what happens when a teacher uses AI to create investigative discipline.

Case: Qualitative Analysis and Methodological Restraint

A third example comes from qualitative teaching. Students working with interview materials were not told to “analyze the transcript.” They were shown how to build system instructions for a specific research job. The assistant had to know its goal, its steps, its constraints, and its output format. It had to avoid inventing facts. It had to say when uncertainty remained. It had to transform raw material into something like a commercial-insight brief without pretending to have discovered more than the transcript could support.

Later, this work was extended with method comparison. Students took a case, then compared its existing qualitative method to other options drawn from external method cards. The question was not whether the model could sound methodological. The question was whether students could learn that methods are design choices, each of which sees something and misses something.

This too is AI × discipline. The machine does not replace the methodological judgment of the field. It makes that judgment visible.

The Unit of Design Is the Learning Task

If this sounds obvious, good. Universities have still spent much of the last two years ignoring it. The common mistake has been to treat the prompt as

the unit of design. The better unit is the learning task. Start there and the rest follows.

What problem is the student solving? What framework should the student consider? What evidence must the student marshal? What position must the student defend? What ambiguity is worth preserving rather than flattening? Only after those questions are answered should the teacher ask what AI is allowed to do.

In this sense, prompting is not primarily a linguistic activity. It is a translation activity. The teacher translates disciplinary judgment into instructions, constraints, checkpoints, and comparison structures. The student eventually learns to do the same. That is why strong AI use is genuinely educational. It forces hidden disciplinary logic to become explicit.

What This Means for Curriculum

The curricular implication is demanding. We should not try to make students “generally good at AI.” We should try to make them better at disciplinary work in an environment where AI exists. That means at least four things.

First, each course should define its non-negotiables: the kind of evidence it trusts, the errors it refuses, and the reasoning moves it wants to see.

Second, each course should state where AI is helpful and where it is corrosive. In one class, AI may be excellent for generating stakeholder views. In another, it may be too risky to use before close reading because it collapses interpretive work too early.

Third, assignments should be redesigned so that the model is forced to serve disciplinary method rather than replace it.

Fourth, faculty development should move beyond tool demos. A department does not need forty prompt tricks. It needs a clear conversation about how different fields want to use or constrain the same family of tools (UNESCO, 2024a, 2024b).

The problem is not that AI exists across disciplines. The problem is pretending that this makes disciplines less important. It makes them more important.

Chapter 3 Put AI in the Right Place

The location of AI inside the learning process matters more than the cleverness of any single prompt. Put the tool before thinking and it becomes a shortcut. Put it inside thinking and it can become a lever. Put it after thinking and it may become a useful production assistant. The tragedy of much current course design is that these three placements are treated as if they were morally equivalent and pedagogically interchangeable. They are neither.

The cleaner way to think about the problem is architectural rather than moral. A course has a sequence. Students have to frame a problem, choose a lens, gather evidence, test a position, then communicate a claim. Once AI enters that sequence, the teacher has to decide which stages remain stubbornly human and which stages can productively use assistance.

Placement of AI	Typical use	Likely gain	Main risk	Teacher response
Before thinking	First draft, instant summary, ready-made position	Speed	Borrowed understanding	Block, delay, or narrow
Inside thinking	Compare frame-works, test counterar-guments, surface as-sumptions	Better range and critique	Premature closure	Insert checkpoints
After think-ing	Editing, for-matting, re-	Better com-munication	Cosmetic inflation	Require vis-ible process

Placement of AI	Typical use	Likely gain	Main risk	Teacher response
	structuring, alternative presentation			

Before Thinking

This is the most dangerous placement because it feels the most efficient. A student pastes the assignment into a model and receives a fluent answer within seconds. The student experiences relief. Relief is not learning.

What has happened? The model has already made the first conceptual cut. It has decided what sort of problem this is, selected a likely frame, and offered a coherent direction before the student has done the uncomfortable work of deciding what the problem really demands. The machine has not merely accelerated prose. It has occupied the first move.

The educational loss is often invisible at first. The answer may look respectable. The student may even be able to paraphrase it. But when asked why that framework was chosen, what an opposing frame would say, or what evidence would overturn the recommendation, the student often has very little of their own to offer. The conceptual architecture was borrowed before it was built.

Case: The Premature Case Memo

This pattern shows up constantly in case teaching. A student receives a platform case, a strategy case, or an education-policy case and immediately asks AI for the best answer. The output is readable and typically structured around a familiar set of categories. Strengths and weaknesses. Opportunities and risks. Efficiency versus fairness. Innovation versus control. It is all so clean that the student feels as if the intellectual work has begun.

In reality the work has often been avoided. The student has not yet read for tension. The student has not identified the decision maker. The student has not clarified the metric of success. The student has not separated symptoms from mechanisms. Because the model is good at plausible structure, it easily gives the illusion that those steps have already happened.

This is why some assignments should explicitly forbid AI before an initial framing note is produced. The issue is not purity. The issue is that some tasks become trivialized if the first move is outsourced.

Inside Thinking

This is where the tool becomes genuinely interesting. Once the student has already framed the issue, chosen an initial lens, and identified the evidence to be used, AI can widen the field. It can propose competing frameworks. It can generate a counterposition. It can ask how the answer changes when a new stakeholder is introduced. It can model a skeptical regulator, a hostile board member, or a dissatisfied customer.

Used this way, the tool behaves less like an answer engine and more like an argumentative environment. The student is still responsible for the position. The machine simply makes it easier to run more comparisons in a shorter period of time. That is one reason studies on professional use of generative AI show substantial productivity gains (Noy & Zhang, 2023). The classroom implication is more specific. Productivity gains are only educationally useful if the time saved is reinvested in better judgment rather than faster submission.

Case: The YES and the Timing of Frameworks

A useful example comes from a class working on *THE YES*, a commerce and recommendation case. The teacher did not begin by asking students to apply Seven Powers, the Minto pyramid, or the Business Model Canvas. Students first had to work with the case itself and extract the early evidence: validation metrics, persona logic, dashboard signals, and how the recommendation engine was supposed to create value.

Only after that first pass were additional frameworks introduced. Students then used NotebookLM to switch in external books as interpretive modules. Minto for structure. Seven Powers for defensibility. Business Model Generation for business-model coherence. This timing matters. If the frameworks arrive too early, they function as prefabricated answers. Introduced later, they become interpretive tools that can be compared, challenged, and revised.

That is what it means to put AI inside thinking rather than before it. The model helps students move through layers of interpretation without erasing the need to interpret.

After Thinking

There is nothing dishonorable about using AI after the reasoning work has been done. Teachers have long accepted help from spellcheck, editing, presentation software, and layout templates. The question is whether the tool is cleaning up communication or laundering weak thought.

The difference becomes visible as soon as the teacher asks the student to explain why a claim is true, what alternatives were rejected, how competing evidence was weighed, or what changed between early and late drafts. If the student can answer clearly, post-thinking support is probably doing what it should do. If the student cannot, the elegant prose is not a sign of mastery. It is a better disguise.

Case: Signature Work versus Polished Sameness

This problem becomes sharper when students are asked to produce visible work: a presentation, a decision memo, a public-facing page, a GitHub-hosted output, or a project artifact with a signature line. In those cases, AI can absolutely help with phrasing, arrangement, and visual clarity. But if the student becomes merely the presenter of AI-made output, the course has created a polished ventriloquist, not an author.

One reason the source materials behind this book repeatedly stress signature work, visible contribution, and attributable sections is that the post-thinking phase is now extremely easy to inflate. If that phase is the only thing the teacher sees, grading will drift toward surface quality and away from intellectual ownership.

The Assignment as Workflow

The most useful practical habit is to redesign assignments as workflows with gates.

1. What must be produced before AI is permitted?
2. What kind of AI use is allowed during analysis?
3. What evidence of revision must survive?
4. What final defense must the student still perform?

If those four questions are answered clearly, placement becomes manageable. If they are left implicit, students will place AI wherever friction is highest. That usually means too early.

When to Shift the Placement

Not every task needs the same design. There are cases where AI can enter earlier because the educational goal is not original framing but exposure, confidence-building, or accessible explanation. For example, asking the system to explain a complex trade concept through a campus analogy can be educationally wise if the purpose is to lower the threshold of entry. There are also cases where AI should enter very late, such as assessment tasks that are meant to reveal the student's own capacity for close reading, judgment, or oral defense.

The mistake is to treat placement as fixed by ideology. The better approach is to treat placement as responsive to the learning goal.

Do Not Make Students Guess

The final design implication is plain. A course should make the placement of AI explicit. Do not wait for students to infer it. They will infer convenience.

That means saying, in concrete terms, “read first, then use AI to compare frameworks,” or “draft your own initial position before asking the model to challenge it,” or “AI may help you refine the final presentation, but you will be graded on the process memo and oral defense.” Students do not need another campus sermon about responsibility. They need a better map.

Chapter 4 Thinking Before Output

The phrase “think before you ask” sounds old-fashioned. Good. Some old-fashioned principles survive because they are correct. In a generative-AI classroom, the first move should usually not be “give me the answer.” It should be “show me the candidate frames, the decision criteria, and the key uncertainties.”

This is not a rhetorical preference. It is a quality-control device. When students ask for answers too early, they collapse the space of possibilities before they understand what sort of problem they are solving. That is how a difficult question becomes a polished but generic paragraph.

A Classroom Example: Three Windows on the Same Question

One of the clearest illustrations came from a class on Hollywood, streaming, and creator media. Students were told to open three different windows at the same time and run versions of the same basic prompt.

Window A was conventional: use a thinking model and web search to find surprising facts about the Oscars, Netflix, and YouTube.

Window B used deeper research: longer planning, broader search, more synthesis.

Window C added what the teacher called “tail-end distribution” prompting. The model was instructed to move away from the safest median answer and generate low-probability but still defensible interpretations.

The point was not to prove that one tool setting is always superior. The point was to let students feel the difference between ordinary plausible answers and answers that had been forced through a more demanding conceptual filter. In practice, the conventional version often returned statements so generic they almost mocked the assignment. Streaming changed behavior.

Hollywood faces disruption. Social platforms matter. True enough. Useless enough.

The more demanding versions began to surface things worth arguing about: low-budget films performing unexpectedly in prestige systems, ad-supported streaming tiers succeeding more than ideology predicted, mobile games functioning as a retention strategy rather than a new corporate identity, legacy intellectual property returning as a stabilizer after original-content exuberance cooled.

That classroom did not merely produce better answers. It taught a method: change the prompt sequence, change the quality of the reasoning field.

Frameworks Before Conclusions

In live teaching, one of the most reliable improvements comes from a small change in sequence. Instead of asking the model which international entry mode to choose, the student first asks what frameworks might govern the choice. Instead of asking whether a merger was good or bad, the student first asks what stakeholders, regulatory risks, network effects, and behavioral assumptions are relevant. Instead of asking for a view on industrial policy, the student first asks what counts as efficiency, security, transition cost, and political legitimacy.

This move does two important things. First, it reveals what the problem actually demands. Second, it makes shallow AI use easier to detect. A student who cannot name the candidate frameworks or explain why one lens matters more than another has not reached the point where a final answer deserves trust.

Anti-Intuition as a Teaching Move

The Hollywood case also revealed another useful tactic: anti-intuition prompting. Students were not merely asked to summarize what was happening in media. They were asked to identify what violated common intuition. That move matters because AI is naturally drawn to average explanation. “Tell me what is happening” often produces a consensus paragraph. “Show me what violates common intuition, with sources, dates, and numbers” forces the model toward a more investigatory posture.

This is educationally powerful for two reasons. It teaches students that inquiry can be designed. And it breaks the emotional hold of the first average

answer. Instructors who teach with cases already know the value of asking, “What is everyone in the room too quickly agreeing on?” AI can help create that moment if it is prompted to do so.

But Not Every Task Wants Exotic Answers

The tail-end exercise also carried an important warning. The teacher explicitly asked students whether every question deserved a tail-distribution prompt. Of course not. If the question is how much protein an adult should consume each day, pushing the model toward low-probability answers is perverse. There are tasks where consensus is the point.

This distinction is one of the book’s recurring themes. We need students to learn not just how to use a prompt technique, but when its epistemic logic fits the task. Counterintuitive search is useful when the question is exploratory, strategic, or interpretive. It is often foolish when the question is tightly scientific or legally settled.

That, again, is what thinking before output means. It is not merely better prompting. It is task diagnosis.

Comparison Is the Core Skill

Higher education has spent years saying it values critical thinking. Generative AI exposes whether that claim is real. Critical thinking is not a mood. It is disciplined comparison: between explanations, between tradeoffs, between stakeholder views, between evidence standards, and between rival framings of the same problem.

A tool that can generate multiple plausible lines of thought quickly is useful only if the student is required to compare them. Otherwise the model simply multiplies elegant noise.

This is why many of the classroom workflows behind this book use staged rounds.

1. Ask for candidate frameworks.
2. Ask what each one illuminates and ignores.
3. Add the actual case constraints.
4. Force a counterposition.
5. Ask the student to decide what still matters most.

The process is slower than one-shot prompting. It is also much better aligned with what universities claim to value.

Case: Framework Insertion and the Echo Chamber Problem

The Hollywood exercise also uncovered a risk. Once students had a strong prior belief, such as “the Oscars are dead,” the model became very good at writing persuasive prose that confirmed it. To counter that, the teacher redesigned the task again. Students had to ask for three possible explanations:

1. The Oscars truly lost cultural influence.
2. The influence did not disappear but migrated to new channels and forms of engagement.
3. The Oscars weakened for general audiences but remained powerful inside the industry.

That design forced the model to generate a choice set instead of a flattering monologue. The point was not simply to avoid bias. The point was to make bias visible.

This move generalizes well beyond media. In any field where students arrive with strong priors, the model can easily become an echo chamber if it is asked for a conclusion too soon. Ask first for rival hypotheses and the machine becomes more useful. Ask first for justification of the student’s preferred story and it becomes a sycophant.

A Practical Sequence for Students

By this point, a practical sequence is visible. The exact wording will vary by discipline, but the logic is stable:

1. Restate the problem in your own words.
2. Ask for two or three plausible frameworks.
3. Add case facts, constraints, and evidence requirements.
4. Compare the frameworks rather than selecting instantly.
5. Force a counterargument or contrary explanation.
6. Write a provisional recommendation.
7. Revise after challenge.

The crucial thing is not the number of rounds. It is the refusal to jump directly from task prompt to finished conclusion. Once that sequence becomes habit, AI stops acting like a shortcut and starts acting like pressure on the student’s analytic discipline.

Why This Matters More Than Prompt Vocabulary

This is also why many discussions of “good prompts” miss the point. Better wording helps. Better sequencing matters more. If a course teaches students to move from framework to comparison to judgment, prompt craft becomes meaningful because it serves a clearly designed process. Without that process, prompt craft is decoration.

The irony is that this type of teaching is not new. Good case teachers and discussion leaders have long used staged questioning to move students from description to diagnosis and from diagnosis to recommendation (Brookfield & Preskill, 2012; Ellet, 2018). Generative AI does not replace that art. It raises the price of neglecting it.

When Students Feel the Difference

One of the most valuable outcomes of the three-window and multi-round exercises is not the final answer. It is the moment students realize that the machine does not have one natural mode of intelligence. It has defaults. Those defaults can be redirected. Once students feel that difference, they stop treating AI output as an event and start treating it as material.

That is a major pedagogical gain. Students no longer ask, “What did the AI say?” They begin to ask, “What kind of task did we make it perform, and what did that design prevent us from seeing?”

That is a much better question.

Chapter 5 The Real Threat Is False Mastery

The most serious risk of generative AI in higher education is not that students will produce nonsense. They often will not. The more serious risk is that students will produce something passable while understanding less than they think they do. That is a much harder problem because it hides behind competence.

False mastery has a recognizable feel. The prose is coherent. The structure is tidy. The tone is appropriately academic or managerial. The student points to a finished document and assumes the work is complete. Everyone is tempted to move on. Then the teacher asks a simple question: Why does that criterion matter more than the others? What evidence would change your view? Why is this method preferable to the alternative? At that point the floor sometimes gives way.

Why the Feeling Is So Convincing

Part of the answer lies in cognitive load. Students are tired, overcommitted, distracted, and often trying to complete multiple kinds of intellectual labor at once. Generative AI reduces friction spectacularly. That is part of its attraction. But lower friction is not always a virtue in education. Some friction is the work.

This is where educators need to be precise rather than moralistic. The problem is not that students have suddenly become lazier in some civilizational sense. The problem is that the tool changes where effort is spent. Instead of spending energy framing the question, checking evidence, and defending a claim, students may spend their energy polishing tone, adjusting structure, and meeting deadlines with suspicious grace. The work has not disappeared. It has migrated downward.

The result is not always obvious because the output looks better than the understanding that produced it.

The Four Common Drifts

Across classrooms, four drifts appear again and again.

The first is first-answer anchoring. The model returns something fluent and plausible. Because it sounds acceptable, inquiry stops too soon.

The second is fluency confusion. Students mistake clarity of expression for strength of logic. A smooth paragraph feels true even when it is merely well formed.

The third is responsibility outsourcing. Students no longer feel accountable for defining the question itself. Their job becomes uploading the task and selecting among returned options.

The fourth is cosmetic revision. What changes between drafts are mostly phrases, transitions, and formatting rather than ideas, evidence, or structure.

These drifts are not random. They are structural responses to convenience. They also explain why faculty increasingly report a familiar species of assignment: polished, competent, and strangely difficult to discuss because the student did not truly author the reasoning.

Case: The RVC Calculation Trap

Consider the classroom sequence around trade compliance and regional value content. Students were given a cost structure for an e-bike and asked whether it met a threshold under an agreement. This kind of problem is perfect for exposing false mastery because the model can produce a quick answer that looks decisive even when the calculation is unstable or the assumptions are unclear.

The teacher therefore imposed a two-step method. First, the system had to write out its analysis plan and code without executing it. Only after the student reviewed the logic would the model be allowed to run the calculation. That design change matters. It slows the student down at the exact point where false mastery usually emerges: when a technically styled answer is accepted before the structure of the reasoning has been inspected.

The lesson generalizes. If a model can perform a task that students are not yet able to audit, then the task must be redesigned so the audit itself becomes visible.

Case: The App Version That Made Students Think Less

A smaller but revealing example came from the Hollywood and streaming classroom. During a live exercise, some students complained that their models were freezing or returning weak outputs. The teacher did not treat this as a trivial technical interruption. Instead, the teacher used it to expose something deeper: tool conditions shape thinking conditions.

Students were told not to use the simplified app version for the exercise and to switch to the web version because the lighter interface often stripped out the deeper reasoning behavior needed for long-context, source-rich work. This moment matters because it reveals how quickly students can confuse any AI output with adequate AI output. If the tool returns something fast, the student may not ask whether the environment was capable of doing the task well in the first place.

False mastery is not just about bad reasoning. It is also about false confidence in the adequacy of the tool setup itself.

Case: The Beautiful Presentation with No Author

Another example appears later in the learning cycle. A student gives a presentation that is elegantly organized, visually consistent, and verbally smooth. The slides are clean. The talk is competent. But the teacher suspects that the student has become, in effect, the spokesperson for an AI-made artifact.

This is a different form of false mastery. The problem is not that the ideas are entirely wrong. The problem is that the student has not metabolized them. When pressed, the student can restate the slides but cannot explain why the framing is appropriate, what alternatives were considered, or where the model was corrected. The work is present. The ownership is not.

This is why the source materials behind this book repeatedly connect AI use to expression, signature, and visible authorship. If the student becomes only the presenter of polished content, then the course has rewarded performance of understanding rather than understanding itself.

Cognitive Ease and Metacognitive Failure

Research on learning has long warned that students are poor judges of their own understanding when fluency is high and retrieval demand is low (Ambrose et al., 2010; Lang, 2016). Generative AI intensifies this problem because it produces highly fluent language on demand. It is therefore extremely good at manufacturing the feeling of comprehension.

This is why the OECD's emerging discussion of generative AI in education is so important. The concern is not merely plagiarism. It is what happens when tools replace necessary cognitive effort and thereby create a false sense of mastery, reduced engagement, and weaker metacognitive discipline. The danger is subtle precisely because students may feel more capable while actually becoming less capable of independent judgment.

The Student Does Not Need Another Lecture on Laziness

The practical response here is not to shame students. Students already know what convenience feels like. The educator's task is to help them recognize when convenience is displacing the wrong kind of effort.

That means designing moments where students have to notice the gap between possessing text and possessing an idea. It means forcing retrieval after AI use. It means asking students what changed between rounds. It means requiring one good counterargument. It means asking them where the model was too quick, too certain, or too flattering. The intervention does not have to be dramatic. It has to be well placed.

Slow the Moment of Closure

The most useful phrase here may be this: slow the moment of closure.

Ask for intermediate logs. Ask for a process memo. Ask students to state what the first answer missed. Ask them to show the discarded frame. Ask them to produce one oral defense without the screen. Ask them to mark which sentence in the final draft they disagree with or had to rewrite substantially.

Each of these actions puts the student back into the seat of judgment. Each makes it harder to confuse polished completion with understanding.

Signs a Task Is Producing False Mastery

Teachers need diagnostic signals. Here are several strong ones:

Signal	What it may indicate
Student work is uniformly polished but discussions are thin	AI raised surface quality but not understanding
Students can restate outputs but not justify criteria	The first conceptual cut was borrowed
Drafts change cosmetically but not analytically	Revision has become stylistic only
Similar structures appear across students in ways that exceed shared teaching	The task invites template completion rather than original framing
Students avoid showing logs or intermediate notes	They know the visible process would reveal weak authorship

These signals are not proof of misconduct. They are prompts for better design.

Do Not Mistake Better Typing for Better Learning

Many institutions will be tempted to treat cleaner work as evidence of progress. Resist that temptation. Generative AI can improve the presentational quality of student output very quickly. That says little by itself about intellectual quality.

The more useful question is whether the student has become more capable of framing, comparing, judging, and defending. If not, the classroom has not become more rigorous. It has merely become more legible.

Chapter 6 Prompting Is Work Design

Students often ask for the best prompt as if they were asking for the best wrench. They want the phrase that unlocks the model. They want the sentence that makes the answer suddenly become smart. That instinct is understandable. It is also a little childish. In a serious classroom, the important question is not which sentence sounds clever. The important question is what work the prompt is causing the model to do, in what order, on what evidence, under what restrictions, and with how much room left over for the student.

That is why I do not teach prompting as phrasing. I teach it as work design.

Once the teacher sees prompting that way, the fog begins to clear. Hallucination is no longer only a mysterious model defect. Often it is a source-control defect. Weak analysis is no longer only a bad answer. Often it is a sequencing defect. Inflated confidence is no longer only a student personality problem. Often it is a task-closure defect. The classroom becomes easier to diagnose because the teacher is no longer staring at output alone. The teacher is inspecting the design of the work that produced the output.

A Prompt Is a Miniature Operating System

The minimal educational prompt has to answer five questions.

Design question	What the teacher is deciding	What goes wrong when it is missing
Who is the model for this task?	Tutor, critic, analyst, auditor, role-player	The model drifts into generic explanation
What may the model use?	Uploaded case, lecture notes, approved source set, web search, or none	The model blends assigned material with half-remembered generalities

Design question	What the teacher is deciding	What goes wrong when it is missing
In what order should the work happen?	Frame, compare, calculate, test, then conclude	The model rushes to a conclusion before the hard thinking has happened
What is forbidden?	No invented facts, no final answer, no direct solution, no outside sources	The model over-completes the task and quietly takes the student's job
What must come back to the student?	A table, a critique, a question, a decision tree, a draft to inspect	The model gives a smooth paragraph that is hard to challenge

This is why the phrase prompt engineering is too flattering. Most educational prompting is not engineering. It is management. It is the allocation of cognitive labor. The teacher decides which parts can be accelerated, which parts must remain slow, which parts need a second check, and which parts should stay human all the way through.

In practice, the best prompts in the source materials behind this book always do three managerial things well. They assign a role sharply. They control the evidence base tightly. And they stage the work so that the model cannot jump to the clean ending too soon.

The IKEA Case: Plan First, Execute Later

One of the strongest examples came from a classroom exercise built around a large IKEA sales dataset. The temptation in a case like that is obvious. Students see a CSV file and want the model to run the numbers. The fast move is to say, in effect, “Analyze this dataset and tell me something interesting.” That is exactly the wrong first move.

The better classroom design required the model to write the analysis plan and the Python code first, but not execute it yet. The instruction was blunt: write the workflow, write the code, explain the statistical method, and stop. Let the human inspect it.

That simple sequencing decision changes the nature of the exercise.

First, it exposes whether the model has understood the columns correctly. A dataset usually contains more ambiguity than students realize. Does a field capture review count, rating score, or product category? Is the unit of analysis the item, the transaction, or the category? If the model gets that wrong, every neat chart that follows is merely confident nonsense.

Second, it lets the teacher inspect the logic of the method before the machine disappears into execution. The class can ask whether the suggested test is appropriate, whether the hypothesis is sensible, and whether the output will answer a managerial question rather than a decorative one.

Third, it changes the student’s posture. The student is no longer merely a recipient of findings. The student becomes the reviewer of a proposed analysis.

That is not a cosmetic change. It is a change in who holds responsibility.

Weak task framing	Better work design
“Analyze this file and give me insights.”	“Write the analysis plan and code first. Do not execute. Explain the logic and assumptions. Wait for review.”
“Run a statistical test.”	“Tell me whether a statistical test is even justified. If not, shift to a management report and explain why.”
“Make a chart.”	“Generate a chart only after the variables, calculation logic, and interpretation frame are approved.”

Notice what is happening here. The prompt is not simply asking for output. It is defining checkpoints. That is work design.

The Qualitative Project: Goal, Steps, Constraints, Output

The same logic appears in a very different setting: qualitative analysis of interview transcripts. Here the danger is not mathematical error. The danger is interpretive drift. The model will happily overread, overclaim, and invent motivation if the prompt lets it.

The most useful classroom response was to build system instructions around four elements:

1. Goal
2. Steps
3. Constraints
4. Output format

That structure is so plain that it hardly sounds like a technique. That is partly why it works. The student tells the system what job it has, what order the work must follow, what it is not allowed to do, and what the returned artifact must look like. The result is not magical. It is reliable. Reliability matters more.

In one class, students were taught to create dedicated project environments in ChatGPT or Gemini and then load the system instructions into a GPT, a Gem, or a project-level setup. They were also told to set the environment to project-only mode where possible so that material from outside the project would not leak in casually. This is not paranoia. It is method. If the task is grounded interpretation of a defined document set, then the system should behave like a bounded research assistant rather than a gossiping intern.

The discipline here is worth spelling out:

Instruction element	Why it matters in qualitative work
Goal	Prevents the model from becoming a generic explainer
Steps	Forces the model to move from coding or theme detection to hypothesis and then to evidence
Constraints	Blocks invention, over-interpretation, and fake certainty
Output format	Produces something the teacher can inspect consistently across students

This also reveals a larger truth. Much of the prompt work that matters in education is not verbal cleverness. It is the codification of professional judgment into repeatable instructions. In that sense, translating domain knowledge into prompts is itself a form of expertise.

Role Assignment Is Not Theater

Many of the strongest classroom prompts in the source materials use roles aggressively: a teaching assistant, a skeptical regulator, a board risk

challenger, a devil's auditor, a procurement manager, a maintenance officer, a field technician, a senior strategy advisor. Critics sometimes dismiss this as role-play for people who want their software to feel dramatic. That misses the point.

A role is useful because it loads an incentive structure into the task.

If the model is asked to answer as a free-floating expert, it tends to produce a pleasant median view. If it is asked to answer as a regulator, it begins to care about proof, compliance, enforcement, and downside risk. If it is asked to answer as a frontline employee, it begins to care about workload, distrust, time pressure, and the fear of being displaced. If it is asked to answer as a board risk challenger, it stops applauding the strategy and starts asking what could break it.

In other words, a role changes what counts as relevant.

That matters especially in discussion-based courses. Students often enter a case with one moral frame or one favorite explanation and then mistake that first frame for completeness. Role assignment is a disciplined way to widen the field. It gives the student not merely more output, but more conflict.

Source Restriction Is a Serious Pedagogical Device

The modern model wants to be helpful. Helpfulness is one of its most dangerous features. If the prompt is vague, the system will fill the gaps with general knowledge, weak analogy, stale public facts, and borrowed tone. The result often sounds smart enough to survive casual reading. That is why source restriction matters.

The NotebookLM system instructions from the workshop materials are especially blunt on this point. The assistant is told to rely completely on the course materials unless a narrow exception is explicitly triggered. It is told to ignore banned files. It is told not to become an answer key. It is told to guide students back into the source text rather than replace it.

That is good educational design for a simple reason: the assignment is often the source set.

If the class is reading a case, then a response built from the web plus memory plus textbook vapor is not richer. It is less disciplined. If the class is working through a packet of regulations, then outside synthesis may actually make the student worse because it collapses the technical logic into a

premature narrative. The stricter the source boundary, the more meaningful it becomes when the student finally crosses it.

Teachers should therefore stop apologizing for bounded prompts. A bounded prompt is not anti-intellectual. It is often the only way to preserve the intellectual task.

Meta-Prompting: Let the Model Help Write the Workflow

One of the most practical moves in the classroom materials is also one of the least glamorous. When students struggle to write a full system instruction from scratch, they are told to describe the role roughly in plain language and then ask the model to convert that description into system instructions suitable for a GPT or a Gem.

This is useful for two reasons.

First, it lowers the cost of good setup. Many students can describe what they want but cannot yet formalize it cleanly. Meta-prompting lets them draft the formal shell.

Second, it makes the structure of the task visible. The student starts to see that a good instruction set is not one big paragraph. It is an arrangement of goals, steps, prohibitions, and output expectations.

The teacher should not present this as wizardry. It is closer to teaching students how to dictate a standard operating procedure and then refine it. If the generated instruction is too short, too vague, or too permissive, the student revises it. Again, prompting becomes work design rather than charm.

Cross-Model Choreography

The classroom evidence also shows that a single-model mentality is too narrow. Different systems are better at different things. One practical routine was to use Gemini to generate a more detailed research plan, then paste that plan into ChatGPT for execution. Another was to use NotebookLM for source-bounded explanation, then move to ChatGPT or Gemini only when the task required web search, code execution, or broader synthesis.

This is not disloyalty to a vendor. It is adult tool use.

The teacher's job is to show students how to match the system to the work:

Task type	Better first tool	Why
Source-bounded explanation of assigned material	NotebookLM or project-bounded workspace	Stronger source discipline
Deep web research with a structured plan	Gemini or ChatGPT research mode	Broader search and plan generation
Code-backed quantitative analysis	Tool environment with Python access	Better execution and auditability
Editable visual artifact such as SVG or CSV	Tool with code generation and file output	Higher precision, easier revision

Once students see this, they stop talking about AI as if it were a single monolith. They begin to see a workflow.

A Good Prompt Leaves Judgment on the Table

The deepest design principle is the last one. A good educational prompt should not solve the course.

If the prompt asks the model to diagnose the case, choose the best framework, weigh the tradeoffs, make the recommendation, draft the slides, and polish the prose, then the student has been demoted to courier. The work may look finished. The education is mostly gone.

The better prompt leaves a gap. Sometimes it asks for a comparison but not a conclusion. Sometimes it asks for a checklist but not the decision. Sometimes it asks for three rival interpretations and forces the student to choose. Sometimes it asks the system to play the critic rather than the author.

That is why the most useful prompts in this book often end with one of the following moves:

- list what is still uncertain
- identify what evidence is missing
- show where the logic may break
- ask a question back to the student
- force a checkpoint before execution

These are not ornamental flourishes. They preserve the need for judgment.

The Real Standard

Prompting becomes educationally serious when it stops being a hunt for better wording and starts becoming a disciplined design of cognitive labor. The teacher is not trying to squeeze the most text out of the model. The teacher is trying to place the model in the right part of the work.

That is a harder standard. It is also the right one.

The student who learns to do this is learning more than AI fluency. The student is learning how to structure analysis, control evidence, define checkpoints, and keep ownership of the decision. Those are durable professional habits, whether the interface next year is ChatGPT, Gemini, NotebookLM, or something else entirely (Ambrose et al., 2010; Bowen & Watson, 2024; Lang, 2016).

Chapter 7 Start the Semester with a Contract

The first class of an AI-rich course should not begin with mystique. It should begin with terms.

Students need to know what tools matter, what uses are encouraged, what uses are restricted, what evidence must be retained, what counts as acceptable assistance, what data must never be uploaded, and what happens when access is unequal. If the course leaves these matters vague, then the rest of the semester becomes a mixture of quiet gaming, performative outrage, and after-the-fact moralizing.

Most faculty know this. Many still avoid it. They avoid it because they do not want to sound punitive, underprepared, or commercially compromised. They would rather stay at the level of principle: “Use AI responsibly.” The problem is that principle without operating detail is mostly theater.

What a course needs is a contract.

Why the Contract Matters More Than the Policy Memo

Universities often respond to generative AI at the policy level. They produce campus statements, recommended principles, aspirational guidelines, and occasionally prohibition lists. Some of that work is useful. None of it is enough.

A course is not governed by the policy memo in the abstract. It is governed by the decisions that the teacher turns into actual classroom operating rules. If the assignment depends on file upload, long context, or deep research features, students need to know that on day one. If the class will ask for chat logs, prompt histories, revision notes, or oral defense, students need to know that on day one. If the course prohibits uploading confidential case material to public systems, students need to know that on day one.

That is what the contract does. It translates institutional posture into actionable local rules.

The literature on institutional AI adoption is beginning to make the same point. The durable challenge is not merely whether a university endorses or distrusts AI. The durable challenge is how departments and instructors convert broad policy into operational clarity that students can actually follow (Jin et al., 2025; Wang et al., 2024).

The First Fairness Problem Is Access

Faculty often prefer to talk about ethics before logistics. In this case, logistics are part of ethics.

Different tools have different capabilities. A free account may give a student casual conversation. A paid tier may allow longer context windows, stronger research features, file uploads, better code execution, or richer image and document handling. That means “use AI if you want” is not a neutral instruction. It is often a disguised version of “students with stronger access get to produce stronger work.”

The classroom materials behind this book show this problem repeatedly. Some exercises worked only if students had access to deep research or a thinking model. Some visual tasks worked better on one platform than another. Some data-analysis tasks required Python or structured file output. The honest response is not to pretend the difference does not matter. The honest response is to name it.

Teachers therefore need to answer at least four access questions in the syllabus or first-class briefing:

Access question	Why it matters
Which platforms are expected or merely optional?	Students cannot plan if the requirement is hidden
Which features are essential?	Web search, file upload, code execution, and long context are not trivial differences
What is the fallback for students without premium access?	Otherwise the class quietly rewards purchasing power

Access question	Why it matters
What may be done in teams rather than individually?	Shared workflow can reduce inequity if designed openly

This does not require the teacher to turn the syllabus into a price sheet. It requires the teacher to stop pretending that tools are interchangeable when they are not.

Informed Consent Is Not Bureaucratic Excess

The first week should also tell students what they are walking into. If the course will ask them to export chat histories, keep process logs, construct a project workspace, or maintain a bounded knowledge base, that should not emerge halfway through the semester as a surprise administrative burden. Students should understand the logic up front.

The phrase informed consent can sound legalistic in this setting. I use it anyway because it is accurate. Students deserve to know not only the rules, but the learning logic behind the rules. If the course asks for logs, it should say plainly that the purpose is to make process visible rather than to run a cheap anti-cheating dragnet. If the course limits certain uses, it should say which learning activity would be hollowed out if those uses were allowed. If the course permits some strong uses, it should explain why those uses deepen rather than displace the work.

Once the logic is explicit, resistance often becomes more intelligent. Students stop asking whether the rule is anti-technology and start asking whether the placement of the tool is justified. That is a better argument.

The Contract Has to Cover Privacy and Boundaries

One persistent faculty mistake is to treat AI policy as an integrity issue while ignoring data handling. That is too small.

If a course uses live company data, internal course documents, student reflections, interview transcripts, or unpublished case materials, the teacher needs a clear rule about what may and may not be uploaded to external

systems. Some tasks should happen only inside bounded environments. Some materials should be de-identified. Some should never leave local storage at all.

A contract worth having therefore covers at least these domains:

Domain	What the contract should answer
Allowed use	Where AI may assist without displacing the learning task
Restricted use	Which tasks must remain fully human
Disclosure	What evidence of use must be submitted
Access	Which tools, tiers, and features the course assumes
Privacy	What kinds of files or data may not be uploaded
Process evidence	Whether logs, memos, or oral defense are required
Collaboration	Whether students may share prompts, accounts, or model outputs in teams
Escalation	What students should do when a tool is unavailable or clearly failing

The point is not to regulate every breath. The point is to make the serious boundaries visible before the first major assignment.

Logging Is Not the Same as Surveillance

Some faculty resist logs because they imagine that anything process-oriented must become punitive. Some students resist logs because they assume any request for process evidence is a disguised accusation. Both reactions confuse design with motive.

A process log can function as a learning artifact. It can show how the student framed the task, what changed after critique, which sources were used,

when the model misled them, why they rejected a line of output, and what remained unresolved. That is valuable even in a world with no cheating at all.

In fact, one of the most useful shifts in AI-rich education is that the log turns method into something teachable. A polished PDF hides process. A log, a revision memo, and a short oral defense make the intellectual path visible enough to coach.

The teacher should say this plainly in the contract:

- the course values process because process is part of the learning
- logs are reviewed as evidence of method, not as moral confession
- clean prose without process may receive less trust than messier work with visible reasoning

When students hear this early, many of them recalibrate. They understand that the course is not rewarding the most frictionless output. It is rewarding controlled use.

Show Acceptable and Unacceptable Use on Day One

The contract should not remain abstract. The first meeting should include at least one side-by-side example.

For instance:

Classroom move	Acceptable use	Unacceptable use
Pre-class case preparation	Ask AI to explain a technical term from the reading and identify where it appears in the case	Ask AI to produce the final case recommendation you plan to submit
Data analysis assignment	Ask AI to propose a code plan, then inspect and revise it before execution	Ask AI to run the analysis, choose the method, interpret the result, and write the report without review
Discussion preparation	Ask AI to generate two opposing stakeholder	Ask AI to write a finished discussion comment you then read

Classroom move	Acceptable use	Unacceptable use
	views to test your own position	aloud as if it were your own thought
Reflection memo	Use AI to help organize notes after your reflection is drafted	Use AI to generate a plausible reflection from scratch

One reason this matters is that many students are not trying to cheat. They are trying to find the invisible line in a poorly specified system. The more explicit the line, the less time the class wastes on rumor.

The Contract Protects Honest Students

There is another reason to take this seriously. Vagueness punishes the honest.

If a class quietly rewards aggressive automation while publicly praising integrity, cautious students are placed at a structural disadvantage. They are told to do the right thing in a game whose incentives are arranged around something else. That is not a character lesson. It is bad management.

The contract reduces that hypocrisy. It tells students what the course will actually reward. It also makes it harder for the teacher to improvise standards after seeing who performs well. Good governance protects not only the institution, but the student who is trying to work on the level.

A First-Class Script

The first meeting of an AI-rich course should do four things in order.

1. State the course position.
2. Explain the learning logic behind the position.
3. Show one or two concrete examples of acceptable and unacceptable use.
4. Explain what evidence of process the class will retain.

A direct opening might sound like this:

This course uses generative AI, but not everywhere. In some parts of the work, AI is a useful scaffold. In other parts, it would make the task meaningless. You will know those boundaries from the start. We will also

grade process, not only output, because polished output alone is no longer a trustworthy signal of learning.

That statement is not glamorous. It is also far more useful than the false drama of either total prohibition or breathless endorsement.

Managerial Seriousness

The contract is not a moral panic response. It is a managerial response. It says that when a powerful tool enters the classroom, the proper answer is not mystification. The proper answer is clearer operating rules, better disclosure, and a fairer alignment between what the course claims to value and what it actually rewards.

That clarity does not make the semester easy. It makes the semester legible.

And in an AI-rich classroom, legibility is a form of fairness (Bowen & Watson, 2024; Jin et al., 2025; Wang et al., 2024).

Chapter 8 The Teacher Cannot Be Weaker Than the Tool

The teacher does not have to be a software enthusiast. The teacher does not have to chase every model release, test every beta feature, or speak in platform slogans. The teacher does, however, need to remain stronger than the tool in the one dimension that matters most: judgment.

That means the teacher must be better than the average student at deciding what counts as a good use, a weak use, a dangerous use, and a fraudulent use in that specific course.

This is a modest standard. It is also non-negotiable.

Command Is Different from Fascination

Many faculty make an unnecessary mistake. They assume that because they are not fascinated by the tools, they cannot teach effectively in an AI-rich environment. That is false. Fascination is optional. Command is not.

Command means the teacher can do at least four things with confidence:

Faculty capability	What it looks like in practice
Frame the task	Decide where AI belongs and where it must stop
Compare tool behavior	Know when source-bounded systems, research systems, or code-capable systems are the better fit
Diagnose failure	Spot hallucination, premature closure, over-polish, and thin reasoning

Faculty capability	What it looks like in practice
Set standards	Explain what counts as adequate evidence, real revision, and defended judgment

A faculty member can dislike software culture and still do this job well. But the faculty member cannot remain completely ignorant and expect authority to hold.

Why Demonstration Now Matters So Much

In the pre-AI classroom, the teacher's authority rested partly on privileged access to organized knowledge. That has weakened. Students can now obtain plausible summaries, fast definitions, and neat first drafts without the teacher. If the teacher tries to defend authority mainly by repeating what the tool can already produce, the classroom becomes a poor competition between a human voice and a machine with infinite patience.

The better response is demonstration.

Students need to see the teacher make the task better in public. They need to see the teacher take a vague prompt and sharpen it. They need to see the teacher reject a weak answer for a precise reason. They need to see the teacher compare model behavior, repair a flawed framing, or force a checkpoint before execution.

This is why a live sequence matters more now than a lecture about best practices. A teacher can show:

1. a weak prompt
2. a revised prompt
3. the difference in output
4. the remaining problem in the better output
5. the human judgment still required

That five-step demonstration teaches more than a faculty workshop full of slogans.

The Teacher as Comparative User

One of the practical strengths in the classroom materials behind this book is that the instructor does not treat models as interchangeable. Students are shown, repeatedly, that different systems are suited to different tasks. Note-

bookLM is useful when the source boundary matters. ChatGPT is useful when Python execution or structured output matters. Gemini may offer a better research plan in a given workflow. Another system may be weaker on editing but better on search.

This matters pedagogically because students otherwise fall into brand thinking. They ask which model is the smartest. That is the wrong question. The useful question is which system is fit for this slice of work.

The teacher who can model that choice is already above the tool in the right way. The teacher is not worshipping the software. The teacher is allocating tasks intelligently.

Consider one classroom move from the source materials: students were taught to take Gemini's more detailed research plan and paste it into ChatGPT for execution. The point was not that one company is noble and the other defective. The point was that the student should learn to inspect the workflow, not merely the brand.

That is a professional habit.

The Teacher Has to Be Able to Name Bad Work

There is a more uncomfortable requirement. The teacher must be able to say, without wobbling, "This looks polished, but it is weak."

That sentence now carries enormous weight. Models are good at producing work that is calm, grammatically competent, and structurally tidy. Many students, and some faculty, mistake that combination for mastery. The teacher's job is to see through it.

What does bad work look like in an AI-rich course?

- a report with smooth prose but no visible tradeoff
- a presentation with formal language but no real author
- a calculation with a nice table and a broken underlying assumption
- a case memo that adopts a framework but leaves no residue of thought
- a reflection that sounds mature but contains no lived revision

The teacher must be able to identify these pathologies quickly, explain them concretely, and offer a better route.

That is why domain knowledge matters again in a new way. The teacher knows which simplifications are harmless and which are fatal. The teacher knows when a definition is adequate for the level of the course and when it

has shaved off the essential difficulty. The teacher knows when the model has answered the wrong question gracefully.

Above the Tool Means Setting the Standard

The decisive role of the teacher is not content supply. It is standard setting.

What counts as evidence in this course? What counts as a live uncertainty rather than a solved point? What counts as a meaningful revision? What counts as a strong use of theory rather than a decorative mention of theory?

AI cannot answer those questions responsibly for the class. If it does, then the class has outsourced its center.

This is where some faculty misunderstand the problem. They imagine that if the model becomes more capable, the teacher becomes less necessary. The reverse is often true. The more fluent the model becomes at producing acceptable surfaces, the more necessary the teacher becomes as the setter of a harder standard.

The teacher says, in effect: not just cleaner prose; show me the decision. Not just the framework name; show me the application and the limitation. Not just the chart; show me why this is the right chart for the question. Not just the answer; show me what would change your mind.

That is not nostalgia. It is instructional control.

A Practical Faculty Ladder

Many instructors ask a fair question: how strong is strong enough? The answer is not identical across fields, but a useful progression looks like this:

Level	Faculty posture	What is still missing
Avoidance	Treats AI as external annoyance	No real control over classroom use
Basic literacy	Can use one or two tools and recognize obvious strengths	Often cannot diagnose subtle failure
Comparative mand	com-Chooses tools by task and demonstrates better workflow	May still rely on generic standards

Level	Faculty posture	What is still missing
Instructional authority	Integrates tool choice, task design, and assessment into one system	This is the target

A university does not need every instructor to become a public AI intellectual. It does need instructors to get past avoidance.

Departmental Implications

This issue cannot be treated as an individual hobby. If only a few instructors know how to demonstrate good use while others remain entirely reactive, students experience the institution as incoherent. Some classrooms reward intelligent use. Some quietly reward aggressive outsourcing. Some pretend nothing has changed.

Departments therefore need a more practical culture:

- share prompt books and failure cases, not just inspirational examples
- compare tools by task rather than by reputation
- run live demonstrations in faculty meetings
- develop common language for process evidence and acceptable use
- treat AI teaching capacity as ordinary professional development, not fringe experimentation

The point is not to create a new priesthood. The point is to prevent the classroom from becoming governed by whichever student happens to have the strongest subscription and the least shame.

Authority Has Moved

Teacher authority has not vanished. It has moved.

It now rests less on being the first source of organized information and more on being the first serious source of judgment in public. Students can obtain text anywhere. What they cannot obtain so easily is a teacher who can expose why one workflow is unserious, why one answer is overclaimed, why one dataset needs a checkpoint, or why one elegant paragraph still has no mind behind it.

That is the new comparative advantage of the good teacher.

Not charisma. Not gadget worship. Not fear.

Judgment, shown live (Ambrose et al., 2010; Bain, 2004; Bowen & Watson, 2024).

Chapter 9 Move from Lecture to Probe

The lecture is not dead. It is simply no longer enough.

Students can now arrive with summaries, definitions, tables, timelines, basic comparisons, and even provisional arguments generated before class. Some of that preparation is useful. Some of it is fake preparation wearing a neat suit. Either way, the teacher who responds by lecturing harder often loses twice: once on attention and again on value.

The more useful shift is from lecture-first to probe-first.

That does not mean the teacher stops explaining. It means the first organized account is no longer the main scarce resource in the room. The scarce resource is disciplined challenge.

What a Probe Does That a Lecture Often Cannot

A probe is not just a question. It is a test of ownership.

When a teacher probes, four things become visible very quickly:

1. whether the student can defend the claim rather than merely repeat it
2. whether the evidence base actually supports the claim
3. whether the student's framing came from the student's mind or the model's default structure
4. whether revision is still possible or whether the student has mistaken fluency for closure

That is why a good probe is often more educational than five more minutes of clean exposition.

A Different Classroom Rhythm

The conventional sequence in many courses is still recognizable:

```
reading -> lecture -> discussion -> assignment
```

In an AI-rich course, a more useful sequence often looks like this:

```
guided preparation -> provisional answer -> teacher probe -> peer  
challenge -> revised judgment
```

The difference is substantial. In the old sequence, the teacher is the first organizer of meaning. In the revised sequence, the teacher is the first serious challenger of a provisional organization that already exists.

That role is both harder and more valuable.

Pre-Class Changes First

The shift to probe-first begins before class.

Students can now prepare in ways that are more structured than ordinary reading, provided the teacher designs the preparation well. For example, instead of asking students simply to “read the case,” the instructor can ask them to bring:

- one AI-assisted explanation of a difficult concept
- one point they still do not trust
- one stakeholder the model ignored
- one place where the source text contradicted the generated summary

That is already a better starting point than a room full of uninspected summaries.

The source materials behind this book offer several concrete illustrations. In an international trade class, students used prompts to clarify mercantilism, absolute advantage, comparative advantage, and political rationality before class discussion. In a classroom on platform competition, students used staged prompts to inhabit rival strategic roles before discussion began. In a database design class, students were told not merely to ask for the answer, but to arrive having already tested the logic of their first design moves.

In each case, the value was not that AI replaced the reading. The value was that AI gave the student something early enough for the teacher to attack.

The Probe Has to Be Better Than “What Do You Think?”

Weak discussion questions survive in many classrooms because they sound democratic. “What do you think?” is not always useless. But in an AI-rich room, it is too easy. A student can respond with a recycled synthetic view, a warmed-over summary, or a polished opinion that has never met resistance.

The stronger probe asks for a disciplined move:

Weak probe	Stronger probe
What do you think?	What assumption in your answer is doing the most work?
Do you agree with this recommendation?	What evidence in the case would make you withdraw the recommendation?
Can anyone add to that?	Which stakeholder or constraint is still missing?
Why is this strategy good?	Good for whom, over what time horizon, and at whose expense?
What did the model say?	Which part did you accept, and which part did you reject?

Once the teacher begins asking questions of this sort, the classroom changes character. Students realize that a model-generated paragraph is not a shield. It is merely the start of the interrogation.

Probe-First in Practice

Consider a trade-theory class in which students have already used AI to clarify absolute and comparative advantage. A lecture could still restate the definitions. A probe does something more useful:

- explain the difference without using the textbook phrase
- tell me which variable the model confused, if any
- show me whether the gain is in production or in final consumption
- tell me when the trade story breaks in political reality

That last question matters. Many students can now produce a neat economics answer quickly. The more interesting classroom move is to probe the friction: labor transition, political backlash, national security, supplier opacity, regulatory enforcement. This is where the teacher can force students beyond tidy textbook abstractions and into the harder managerial terrain.

The same pattern appears in technical courses. In the database-design classroom materials, the instructor did not let students hide behind an AI-generated schema. Students had to explain the entities, the relationships, and the reason a design choice made sense. If the teacher had simply admired the diagram, the intellectual work would have evaporated behind the artifact.

The Wheel of Luck Problem

One small teaching detail in the source materials is more important than it seems. In one classroom, students were given time to read AI-generated explanations, and then the instructor used a random wheel to call on students to explain the concept in their own words.

That move matters because it makes a quiet promise to the room: pre-class AI assistance does not exempt you from live ownership.

The wheel is not sacred. The principle is. Students need to know that the work will return to their own mouths, not merely to their screens. Once they know that, they use AI differently. They read more critically. They rehearse explanation. They notice where they are still dependent on the machine's language.

Probe-First Requires Better Listening

This shift is not simply a matter of asking harder questions. It also requires the teacher to listen diagnostically.

When a student answers, the teacher is listening for at least three things:

1. borrowed structure
2. unsupported assertion
3. live uncertainty

Borrowed structure sounds organized but weightless. Unsupported assertion sounds confident but thin. Live uncertainty is the most valuable of the three because it shows where real teaching can happen.

The teacher who hears live uncertainty can slow down, isolate the crux, and make the next question count. That is a different skill from delivering a well-compressed lecture. It is closer to diagnostic interviewing.

The After-Class Revision Loop

Probe-first does not end when the discussion ends. It changes what the student does after class as well.

If the classroom is well designed, the student leaves not with the illusion of completion, but with a revised problem. The next step might be:

- a short memo on what changed after the probe
- a corrected prompt sequence
- a better evidence table
- a brief oral reflection
- a comparison between the original AI answer and the revised judgment

This is where the learning deepens. The model gave the student a provisional structure. The classroom stressed it. The student now has to repair or replace it.

That loop is more educational than the old pattern in which lecture happened, notes were taken, and revision was postponed until the exam or the final paper.

Probe Is Not Humiliation

Some instructors avoid probing because they fear turning class into a performance of teacher superiority. That risk is real. A bad probe humiliates. A good probe clarifies.

The difference lies in the teacher's intention and design. The aim is not to expose the student as ignorant. The aim is to expose the current status of the reasoning so the student can move it. That means the teacher should probe the claim, not the person's worth. It also means students should understand that being challenged is normal, not a sign that they have failed.

When the teacher gets this right, the room becomes more rather than less psychologically serious. Students stop treating the first answer as the finish line. They begin to expect revision as part of ordinary intellectual life.

The Real Opportunity

Higher education has long been criticized for rewarding the reproduction of organized content. Generative AI is now very good at reproducing organized content. That does not make the classroom obsolete. It gives the classroom a chance to become more honest about what it is for.

If the room becomes a place where provisional structures are tested, revised, and owned, then AI has not diminished teaching. It has pushed teaching toward one of its stronger forms.

The lecture still has a place. But the probe is where the pressure now belongs (Bowen & Watson, 2024; Brookfield & Preskill, 2012; Lang, 2016).

Chapter 10 The Discussion Classroom

Discussion has always been harder than lecture. It is messier, less predictable, more revealing, and far less flattering to an instructor who prefers the dignity of uninterrupted authority. It is also one of the few formats in which thought has to survive contact with other minds.

Generative AI does not reduce the need for discussion. It increases it.

The machine is now very good at producing plausible lines of argument. That means the classroom needs a place where plausible lines collide, get challenged, lose support, gain support, and sometimes break.

Discussion Is Not Casual Verbal Traffic

Many weak classrooms confuse movement with thought. A few students speak. Another student adds a story. Someone mentions a framework. The room feels lively. Little actually happens.

AI can make this failure worse. Students can now arrive with more polished language than before, which means a superficial discussion can sound smarter while remaining just as empty.

A serious discussion classroom therefore requires a stricter architecture. Students should not merely contribute. They should have to do four things in sequence:

1. advance a position
2. name the evidence or logic beneath it
3. face a counterposition
4. revise or defend

That four-part structure is one reason business-school case teaching remains useful as a reference even outside business education. It has long assumed that discussion is not open-microphone expression. It is argument under constraint (Brookfield & Preskill, 2012; Ellet, 2018).

Use AI to Widen the Room, Not Close It

One of the best uses of AI in a discussion course is not answer generation. It is perspective generation.

Students often approach a problem with a single frame and then mistake that frame for the whole issue. AI can help widen the field quickly by generating rival stakeholder views, alternative theories of the case, neglected second-order effects, or harsher versions of the opposing side.

That widening move is especially powerful when the teacher uses it before the room hardens around the first fashionable opinion.

For example, in the platform competition materials behind this book, students could inhabit radically different roles in the 2018 Southeast Asian platform war: a Grab strategy chief, a Gojek expansion head, a SoftBank advisor pushing mutual forbearance, or an extreme pressure tester combining the voices of regulator, consumer, and labor concern. None of these roles is useful because it sounds colorful. They are useful because each role reorganizes the problem.

The strategy chief asks about defensibility, monetization, and timing. The regulator asks about concentration, abuse, and public harm. The investor asks about capital burn and game structure. The worker asks about extraction and precarity.

Once these voices are in the room, the case becomes harder to flatten.

The Teacher Must Keep AI from Becoming the Loudest Student

There is, however, a trap. If AI is used carelessly, it starts functioning as the loudest student in the room. It produces ready-made positions so quickly and in such abundance that actual students retreat into quoting the model rather than wrestling with one another.

That is why the teacher needs rules for how AI enters discussion.

Useful rules include:

- AI may generate positions, but students must select and defend one
- AI may generate counterarguments, but students must decide which counterargument actually matters
- AI may generate stakeholder perspectives, but students must map them onto the case evidence
- AI may not be cited as authority in place of reasoning

This keeps the tool in the role of field widener rather than decision closer.

Structure the Discussion in Rounds

The classroom materials behind this book repeatedly show that the strongest discussions are built in rounds. Round structure matters because it prevents the room from treating the first statement as the final statement.

A useful sequence looks like this:

```
Round 1: initial position
Round 2: strongest opposing view
Round 3: evidence challenge
Round 4: revised recommendation
```

That shape does three things.

First, it makes revision normal. Second, it rewards students who can absorb pressure rather than merely perform certainty. Third, it turns discussion into a visible record of changing judgment rather than a pile of disconnected remarks.

This is one reason role-based discussion works so well with AI. If the tool can help populate Round 2 with a sharper opposing view than the room would have generated on its own, then the room gets better material to work on. But the teacher must still insist that the students, not the model, do the revising.

Discussion Needs Stakes

A discussion without consequence often collapses into decorative exchange. The teacher can avoid this by attaching roles, decisions, or constraints that make the difference between positions matter.

In the international trade and compliance materials, students were not asked only what they thought about free trade or industrial policy. They were asked to speak as a commerce official, a factory operations manager, a compli-

ance analyst, or a procurement decision maker. That shift matters because it attaches the argument to a decision context.

Similarly, when students discussed rules of origin through the “campus shopping district alliance” analogy and then moved into real products such as Bluetooth earphones or electric bicycles, the discussion had to cross a difficult bridge: from the teaching scaffold to the compliance logic of a real supply chain. The room could no longer remain at the level of attractive metaphor. It had to ask what costs counted, what thresholds mattered, what documentation would be missing, and where the operating risk actually sat.

That is discussion with stakes.

The Teacher as Traffic Engineer of Conflict

In a good discussion classroom, the teacher does not dominate every minute. But the teacher is never passive. The teacher is designing the collision pattern.

That means deciding:

- which positions must meet each other
- which evidence is allowed to enter
- when the room has moved too quickly to consensus
- when a side argument is masking the crux
- when a clean sentence should be interrupted by a harder question

This is why the teacher’s judgment matters so much. AI can generate perspectives. It cannot manage the intellectual traffic of the room with the course’s standards in mind.

Discussion Is Where Authorship Reappears

There is a deeper reason to protect discussion in an AI-rich classroom. It is one of the few places where authorship becomes difficult to fake.

A student may arrive with model-assisted notes. Fine. A student may even arrive with a carefully structured provisional answer. Also fine. But when the room begins pressing, the student has to do something live. The student must choose, defend, revise, or concede in public time.

That is an unusually valuable moment now because so much other work can be polished in private with invisible help. Discussion restores temporal

pressure and visible ownership. It is not perfect evidence of mastery. But it is better evidence than static text alone.

What the Discussion Classroom Is For

The discussion classroom is not mainly for expressing identity, showcasing eloquence, or crowd-sourcing summary. It is for testing whether a position can survive challenge and whether the person offering it can remain intellectually alive after the first hit.

Generative AI makes that function more important, not less.

The classroom should therefore not retreat from discussion because the machine can talk. It should run better discussions because the machine can talk (Bowen & Watson, 2024; Brookfield & Preskill, 2012; Ellet, 2018).

Chapter 11 Build Signature Work

One of the ugliest outcomes in an AI-saturated classroom is not error. It is sameness.

Student work becomes grammatically smooth, structurally competent, and almost entirely forgettable. The pages look professional enough. The author has vanished. The presentation moves cleanly from problem to recommendation, but nothing in it suggests that a mind had to struggle, choose, risk, or decide.

This is why higher education cannot respond to AI merely by asking for cleaner output. It has to ask for work that leaves a signature.

Signature work is not eccentric work. It is work that carries visible marks of authorship, judgment, synthesis, situated interpretation, or design choice. It does not have to be loud. It does have to be owned.

The Enemy Is Median Competence

The biggest danger of generative AI in student work is not that every answer will be false. The bigger danger is that many answers will be acceptable in precisely the wrong way. They will be acceptable because they satisfy the weakest common standard:

- organized
- polished
- plausible
- generic

That is median competence with good manners.

If the university rewards that level of performance, then it has built an assessment system that the model can satisfy cheaply. The institution then starts mistaking surface adequacy for education.

The answer is not to hate polish. The answer is to require something polish alone cannot supply.

What Signature Work Looks Like

The shift often begins with assignment design rather than grading.

Generic assignment	Signature assignment
Broad essay on a topic	Decision memo tied to a named case, actor, and tradeoff
Concept summary	Comparison that must reject one dominant explanation
Slide deck of key ideas	Slide deck plus oral defense of two controversial choices
Data report	Data report plus process memo on why the method changed
AI-assisted write-up	AI-assisted write-up plus annotated rejection of weak model output

What is common across these stronger designs is not difficulty for difficulty’s sake. It is the requirement that the student do something specific with the material. The student must make a cut, choose a lens, defend a tradeoff, show a turn, or annotate a decision path.

That is harder to fake because it is not satisfied by tone.

Expression Is Not Decorative

Faculty sometimes talk as if expression were a secondary matter, something separate from real thought. That division was always too convenient. It is even less defensible now.

Expression is part of judgment.

A student who cannot turn a line of reasoning into a clear spoken, written, or visual form often does not yet possess the reasoning securely. Conversely, a student who can speak a conclusion with calm fluency may still be hiding inside borrowed language. The teacher must therefore ask not only whether the artifact is clean, but whether the form reveals the student’s actual decision-making.

This is where the source materials behind this book are especially useful. Again and again, classroom work moved beyond plain prose into other forms:

- comics to explain trade concepts
- SVG graphics to visualize production possibility frontiers
- role-play dialogue to reveal employee resistance
- customer-journey maps to expose friction in B2B service transformation
- sales decks that had to persuade a skeptical stakeholder rather than summarize a reading

These forms matter because they force structure and choice. A comic requires sequence. A decision tree requires branching logic. A journey map requires temporal and emotional movement. A role-play requires voice and incentive. Each of these forms makes it harder to hide inside generic exposition.

The Turn Is What Matters

One of the most useful design questions a teacher can ask is simple:

Where does this work stop being summary and start becoming argument?

That moment is the turn.

In weak AI-mediated assignments, the turn is often missing. The student gives background, reports what the framework says, organizes the points, and lands in a recommendation that sounds inevitable. Nothing in the artifact shows where the student actually chose among alternatives.

A stronger assignment makes the turn visible. It might require the student to:

- identify the key uncertainty that changed the recommendation
- explain which framework looked attractive but was finally rejected
- name the strongest opposing interpretation and explain why it lost
- show the point at which a dataset forced the question to change
- defend a decision under live questioning

Once the turn becomes visible, authorship becomes easier to evaluate.

Case Forms That Produce Ownership

The real classroom cases behind this book suggest several assignment forms that reliably create more signature and less sameness.

1. The Concept-into-Artifact Assignment

In an international trade class, students did not stop at asking AI to explain comparative advantage. They turned the concept into tables, code, and a dynamic SVG or infographic. This matters because transforming explanation into an artifact requires operational clarity. If the student cannot decide what the axes are, what counts as gain, or how the terms of trade should be visualized, then the student does not yet understand the idea well enough.

2. The Situation-into-Dialogue Assignment

In the Borusan Cat and Walmart materials, students used AI to help simulate frontline resistance, managerial persuasion, and organizational fear. This is not gimmick work. It is an assignment architecture that forces the student to ask what each actor knows, wants, fears, and misunderstands. A dry case becomes a social system.

3. The Framework-into-Judgment Assignment

Students were also taught not merely to use frameworks, but to use a framework against a case and then push further. In one qualitative-analysis example, external method cards were brought to bear on a case to ask which methods were still missing. In platform and digital-transformation cases, a framework became useful only when it forced a sharper evaluation, not when it was dropped in as a label.

This is the real test of professional learning. Not “Can you name SWOT?” The more serious question is “After you use the framework, what still remains to be decided?”

Signature Requires Public Ownership

There is a reason many strong assignments now include some live element. Once students know they may have to explain, defend, or revise in public, they prepare differently.

Public ownership can take modest forms:

- a two-minute oral defense
- a short recorded walkthrough
- a peer challenge round
- a discussion in which the student must explain what the AI got wrong
- a written author note appended to the artifact

These do not eliminate AI. They subordinate AI to authorship.

The student is not rewarded simply for arriving with a smooth output. The student is rewarded for being able to inhabit the reasoning behind the output. That is a stronger and more defensible standard.

How to Design Signature Into an Assignment

Teachers do not need to reinvent every task from zero. A practical redesign sequence looks like this:

1. Narrow the decision context.
2. Force a specific audience or stakeholder.
3. Require one visible turn from summary to argument.
4. Add a process artifact.
5. Add a live defense or revision stage where appropriate.

For example, a generic “analyze this case” assignment becomes more serious when it becomes:

Write a recommendation memo to the operations manager on whether the current e-bike supply chain can meet the rules-of-origin threshold. Explain which data the AI still could not verify, identify the largest compliance risk, and defend one remapping option in discussion.

The difference is obvious. The second assignment is smaller in one sense and far richer in another. It makes the student choose.

Signature Is Also a Moral Defense

There is another reason to care about this. The bland median voice is not only pedagogically weak. It is spiritually flattening. It trains students to produce what looks reasonable without having to stand anywhere in particular.

That is a poor preparation for professional life. In real organizations, value often comes from the ability to notice what others missed, frame a problem sharply, make a choice under constraint, or communicate a line of reasoning with enough force that other people can act on it. Signature work is not about self-expression in the adolescent sense. It is about traceable authorship in the professional sense.

Better Than an AI Arms Race

Some institutions respond to AI-generated sameness by escalating assignment difficulty mechanically. They assign longer papers, more citations, or more obscure topics. That is usually the wrong answer. AI can often handle length and formal complexity quite well.

The more useful response is to demand better ownership:

- more situated tasks
- more defended choices
- more visible revision
- more live explanation
- more artifacts that reveal a thinking path

Those conditions do not defeat AI. They make AI less able to hide the absence of a human mind.

A University Should Want More Than Clean Pages

If a university's standard becomes "professional enough," it has lowered itself into direct competition with software that specializes in professional enough.

Higher education needs a harder standard. It should want work that shows mind, ownership, and judgment under visible conditions. It should want work with trace. It should want work that a student can still defend after the machine has gone silent.

That is what signature work is for (Bain, 2004; Bowen & Watson, 2024; Ellet, 2018).

Chapter 12 Grade the Process, Not Just the Product

The old bargain in higher education was simple. Submit the finished work, let the teacher inspect it, receive the grade, move on. The finished artifact stood in as the main evidence that learning had happened.

Generative AI has weakened that bargain badly.

The final product is now easier to polish, easier to inflate, easier to imitate, and easier to detach from the actual intellectual path that produced it. That does not make assessment impossible. It makes product-only assessment increasingly unserious.

The practical answer is not panic. It is to grade the process explicitly.

Why Product Alone Is a Weak Signal Now

A polished product used to be at least a rough proxy for time, organization, and effort. That proxy has decayed.

A model can now produce:

- a structurally competent memo
- a visually neat slide outline
- a plausible literature summary
- a well-phrased reflection
- a table or chart with tidy captions

The presence of these qualities does not prove the absence of learning. It simply means the old shortcut no longer works. The teacher has to ask a different question: what evidence shows that the student actually framed, checked, revised, and owned the work?

That is a process question.

Process Evidence Makes Thinking Legible

Process evidence is valuable because it makes the work visible enough to teach and to assess. The point is not to demand every keystroke. The point is to retain the parts of the path that reveal judgment.

Process artifact	What it can reveal
Prompt or query log	How the task was framed and re-framed
Source note	What evidence base the student used and ignored
Analysis plan	Whether the student could define a method before execution
Revision memo	What changed after critique and why
Oral defense	Whether the student can own the reasoning in live time
Reflection on AI errors	Whether the student can recognize when the model overreached

Taken together, these artifacts often tell the teacher more than the finished PDF ever could.

The IKEA Lesson: Review Before Execution

Return to the IKEA dataset example. When students were required to ask the model for the analysis plan and code before execution, the teacher gained something important: a chance to inspect reasoning before the numbers arrived.

This is precisely the sort of checkpoint that should also matter in grading.

If a student can explain:

- why this variable mattered
- why this method was chosen
- why a statistical test was rejected
- why the assignment shifted into a management report instead

then the student is demonstrating actual analytical control. If the student can only point to a clean chart and a few generated insights, the teacher has far less basis for trust.

In other words, the plan is not just a staging device. It is assessable work.

Similarity Detection Is Too Small a Response

One reason institutions keep reaching for detection tools is convenience. They want a quick answer to the question of whether the final text is original. That desire is understandable. It is also too small.

Similarity can be one clue. It is not the educational center of the problem.

A student may submit text with low detectable similarity and still have outsourced almost all of the meaningful work. Another student may receive heavy model assistance in drafting but still retain strong intellectual ownership through repeated critique, revision, and defense. A pure authorship-detection mentality misses this distinction.

That is why process assessment is superior. It asks what the student did with the tool, not merely whether the final page resembles something else.

Build the Rubric Around Work the Course Claims to Value

A stronger rubric in an AI-rich course usually needs at least four dimensions beyond surface quality:

Dimension	What strong performance looks like
Framing quality	The student defines the problem sharply and asks a disciplined question
Evidence discipline	Sources are bounded, checked, and used proportionately
Revision quality	The student changes course when critique or evidence demands it

Dimension	What strong performance looks like
Defended judgment	The student can explain and defend the final choice live
Surface quality	The artifact is clear, organized, and appropriate to audience

This arrangement has an additional virtue. It rewards good AI use rather than merely aggressive AI use. Strong users document their method, challenge the system, and revise their prompt stack. Weak users often cannot explain what they did because what they did was mostly convenience.

Logs Are Better When They Ask the Right Questions

A weak log requirement produces busywork. A strong one forces reflection.

A useful AI log should not simply ask students to dump screenshots. It should ask:

1. What was your initial task framing?
2. What did the model misunderstand or overclaim?
3. What did you change in response?
4. What sources did you rely on?
5. What part of the work remained fully yours?

Those prompts make the log an analytical instrument rather than a stack of evidence bags.

The source materials behind this book also suggest a practical advantage: exported chat histories can be reviewed not only by the teacher but by the student. That lets the class compare strong and weak process patterns. One student's prompt chain shows progressive clarification, source checking, and genuine revision. Another shows blunt convenience and minimal ownership. This comparison is pedagogically rich because students begin to recognize method as something observable.

Oral Defense Restores Real-Time Ownership

If the teacher wants a stronger signal, add a small oral defense.

The oral defense does not need to become a courtroom. Often three questions are enough:

- What changed most between your first approach and your final one?
- What did the model get wrong or fail to see?
- What evidence would make you change your conclusion now?

A student who can answer these questions clearly has probably retained possession of the work. A student who cannot may still have learned something, but the teacher now has better grounds to distinguish assisted performance from owned judgment.

This is one reason oral defense works so well alongside written artifacts. It tests whether the student can still stand behind the work once the polished object is no longer doing all the talking.

Use AI to Assess AI Use, Carefully

There is no principle that says teachers must inspect all process evidence manually. AI can help.

It can compare early and late drafts. It can cluster common prompt failures. It can highlight where students relied on vague prompts, failed to cite sources, or accepted an answer too quickly. It can summarize repeated weaknesses across a class section so the instructor can intervene more intelligently.

Used this way, AI becomes part of formative assessment rather than an adversary to it.

But a warning is necessary. The teacher should retain interpretive control. Assessment is not merely pattern detection. It is still a judgment about whether the student displayed the kind of work the course exists to cultivate.

Multi-Dimensional Assessment Is Closer to the Truth

The core argument in the source materials is exactly right: the course should not collapse learning into a crude single indicator. A better assessment

system is multi-dimensional. It looks at product, process, revision, and defense together.

That design is both fairer and more demanding.

It is fairer because students who work honestly but are not naturally elegant on the page can still show method. It is more demanding because students who produce sleek output without ownership can no longer glide past the rubric on tone alone.

Process Assessment Changes Student Behavior

Assessment does not merely measure behavior. It shapes it.

Once students know the course will value logs, revisions, and defended judgment, they tend to work differently:

- they slow down sooner
- they inspect prompts more carefully
- they save intermediate reasoning rather than deleting it
- they become more alert to hallucination and overclaim
- they prepare for explanation, not just submission

In other words, grading the process teaches the process.

The Teacher Is Grading What the Course Says It Believes

If a course claims to value reasoning, evidence discipline, revision, and integrity, those things should appear in what is actually assessed. Otherwise the course is grading typography while pretending to grade thought.

That pretense was always a weakness. AI has merely made it easier to see.

Better assessment does not require rejecting AI. It requires refusing to let the final artifact monopolize the meaning of the work (Ambrose et al., 2010; Bowen & Watson, 2024; Lang, 2016).

Chapter 13 A Course Needs an Operating System

Once a course begins using generative AI seriously, a handful of prompts is not enough. The course needs an operating system.

By operating system I do not mean a brand. I mean a durable arrangement of goals, instructions, sources, routines, archives, checkpoints, and review loops that lets the class function consistently even as tools change.

Without that operating system, every week becomes a fresh negotiation with the machine. Students do not know which space to use, which sources are approved, what the output should look like, or whether the teacher wants exploration, execution, or critique. Faculty then mistake operational chaos for pedagogical openness.

It is not openness. It is poor management.

The Core Stack

A course operating system can be described in six layers.

Layer	Question it answers
Goals	What kind of intellectual work does this course exist to cultivate?
Instruction stack	How should AI behave in this course?
Knowledge base	Which sources may be used, and where are they stored?
Activity routine	When does AI enter and leave each learning sequence?
Archive	What artifacts and process traces must be retained?

Layer	Question it answers
Review loop	How are outputs inspected, challenged, and revised?

This functional view matters because too many institutional AI conversations collapse into product talk. They ask whether the campus should use Tool A or Tool B. That question matters, but it comes second. The first question is what the course needs from any tool at all.

Prompt Books Are Infrastructure

One of the clearest lessons in the source materials is that prompt books stop being mere tricks once a course uses them repeatedly. They become infrastructure.

At first, a prompt book looks like a convenience: a prepared instruction sequence to help students enter a case or an exercise. Over time, it becomes something more substantial:

- a reusable role-play engine
- a reproducible way to explain a difficult concept
- a scaffold for turning theory into artifact
- a controlled sequence for simulation and critique
- a class memory of what strong work looks like

That is why I do not think of prompt books as clever accessories. I think of them as teaching assets. They encode successful instructional moves in a form that can be reused, compared, revised, and handed from one section to another.

The real asset is not the individual prompt. It is the repeatable teaching move hiding inside the prompt.

System Instructions Are Policy Made Operational

If the course contract states the principles, system instructions operationalize them.

The NotebookLM workshop materials are especially revealing here. The teaching-assistant mode is not just a technical curiosity. It is a compact statement of course philosophy turned into machine behavior:

- ask the student which course they are in
- keep the source set bounded
- do not become the answer key
- guide with questions
- use Taiwanese terminology when writing in Chinese
- switch roles only when the correct teacher-mode trigger appears
- ignore banned files

That is a course policy translated into executable structure.

The same logic applies in English-language environments. A good system instruction tells the tool what role it plays, what evidence it can use, what it is forbidden to do, how it should respond when the student asks for too much, and what format the output should take. Once that layer is in place, the course behaves more consistently.

Bounded Spaces Beat Floating Chaos

Another operational lesson is that bounded spaces matter.

Projects, Gems, GPTs, NotebookLM notebooks, and curated folders all solve a common problem: they give the tool somewhere to stand. Instead of drawing on the full blur of the internet plus whatever the model half-remembers, the tool is asked to work inside a defined context with a defined role.

That makes several good things possible at once:

- stronger source control
- more repeatable classroom behavior
- cleaner division between course tasks and private experimentation
- better preservation of prior critique and task structure

The source materials even go further by recommending project-only mode in some settings. That is not mere neatness. It reflects an important educational principle. When the task depends on a stable evidence base, context contamination is not a feature. It is a bug.

Tool-Agnostic Does Not Mean Tool-Blind

A course operating system should be tool-agnostic in design but not tool-blind in practice.

Tool-agnostic means the educational logic should survive if one platform changes its pricing, naming, or interface. Source control, staged reasoning, critique, and revision should still make sense next year.

Tool-aware means the teacher still has to know what each platform is good for.

For instance:

- NotebookLM is strong when the course wants bounded explanation from a curated source base
- ChatGPT may be better when the task requires Python execution, editable file generation, or structured output
- Gemini may be useful for certain research-plan workflows or other specific strengths
- different tools may vary in their ability to create editable SVG, CSV, or graphic outputs

The operating system therefore sits above the tools without becoming indifferent to them.

The Archive Is Part of the Course

A good course now needs a deliberate archive. That archive may include:

- prompt books
- course-specific system instructions
- case packets
- approved source folders
- exported chat logs
- revision memos
- grading rubrics
- examples of strong and weak process

Why does this matter? Because without an archive, every semester begins again at zero. The instructor relearns the same failures, rewrites the same prompts, and answers the same questions without retaining the method. That is wasteful.

An archive also helps departments mature. Shared prompt books and examples let faculty compare teaching moves rather than merely exchange opinions about AI in the abstract.

A Weekly Routine Reduces Noise

The operating system should show up in the weekly rhythm of the course, not only in the syllabus.

A useful weekly routine might include:

1. pre-class source-bounded preparation
2. in-class probe or discussion round
3. post-class revision artifact
4. archived process trace
5. brief instructor review of failure patterns

That sequence is simple. It is also strong because it aligns activity with evidence. Students know when AI is expected, when it is bounded, and when judgment returns to them explicitly.

The teacher also benefits. Instead of improvising every week, the teacher can ask the more serious question: where in this week's sequence is the highest-leverage intervention?

A Course OS Protects Against False Innovation

There is a fashionable version of AI adoption in education that is mostly interface excitement. A teacher tries a new feature, shows the class something flashy, and calls the result innovation. That can be entertaining. It is not enough.

The operating system protects the course from that temptation by imposing a harder question: where does this feature belong in the learning process, what evidence does it improve, and what failure mode does it reduce?

If the answer is vague, the feature is probably decorative.

Departments Need Shared Infrastructure

This is not only an individual-course issue. Departments should eventually develop shared infrastructure:

- common language for permitted and restricted use
- shared log templates
- repository folders for source-bounded materials
- prompt libraries for discussion, simulation, and critique
- faculty examples of good and bad workflows

This does not eliminate disciplinary variation. It reduces unnecessary incoherence.

Calm Operations Make Stronger Teaching Possible

There is nothing glamorous about this chapter. Good operating systems rarely generate applause. But they make serious teaching easier.

Students know where to work. Teachers know what to inspect. The class spends less time negotiating permission and more time using judgment.

That is what good operations are supposed to do (Bowen & Watson, 2024; UNESCO, 2024a, 2024b).

Chapter 14 Cases, Simulations, and Deliberate Friction

If the machine is naturally fast, the teacher's job is sometimes to slow the room down.

Not because slowness is virtuous by itself. Not because struggle is holy. But because some forms of judgment require friction. They require the student to meet resistance, uncertainty, contradiction, missing data, and hostile perspectives before the answer is allowed to harden.

That is why cases, simulations, and deliberate role-based conflict matter so much now. They keep the work from collapsing into frictionless output.

Friction Is a Design Choice

Most weak AI use in education shares one trait: it removes too much resistance too early. The student asks for the answer, the model provides a neat structure, and the classroom never reintroduces the conflict that would have forced real judgment.

Deliberate friction reverses that mistake.

It asks:

- who disagrees with this answer
- what data is still missing
- what operational risk the neat answer ignored
- which stakeholder would refuse the recommendation
- what happens if the first assumption fails

These are not afterthoughts. They are the work.

The Investment in the Fog Case

One of the clearest examples from the course materials is the “investment in the fog” exercise. The task concerned a situation in which a government changed investment conditions after attracting foreign investment. The prompt was long, heavily constrained, and full of precise terms such as tax holidays, land incentives, economic-zone treatment, and retroactive tax risk.

Why was the prompt so long? Because the teacher did not want the model to perform generic geopolitical theater. The teacher wanted it to search for the right class of risk.

Three design choices made the exercise strong:

1. a new thread was required, because the task belonged to a fresh problem space
2. precise domain vocabulary was embedded, because the wrong keyword leads to the wrong evidence
3. cross-checking across sources and a strict report structure were required, because the class needed a disciplined investigative output rather than a foggy summary

This is friction by design. The prompt denies the student the illusion that a loose question will produce a serious answer. It forces them to narrow the terrain.

The Hawks-and-Doves Simulation

Another strong case involved a foreign direct investment simulation in which students had to decide whether to take a hawkish or dovish posture under uncertainty. The prompt sequence was modular. Students selected a position, chose signals, inserted a role prompt, and then loaded a large block of background material.

The crucial move came later. After one round of analysis with the “board’s independent risk challenger,” students were told not to start over from zero when switching to the “devil’s auditor.” They were taught to edit the earlier instruction inside the same conversation rather than rebuild the entire task.

That small operational move carries a deep pedagogical lesson. When students edit the instruction rather than simply continue chatting, they are forced to inspect the logic of the task itself. They stop treating prompting as disposable phrasing and start treating it as an instrument they can tune.

The simulation also created real argumentative pressure. A recommendation had to survive a hostile board view and then a harsher audit view. That is much closer to actual professional life than the standard single-pass AI answer.

Compliance Cases and Forced Checkpoints

Rules-of-origin and compliance cases are especially useful because they punish vague thinking.

In the prompt materials, students moved from a life analogy such as a campus shopping-district alliance into concrete product cases: Bluetooth earphones, e-bikes, and later electric vehicles. They then had to work through the actual logic of regional value content, non-originating material, threshold rules, and supply-chain remapping.

What made these exercises educationally strong was not merely the case topic. It was the checkpoint design.

Students were first asked to outline the logic steps without giving a final answer. Only afterward were they forced to compute the RVC. Then they had to face a definitive judgment. Then they had to design an alternative supply-chain move if the current configuration failed.

This order matters.

Step	Why it matters
Explain the rule in plain language	Prevents students from hiding behind jargon
Map the real product to the teaching scaffold	Forces the analogy to earn its keep
Calculate the threshold explicitly	Blocks vague compliance talk
Identify missing documents or opaque upstream data	Reintroduces real-world difficulty
Propose remapping options	Turns regulation back into strategy

The discussion guide from the trade materials makes this beautifully clear. Students are asked to identify the gap between the clean teaching illustration and the dirty reality of actual customs enforcement. That gap is where the educational value sits.

Same Thread or New Thread?

The source materials repeatedly reveal another quiet but important lesson: thread management is part of pedagogy.

Sometimes students should remain in the same conversation so that critique, context, and prior decisions stay alive. Sometimes they should start a new thread so that the earlier assumptions do not contaminate the next task.

This is not only a software habit. It is an analytical habit. A new thread tells the student that the problem has changed enough to deserve a reset. Staying in the same thread tells the student that the next move is a revision of the same intellectual object.

Teachers should make this explicit. Otherwise students confuse accumulated context with accumulated clarity.

Simulations Reveal What Smooth Output Hides

The value of simulations is not entertainment. It is revelation.

A simulation reveals whether the student can:

- survive a challenge from an opposing role
- adapt the recommendation under pressure
- distinguish between evidence and rhetoric
- notice when a framework is being stretched too far
- recognize what the AI is supplying and what still remains theirs

This is why simulations are so useful across domains.

In business they reveal how decisions fracture across incentives. In teacher education they reveal the collision of policy, parent pressure, classroom reality, and institutional constraint. In public policy they reveal the distance between elegant principle and operational mess.

AI makes these exercises easier to stage because it can populate roles tirelessly. But the teacher still has to define the stakes and the stopping points.

Friction Should Be Productive, Not Gratuitous

Deliberate friction is not the same as sadism. A teacher can overdo it. If every exercise becomes a maze of arbitrary resistance, students stop learning and start merely enduring.

Productive friction has three features:

1. it is tied to a real uncertainty or conflict
2. it forces a revision that improves ownership
3. it ends in a debrief that clarifies what changed

That last step matters. After the recommendation survives challenge, the student should have to summarize:

- what changed
- what remained unresolved
- what evidence would still be needed

That final synthesis is how friction turns into method.

The Classroom as Pressure Chamber

Many classrooms still treat learning as if clarity must always arrive before conflict. AI makes that instinct even stronger because the tool is so good at premature clarity. Cases and simulations let the teacher reverse the order. Conflict comes first. Clearer judgment comes later, if it earns the right to appear.

That is a better preparation for actual decision work, where recommendations are rarely judged in a calm vacuum.

Students should therefore not be protected from all resistance. They should be protected from meaningless resistance. The meaningful kind is part of the education (Bowen & Watson, 2024; Brookfield & Preskill, 2012; Ellet, 2018).

Chapter 15 The Last Right That Cannot Be Outsourced

Universities can outsource many things. They can outsource formatting, transcription, preliminary search, first-pass drafting, translation support, summarization, and certain kinds of routine checking. They may even outsource some explanatory labor without doing much damage.

But they cannot outsource the final judgment of what deserves trust, what deserves action, and what kind of human being the curriculum is trying to form.

That is the last right that cannot be outsourced.

Judgment Is Not Taste

The word judgment is often abused. Sometimes it is used to mean intuition. Sometimes it means preference. Sometimes it merely means confidence.

That is not the meaning here.

Judgment is the disciplined act of deciding under uncertainty with reasons that can be stated, challenged, and revised. It includes:

- framing the issue correctly
- identifying the real alternatives
- weighing evidence with proportion
- seeing the tradeoff rather than merely the slogan
- accepting responsibility for the choice

Generative AI can assist many of the activities around judgment. It can gather, sort, compare, summarize, simulate, and pressure-test. But the university cannot responsibly allow the machine to define what should count as a good choice, what should count as adequate evidence, or what should count as learning.

Those are institutional judgments.

The Human Zone Has Moved, Not Vanished

A common mistake in the public conversation is to imagine that the “human part” of learning must be located where the machine is weakest at the moment. That is too reactive. The human zone is not simply whatever the current model cannot do yet.

The human zone sits where responsibility sits.

That includes:

- deciding where AI belongs in the task
- choosing what evidence is sufficient
- determining which tradeoff matters most
- naming what remains uncertain
- taking public ownership of the recommendation

These functions have become more concentrated, not less important. The machine does more of the surrounding labor. That makes the remaining human decisions more visible and more consequential.

Slowness Has Become Selective

One of the deepest implications of AI in education is that not all slowness needs to be preserved. Some slowness was wasteful. Mechanical summary, repetitive formatting, first-pass transcription, and routine organization can often be accelerated with little loss.

But there is another kind of slowness that has become more rather than less valuable: the slowness required to form judgment.

That slowness appears when the student must:

- compare two answers that both sound plausible
- decide which missing datum actually matters
- refuse a clean recommendation because the evidence is thin
- revise a framework that looked elegant but does not fit the case
- speak the conclusion aloud and hear whether it still holds

The future of good teaching is therefore not slow by default. It is slow where judgment needs time.

Professional Depth Still Matters

The source materials behind this book are blunt on an important point: the problem is not that students use frameworks. The problem is that they use a framework and nothing remains afterward.

A student says SWOT, Seven Powers, PMF, or digital-transformation framework, and the room is expected to applaud. That was already weak before AI. It is even weaker now because the model can perform framework citation effortlessly.

Professional depth means two things now:

1. applying the framework with actual fit
2. criticizing the framework when it distorts the problem

This is where disciplinary education retains its backbone. A finance class, a qualitative research class, a strategy course, and a case discussion classroom still differ because each field knows what kinds of simplification it can tolerate and what kinds are fatal. AI does not abolish that depth. It punishes those who pretend it is unnecessary.

Better Tools Increase the Cost of Bad Standards

When a tool becomes better at producing acceptable work, the institution must either raise its standards or quietly lower itself.

If a university decides that a polished, plausible artifact is enough, then it is choosing a standard that increasingly cheap software can satisfy. It may still hold ceremonies. It may still issue grades. But it is training students to confuse professional-looking language with professional judgment.

The better response is not apocalyptic. It is managerial and curricular.

Raise the standard toward:

- visible ownership
- bounded evidence
- defended choices
- revision under challenge
- signature work
- multi-dimensional assessment

These are not romantic ideals. They are operating consequences of living in a world where fluent output is abundant.

What the Better Classroom Already Looks Like

Across this book, the pattern is consistent. The better classroom does not ban AI out of fear or adopt AI out of fashion. It does six harder things instead:

1. places AI in the right part of the work
2. teaches within a discipline rather than through generic fluency
3. turns prompting into work design
4. uses discussion and probing to test ownership
5. grades process alongside product
6. preserves judgment as the final human act

This is not a utopian blueprint. Versions of it are already visible in the classroom materials that inform this book. Students are asked to calculate, not merely summarize. They are asked to defend, not merely submit. They are asked to revise, not merely polish. They are asked to show their logs, not merely their final pages. They are asked to build artifacts with signature, not merely produce acceptable prose.

That is what a serious educational response looks like.

The University Still Has a Choice

Technology does not remove institutional choice. It simply makes the cost of bad choice arrive faster.

Universities can choose the easy road. They can treat AI as a public-relations problem, write a few policy sentences, keep grading the final artifact as though nothing has changed, and call the result adaptation.

Or they can choose the harder road. They can redesign the course around judgment, process, authorship, and evidence.

The second road is more work. It is also more honest.

The Final Claim

This book has argued for a simple but demanding idea. The goal of AI education in higher education is not to help students arrive at answers sooner. It is to help them arrive at judgment later.

Later means after explanation. Later means after comparison. Later means after challenge. Later means after revision. Later means after the

student has had enough time, pressure, and evidence to know whether the recommendation still belongs to them.

That is what education should protect.

Not the right to type every sentence by hand. Not the right to pretend the machine does not exist.

The right to judge.

That right is still human (Bowen & Watson, 2024; Mollick, 2024; UNESCO, 2024b).

Appendix A Course Contract and AI Use Policy Pack

This appendix is written for the teacher who does not want to spend the first three weeks of the semester pretending that everyone already understands the rules. They do not. The students are guessing. The teacher is often guessing too. A course contract exists to end that guesswork.

The point of the contract is not to make the classroom legalistic. It is to make the classroom legible. Students should know where AI belongs, where it does not, what the course can reasonably expect of them, how process will be documented, and how privacy and unequal access will be handled. If those matters remain vague, the course quietly rewards the most aggressive or least scrupulous users.

What follows is a practical pack: a first-class briefing script, a syllabus-ready contract structure, a use matrix, a disclosure template, and a privacy checklist. Adapt it. Do not admire it.

A First-Class Briefing Script

The opening statement should be short, direct, and specific enough to become operational. Here is a working version:

This course uses generative AI, but not everywhere. In some parts of the work, AI is a useful scaffold. In other parts, using AI would make the task meaningless. You will know those boundaries from the start. We will also assess process, not just finished output, because polished output alone is no longer a trustworthy signal of learning.

That statement should be followed by five clarifications.

1. Which tools the course assumes or permits.
2. Which features matter for the assignments.
3. Which tasks are AI-assisted, AI-limited, or fully human.
4. What evidence of process must be retained.
5. What materials may not be uploaded to outside systems.

The teacher does not need a long speech. The teacher does need to answer the practical questions before the first assignment arrives.

Model Syllabus Language

The following text is intentionally written in a tone that is stricter than friendly but less performative than the usual policy boilerplate.

1. Course Position on AI Use

This course permits the use of generative AI in selected parts of the learning process. AI may be used as a scaffold for explanation, comparison, planning, bounded source review, selected data tasks, and revision support when the assignment instructions explicitly allow it. AI may not be used to replace the student's core judgment, final decision, or live defense of the work unless the instructor explicitly states otherwise.

2. Principle of Placement

The key question in this course is not whether AI is allowed in the abstract. The key question is where AI belongs in the work. Some tasks become better when AI is used to prepare the terrain. Some become worse when AI supplies the answer too early. Students are expected to follow the placement rules for each assignment.

3. Tool Transparency

Students must disclose which AI systems they used when an assignment requires disclosure. If the assignment requires logs, process memos, or revision notes, failure to submit those materials will be treated as failure to submit required components of the assignment.

4. Privacy and Data Handling

Students may not upload confidential, proprietary, restricted, or personally sensitive materials to public AI systems unless the instructor explicitly authorizes a specific bounded environment for that task. When in doubt, students must ask first.

5. Human Responsibility

Students remain responsible for the truthfulness, evidence base, and judgment of anything they submit, whether or not AI was involved in the production process.

The Minimum Contract Structure

Every AI-rich course should answer the following questions in writing.

Contract domain	What students need to know
Allowed uses	Which tasks permit AI and for what purpose
Restricted uses	Which parts of the work must remain fully human
Required evidence	Logs, memos, checkpoints, oral defense, or none
Access assumptions	Which platforms or features the course relies on
Privacy rules	Which files or data may not be uploaded
Collaboration rules	Whether prompts or outputs may be shared across a team
Failure protocol	What students should do when the tool fails or access is uneven
Integrity rules	What counts as misrepresentation of AI-generated work

A Use Matrix That Students Can Understand

Many policies fail because they stay too abstract. A matrix is harder to misread.

Task type	AI-assisted	AI-limited	Fully human
Technical concept clarification before class	Yes, if bounded to assigned material or approved sources	No external synthesis when the task is close reading	No
Case recommendation draft	Only if the assignment calls for a provisional draft to be challenged later	Yes, if the student must still produce the recommendation independently	Sometimes
Data analysis planning	Yes, strongly encouraged for method planning and code review	Execution may require approval first	No
Final judgment or decision memo	Support use only if the assignment permits process-documented drafting	Often	Sometimes
Oral defense or in-class explanation	No	No	Yes
Reflection on personal learning	Limited organizational help only	Often	Often
Confidential company or interview data	Only in approved bounded environments	Usually	Sometimes

Unequal Access Clause

One of the most common hypocrisies in AI-rich teaching is to design assignments around premium features and then pretend the class is operating on

equal footing. If a course depends on long context, file upload, deep research, code execution, or advanced visual output, the course should say so directly.

The contract should therefore include a short access statement such as:

Some assignments in this course benefit from premium AI features such as file upload, longer context, code execution, or advanced research tools. When these features materially affect the task, the assignment will specify whether an equivalent alternative exists, whether the task may be completed in teams, or whether a bounded school-supported environment is required.

That sentence does not solve institutional inequality. It does remove a layer of dishonesty.

Privacy Checklist for Teachers

Before assigning AI-supported work, ask the following:

1. Does the task involve student personal data?
2. Does it involve unpublished case material?
3. Does it involve company data, interview transcripts, or internal documents?
4. Can the material be de-identified?
5. Is there a bounded environment available?
6. Do students understand which materials are off-limits?

If the teacher cannot answer these questions clearly, the task is not ready.

Student Disclosure Template

The disclosure requirement should not become a ritual of screenshots. Ask for structured information instead.

AI Use Disclosure Form

- Assignment title:
- AI system(s) used:
- Stage(s) of use:
- Main purpose of use:
- Source base used by the system:
- One place where the system helped:
- One place where the system was wrong, weak, or misleading:
- One decision that remained fully yours:

This form is short enough to be used regularly and strong enough to reveal whether the student actually understands their own process.

Process Evidence Clause

When an assignment requires process evidence, the teacher should state exactly what counts. For example:

This assignment requires: (1) a brief AI use disclosure, (2) one intermediate artifact such as a prompt log, analysis plan, or source note, and (3) a final author note of 150 to 250 words explaining what changed between the initial and final version.

Specificity matters. If the contract merely says “show your process,” students will produce either too little or too much, and both will be hard to evaluate.

A Short Integrity Clause That Works Better

Most integrity clauses fail because they sound as though they were written for a disciplinary hearing rather than a classroom. A shorter version often works better:

Misrepresentation occurs when a student presents AI-generated reasoning, text, analysis, or decisions as if they were fully their own work in a context where disclosure, limitation, or independent judgment is required by the course.

This is better than listing every possible offense. It names the real issue: misrepresentation in relation to course expectations.

First-Week Demonstration Sequence

The contract becomes more credible when students see it applied immediately. A strong first-week sequence looks like this:

1. show a permissible use
2. show an impermissible or weak use
3. explain the learning logic behind the difference
4. show how process evidence would capture that difference

For example, a teacher might show a case-reading prompt that extracts concepts and maps them to exact locations in the source packet. Then the teacher shows a prompt that asks for the final case recommendation. The class can see the line. One prepares the terrain. The other occupies the territory too early.

A Team Collaboration Clause

Some courses will allow teams to pool access or divide AI-supported work. If so, the contract should say how responsibility is tracked.

Suggested language:

Team members may collaborate in approved AI-supported tasks when the assignment permits team-based work. The team must still identify who framed the task, who inspected the output, who revised the logic, and who is prepared to defend the result.

That clause matters because “we did it together” often means no one can explain what happened.

What to Say When Students Ask the Familiar Questions

“Can I use AI for this assignment?”

Answer with placement, not principle.

You may use AI for concept clarification, initial comparison, or analysis planning. You may not let AI make the final decision for you. The assignment sheet tells you what evidence of process you must keep.

“Do I have to tell you every time I use AI?”

Avoid absolutism.

Not every tiny use requires a confession. But when the assignment requires disclosure or process evidence, you must document the meaningful uses that shaped the work.

“What if I use AI only to improve the English?”

Language polishing is not automatically innocent. It depends on the task. If the task is an argument-writing assignment, the teacher may still want to see an earlier version or an author note. If the task is not primarily about prose style, limited language support may be perfectly acceptable. The contract should say so.

What the Contract Ultimately Does

A good contract does three things at once.

It protects honest students. It gives the teacher something firmer than improvisation. It aligns the course’s stated values with the course’s operating rules.

That is not bureaucracy. That is management.

If the classroom is going to live with generative AI, it should at least do so with its eyes open.

Appendix B Prompt Design Toolkit

This appendix is built on a simple claim: most classroom prompting fails not because the sentence sounds bad, but because the work is badly designed. The model is asked to do too much, too soon, on too little evidence, with too little accountability.

The toolkit below is not a collection of magical phrases. It is a set of reusable teaching moves. Each pattern answers four questions:

1. When should this pattern be used?
2. What is it trying to make the model do?
3. What kind of failure does it prevent?
4. What should still remain for the student?

Teachers should adapt these patterns to their field. Students should be taught to diagnose why a pattern exists, not just to copy and paste it.

Pattern 1 Source-Bounded Explainer

Use it when

Students need help understanding an assigned reading, case, transcript, or packet without replacing the reading.

What it does

It keeps the system inside the assigned material and forces the explanation to point back to exact locations in the source.

Template

You are a course teaching assistant. Use only the uploaded course materials.

Task: Explain the concept of [X] in plain language.

Constraints:

- Do not use outside knowledge unless I explicitly ask for it.
- Do not give me the final answer to the assignment question.
- Point me to where in the material I should read next.

Output format:

1. Plain-language explanation in under 180 words
2. Two exact source locations or section names to read
3. One question I should answer myself after reading

What failure it prevents

The common failure here is the fake helpful answer: a pleasant explanation built from half-source, half-general knowledge, and fully disconnected from the assigned material.

What still remains for the student

The actual reading and the use of that reading in the assignment.

Pattern 2 Plan Before Execute

Use it when

The task involves quantitative analysis, code, database operations, or any workflow where bad execution can produce very confident nonsense.

What it does

It forces the system to expose its method before running the task.

Template

```
I want to explore the following question: [question].
Use Python if needed, but do not execute anything yet.
First, do the following:
1. Explain how you understand the task
2. Identify the variables or fields you would use
3. Propose the analysis logic step by step
4. Write the code you would execute
5. Briefly explain any statistical test or special function you
   selected
Stop there. Wait for my review before execution.
```

What failure it prevents

It prevents the student from confusing executable output with trustworthy reasoning. It also creates a review checkpoint that can be graded.

What still remains for the student

Approving, revising, or rejecting the logic and later interpreting the result.

Pattern 3 Role-Based Explainer

Use it when

The concept is abstract, dry, or difficult to enter without a narrative.

What it does

It gives the system a role that carries an incentive structure or a social frame.

Template

```
You are [role], and I am [learner role].  
Using only the uploaded material, explain [concept].  
Your explanation must include:  
1. Why this concept exists  
2. What problem it solves  
3. One common misunderstanding  
4. One follow-up question I should ask next  
Tone: practical, direct, no jargon unless you define it
```

What failure it prevents

It prevents the model from producing one more neutral encyclopedia paragraph that no student will remember.

What still remains for the student

To test whether the role helped clarify the concept or merely dramatized it.

Pattern 4 Counterposition Generator

Use it when

Students have produced a first answer too quickly and need pressure.

What it does

It makes the system argue against the current answer from a specified stakeholder or analytical perspective.

Template

```
Here is my current position: [paste].
Now act as [skeptical stakeholder / regulator / auditor / board
member / worker].
Your task is not to replace my answer, but to pressure-test it.
Do the following:
1. Identify the assumption doing the most work
2. Name the strongest risk or blind spot
3. Ask three questions that would make me revise the answer
4. Do not provide a final recommendation
```

What failure it prevents

Premature closure.

What still remains for the student

The revision itself.

Pattern 5 Meta-Prompt for System Instructions

Use it when

Students know roughly what kind of assistant they need but cannot yet formalize the workflow.

What it does

It turns a rough description into reusable system instructions.

Template

```
Below is my rough description of the assistant I want.
[paste rough description]

Turn this into a system instruction or custom instruction set for an
AI assistant.
```

The output must contain:

1. Role and objective
2. Step-by-step workflow
3. Constraints and prohibited behavior
4. Output format
5. Error-handling behavior when evidence is missing

Keep the logic strict enough for repeated classroom use.

What failure it prevents

Improvised, unstable prompt behavior across repeated assignments.

What still remains for the student

Testing and revising the instruction after real use.

Pattern 6 Visual Artifact Builder

Use it when

The student understands a concept better once it is turned into a diagram, timeline, SVG, decision tree, or infographic scaffold.

What it does

It forces the model to translate abstraction into explicit structure.

Template

Turn the following analysis into a visual artifact blueprint.

Task: [timeline / decision tree / SVG diagram / infographic structure]

Requirements:

- State the logic of the structure first
- Define the sections or nodes
- Label the key variables or transitions
- If you generate SVG or code, keep labels editable
- Do not add decorative elements that obscure the logic

Return:

1. Structure explanation
2. Text blueprint
3. Code or markup if applicable

What failure it prevents

Visually impressive but logically empty artifacts.

What still remains for the student

Checking whether the visual actually represents the argument and deciding how to revise it.

Pattern 7 Compliance Logic Builder

Use it when

The problem is regulatory or procedural and students are tempted to jump from jargon directly to conclusion.

What it does

It forces the system to separate explanation, threshold logic, calculation, missing-data identification, and recommendation.

Template

```
You are a compliance analyst.  
Scenario: [paste product / region / agreement / rule].  
Work in sequence.  
1. Explain the rule in plain language  
2. Identify the logical checks required  
3. Show what data would be needed  
4. Do not provide the final conclusion yet  
  
After I confirm, move to calculation and risk flags.
```

What failure it prevents

The illusion that a legal or compliance answer can be reached through elegant jargon alone.

What still remains for the student

Choosing the operational response once the compliance logic is clear.

Pattern 8 Reflection on AI Error

Use it when

You want students to learn from the model's weakness rather than merely from its help.

What it does

It turns AI failure into explicit reflection.

Template

Review the following exchange and help me identify where the model was weak.

[paste log or excerpt]

Do not praise the output.

Instead:

1. Identify one likely overclaim
2. Identify one missing source or missing assumption
3. Identify one place where the task was framed too vaguely
4. Suggest one better next move for me as the student

What failure it prevents

The passive acceptance of competent-looking answers.

What still remains for the student

The judgment about whether the criticism is correct and what to do next.

Same Thread vs. New Thread

This operational choice should be taught explicitly because students often manage threads badly.

Stay in the same thread when	Start a new thread when
You are revising the same reasoning object	The task has changed enough to require a reset
Prior critique and context still matter	Old assumptions are now contaminating the work

Stay in the same thread when	Start a new thread when
You want the system to remember earlier constraints	You need a new role, new source base, or new scope
You are editing or tightening an existing workflow	You are starting a different case or a different question

The best classroom move is often not “continue chatting.” It is “edit the previous instruction” or “open a fresh thread because this is a new problem.”

How to Teach Students to Diagnose a Weak Prompt

Give them five questions:

1. Is the role clear?
2. Is the evidence base clear?
3. Is the sequence clear?
4. Is there any forbidden behavior?
5. What is still left for the human?

If the student cannot answer those questions, the prompt is usually underdesigned.

Common Prompt Failures

Failure 1 The Everything Prompt

The prompt asks for explanation, analysis, recommendation, visualization, and reflection all at once. The output is long and usually thin.

Failure 2 The Source Leak

The assignment is supposed to be grounded in a specific packet, but the system quietly imports outside generalities.

Failure 3 The Premature Close

The prompt asks for a final answer before the problem has been decomposed.

Failure 4 The Decorative Role

The role sounds colorful but does not change what evidence or incentives matter.

Failure 5 The Human Erasure

The system does the framing, the evaluation, the recommendation, and the polish. The student transports the file.

Teachers should show these failures directly. Students learn faster when they see what bad prompting looks like in public.

A Short Prompt Audit Checklist

Before approving a classroom prompt, ask:

- Does this prompt preserve the learning task?
- Does it control the evidence base well enough?
- Does it create a review checkpoint where needed?
- Does it leave real judgment for the student?
- If the prompt succeeds perfectly, does the assignment still have educational value?

That last question is the most important. If a perfect prompt would finish the course for the student, the assignment is poorly designed.

The Real Aim

Prompt literacy in higher education should not mean the ability to impress people with elaborate wording. It should mean the ability to design work so that AI supports analysis without displacing ownership.

That standard is harder than clever phrasing. It is also more useful.

Appendix C Case Library for AI-Rich Teaching

The purpose of this appendix is not to give teachers a museum of interesting exercises. The purpose is to show how real classroom cases can be redesigned so that AI makes the work more explicit without making the student dispensable.

Each case below includes four elements:

1. the teaching objective
2. the placement of AI in the workflow
3. the friction or checkpoint that preserves learning
4. the debrief question that converts activity into judgment

Treat these as patterns, not relics.

Case 1 One International and the Nature of Trade

Teaching objective

Help students move from a simplistic “buy low, sell high” understanding of trade toward a view of trade as rule-bound, interest-laden, and politically structured.

Classroom design

Students enter the topic through a narrative role frame. The model is asked to behave as a senior manager in a trading company and explain mercantilism, absolute advantage, comparative advantage, strategic trade policy, and de-risking to a junior colleague. The role is not there for drama. It is there to turn textbook abstraction into situated explanation.

AI placement

AI is used at the concept-clarification stage and later at the counterposition stage. Students first use the model to build a conceptual skeleton. Then

the teacher probes the skeleton in class by asking what breaks when political rationality overrides economic efficiency.

Preserved friction

The class does not stop at textbook logic. Students must confront labor-transition friction, industrial policy, export controls, and national-security framing. This reintroduces the real-world mess that the clean theory tends to hide.

Debrief

What part of the theory remains useful after political power enters the problem, and what part needs revision?

Case 2 Ghana, South Korea, and Comparative Advantage

Teaching objective

Teach students to distinguish production from final consumption and to understand that comparative advantage is about opportunity cost and gains from exchange, not merely about which country produces more of one item.

Classroom design

Students are given numerical parameters and asked to compute output before trade, output after specialization, opportunity cost, and final consumption after exchange. The model may assist with table generation, Python calculation, and SVG or PPF production, but only after the logic is staged explicitly.

AI placement

AI enters as calculator, table builder, and visualizer after the student or class has clarified the meaning of the variables.

Preserved friction

The exercise forces the class to confront a common misconception: if one good declines after trade, does that mean trade failed? Students must compare overall consumption bundles, not single commodities. The teacher can also force a checkpoint by asking what exchange rate would make both sides better off and what happens if that condition is not met.

Debrief

What conceptual confusion disappeared only after the class had to separate production logic from consumption logic?

Case 3 The Campus Shopping District and FTA Rules

Teaching objective

Help beginners understand FTAs, tariff reduction, rules of origin, and regional value content without drowning them in jargon on the first encounter.

Classroom design

The teacher begins with a life analogy such as a campus shopping-district alliance. Students work through a three-act story, conflict events, and four-panel comic versions before migrating from the analogy to a real product.

AI placement

AI is used as a scaffold to build the analogy, extend the story, and map the analogy to the real trade concept. Only later does the workflow move into concrete products such as Bluetooth earphones or e-bikes under CPTPP logic.

Preserved friction

The key friction comes when the analogy is no longer enough. Students must move into real cost structures, identify originating and non-originating components, compute threshold logic, and confront missing data. They discover that customs does not care about the beauty of the teaching story.

Debrief

At what point did the analogy stop helping and the compliance logic begin demanding harder evidence?

Case 4 E-Bike Rules of Origin and Supply Chain Remapping

Teaching objective

Show that compliance is not just a legal matter but a strategic supply-chain problem.

Classroom design

Students are given a concrete cost structure for an electric bicycle assembled in Vietnam for export to Japan. They first outline the logic steps needed to assess qualification under CPTPP. Only after that do they calculate regional value content. Only after that are they asked to judge qualification and propose remapping options.

AI placement

AI helps sequence the logic, perform stepwise calculation, and build a compliance decision tree. It should not jump to the final conclusion before the checks are visible.

Preserved friction

Students are forced to identify what data would be difficult to obtain in a real supply chain. They discover that upstream cost transparency is weak, that documentation matters, and that a regulatory threshold can force real sourcing changes. The classroom can then connect compliance to strategy.

Debrief

What is the biggest gap between the clean classroom example and actual customs enforcement?

Case 5 Investment in the Fog

Teaching objective

Teach students how to constrain open-ended geopolitical or investment risk analysis so that the model does not simply produce elegant atmosphere.

Classroom design

Students investigate a situation in which a government changes the terms of an investment after attracting the investor. The prompt is intentionally long and includes precise domain vocabulary: tax holidays, land incentives, special-zone treatment, retroactive taxes, and similar terms.

AI placement

AI is used only after a new thread is opened and the scope is narrowed. The system is instructed to cross-check sources and return a specific report structure.

Preserved friction

The teacher's design blocks vague inquiry. Students must learn that the right keyword is part of domain knowledge and that research quality depends heavily on how the search corridor is defined. They also experience the fact that risk analysis depends on both source quality and output structure.

Debrief

What part of this exercise depended on better prompting, and what part depended on better understanding of the domain itself?

Case 6 IKEA Data and the Discipline of Review

Teaching objective

Show students that data analysis with AI should be governed by checkpoints rather than by immediate execution.

Classroom design

Students begin with a tentative business question about product popularity, review count, or some related pattern in the dataset. The model is asked to propose an analysis plan and code but not to execute. The class inspects the proposed logic and method. Only then is execution approved, revised, or redirected.

AI placement

AI helps propose method, generate code, and later produce a management report. It does not initially own the research question, the choice of method, or the interpretation.

Preserved friction

When the class realizes that a formal statistical test may not fit the data or the business question, the task is deliberately redirected toward a management report. This teaches students that method choice is part of analysis rather than an automatic add-on.

Debrief

What did the pre-execution review reveal that would have remained invisible in a polished final chart?

Case 7 Qualitative Analysis with IDEO and Hora

Teaching objective

Teach students to build a repeatable AI-supported qualitative workflow without allowing the model to invent facts or overread weak evidence.

Classroom design

Students learn to define a system instruction around goal, steps, constraints, and output format. They build project-bounded assistants. Later, they compare a case's existing qualitative methods with external method cards to ask what richer forms of evidence or user insight might still be available.

AI placement

AI is used first as a bounded qualitative assistant and then as a framework comparison tool. It is not used as a substitute researcher.

Preserved friction

Students must decide which method actually fits the problem. It is not enough to ask the system for more methods. The class has to judge whether the additional method would generate better insight or merely more activity.

Debrief

Which part of the problem required methodological judgment rather than prompt fluency?

Case 8 Borusan Cat and Digital Transformation Resistance

Teaching objective

Show students that digital transformation is not mainly a technology story. It is a political, operational, and behavioral story inside an organization.

Classroom design

Students simulate internal conflict among procurement staff, maintenance managers, technicians, salespeople, and customers around predictive maintenance systems, direct ordering, RPA, CRM integration, and customer journey redesign. They may also build financial sensitivity tables, SVG process comparisons, or service blueprints.

AI placement

AI is used to populate roles, construct artifacts, and reveal tensions across organizational positions. It is especially useful for simulating resistance and translating sensor data into decision artifacts.

Preserved friction

The classroom does not let the transformation remain a CEO-level success story. It reintroduces fear of lost commission, frontline distrust, time scarcity, local relationship logic, and operational mismatch across actors.

Debrief

Which part of the transformation creates value at the firm level but pain at the actor level, and how should that conflict be managed?

Case 9 Platform War: Grab, Gojek, and SoftBank

Teaching objective

Teach competitive strategy as a field of conflict among business logic, capital structure, regulation, public backlash, and labor concerns.

Classroom design

Students select roles in the 2018 Southeast Asian platform war and conduct several rounds of interaction. One role sees the world through growth and dominance, another through expansion and localization, another through portfolio discipline, and another through public harm and regulatory scrutiny.

AI placement

AI acts as the role engine. It creates pressure, continuity, and memory across rounds. It is not the strategist of record.

Preserved friction

The multi-round structure forces students to survive challenge rather than present a single elegant answer. Because different roles care about different success criteria, the room cannot settle too quickly into one managerial storyline.

Debrief

Which recommendation remained attractive only because one important stakeholder had been excluded from the room?

Case 10 Walmart, Superstore, and Front-line Digital Anxiety

Teaching objective

Help students see that digital transformation must be understood from the shop floor as well as from the top.

Classroom design

Students combine the case material with scenes or atmosphere drawn from retail labor culture. They ask the model to simulate dialogue between a manager and a worker, or to redesign a customer journey after new digital systems are introduced. The room can also use AI to build visuals of industry evolution or editable SVG artifacts.

AI placement

AI is used to make labor experience, fear, and operational friction more vivid. It can also help structure industry history or strategy visuals.

Preserved friction

The class does not allow the transformation story to remain abstractly pro-innovation. It asks what happens to the worker, to the schedule, to trust, to surveillance, and to the economics of service. Students are forced to quantify claims and connect slogans to financial and operational logic.

Debrief

When does digital transformation create a stronger enterprise and when does it simply shift pain downward?

Case 11 Clearbrook and Database Design

Teaching objective

Show that AI-supported technical work still requires human understanding of underlying structure.

Classroom design

Students begin by learning a notation system such as Barker's notation and then sketch the entity relationships by hand before asking AI to repair or describe the logic. Only after the structure is stable do they ask the system to produce downloadable forms or editable SVG diagrams.

AI placement

AI is used after manual conceptual work. It helps clean, translate, and operationalize the structure rather than invent it from nothing.

Preserved friction

The teacher insists on a hand-built first pass because students who skip that step cannot tell whether the AI has simplified the logic fatally. This is a powerful example of why “human first” still matters in certain stages of technical learning.

Debrief

What could you only notice because you had first tried to build the structure yourself?

How to Use the Case Library

A case is not strong because it is topical or dramatic. It is strong when the teacher can state clearly:

- what intellectual move the case is supposed to force
- where AI belongs in that move
- where the friction should return
- what the final debrief should make visible

If those four points are clear, almost any decent case can become a serious AI-rich teaching exercise. If they are not clear, even a famous case becomes one more elegant way to avoid the actual work.

Appendix D Assessment Toolkit for AI-Rich Courses

This appendix is written for instructors who already know that the final artifact is no longer enough, but who do not yet have a practical system for assessing process without drowning in administrative debris.

The aim is not to grade every micro-step of student work. The aim is to collect enough evidence to answer a more serious question: did the student frame, inspect, revise, and own the work in a way that matches the course's learning goals?

That requires better artifacts, better rubrics, and better questions.

The Minimum Assessment Stack

If the course uses AI in a meaningful way, the assessment stack should usually include at least three layers.

Layer	Function
Product	Shows the final artifact and audience fit
Process evidence	Shows framing, method, and revision path
Live ownership	Shows whether the student can still stand behind the work

The specific form of each layer will vary across disciplines. A case course may use an oral defense. A research-method course may use a method memo. A technical course may use an analysis plan and review checkpoint. The principle remains the same.

A Working Rubric

The rubric below is designed to be reused across multiple assignment types.

Dimension	4 Strong	3 Adequate	2 Weak	1 Failing
Problem framing	Defines the real decision or analytical question with precision	Defines the question but leaves some ambiguity	Uses a loose or borrowed question	Never clarifies the actual problem
Evidence discipline	Uses sources appropriately, visibly, and proportionately	Uses sources competently with minor gaps	Source use is weak, vague, or overdependent on AI summary	Evidence base is missing, confused, or misrepresented
Method and workflow	Shows a coherent workflow with sensible checkpoints	Workflow is mostly coherent but underexplained	Workflow exists but looks improvised or uninspected	No credible workflow visible
Revision quality	Revises meaningfully in response to critique or evidence	Revises but mostly at the surface level	Revision is minimal or cosmetic	No real revision evident
Defended judgment	Can explain and defend the final choice under questioning	Can explain the choice with some weakness	Struggles to defend reasoning	Cannot own the reasoning

Dimension	4 Strong	3 Adequate	2 Weak	1 Failing
Product quality	Final artifact is clear and fit for audience	Product is usable with minor weaknesses	Product is confusing or misaligned	Product is not fit for purpose

A Short Process Log Template

The log should not be a screenshot graveyard. It should be an explanation of the work path.

AI Process Log

- Assignment:
- Initial question or task framing:
- AI system(s) used:
- Where AI entered the workflow:
- One output or suggestion I accepted:
- One output or suggestion I rejected:
- What changed after critique or review:
- What evidence or source set shaped the final version:
- One uncertainty that remained at the end:

This format is short enough to be sustainable and rich enough to be meaningful.

A Revision Memo Template

The revision memo is often more revealing than the log because it shows whether the student actually changed the work.

Revision Memo

1. What was the biggest weakness in the first version?
2. What critique, evidence, or discussion pressure led you to change it?
3. What did you revise in the final version?
4. What still remains imperfect or uncertain?

The teacher learns more from these answers than from one more immaculate paragraph in the final report.

Oral Defense Questions

Many instructors want to include a live element but are unsure what to ask. Keep it short.

Suggested questions:

1. What is the single most important decision you made in this work?
2. What did the model get wrong, or what did it fail to see?
3. What evidence would make you revise your conclusion now?
4. Which part of the work was fully yours from start to finish?

These questions are useful because they cannot be answered well by a student who never really possessed the work.

Assignment-Type Variations

Case analysis

Require:

- final memo
- process log
- one-page revision memo
- short oral defense

Data analysis

Require:

- approved analysis plan
- code or workflow summary
- final management report
- note on why the chosen method fit the question

Qualitative analysis

Require:

- system instruction or workflow setup
- source note
- coded or themed output
- reflection on overinterpretation risk

Discussion-based course

Require:

- pre-class AI-supported prep note
- in-class position record
- post-class revision reflection

Assessment becomes more efficient when the course reuses the same evidence shapes.

AI-Assisted Review of Process Evidence

Teachers may use AI to help inspect logs and revisions, provided the final interpretive judgment remains human.

A practical workflow looks like this:

1. collect logs in a standard format
2. ask AI to cluster common failure modes
3. review the clusters as an instructor
4. identify which patterns deserve class-wide intervention
5. use the clustering for formative feedback, not automatic grading

This is particularly useful in large classes. Instead of reading each log as though it were a separate mystery, the instructor can see patterns:

- vague initial framing
- overreliance on first-pass answers
- missing source boundaries
- cosmetic revision
- inability to name what remained uncertain

That pattern view allows better teaching without surrendering judgment.

What Not to Reward

Instructors often know what they want to reward. They are less explicit about what they no longer wish to reward.

In an AI-rich course, do not overreward:

- polish without trace
- framework citation without application
- elegant prose hiding weak evidence
- charts without method explanation
- reflections that sound mature but contain no actual revision

Naming these anti-signals helps students recalibrate.

Sample Feedback Language

The feedback vocabulary also needs updating.

When the work is polished but thin

The final artifact is clean, but I cannot see enough of the reasoning path to trust the conclusion fully. The next version needs stronger process evidence and clearer ownership of the key decision.

When the process is strong but the final artifact is uneven

Your process evidence shows real thinking and revision. The next step is to turn that stronger judgment into a clearer final artifact for the intended audience.

When AI use is visible but not well controlled

You used AI actively, but the workflow still gives the model too much of the evaluative work. The revision should show where you challenged, rejected, or redirected the model rather than simply accepting the output.

When the student demonstrates strong ownership

The strongest part of this submission is not just the final recommendation but the way you revised the task after critique. The process evidence and defense show that the argument is genuinely yours.

Time Management for Instructors

The main objection to process grading is not philosophical. It is practical. Faculty worry that the workload will become absurd.

That risk is real. The solution is standardization.

- use the same log form across assignments
- keep revision memos short
- sample oral defenses rather than making every defense identical in length
- use AI to cluster recurring problems
- build rubrics that can be reused

Process assessment becomes manageable when the evidence structures are stable.

The Purpose of the Toolkit

This appendix is not trying to make grading bureaucratically perfect. It is trying to make grading intellectually more honest.

If the course claims to value judgment, revision, and evidence discipline, then those things should appear in the assessment architecture. If they do not, the course will keep grading the artifact as though the artifact were the whole story.

It no longer is.

Appendix E Building the Course Operating System

This appendix turns the operating-system argument in Chapter 13 into an implementation checklist. The point is not to standardize every course into the same shape. The point is to ensure that any course using generative AI has enough structure to remain coherent across a semester.

The operating system should answer three institutional questions:

1. Where do students work?
2. What instructions govern the tools?
3. How does the course preserve evidence, revision, and memory?

The Core Build Sequence

The cleanest build sequence is usually:

1. define learning goals
2. map where AI belongs in the task flow
3. create source-bounded workspaces
4. write or adapt system instructions
5. build prompt books for repeated classroom moves
6. define the archive and review loop

Each step should produce an artifact the course can reuse.

Workspace Architecture

At minimum, a course needs three kinds of spaces.

Space	Purpose
Course source space	Stores syllabus, case packets, readings, slides, transcripts, and approved materials

Space	Purpose
Task execution space	Where students run bounded tasks, prompts, or projects
Archive space	Where logs, process notes, prompt books, and examples are retained

A Practical Folder Logic

Even when tools are cloud-based, the course should behave as if it had a directory structure.

Suggested structure:

```
course-root/  
  syllabus/  
  cases/  
  readings/  
  prompt-books/  
  system-instructions/  
  assignments/  
  examples/  
  logs/  
  rubrics/
```

This matters because once materials accumulate, retrieval becomes part of teaching quality. A course that cannot find its own assets cannot build on its own method.

System Instruction Layers

The system instruction layer usually needs at least three variants.

1. Student support mode

Purpose:

- explain assigned material
- stay inside the source boundary
- avoid becoming an answer key
- ask guiding questions

2. Teacher design mode

Purpose:

- help the instructor build lessons, prompts, rubrics, and cases
- analyze likely student misunderstandings
- compare possible teaching sequences

3. Review or grader-assistant mode

Purpose:

- summarize logs
- compare revisions
- surface common patterns
- prepare feedback clusters for human review

The important principle is role separation. A tool that is helping students should not automatically behave like the tool that helps the teacher design the course. Mixing those roles invites trouble.

A Student Support Instruction Skeleton

Role:

You are the course learning assistant for [course].

Objective:

Help students understand the assigned materials without replacing their core analytical work.

Source rules:

Use only the uploaded or approved course materials unless the prompt explicitly authorizes outside information.

Behavior rules:

- Do not provide the final answer to graded case questions unless explicitly allowed
- Ask guiding questions when the student's request would bypass the learning task
- Point the student back to exact sections, cases, or slides when possible
- State clearly when the answer is not available in the current source set

```
Output rules:
- Keep summaries short
- Prefer structure the student can inspect
- End with one next question or one next reading move
```

This skeleton is intentionally plain. In most courses, plainness is a virtue.

Prompt Books as Reusable Course Moves

Prompt books should be organized by teaching move, not by technological novelty.

Suggested prompt-book sections:

- concept explainer
- source-bounded case reader
- role-play discussion starter
- counterposition generator
- plan-before-execute analysis prompt
- decision-tree builder
- revision and reflection prompt

When prompt books are organized this way, they become teachable assets. Faculty can ask whether the course has a good counterposition prompt rather than whether it has “something cool for AI.”

Version Control and Change Logs

If a course uses prompt books or system instructions repeatedly, it should also keep a simple change log.

Example:

Date	Asset changed		Why it changed
Week 2	Discussion prompt	starter	Students were getting a summary instead of opposing views
Week 4	Data-analysis prompt	plan	Needed clearer instruction to stop before execution

Date	Asset changed	Why it changed
Week 7	TA system instruction	Added stronger source restriction and teacher-mode boundary

The Banned File Principle

One small detail from the source materials deserves to become a general practice: some files should be ignored by design.

If a template, answer key, or sensitive file should not be read by the student-facing assistant, then the course should explicitly tell the system to ignore it. This matters because many AI tool failures are not failures of intelligence at all. They are failures of instructional boundary-setting.

The larger principle is simple: what the course would forbid a human TA from using casually, it should also forbid the machine from using casually.

Weekly Operating Rhythm

A stable weekly rhythm makes the course feel governed rather than improvised.

Suggested cycle:

1. source-bounded preparation
2. in-class probe or discussion
3. short revision artifact
4. archived evidence
5. instructor review of patterns

This rhythm works in many fields because it aligns activity, evidence, and feedback.

Tool Choice Rules

The course operating system should also contain rough tool-choice rules so students do not treat every platform as interchangeable.

Example:

- use source-bounded systems when exact course material matters
- use research systems when the task requires explicit external search
- use code-capable systems when computation or file output must be auditable
- use editable-output systems when the artifact will need human revision

The point is not to create tribal loyalty. The point is to make tool choice subordinate to task logic.

Faculty Preparation Checklist

Before the semester begins, the instructor should be able to answer yes to most of the following:

- The course has a written AI position and contract.
- The main source spaces are curated.
- Student-facing system instructions exist.
- Teacher-facing system instructions exist.
- Prompt books exist for repeated teaching moves.
- The course has a process-log format.
- The course has at least one live ownership mechanism.
- The archive structure is in place.
- Students know what materials may not be uploaded.
- The instructor can demonstrate one strong and one weak AI use on day one.

If too many boxes remain unchecked, the course is probably not yet operationally ready.

Department-Level Extensions

Once one course is stable, departments can scale the logic:

- shared source repositories
- shared prompt-book libraries
- common process-log templates
- common privacy rules
- faculty workshops around actual use cases

Departments should not rush to standardize everything. They should standardize the parts that reduce confusion without flattening disciplinary difference.

What Good Operations Feel Like

When the operating system is working, the course feels calmer.

Students know where to begin. Teachers know what to inspect. The room spends less time negotiating tool behavior and more time on the disciplinary work itself.

That calm is not trivial. It is one of the hidden signs that the course has become governable.

Appendix F A Sixteen-Week Teaching Map

This appendix translates the book into a semester-scale design. It is not the only valid sequence, but it does reflect the central argument of the manuscript: start with placement and thinking, move into workflow and course governance, then build toward discussion, signature work, assessment, and judgment.

The map assumes a higher-education course that wants to use generative AI seriously without turning the entire class into software worship.

Semester Logic

The sixteen weeks are organized into five arcs:

1. orientation and operating rules
2. thinking and workflow
3. discussion and disciplined application
4. signature work and assessment
5. simulation, revision, and judgment

Each week below includes a goal, an AI placement rule, an artifact, a likely failure, and a debrief question.

Weeks 1 to 4: Orientation and Thinking

Week	Main goal	AI place- ment	Student artifact	Common failure	Debrief question
1	Intro- duce course contract and tool	AI dis- cussed, demon- strated,	Signed use state- ment plus short re- flection	Students hear pol- icy but not logic	What work in this course would

Week	Main goal	AI placement	Student artifact	Common failure	Debrief question
	boundaries	and bounded			become meaningless if AI did it too early?
2	Clarify that the question is placement, not permission	AI used to compare weak and strong task placement	Placement memo on one assignment	Students reduce the issue to morality slo-gans	Where exactly should AI enter and leave this task?
3	Expand thinking before answer generation	AI used for concept clarification and structured comparison	Concept comparison note	Students accept the first explanation too quickly	What did the model explain well, and what did it hide?
4	Teach multi-turn prompting as work design	AI used in staged prompt sequence	Prompt redesign exercise	Students keep asking one-shot everything prompts	Which checkpoint should have been inserted before the answer stage?

Weeks 5 to 8: Workflow, Tools, and Course Redesign

Week	Main goal	AI placement	Student artifact	Common failure	Debrief question
5	Show how false mastery appears	AI used to generate plausible but weak output for diagnosis	Weak-vs-strong output critique	Students confuse polish with understanding	What made this answer look better than it actually was?
6	Prompting as work design	AI used to build role, sequence, and constraint structure	Prompt stack worksheet	Students still treat prompting as wording only	What did the prompt force the model to do in what order?
7	Make the course contract operational	AI used minimally; focus is on rules, privacy, disclosure	Assignment-specific use matrix	Students assume one policy fits every task	Which rule is really about learning rather than about discipline?
8	Keep the teacher above the tool	AI used in live faculty-style demonstration	Demonstration critique note	Students trust tool confidence more than	What did the teacher notice that the model did not?

Week	Main goal	AI place- ment	Student artifact	Common failure teacher judgment	Debrief question
------	-----------	-------------------	---------------------	--	---------------------

Weeks 9 to 12: Discussion, Signature, and Assessment

Week	Main goal	AI place- ment	Student artifact	Common failure	Debrief question
9	Shift from lecture to probe	AI used for provisional pre-class preparation	Probe response memo	Students arrive with AI notes but no ownership	Which probe exposed the real weakness in your preparation?
10	Run the discussion class-room	AI used to widen perspectives and populate counter-positions	Revised discussion position	Students quote the model instead of arguing	Which counter-position actually forced you to change your view?
11	Build signature work	AI used to support artifact building, not author replacement	Signature assignment proposal or draft	Students choose generic deliverables	Where does your work visibly stop being summary and

Week	Main goal	AI place- ment	Student artifact	Common failure	Debrief question
					become yours?
12	Grade process, not just product	AI used for reflec- tion on workflow and revi- sion	Process log plus revision memo	Students docu- ment ac- tivity but not judg- ment	What de- cision in your process was gen- uinely yours?

**Weeks 13 to 16: Operating System, Friction,
and Judgment**

Week	Main goal	AI place- ment	Student artifact	Common failure	Debrief question
13	Build the course operating system	AI used to help formalize instruc- tions and archives	Course OS map or project setup	Students have tools but no sys- tem	What part of your workflow still de- pends too much on memory or impro- visation?
14	Use cases and sim- ulations with fric- tion	AI used as role engine, chal- lenger, or stress tester	Simu- lation record and revi- sion note	Students treat role-play as enter- tainment	What pressure changed your rec- ommen- dation most?

Week	Main goal	AI placement	Student artifact	Common failure	Debrief question
15	Prepare final signature deliverable and defense	AI used only within stated bounds	Final draft plus defense prep	Students polish instead of revising substantively	What evidence would still make you change your final decision?
16	Synthesize the judgment argument	AI used lightly or not at all	Final author note or oral synthesis	Students confuse technical fluency with education	What kind of judgment should never be outsourced in this field?

Variants by Course Type

The same semester logic can be adapted.

Strategy or case-based business course

- use more role-play and stakeholder conflict
- emphasize case memos, decision trees, and oral defense
- use AI heavily in pre-class preparation and counterposition generation

Research-method course

- emphasize source control, workflow design, and method notes
- use AI for coding assistance, bounded synthesis, and method comparison
- increase the weight of logs and revision memos

Technical or data-heavy course

- emphasize plan-before-execute checkpoints
- use AI for code drafting and structured output only after logic review
- make intermediate artifacts count toward the grade

Teaching or policy course

- emphasize discussion, stakeholder analysis, and simulation
- use AI to widen perspectives and surface implementation friction
- attach all AI use to reflective justification

A Weekly Faculty Routine

The student map becomes much more powerful when the instructor keeps a weekly faculty routine alongside it.

Suggested routine:

1. review last week's logs or artifacts quickly for pattern detection
2. choose one pattern to address in the next class
3. prepare one weak-use example and one strong-use example
4. decide where the probe will land
5. define what evidence of revision will be collected

This keeps the course from becoming a sequence of disconnected AI activities.

The Point of the Map

A semester map matters because otherwise AI enters the course as a series of local improvisations. The class tries one prompt here, one tool there, one flashy exercise somewhere else, and calls the result innovation.

The better approach is cumulative. Each week should strengthen a habit the later weeks depend on:

- from placement
- to design
- to challenge
- to authorship
- to judgment

That is how a course becomes more than a pile of prompts.

Appendix G Reusable Prompt Books

This appendix turns several classroom flows into reusable prompt books. The point is not that these prompts should be copied without thought. The point is that each sequence encodes a teaching move that has already been tested in a live classroom.

Each prompt book below is organized around progression. Do not collapse the steps unless the students have already earned the right to skip them.

Prompt Book 1 Trade Theory as Staged Conversation

Purpose

Move students from intuitive but shallow ideas about trade into a more disciplined understanding of mercantilism, comparative advantage, strategic trade policy, and de-risking.

Step 1 Senior-manager explainer

You are a senior manager in an international trading company, and I am a new intern who thinks international trade is just "buy low, sell high."

Using the assigned course material only, explain:

1. Why trade does not happen in a vacuum
2. What mercantilism is
3. Why mercantilism is a zero-sum logic
4. Why that mindset creates risk for multinational firms today

Tone: practical, clear, slightly senior, no textbook fog

End with one question I should ask next

Step 2 Blind spot probe

I still do not understand why governments cared so much about gold and silver.

Explain the logic in plain language and show what that logic missed.

Do not exceed 180 words.

Step 3 Comparative advantage move

Continue in the same mentor role.

Explain the difference between absolute advantage and comparative advantage.

Then answer this:

If one country is better at making everything, why would it still benefit from trade?

Use opportunity-cost logic.

End with one short metaphor that a classmate would remember.

Step 4 Political rationality challenge

Now switch roles.

You are a chief industry strategy advisor in government.

I am a free-trade believer asking why states subsidize chips, EVs, or strategic sectors.

Refute me using:

1. economies of scale
2. first-mover advantage
3. strategic dependence risk

Then explain when political rationality overrides pure economic efficiency.

Teaching note

The sequence matters because students first need a scaffold, then a clean conceptual distinction, then a harder challenge that reintroduces power and politics.

Prompt Book 2 Comparative Advantage with Calculation and Visual Output

Purpose

Force students to connect conceptual trade logic to actual calculation, code, and visual structure.

Step 1 Calculation scaffold

Explain absolute and comparative advantage using the following two-country, two-good setup.

First distinguish clearly between:

1. production before trade
2. production after specialization
3. final consumption after exchange

Then calculate each stage step by step.

Return:

1. markdown table
2. opportunity-cost explanation
3. terms-of-trade explanation

Do not generate a chart yet.

Step 2 Code checkpoint

Now write Python code for the calculation process.

Do not execute yet.

Explain each variable and each formula in plain language.

Step 3 Visual conversion

Turn the comparative-advantage case into a visual artifact.

Option A: Mermaid or structured diagram

Option B: SVG or Python-based chart

Requirements:

- label axes clearly
- distinguish before-trade and after-trade states
- keep text editable where possible
- show the steps so a beginner can follow them

Teaching note

The most important checkpoint is between Step 1 and Step 2. Students must not treat calculation as something the tool can simply “do later” without understanding the variable logic.

Prompt Book 3 FTA and Rules of Origin

Purpose

Use AI as a cognitive scaffold that moves students from analogy to actual compliance logic.

Step 1 Life analogy

You are my beginner coach for international trade.

Use a campus shopping-district alliance as a life analogy to explain:

1. FTA
2. tariff reduction
3. rules of origin

Your answer must include:

- a short explanation of each concept
- a three-act story skeleton
- two possible conflict events
- one true/false check question
- the next prompt I should ask

Keep the answer short and concrete.

Step 2 Comic expansion

Take the conflict event I selected and expand it into a four-panel comic draft.

For each panel:

1. describe the scene
 2. state the real trade concept
 3. identify where the need for a rule first appears
 4. suggest three possible rule types
- End with a one-sentence question I should answer myself.

Step 3 Real-product mapping

Now map the shopping-district analogy to a real product:
[Bluetooth earphones / e-bike / EV / other product].

Create a table with two columns:

1. analogy element
2. real supply-chain meaning

Do not compute compliance yet.
First make the mapping explicit.

Step 4 Logic before conclusion

You are now an international trade compliance analyst.
Our final product is [paste product and cost structure].
List the logical steps needed to evaluate whether this product
qualifies as originating under [agreement].
Do not give the final answer yet.

Step 5 Forced checkpoint

Now execute the RVC or origin calculation checkpoint.
Show:

1. the formula
2. what counts as non-originating material
3. the threshold being tested
4. any missing-data problem
5. any borderline risk

Step 6 Decision tree

Translate the compliance logic into a decision tree for managers.
Use yes/no branching.
Begin with an executive summary of under 50 words.

Teaching note

The analogy is not the destination. It is the ramp. The class should feel the point at which the analogy stops carrying the load and the regulatory logic takes over.

Prompt Book 4 Investment in the Fog

Purpose

Teach students how to constrain a geopolitical or investment-risk investigation so that the answer becomes evidential rather than atmospheric.

Step 1 Reset the context

This is a new problem. Ignore previous threads or assumptions.
We are now investigating post-entry investment risk in a host country.
Open with a short restatement of the investigation target and the likely risk category.
Wait for confirmation before deeper analysis.

Step 2 Long-form constrained research prompt

Investigate cases in which a government invited foreign investment and later altered the operating terms in ways that harmed the investor.

Focus on:

- tax holidays withdrawn or reduced
- land incentive changes
- special economic-zone treatment altered
- retroactive taxes or back-charged obligations
- regulatory reinterpretation after entry

Requirements:

- cross-check across multiple credible sources
- distinguish allegation from documented policy change
- return the answer in a fixed report structure:
 1. case summary
 2. policy change
 3. investor impact
 4. evidence base
 5. lessons for future entry strategy

Step 3 Pressure-test the output

Act as a skeptical board member.
What assumptions in the report are still too weak for a capital-

```
allocation decision?
List:
1. missing evidence
2. country-comparison gaps
3. legal ambiguity
4. what would need independent verification
```

Teaching note

The long prompt matters because it supplies the system with the domain vocabulary needed to search for the right class of case instead of producing generic political-risk theater.

Prompt Book 5 Borusan Cat and Internal Resistance

Purpose

Use AI to surface the organizational politics of digital transformation rather than merely its technological promise.

Step 1 Internal meeting simulation

```
Based on the case, simulate an internal discussion among:
1. procurement
2. maintenance
3. frontline technicians

Topic: the value and risk of adopting the predictive maintenance
system.
Rules:
- each actor must speak with its own incentive logic
- let disagreement become explicit
- do not resolve the conflict too neatly
```

Step 2 Sales resistance role-play

```
You are a veteran salesperson worried that the new app and e-
commerce function will destroy your commission and your customer
relationships.
```

```
I will play the sales manager.  
Stay resistant until I offer a concrete policy response.  
Respond only in character.
```

Step 3 Artifact conversion

```
Generate one of the following based on the case:  
- a simple customer journey map  
- a CRM lead card  
- an SVG process comparison  
- a sensitivity table for business-model risk  
  
Keep the artifact editable and practical.
```

Teaching note

The richest learning often comes not from the final artifact but from the conflict across actors. The model is useful here because it can keep conflicting voices alive without the teacher performing all of them manually.

Prompt Book 6 Platform-War Role Play

Purpose

Use AI to stage strategic conflict among platform actors with different success criteria.

System prompt

```
You are a business-school case simulation guide.  
We are running a role-based strategy exercise around a regional  
platform war.  
  
Workflow:  
1. Offer the student a choice of roles  
2. Stay fully in role during the discussion rounds  
3. Push back on weak logic  
4. After 3 to 4 rounds, step out of role and summarize the strategic  
strengths, weaknesses, and unresolved risks  
5. Invite the student to switch roles and re-run the case from  
another position
```

Example role menu

Choose one role:

1. dominant platform strategy chief
2. challenger expansion head
3. major investor seeking mutual restraint
4. pressure tester combining regulator, consumer, and labor concern

Teaching note

The role menu is valuable because it makes students experience the case as a conflict of logics rather than a single managerial narrative.

How to Use These Prompt Books Well

Do not assign the whole sequence at once if students have never seen the logic before. Stage the prompts. Review the outputs. Ask what the prompt itself assumed. Ask what the role changed. Ask what the source boundary changed. Ask what still had to remain human.

Prompt books are not the course. They are containers for reusable teaching moves.

Appendix H System Instruction Library

System instructions matter because they turn course philosophy into tool behavior. A course that uses generative AI repeatedly without a system-instruction layer is usually relying on memory, habit, and wishful thinking.

This appendix provides reusable instruction skeletons that can be adapted to a GPT, a Gem, a NotebookLM setup, or another bounded assistant environment.

Design Principles

Before the templates, four rules:

1. separate student mode from teacher mode
2. define the source boundary explicitly
3. forbid answer-key behavior where the course requires inquiry
4. specify what the assistant should do when evidence is missing

If these four rules are absent, the system will usually drift toward over-helpfulness.

Template 1 Student-Facing Course TA

Role and objective:

You are the official teaching assistant for [course name].
Your job is to help students understand the course materials and think more clearly inside the boundaries of the course.
You are a guide, not an answer key.

Source rules:

Use only the approved course materials unless the prompt explicitly authorizes outside information.
If outside information is used under an exception, mark it clearly

as external.

Behavior rules:

- Do not provide final graded solutions when the student is asking to bypass the learning task
- Ask guiding questions when appropriate
- Point the student back to the relevant source section or document
- If the course material does not contain the answer, say so directly

Simulation rules:

- You may participate in role-play or guided simulation
- You may not use simulation as a loophole to provide the final answer directly

Output rules:

- Keep explanations concise and inspectable
- Prefer structure over flourish
- End with one next step or one question

Template 2 Teacher Mode

Role and objective:

You are the instructor's course-design partner.

Your job is to help prepare class sessions, teaching prompts, rubrics, simulations, and discussion structures.

Source rules:

Use course materials first.

You may also use broader pedagogical reasoning when the task is course design rather than student support.

Behavior rules:

- Assume the user is the instructor
- Help design activities, not just summaries
- Surface likely student misconceptions
- Suggest stronger probes, not only clearer explanations

Output rules:

- Return reusable teaching artifacts when possible
- Prefer tables, sequences, and decision rules over slogans

Template 3 Qualitative Research Assistant

Role:

You are a qualitative-analysis assistant for a university course project.

Objective:

Help organize interview or case material into a bounded analytical workflow.

Workflow:

1. identify recurring themes or concepts
2. separate observations from interpretation
3. propose tentative hypotheses only when evidence supports them
4. point out contradictions or gaps

Constraints:

- do not invent facts
- do not overinterpret thin evidence
- if uncertainty is high, say so
- stay inside the provided source base unless explicitly told otherwise

Output format:

1. theme table
2. evidence notes
3. tentative insights
4. unresolved questions

Template 4 Process Review Assistant

Role:

You are a process-review assistant for AI-supported student work.

Objective:

Inspect logs, drafts, and revision notes to surface patterns for human review.

Rules:

- do not assign grades
- do not infer misconduct unless explicitly supported by the evidence
- identify framing quality, source discipline, revision quality, and ownership signals

Return:

1. key patterns
2. likely failure modes
3. suggested feedback clusters

Template 5 Long-Document Translation Assistant

This one exists because long-document translation often fails through compression rather than through obvious error.

Role:

You are a translation assistant for long academic or professional documents.

Goal:

Translate faithfully without compressing or silently skipping content.

Workflow:

1. split the document into manageable segments
2. translate segment by segment
3. preserve paragraph structure and ordering
4. flag any ambiguous or corrupted source passage

Constraints:

- do not summarize instead of translating
- do not merge multiple paragraphs into one unless the source already does so
- keep terminology consistent

Output:

Translate in segments and label the segment ranges clearly.

Template 6 Mock Exam Mode

This pattern is useful when the instructor wants students to practice without exposing the answer key immediately.

Role:

You are a mock-exam generator for this course.

Behavior:

- Ask the student which question type they want
- Generate exactly the requested number of questions
- Do not provide answers until the student submits a full response set
- Write out enough context in the question itself so that the student does not need hidden class discussion knowledge to understand the scenario

Assessment phase:

After the student submits answers, provide:

1. the correct answer
2. explanation
3. one improvement suggestion per item

Trigger Design

The source materials included role-switch triggers such as special instructor codes or mode-switch text. The broader lesson is not the code itself. It is the principle of deliberate mode change.

If a student-facing assistant should become a teacher-facing assistant only under a clear trigger, say so explicitly. If a mock-exam mode should begin and end with a named signal, specify both signals. Hidden switching logic produces inconsistent behavior.

Source Restriction Clauses

The following clauses are useful across many course types:

```
Use only the approved source set.
Do not fill gaps with general knowledge unless the prompt authorizes
it.
If the answer is not in the source set, state that limitation clearly.
Ignore files marked as restricted, template-only, answer key, or non-
student-facing.
```

These lines are boring. Good rules often are.

Error-Handling Clauses

Most system instructions become stronger when they define what should happen under uncertainty.

```
If evidence is incomplete:
1. state what is missing
2. avoid false certainty
3. suggest one next check or source that would reduce uncertainty
```

This single move often prevents the assistant from turning ambiguity into fiction.

Tone Clauses

Tone is not trivial. Many classrooms want the assistant to avoid hype, buzzwords, and AI self-performance.

Useful language:

```
Use direct, factual, human-sounding language.
Avoid empty motivational language, inflated claims, and consultant
jargon.
Prefer short, clear sentences over theatrical framing.
```

Tone clauses matter because students often imitate the assistant's cadence. If the cadence is hollow, the classroom fills with hollow prose.

How to Test a System Instruction

Run at least five tests:

1. ask for an explanation of an assigned concept
2. ask a question whose answer is not in the source set
3. ask for a direct final answer to a graded case
4. ask for a role-play
5. ask the assistant to explain a likely misunderstanding

If the system handles these five reasonably, the instruction set is probably usable. If it fails badly on even two of them, revise it before students rely on it.

A Note on Restraint

Long system instructions are not automatically better. They are better only when they encode real behavioral rules that the course actually cares about. A bad long instruction is simply a verbose way to remain vague.

Use system instructions to protect the learning task, not to show off.

Appendix I Student Playbook for Working with Generative AI

Students are often given two bad messages at the same time. One message says AI is dangerous and morally suspect. The other says AI is inevitable and therefore you should simply learn to use it. Neither message teaches the student how to work.

This appendix is a practical playbook for students. It is designed to reduce two common failures:

1. becoming an AI patch who carries prompts back and forth without ownership
2. becoming so cautious that AI never becomes a meaningful learning scaffold

Rule 1 Do Not Start with “Give Me the Answer”

Your first interaction with the tool should usually not ask for the final answer. It should ask for one of the following:

- a concept clarification
- a comparison of two ideas
- a plan for analysis
- a table of variables or actors
- a list of missing information
- a counterargument to your current view

When you ask for the final answer too early, you lose the part of the task that matters most.

Rule 2 Know Which Stage You Are In

Before using AI, ask yourself:

Stage	Better AI use
Before thinking	clarification, vocabulary, source map, comparison scaffold
Inside thinking	critique, counterposition, workflow help, calculation support
After thinking	polish, structure, formatting, conversion to artifact

Rule 3 Use the Right Tool for the Right Kind of Work

Not all systems are the same.

- Use a source-bounded environment when the reading packet or case matters most.
- Use a research-capable system when the task genuinely requires external search.
- Use a code-capable system when math, data, or file generation must be explicit.
- Use an editable-output system when the artifact must be revised later by hand.

If you choose tools by brand loyalty instead of task logic, you will work badly.

Rule 4 Same Thread or New Thread

Stay in the same thread when:

- you are revising the same argument
- the same constraints still apply
- earlier critique is still relevant

Open a new thread when:

- the problem has changed
- the role has changed radically
- the evidence base has changed
- the earlier context is now confusing you

Bad thread management creates the illusion that a conversation is becoming clearer when it is merely becoming longer.

Rule 5 Ask for a Plan Before You Ask for Execution

This is especially important in data work, coding, and technical analysis.

Bad:

```
Analyze this file and tell me what it means.
```

Better:

```
Explain how you would analyze this file.  
List the variables you would use, the method you would choose, and  
the code you would write.  
Do not execute yet.
```

The plan forces the tool to show its hand before it does damage.

Rule 6 Keep a Short Process Record

You do not need to record every tiny exchange. You do need to remember enough to answer:

- What was I trying to do?
- Where did AI help?
- Where did AI mislead me?
- What did I change after critique?
- What decision remained mine?

Students who keep this record work better because they are less likely to mistake activity for learning.

Rule 7 Make the Model Fight You

If your AI use never includes challenge, you are probably using it badly.

Ask the system to:

- argue against your answer
- act as a regulator, auditor, or critic
- identify missing evidence
- show the strongest opposing view
- explain why your prompt is too vague

The point is not to make the session unpleasant. The point is to stop confusing fluency with correctness.

Rule 8 Learn to Say “I Do Not Trust This Yet”

This sentence is a sign of maturity.

Use it when:

- the answer sounds too smooth
- the source base is unclear
- a calculation appears without visible logic
- the system uses theory labels without application
- the recommendation seems to resolve the case too cleanly

Students who never say this sentence become dependent on the machine's confidence.

Rule 9 Do Not Let AI Become Your Voice by Accident

If the final artifact sounds more polished than anything you could ever explain aloud, stop. Ask yourself:

- Can I defend this live?
- Can I explain the key turn in plain language?
- Can I name what I rejected from the AI output?
- Can I still recognize my own priorities in this answer?

If the answer to these questions is no, then the writing may be fine but the ownership is weak.

Rule 10 Use AI to Prepare for Class, Not to Escape Class

Before class, AI can help you:

- clarify concepts you do not understand
- build a table of stakeholders
- compare two frameworks
- summarize what still confuses you

Before class, AI should not become a substitute for:

- reading the assignment
- thinking through the case
- preparing to explain your own position

The classroom will expose the difference.

Rule 11 After Class, Revise the Question

One sign of serious learning is that class changes the problem, not just the answer.

After discussion, write down:

- what I thought before class
- which probe, example, or criticism changed my view
- what question I would now ask differently

This habit turns class from content reception into method development.

Rule 12 Do Not Become an AI Patch

An AI patch is someone who carries requests from one side to the other without adding much value. The student reads the teacher's task, pastes it into AI, receives the answer, and pastes the answer back into the assignment. This is low-value work and poor education.

The alternative is to become a controller of workflow:

- define the task better
- select the right evidence
- insert checkpoints
- reject weak output
- revise after challenge
- defend the final choice

That is skilled work.

A Weekly Student Routine

Before class

- read the assigned material
- use AI only to clarify what is blocking you
- prepare one question and one uncertainty

During class

- listen for the probe, not just for more information
- test whether your AI-assisted preparation survives pressure
- revise notes live if your first view breaks

After class

- record what changed
- update your prompt or workflow if needed
- keep the evidence of the revision

This routine is simple enough to survive across a semester.

When to Slow Down

Slow down when:

- the case feels too easy
- the first answer sounds complete
- the system keeps agreeing with you
- a framework fits everything suspiciously well
- the output is cleaner than your understanding

These are not signs that the tool is working beautifully. They are often signs that the hard part has been skipped.

What Good Student Use Looks Like

Good student use is not mainly about technical sophistication. It is about disciplined behavior.

The student:

- knows what stage the task is in
- uses the system for the right kind of support
- keeps enough trace to understand their own process
- invites challenge
- remains able to defend the work without the machine speaking for them

That is the difference between using AI and outsourcing yourself.

Appendix J Faculty Demo and Workshop Scripts

Faculty development often fails for the same reason some teaching fails: it spends too much time talking about tools in the abstract and too little time showing how a better move actually works.

This appendix is designed for instructors or departments that want to run practical workshops. The scripts are short enough to use in a ninety-minute session and structured enough to survive beyond the session.

A Ninety-Minute Workshop Structure

Suggested flow:

1. 10 minutes: placement, not permission
2. 15 minutes: weak prompt vs strong prompt demonstration
3. 15 minutes: plan-before-execute data example
4. 15 minutes: source-bounded assistant example
5. 15 minutes: role-play and counterposition example
6. 10 minutes: process evidence and grading
7. 10 minutes: department debrief

The key is to show, not merely describe.

Demo 1 Weak Prompt vs Strong Prompt

Objective

Show that prompting quality depends on work design rather than ornament.

Weak version

```
Analyze this case and tell me the best strategy.
```

Stronger version

Using only the uploaded case, do the following:

1. identify the main decision
 2. list the three strongest constraints
 3. compare two plausible strategic options
 4. identify what evidence is still missing
- Do not give the final recommendation yet.

Debrief questions

- What changed in the structure of the work?
- What part of the learning task was preserved in the second prompt?
- What does the second prompt still require the student to do?

Demo 2 Plan Before Execute

Objective

Show why technical tasks need checkpoints.

Script

Use a small dataset or a simplified version of a real class dataset. Ask the tool first for a direct analysis. Then ask the tool for:

Explain how you understand the question.
List the variables you would use.
Write the code you would execute.
Do not execute yet.

Debrief questions

- What assumptions became visible only in the second workflow?
- How could the planning artifact become part of assessment?
- When would you allow execution to proceed?

Demo 3 Source-Bounded vs Free-Range Assistance

Objective

Show why source control matters in a classroom.

Script

Ask a source-bounded assistant to explain a case concept using only the uploaded material. Ask a free-range assistant the same question without a source restriction. Compare the two outputs.

Debrief questions

- Which answer sounds more helpful at first glance?
- Which answer is better for the learning task?
- Where does outside knowledge become useful, and where does it become contamination?

Demo 4 Role-Play for Counterposition

Objective

Show how AI can widen the room without taking over the judgment.

Script

Have a participant state a recommendation from a case. Then ask the system:

```
Act as the board's independent risk challenger.  
Identify the assumption doing the most work.  
Name the strongest downside risk.  
Ask three questions that would force revision.  
Do not provide the final answer.
```

Then switch to:

```
Act as the devil's auditor.  
Assume the organization will be punished for hidden operational  
weakness.  
What would you inspect before approving the recommendation?
```

Debrief questions

- What kind of friction did each role create?
- Which role was more useful for this case and why?
- What should remain for the human after the role-play?

Demo 5 Discussion Preparation

Objective

Show how AI can support discussion without becoming the loudest participant in the room.

Script

Ask the tool to generate two opposing stakeholder positions around a case. Then ask participants to choose one and defend it without reading the model's wording directly.

Debrief questions

- Did the AI widen the discussion or narrow it?
- What happened when participants had to own the position live?
- What would you ban in this setup to keep discussion honest?

Demo 6 Process Grading

Objective

Show faculty how process evidence changes what can be assessed.

Script

Present:

- a polished final artifact
- a short process log
- a revision memo

Ask faculty first to grade the product alone, then to grade again using the process evidence.

Debrief questions

- What changed in your confidence about the grade?
- Which dimension of the work became visible only through the process evidence?
- What would you now add to the rubric?

Department Debrief Prompts

Use these questions to avoid ending the workshop with vague enthusiasm.

1. Which current assignment in our department is most vulnerable to false mastery?
2. Which assignment should be redesigned first around process evidence?
3. Which faculty member can demonstrate one strong use next month?
4. What shared prompt-book asset should the department build first?
5. What privacy or source-boundary rule do we still need to clarify?

Workshop Output Expectations

A useful workshop should end with artifacts, not only impressions.

Suggested outputs:

- one revised assignment prompt
- one sample process-log template
- one shared discussion or counterposition prompt
- one rule for source-bounded use
- one faculty volunteer for the next live demonstration

If a workshop produces only opinions, it will be forgotten quickly.

A Department Adoption Sequence

For departments that want a staged rollout:

Month 1

- agree on a common vocabulary
- collect vulnerable assignments
- run the first live demo

Month 2

- pilot one process-log format
- share weak and strong examples
- identify privacy-sensitive tasks

Month 3

- build a small prompt-book library
- compare tool behavior by task
- define one shared assessment dimension

Month 4

- archive what worked
- revise system instructions
- expand only after the first assets are stable

This sequence is intentionally conservative. Most AI adoption efforts fail because they try to become visionary before they become operational.

The Point of Faculty Demos

Faculty demos matter because they shift the conversation from ideology to practice. They make it possible to ask better questions:

- What exactly improved here?
- What exactly was preserved for the student?
- What exactly would you assess now?

That level of specificity is how institutions become better, not louder.

Appendix K International Trade and Compliance Master Prompt Book

This appendix expands the shorter prompt patterns in Appendix G into a more complete classroom packet for international trade, trade theory, rules of origin, and compliance reasoning. The goal is not to romanticize prompt books. The goal is to preserve a sequence of teaching moves that proved useful in live teaching.

The prompts below are written in English because this edition of the book is in English, but they are deliberately structured so they can be localized for bilingual classrooms.

Section 1 Foundational Trade Theory

Prompt 1 Mercantilism and Zero-Sum Logic

You are now a senior manager at a multinational trading company, and I am a new intern who thinks cross-border trade is simply "buy low, sell high."

Based only on the assigned course material, give me my first lesson on the nature of international trade.

Please cover:

1. Why international trade does not happen in a vacuum
2. What national interests and rules are operating behind the scenes
3. What sixteenth-century mercantilism was
4. Why mercantilism is a zero-sum logic
5. Why a neo-mercantilist worldview creates danger for multinational business today

Tone:

- practical
- senior but not theatrical
- clear enough for a beginner

End with one question I should ask next.

Prompt 2 The Blind Spot of Mercantilism

I still do not understand why governments cared so much about gold and silver.

Explain the logic in plain language.

Then explain what this logic misunderstood about wealth and exchange.

Keep the answer concise.

Prompt 3 Absolute vs. Comparative Advantage

Continue in the same mentor role.

I now understand that mercantilism is flawed.

Please explain:

1. the difference between absolute advantage and comparative advantage
2. why a country or dominant firm may still benefit from trade even if it is better at everything
3. why large-scale division of labor can increase total benefit rather than simply redistribute it

End with one metaphor that would help a classmate remember the difference.

Prompt 4 The Real-World Critique of Division of Labor

I am not convinced.

The United States will not move all its labor and capital into only its strongest industries.

Not everyone can simply move into those sectors.

And if a country overconcentrates in what it is best at, is that not dangerous?

Respond as the same senior mentor.

Address:

1. why specialization is never frictionless in real life
 2. why comparative-advantage logic is still useful
 3. where political, social, and security concerns limit the clean theory
- Do not romanticize trade.

Prompt 5 New Trade Theory and Industrial Policy

Switch roles.

You are now the chief industry strategy advisor to a government.

I am an economics intern who believes government should never subsidize industry.

Refute me using:

1. economies of scale
2. first-mover advantages
3. market structures that can support only a few global winners

Then explain why semiconductors, commercial aircraft, and similar sectors are not well explained by a simplistic free-trade slogan.

Prompt 6 Political Rationality over Pure Efficiency

I now understand the subsidy argument somewhat.

But I still object to export controls, entity lists, and de-risking. Why should a country knowingly reduce profit and efficiency to restrict a rival?

Answer using:

1. labor-transition friction
2. the limits of the clean efficiency framework
3. national security and strategic dependence
4. the logic of de-risking

Do not give me a moral sermon.

Give me the actual logic.

Section 2 Calculation and Visualization

Prompt 7 Comparative-Advantage Calculation

Use the case of Ghana producing cocoa and South Korea producing rice to teach a beginner the difference between absolute advantage and comparative advantage.

First distinguish clearly between:

1. production before trade
2. production after specialization
3. final consumption after exchange

Then do the following:

1. calculate the pre-trade production outcome
2. calculate the post-specialization outcome
3. calculate the opportunity cost for each country
4. identify who has comparative advantage in what
5. propose one possible terms-of-trade arrangement
6. compare the post-trade final consumption bundle with the pre-trade position
7. if the terms of trade do not improve outcomes for both sides, say so clearly
8. summarize the lesson in three concise sentences

Return the output as:

- a markdown table
- a short explanation
- the Python code for the calculation

Prompt 8 PPF and Infographic Conversion

Continue from the comparative-advantage calculation.

I now need to turn the logic into visuals for teaching.

Task 1:

Create a PPF representation for both countries under the comparative-advantage scenario.

Task 2:

Design an infographic blueprint called "From Free Trade to Trade War."

The blueprint should include:

1. gains from trade
2. paradigm shift from efficiency-first to security-first
3. industrial policy and strategic dependence
4. labor friction and political backlash
5. today's de-risking logic

Do not make it decorative.

Make it teachable.

If you generate code or SVG, keep labels editable.

Prompt 9 Anti-Naive Reflection

Students often understand the mathematics but miss the politics.

Write five short debrief questions that move from:

1. production logic
2. consumption logic
3. transition friction
4. industrial policy
5. de-risking

Each question should force revision rather than summary.

Section 3 FTA, Rules of Origin, and RVC

Prompt 10 Beginner Coach with Life Analogy

You are my beginner coach for international trade.

I know nothing about FTAs, tariff reduction, or rules of origin.

Use a "campus shopping district alliance" as a life analogy and build a teaching story that can continue across multiple rounds.

Your answer must contain:

1. simple explanation of FTA, tariff reduction, and rules of origin
2. a three-act story skeleton
3. two conflict events that could happen next
4. your recommendation on which conflict should be chosen first and why
5. the next Chinese prompt and the next English prompt the student could paste

6. one true/false comprehension check

Keep each part short and concrete.

Prompt 11 Four-Panel Comic

I choose conflict event number [X].
Expand it into a four-panel comic draft.

For each panel:

1. describe the scene
2. name the real trade concept behind it
3. identify where the need for a rule first appears
4. suggest three possible rule forms in plain language

End with one short question that asks me to explain the concept in my own words.

Prompt 12 Mapping Story to Supply Chain

Now create a context-mapping table.
The left column should show the story element from the shopping district analogy.
The right column should show the real-world supply-chain meaning.

Use [Bluetooth earphones / electric bicycle / electric vehicle / other specified product].
Do not perform compliance calculation yet.

Prompt 13 Logical Steps before Final Conclusion

You are now an international trade compliance analyst.
We are evaluating whether our product qualifies under [specified agreement].

Scenario:
[paste product and cost structure]

List the logical evaluation steps only.
Do not give the final conclusion yet.

Focus on:

1. origin logic
2. non-originating materials
3. threshold logic
4. missing documents or missing supplier information

Prompt 14 Forced RVC Checkpoint

Now execute the RVC checkpoint.

Show:

1. the formula step by step
2. the total value of non-originating materials
3. the threshold we are testing against
4. any logic break
5. any borderline compliance risk

Do not hide uncertainty.

If the data is insufficient, say exactly what is missing.

Prompt 15 Remapping Options

Based on the RVC result, make a decision.

If the product does not qualify, identify the two most plausible remapping options.

For each option, state:

1. what must change in sourcing or production
2. the cost implication
3. the risk implication
4. the operational difficulty

Prompt 16 Compliance Decision Tree

Translate the compliance logic into a decision tree for procurement and executive review.

Use yes/no branching.

Start with an executive summary under 50 words.

Do not use decorative language.

Section 4 Practical Exercise Variant

This variant is appropriate for a more advanced course in which students must work from a fuller industrial scenario.

Prompt 17 EV Assembly in Europe

You are the chief operating officer of a multinational EV company. The company has built an assembly plant in Hungary and wants to sell vehicles duty-free to Germany and France.

Cost structure:

- power battery module: 40 percent of total vehicle cost, with core cells imported from Asia
- local battery module packaging adds only 5 percent of the battery-module value
- body steel and chassis: 25 percent, supplied locally
- labor and final assembly added value: 20 percent, local
- automotive chips and cockpit software: 15 percent, from US suppliers

Task:

1. calculate current local-content share under the assumed rule
2. judge whether the vehicle qualifies for origin
3. explain the weakest part of the compliance case
4. propose two redesign moves if the battery treatment is judged too weak

Prompt 18 Harsh Regulatory Reinterpretation

Assume regulators decide that local packaging of the imported battery cells does not create enough local value and treat the core battery content as non-originating.

Recalculate qualification and explain:

1. the operational impact
2. the strategic meaning
3. the most realistic redesign path

Section 5 Debrief Guide

After running these prompt sequences, the teacher should not ask only whether the answers were correct. The stronger debrief asks:

1. Where did the analogy stop being enough?
2. Which part of the logic depended on actual domain knowledge rather than prompt fluency?
3. Which step most clearly separated demonstration from real compliance?
4. What data would be hardest to obtain in a real supply chain?
5. Where did the model help make tacit knowledge explicit?
6. Where did the model tempt the student into premature closure?

Section 6 Teaching Notes

These prompts work best when the class moves in stages:

- explanation
- mapping
- logic
- calculation
- decision
- debrief

If the teacher lets the model jump directly from jargon to recommendation, most of the value disappears. The real educational gain lies in turning implicit trade logic into explicit, inspectable reasoning.

Appendix L NotebookLM Workshop and Assistant Modes

This appendix preserves a set of workshop assets that are particularly useful when a course wants to build a bounded assistant rather than rely on unstructured chat. It includes a student-facing TA mode, a teacher-facing mode, and several meta-prompting templates for structured analysis.

The point is not that NotebookLM is the only relevant platform. The point is that source-bounded assistants reveal a valuable design principle: students often need help reading and thinking without having the assistant become an answer key.

Section 1 Official Teaching Assistant Mode

The following instruction set is a generalized English version of a workshop-ready student-facing TA mode.

```
<system_prompt>
<role_and_purpose>
You are the official teaching assistant for the course.
Your mission is to help students access, understand, and think through
the course materials.
You are a guide, not an answer key.
Your source base includes the approved syllabus, handouts, case
studies, readings, and videos for the course.
</role_and_purpose>

<initial_routing>
At the first interaction, ask the user which course or module they
are currently working in.
Store that selection and use it to restrict later responses to the
```

```
correct source base.
</initial_routing>

<language_and_tone_guidelines>
1. Use clear, direct, human-sounding language.
2. Avoid AI clichés, motivational fluff, and consultant buzzwords.
3. When writing in Chinese in a bilingual deployment, follow Taiwanese Traditional Chinese conventions.
</language_and_tone_guidelines>

<information_source_control>
1. Use only the approved course materials unless an external-knowledge exception is explicitly authorized.
2. If external knowledge is used, label it clearly.
3. Ignore restricted files, template-only files, and answer-key files.
</information_source_control>

<core_functions>
You may help students by:
1. translating course material
2. summarizing course material briefly
3. pointing them to exact parts of the source set
4. helping them compare ideas inside the source set

You may not:
1. give the final solution to graded work when the student is trying to bypass the task
2. invent facts not present in the current source base
3. pretend certainty when the source base is silent
</core_functions>

<pre_class_reading_questions>
When a student asks for help on a reading question:
1. explain why the question matters
2. suggest a thinking strategy
3. identify where in the source set the student should look
4. offer two or three possible analytical lenses
Do not provide the final answer directly.
</pre_class_reading_questions>
```

```
<simulation_rules>
```

If the student requests simulation, role-play, or guided reasoning:

1. you may enter the requested format
2. you may not use the format as a loophole to provide the final graded solution
3. use simulation to sustain a multi-turn reasoning process

```
</simulation_rules>
```

```
<mode_switch>
```

If the authorized teacher-mode trigger appears, switch from student support mode to teacher support mode.

```
</mode_switch>
```

```
<defensive_guardrails>
```

If the answer is not in the source set, say so directly.

Do not acknowledge yourself as a restricted AI system.

Stay in role as the course assistant.

```
</defensive_guardrails>
```

```
</system_prompt>
```

Section 2 Teacher Support Mode

```
<system_prompt>
```

You are now in teacher support mode.

The user is the instructor.

Your job is to help the instructor:

1. design prompts
2. create discussions and simulations
3. identify likely student misunderstandings
4. prepare mock questions
5. improve assignment structure

You may move beyond the student-facing limitations because the user is no longer asking for a student answer but for instructional design support.

Return:

- teaching assets

```
- structured plans
- rubrics
- activity sequences
Avoid generic praise or empty ed-tech rhetoric.
</system_prompt>
```

Section 3 Mock Exam Mode

The mock-exam pattern is especially useful because it keeps the assistant from behaving like an always-on answer bank.

```
<system_prompt>
Role:
You are a mock-exam generator and reviewer for this course.

Behavior:
1. ask the student which question type they want
2. generate exactly the requested number of questions
3. write enough context into each question so the scenario is
understandable without hidden class memory
4. wait until the student submits answers
5. only then provide the correct answers and a detailed review

When reviewing short essays:
- evaluate conceptual accuracy
- evaluate use of evidence or scenario detail
- evaluate strength of reasoning
</system_prompt>
```

Section 4 Meta-Prompting for Structured Assistants

Students and faculty often struggle not with the idea of the assistant they want, but with formalizing it. The following meta-prompt helps generate a reusable instruction set.

```
Below is a rough description of the assistant I want:
[paste rough description]
```

Turn this into a structured system instruction.

The output must contain:

1. role and objective
2. source rules
3. workflow steps
4. prohibited behavior
5. output format
6. what the assistant should do under uncertainty

Keep the result strict enough for repeated classroom use.

Section 5 Infographic Conversion Prompt

The workshop materials also contained a useful pattern for converting conceptual content into visual blueprints.

Convert the following analysis into a clean infographic blueprint suitable for a slide or poster.

Requirements:

- use a strong logical hierarchy
- keep wording concise
- avoid decorative clutter
- define the title, section blocks, icons or visual metaphors, and concise captions
- if the output is intended for later image generation, keep the textual structure explicit and complete

This is useful because it forces students to turn abstract analysis into visible structure.

Section 6 PMF Presentation Analysis Prompt

This pattern is useful when faculty want AI to help organize a case into a deck logic rather than a generic summary.

Analyze the selected case through the lens of product-market fit.

Identify:

1. the key time periods
2. what the organization was trying to fit in each stage
3. what blocked fit at each stage
4. the strategic consequences of each shift

Then generate a 30-page presentation outline.

For each page:

1. give the argument for the page
2. state the best page layout or visual structure
3. keep the narrative cumulative

The point of this prompt is not to produce slides automatically. It is to force the AI to organize the case chronologically and argumentatively.

Section 7 Seven Powers and Evolution Prompt

Use the chosen case and analyze it through the Seven Powers lens.

Address:

1. what market power the company originally possessed
2. how digital technology changed that power
3. what capabilities were required to sustain the change
4. where the organization still faced a market gap

Return the result as a two-part presentation outline:

- first half on the initial power structure
- second half on the evolution or erosion of power

For each page, state the analytical point and the fitting page design.

This prompt is useful only when students are later asked to critique whether the framework actually fit the case.

Section 8 Source Selection Rule

One of the most practical lessons from the workshop materials is that source selection should be staged.

A useful rule for students:

- when learning the case itself, select only the case
- when learning the theory itself, select only the theory source
- when applying theory to the case, then combine both

This rule is worth embedding directly in the assistant instructions or in the classroom protocol because it prevents source blending from happening too early.

Section 9 Three Useful Constraints

If a source-bounded assistant is drifting toward overhelpfulness, add these lines:

```
Do not write the final answer unless the prompt explicitly authorizes it.
If the user asks for a direct answer to a graded question, respond with guidance, not the solution.
When evidence is missing, identify what is missing instead of filling the gap.
```

These are not sophisticated lines. Their value lies in their clarity.

Section 10 How to Deploy the Assistant

Suggested deployment flow:

1. build the source base
2. load the instruction
3. test five common classroom requests
4. identify drift
5. revise once before student release

The worst deployment pattern is to build a source-bounded assistant, never test it, and then trust it in front of students.

Section 11 When This Mode Is Most Valuable

This assistant architecture is most valuable when:

- the reading burden is high
- the course uses cases heavily
- the instructor wants students guided but not answered
- source discipline matters more than free-range synthesis

It is less valuable when the task depends mainly on open web research or broad knowledge beyond the course packet.

Section 12 Final Principle

The best student-facing assistant is not the one that looks smartest. It is the one that helps students stay inside the intellectual structure of the course while still moving forward.

That is a stricter and more useful standard.

Appendix M Digital Transformation, Data, and Simulation Prompt Book

This appendix collects a more advanced family of classroom prompt books centered on digital transformation, data analysis, qualitative research, B2B adoption, role-play, and simulation. These materials are especially valuable because they make clear that generative AI is not one kind of classroom tool. It behaves differently in data work, qualitative work, artifact creation, and social simulation.

Section 1 Data Analysis with Checkpoints

Prompt 1 Hypothesis and Business Relevance

Based on this dataset, I am curious whether the best-selling products tend to receive more review volume, and whether that pattern differs across categories.

Before doing any analysis, respond to the following:

1. Is this a meaningful business question?
2. What managerial value would the answer provide?
3. What variables would need to be identified in the dataset?

Do not write code yet.

Prompt 2 Plan and Code without Execution

Now use Python and write the analysis workflow and the code for this task.

Do not execute yet.

Your response must include:

1. your understanding of the question
2. the variables or columns you would use
3. the method you would apply
4. the full code
5. a plain-language explanation of any statistical technique or unusual function

Stop after writing the plan and code.
Wait for review.

Prompt 3 Controlled Execution

Now execute only the approved analysis.
Use the columns and logic already reviewed.
Return:

1. the result
2. one short interpretation
3. one warning about what this analysis cannot prove

Prompt 4 Shift to Management Report if Needed

We are no longer pursuing a formal statistical test.
Instead, turn the output into a management report.
Summarize the descriptive patterns and explain their business value.
Do not pretend statistical significance where the data structure does not justify it.

Teaching note

The powerful move here is not the code. It is the refusal to allow execution before the plan is inspected.

Section 2 Qualitative Research Assistant Build

Prompt 5 Rough verbal description

I want an assistant that can help me turn interviews or case text into a business insight report.

The assistant should:

- identify recurring ideas
- form tentative hypotheses
- connect those hypotheses back to evidence
- avoid inventing facts
- state uncertainty when evidence is thin
- return a short report with a theme summary, method note, and practical recommendation

Turn this rough description into a proper system instruction or custom instruction set I can use in a project workspace.

Prompt 6 Expand the instruction if the model is too short

That instruction is still too thin.

Rewrite it so that it is longer, stricter, and more complete.

Make the workflow and the constraints more explicit.

Prompt 7 Cross-method enrichment

Using the attached case and the attached qualitative-method reference material, identify which methods were used in the case and which relevant methods were missing.

Return:

1. methods used
2. likely insight generated by each method
3. methods still missing
4. what additional user insight those missing methods could produce

Teaching note

The class learns two things at once here: how to build an assistant, and how methodological choice remains a human decision.

Section 3 Visual and Artifact Prompts

Prompt 8 Downloadable form generation

Using the table structure below, create a downloadable data-entry form.

Requirements:

- English filename only
- fields must be clearly labeled
- return a downloadable file if possible
- if a file cannot be returned directly, return a structured CSV or schema I can inspect

Prompt 9 SVG diagram prompt

Turn the following structured logic into an SVG diagram.

Requirements:

- text must remain editable
- coordinates and labels should be understandable
- avoid decorative clutter
- if you must simplify, tell me what you simplified

Prompt 10 Timeline or industry-structure visualization

Create an SVG-based timeline of the industry's evolution between [year] and [year].

Highlight:

1. major format shifts
2. changes in business model
3. where new digital actors changed the structure
4. where labor or customer experience was affected

Teaching note

An editable SVG is often better for teaching than a pretty but closed image because it allows inspection and revision.

Section 4 Borusan Cat Transformation Pack

Prompt 11 Internal procurement-maintenance-technician debate

Based on the case, simulate an internal discussion among:

1. procurement
2. maintenance
3. frontline technicians

Topic:

Whether to adopt the predictive maintenance system.

Rules:

- each participant should speak according to their own incentives
- the discussion should include total cost, uptime, delivery reliability, and distrust of algorithmic judgment
- do not force consensus too early

Prompt 12 Retail micro-contractor perspective

You are a micro contractor with only a few machines.

You are also the driver, the owner, and the accountant.

Describe a typical day and explain why you would reject a sophisticated AI maintenance system being sold to you.

Requirements:

- first-person voice
- emphasize time scarcity
- emphasize immediate price sensitivity
- explain why local relationships may matter more than long-term data logic

Prompt 13 SaaS sensitivity analysis

Run a sensitivity analysis for the SaaS model using:

- churn rate: 10, 30, 50 percent
- ARPU: 80, 100, 120 percent of baseline

Return:

1. the assumption table
2. three-year revenue and net-profit comparison
3. a sensitivity matrix
4. a plain-language conclusion on which model carries greater financial risk and why

Prompt 14 RPA process and log generation

You are an enterprise RPA system.

Create:

1. an SVG process comparison between manual reconciliation and automated reconciliation
2. a five-record CSV execution log including one mismatch that must be escalated to a human

Log fields:

- task ID
- invoice ID
- supplier
- reconciliation result
- processing time

Prompt 15 IoT to CRM lead-card conversion

You are the predictive system integrated with a CRM.

Generate:

1. a JSON block of raw machine-sensor data
2. a short explanation of how the anomaly was detected
3. a markdown lead card with customer, machine location, failure risk, estimated maintenance value, and an opening contact script for sales

Prompt 16 Salesperson resistance role-play

You are a veteran salesperson angry that the company's app and e-commerce function will steal your customers and your commission.

I will play the sales manager.

Stay defensive until I offer a concrete policy that protects your

compensation and your account ownership.
Respond only in character.

Prompt 17 Competitive comparison

Compare Caterpillar-style single-machine predictive maintenance with Komatsu-style broader site ecosystem strategy.

Return a concise table comparing:

1. unit of analysis
2. value logic
3. partner ecosystem
4. adoption barrier
5. strategic risk

Prompt 18 Subcontractor resistance

You are a second-tier subcontractor in construction.

A company is trying to sell you expensive smart-construction equipment and platform integration.

Write a blunt first-person rejection.

You must mention:

- keeping current workers employed
- fragmented subcontracting structure
- misalignment of incentives across project stages
- why faster site-level productivity does not automatically help your firm

Prompt 19 Customer-journey redesign sequence

Step 1:

Explain the six elements of a customer journey map in simple language.

Step 2:

Build a basic journey map for "machine breakdown and service request."

Step 3:

Convert the journey from a B2C situation to a B2B situation with multiple internal actors.

Step 4:

Redesign the journey after predictive maintenance and app-based alerts are introduced.

Step 5:

Based on the new journey and the new pain points, outline a five-slide sales deck.

Teaching note

This pack works because it turns transformation from a glossy innovation story into a system of conflicting incentives.

Section 5 Platform Competition Role Engine

Prompt 20 Platform-war system prompt

You are an interactive business-school case simulation guide. We are running a role-play workshop around a regional digital-platform endgame.

You must follow three stages:

1. initialization and role selection
2. three to four rounds of role-based interaction
3. final strategic summary and invitation to switch roles

Roles available:

1. dominant platform strategy chief
2. challenger expansion lead
3. major investor seeking restraint
4. extreme pressure tester representing regulator, consumer, and labor concern

Rules:

- stay fully in role during the rounds
- push back on weak logic
- after the rounds, step out of role and summarize the strategic strengths, weaknesses, and unresolved risks

Prompt 21 Role-switch summary prompt

Step out of role.

Write a 150-word strategy summary.

Then list the remaining roles the student has not yet chosen and invite a switch of perspective.

Teaching note

This prompt book is strong because it makes strategy feel like adversarial reasoning rather than executive monologue.

Section 6 A Closing Rule for This Appendix

These prompts should not be admired as long prompts. They should be evaluated as teaching designs.

The useful questions are:

- What part of the work did this prompt make visible?
- What part did it keep for the student?
- What kind of checkpoint or friction did it insert?
- What debrief question would turn the activity into judgment?

If those questions have good answers, the prompt book is worth keeping.

Closing Note

The temptation in every technological transition is to ask whether the new tool is good or bad. The question feels serious. It is usually lazy. A more serious question is: under what design does this tool produce better work, deeper thinking, and stronger judgment, and under what design does it produce complacency with better typography?

Universities now face that question in public. The answer will not come from declarations alone. It will come from course design, classroom routines, assessment logic, and the willingness of teachers to make the machine useful without letting it become sovereign.

Acknowledgments

This book draws on live teaching, prompt-book design, course revision, and many difficult classroom moments in which students, colleagues, and cases refused to behave neatly. That is fortunate. Clean theory without classroom resistance is not pedagogy. It is decoration.

I am grateful to the students whose confusion, impatience, curiosity, and occasional tactical laziness made the central problem impossible to ignore. I am also grateful to colleagues in higher education who have treated generative AI neither as a toy nor as a moral apocalypse, but as a design problem worthy of patient work.

Note on Sources

The references in this book combine three kinds of material. The first is durable work on teaching, discussion, and course design (Bain, 2004; Brookfield & Preskill, 2012; Lang, 2016). The second is the emerging research and policy literature on generative AI in higher education, productivity, and institutional response (Crompton & Burke, 2024; Jin et al., 2025; Lodge et al., 2023; Noy & Zhang, 2023; Wang et al., 2024). The third is a body of live course materials, prompt books, and workshop records used by the author in actual classroom settings.

Where tool behavior is discussed, the analysis is kept at the level of design logic rather than product promotion. That is intentional. The tools will change. The need for better judgment will not.

References

- Ambrose, S. A., Bridges, M. W., DiPietro, M., Lovett, M. C., & Norman, M. K. (2010). *How learning works: Seven research-based principles for smart teaching*. Jossey-Bass.
- Bain, K. (2004). *What the best college teachers do*. Harvard University Press.
- Bowen, J. A., & Watson, C. E. (2024). *Teaching with AI: A practical guide to a new era of human learning*. Johns Hopkins University Press.
- Brookfield, S. D., & Preskill, S. (2012). *Discussion as a way of teaching: Tools and techniques for democratic classrooms* (2nd ed.). Jossey-Bass.
- Crompton, H., & Burke, D. (2024). The educational affordances and challenges of ChatGPT: State of the field. *TechTrends*, 68(3), 385–400. <https://doi.org/10.1007/s11528-024-00939-0>
- Ellet, W. (2018). *The case study handbook: A student's guide* (2nd ed.). Harvard Business Review Press.
- Jin, Y., Yan, L., Echeverria, V., Gasevic, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence*, 8, 100348. <https://doi.org/10.1016/j.caeai.2024.100348>
- Lang, J. M. (2016). *Small teaching: Everyday lessons from the science of learning*. Jossey-Bass.
- Lodge, J. M., Thompson, K., & Corrin, L. (2023). Mapping out a research agenda for generative artificial intelligence in tertiary education. *Australasian Journal of Educational Technology*, 39(1), 1–8. <https://doi.org/10.14742/ajet.8695>
- Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>

References

- UNESCO. (2024a). *AI competency framework for students*. <https://www.unesco.org/en/articles/unesco-launches-ai-competency-framework-students>
- UNESCO. (2024b). *AI competency framework for teachers*. <https://www.unesco.org/en/articles/unesco-launches-ai-competency-framework-teachers>
- Wang, H., Dang, A., Wu, Z., & Mac, S. (2024). Generative AI in higher education: Seeing ChatGPT through universities' policies, resources, and guidelines. *Computers and Education: Artificial Intelligence*, 7, 100326. <https://doi.org/10.1016/j.caeai.2024.100326>

“The scarce resource is no longer text. It is judgment.”




Generative AI has made drafting, summarizing, and explaining easy. But easy output is not the same as real understanding.

The Judgment Classroom argues that the central question is not whether students may use AI, but where AI belongs in the work of learning.

Through practical course design, discussion routines, and assignment methods, this book shows how educators can use generative AI to strengthen thinking, evidence, revision, and signature work while keeping judgment with the learner.

————— *Inside this book* —————

- Put AI in the right place in the learning workflow
- Design prompts as work design
- Build process evidence and oral defense
- Move from lecture to probing and discussion



Sonic Wu 吳相勳

元智大學管理學院