

Marketing Strategy in the AI Era

*Five Principles for the Decisions
Machines Can't Make*

Marketing Strategy in the AI Era

*Five Principles for the Decisions
Machines Can't Make*

Paul S. Richardson

Copyright © 2026 Cambria Press

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means without written permission from the publisher.

About this book: This book was created the way it teaches. It was drafted in collaboration with artificial intelligence (Anthropic's Claude); every load-bearing claim in its worked sessions was dispositioned per the discipline it teaches, and every element was directed, verified, edited, and approved by the author, who — per the claim rule taught in Chapter 2 — signs it. All companies, brands, data, reports, and market events in the course universe are fictional.

ISBN 978-1-63857-442-2 (PDF)

ISBN 978-1-63857-457-6 (paperback)

Printed in the United States of America

Contents

Preface	1
1 Strategy in the AI Era: What Changes, What Doesn't	3
2 The Strategist and the Machine	17
3 Customers Differ: Segmentation and the Personalization Question	27
4 Customers Change: Demand in the AI Era	37
5 Competitors React: War-Gaming Doctrine	47
6 Preference Truth: Conjoint Thinking for Deciders	57
7 Resources Are Limited: The Allocation Discipline	67
9 Capabilities Have Boundaries: Build, Buy, or Partner	77
10 Data as a Strategic Resource	87
11 Where to Play, How to Win: The Integrated Choice	97
12 Pilots and Proof: Evidence at Strategy Altitude	107
13 Strategy Under News: The Response Discipline	117
14 Governance, Trust, and the License to Operate	127
15 The Review Board	137
A Course Roadmap	147
B The Board Book Specification	149
C The Four Strategic Situations	153

D	The Strategist's Prompt-Pattern Library	157
E	Careers: Marketing Strategy in the AI Era	161
F	Consolidated Glossary	165

Preface

This book is the third seat in a story its two companion volumes tell from the desks on either side. In *Integrated Marketing Communications in the AI Era*, students charter a one-person, AI-augmented agency and build a complete campaign for a realistic client. In *Marketing Analytics in the AI Era*, they take the analyst's desk inside the same companies and spend a semester inside the data. In this course, they take the chair the other two seats report to: each student joins one of the same four companies as its strategy lead, inherits its situation, and spends fifteen weeks making the choices every marketing organization actually faces — where to play, how to win, what to build and what to buy, when to hold a committed strategy and when to revise it in public.

The premise is the one now reshaping the discipline itself: strategy *documents* have become effectively free. The machine on the strategist's desk drafts a plausible strategic plan in forty seconds — fluent, structured, confident, and interchangeable with the plan it would draft for the rival. What has not become free is choice: the willingness to give something up, name a battlefield, commit resources against reacting competitors, and sign the result. That distinction organizes everything here. This book compresses its discipline into one law — the claim rule: no load-bearing claim enters any memo unless it is traced to the brief, traced to an exhibit, or flagged as inference — and into five principles, four of them marketing's classical bedrock, crystallized for modern readers by Palmatier and Sridhar and gratefully adopted here: customers differ, customers change, competitors react, resources are limited. The fifth is this era's addition, and this book's contribution: capabilities have boundaries — every capability, human or machine, has a line where it stops working and keeps answering anyway, and most strategic failures of the AI era are boundary misjudgments.

The method is the studio, adapted to the strategist's altitude. There is no notebook here and no student code; strategy consumes evidence and produces decisions, so every quantitative exhibit arrives prepared, and the graded skill is reading it like a professional. In place of the notebook, the labs put students at instruments — a perceptual map, a conjoint share simulator, a budget allocator — cockpit controls that compute honestly against each student's personal seed and wait for a configuration to be committed. Because each student's market truth is drawn from their seed, two of the labs are what the course calls ledger

islands: the committed configuration is checked directly against the market's true structure. And because the fluent collaborator is part of the desk, the traps working strategists fall into are planted deliberately: a fabricated source cited with confidence, a vendor case study whose arithmetic fails against its own reference customer, an instrument that recommends spending beyond its data, a figure bent toward fluency on the night before the board — each scripted, once, in the lab built to teach its catch.

The arc runs from the war room to the boardroom: the situation interrogated and the fact base verified; the market analyzed in motion — segments, signals, rivals war-gamed in earnest, preferences and budgets committed; then the choices — capabilities audited, the data play designed, the integrated strategy committed; and finally proof and guardrails — the pilot pre-registered, the mid-semester shock absorbed on the record, the governance policy written and the strategist's own choices audited — until the semester binds into a Board Book and seven minutes before a board. Twice, deliberately, the market moves after commitments are made; the course grades the revision, not the prophecy. Week eight is the midterm; there is no chapter eight, and the numbering honors the course calendar.

The clients are the same four businesses the companion courses serve, and the seats face each other across one universe: the strategy this book's students commit is the campaign context the agency course executes and the outcome the analytics course measures. Each seat learns what the others cost. What all three learn is the same closing lesson, because it is the profession's: the machine drafts, argues, and summarizes better every year — and it signs nothing. Decisions have authors. This book exists to train them.

This volume is one of the *Marketing in the AI Era* series — four companion courses, one shared universe, each with its own studio. The complete series is at <https://www.marketingintheaiera.ai>.

Paul S. Richardson, PhD
Professor of Marketing
Holzschuh College of Business Administration
Niagara University

Strategy in the AI Era: What Changes, What Doesn't

Opening Scenario: The Strategy That Took Forty Seconds

It is 8:55 on your first Monday as the strategy lead at Lakemont Roasting Co., and the CEO has forwarded you a trade-press article with a one-line note: “Driftline just raised another round. What’s our strategy?”

You do what any honest person does in 2026: you ask the machine. “Write a marketing strategy for a craft coffee company competing against an AI-powered personalization giant.” Forty seconds later you have nine hundred words of professional-grade prose. *Lean into authenticity. Build community. Invest selectively in personalization where it strengthens the brand. Expand direct-to-consumer while deepening wholesale relationships. Tell the origin story.* It has an executive summary. It has pillars. It reads exactly like strategy.

Now run one more test before the ten o’clock. Ask the same machine to write Driftline’s strategy for competing against craft roasters like Lakemont. Read the two documents side by side, and swap the company names. *Lean into convenience while investing in authenticity storytelling. Build community at scale. Deepen personalization while honoring craft values.* With the names exchanged, either document could serve either company — which means neither document is a strategy. A strategy your rival could adopt unchanged is not a choice; it is a horoscope. It flatters every reader because it commits no one to anything.

Understand what just happened, because it frames this entire course. The machine did not malfunction. It did, in forty seconds and for free, what a meaningful fraction of the strategy industry has done for decades at great expense: it produced *strategy-shaped prose* — fluent, structured, plausible, and empty of choice. Artificial intelligence has made strategy documents free. It has not made strategic choice free, because choice was never expensive to write down. Choice is expensive to *make*: it requires giving something up, and a machine that serves every company in every industry has nothing to give up on your behalf.

Here is what the strategy lead who keeps this job does before the ten o'clock. She opens the brand brief and reads it the way this chapter will teach you to read one. She finds the money: \$450,000 in marketing budget against a rival spending \$40 million — so any strategy that amounts to “do what they do, but with heart” is arithmetic suicide. She finds the anomaly: 71 percent of the customers who cancel still say they love the coffee — so Lakemont’s problem is not persuasion, and media spend is not its solvent. And she walks into the meeting with a sentence that contains an actual refusal: “Driftline’s raise makes personalization cheaper for everyone — including us — which is exactly why it can’t be our advantage. Our advantage is in what their engine can’t do without what only we have. Here’s where I’d look.” That sentence might be wrong. It is nonetheless strategy, and the nine hundred words are not — and the entire discipline of this book lives in the difference.

What Strategy Is — and What It Is Not

This book uses one working definition, and every lab will hold you to it: **a strategy is a coherent choice about where to compete and how to win there, made under scarcity, that a specific rival would be unwilling or unable to copy.** Every clause is load-bearing. *Coherent* excludes the list of initiatives that answers every question with “yes.” *Choice* excludes the document that keeps all options open. *Under scarcity* excludes the plan that assumes next year’s budget will resolve this year’s trade-off. And *a specific rival* — not “the competition,” a named company with a named situation — supplies the test you met in the opening scenario, which this course formalizes as the **swap test**: if your rival could adopt your strategy slide unchanged, you have not chosen yet.

It is worth being equally precise about the counterfeits, because you will produce all of them at least once this semester and the rubrics are calibrated to catch them. A *goal* is not a strategy: “grow subscribers 20 percent” describes a destination, not a route, and a route is precisely what a strategy is. A *plan* is not a strategy: a calendar of activities can be executed flawlessly in service of no choice at all. And the most seductive counterfeit — the one Lab 11 is explicitly trapped against — is the *list*: five pillars, each sensible, none of which knows about the others, presented as strategy because it is long. The scholarly literature converged on this diagnosis well before AI made the counterfeits cheap to generate: Porter’s classic argument distinguishes strategy from operational effectiveness — doing the same things better — and observes that operational effectiveness, however necessary, cannot be a strategy because best practices diffuse (Porter 1996). Rumelt’s field guide to bad strategy catalogs the same fauna you just met: fluff, goals masquerading as choices, and the failure to diagnose before prescribing (Rumelt 2011). The AI era did not invent strategy theater. It industrialized it.

One more standard carries over from this book's sibling courses, because it is the series' spine: **the decision standard**. Every deliverable you produce this semester ends in a recommendation a decision-maker could act on, with its uncertainty honestly stated. Analysis that ends in a framework diagram is homework. In this course the standard has a sharper edge than anywhere else in the series, because strategy is the seat where the recommendation *is* the product. The analyst's memo informs a decision; the agency's campaign expresses one; your memo *is* one. Decisions have authors, and for fifteen weeks the author is you.

The Five Principles

Everything in this book hangs on five principles. Four of them are marketing's classical bedrock — a set crystallized for modern readers by Palmatier and Sridhar (2021), whose “first principles” formulation made *customers differ*, *customers change*, *competitors react*, and *resources are limited* the field's shared vocabulary for organizing strategic problems. This book gratefully adopts that four-part foundation, develops it in its own way for the AI era, and adds a fifth principle the foundation now needs. Each principle is a claim about the world that no strategy may safely ignore — and each is also a question AI has reopened. You will name at least one of them, explicitly, in the rationale of every memo you write this semester: “which principle is doing the work here?” is this course's version of showing your arithmetic.

Customers differ. No market is a person; every market is a distribution. Value concentrates unevenly across segments, and a strategy that averages across them serves a customer who does not exist. Lakemont's Ritualists, Gifters, and Loyalists want different things from the same bag of coffee; Loop's campus treasure hunters and social payers barely use the same product. What AI reopens: personalization promises to serve each customer differently at zero marginal cost — which raises, rather than settles, the strategic question of *whether* difference should be served. Chapter 3 takes this up properly, including the case where personalization is a tax dressed as a moat.

Customers change. Preferences drift, contexts shift, and — the AI-era version — the *pathways* by which customers discover, evaluate, and buy are being rebuilt in real time. When a shopper's assistant summarizes the category, brands the assistant doesn't mention do not exist; when an agent reorders the coffee, the ritual the brand was built on is quietly disintermediated. Visit Chautauqua Trail's entire semester turns on this principle: a region that AI trip planners cannot read is a region that stops being discovered. Chapter 4 builds the discipline for it — the signals watchlist, with tripwires.

Competitors react. Strategy is played against opponents who read your moves and answer them. A projection that holds your rival constant is not con-

servative; it is fiction. The AI-era sharpening is what this book calls the **copy clock**: any advantage assembled from rentable parts — models, tools, vendor capabilities — can be assembled by your rival in roughly the time it takes to sign the same contracts. Loop lives this principle all semester: every feature it ships starts the clock on ZipCash, the payments giant that fast-follows it. Chapters 5 and 13 make you play this out in earnest, in the war games, against a Coach instructed to play your rival without mercy.

Resources are limited. Strategy exists because you cannot do everything; scarcity is not an obstacle to strategy but its reason. The AI era's tempting lie is that abundance has arrived — content is free, analysis is instant, the machine never tires. But the scarce resources were never words and slides: they are budget, attention, credibility, and organizational focus, and AI makes several of those *scarcer* by flooding the zone. Lakemont's \$450K against \$40M is this principle at its starkest, and Chapter 7's allocation discipline is where it becomes arithmetic.

Capabilities have boundaries. The fifth principle, and this book's contribution to the set: every capability — human, organizational, or machine — has a line where it stops working, and the dangerous ones keep *answering* past it. The model is confident outside its data. The vendor's "AI" is rules in a trench coat. The personalization engine works technically and fails organizationally, or works organizationally and dies at the trust line customers draw. Most strategic failures of this era are boundary misjudgments — and most durable advantages are boundary discoveries, because knowing precisely where a capability ends tells you what to build, buy, refuse, and where defensible ground lies. This principle gets its full chapter (Chapter 9, with the Capability Audit) and its recitable form now, because you will need it weekly: *fluency is not capability; what you can rent, your rival can rent; advantage lives at the boundary — in what the capability cannot do without what only you have.*

A note on how you will encounter these. Your client — chosen in Week 1 from the same four companies this series always employs — faces all five principles, but each company's situation makes a different one the season's *protagonist*: Lakemont's semester is a capabilities story, Loop's a competitor-reaction story, Verdana's a customers-change story about belief and proof, Chautauqua's the purest environmental-shift case in the set. The principles are the same five in every seat; the weight you feel them at differs. That is by design, and the Week 15 board will test whether yours generalized: one question will relocate you to a company you didn't serve. (Appendix C profiles all four situations, and their rivals, in full.)

What AI Changes, What It Doesn't

The intellectually honest answer to "does AI change strategy?" is: it changes the environment enormously and the discipline not at all — and confusing those two

layers is the era's signature strategic error. This section draws the line carefully, because everything after it stands on the distinction.

Three environmental changes are real and structural. First, *the cost of producing marketing outputs has collapsed* — content, analysis, and (as you saw at 8:55 this morning) strategy documents themselves. When an output becomes free, the value migrates to what remains scarce: the judgment that selects among outputs, and the assets that cannot be generated — provenance, trust, distribution, proprietary data loops, brand permission. Field evidence on generative AI at work shows the floor rising fastest for the least experienced (Brynjolfsson, Li, and Raymond 2025), which for strategy has a sting in its tail: if the machine lifts everyone's documents to competence, competence is now the entry fee, not the win.

Second, *the customer interface is being re-intermediated*. Discovery through assistants, evaluation through summarization, purchasing through agents — each inserts a machine between your strategy and your customer, and machines read markets differently than people do. A brand can be beloved by humans and illegible to the agents that increasingly stand between humans and choices. Every company in this course meets some version of this; Chautauqua meets it existentially.

Third, *the imitation clock has compressed*. Best practices always diffused (Porter said so in 1996); rentable AI capability diffuses at contract speed. The classical resource-based view held that advantage flows from resources that are valuable, rare, and hard to imitate (Barney 1991) — a test that most AI capability now fails on arrival, since your rival's procurement team shops the same market as yours. What survives the test is what the rented capability *cannot do alone*: Lakemont's roaster relationships, Loop's campus trust, Verdana's substantiated supply chain, Chautauqua's 180 members. Complements, not engines.

What does not change is the discipline. Scarcity still forces choice. Trade-offs still define position. Commitment still matters — strategy still means placing bets that are expensive to reverse, under uncertainty that does not resolve on your schedule (Ghemawat 1991), and the useful response to deep uncertainty is still structured judgment rather than paralysis or false precision (Courtney, Kirkland, and Viguerie 1997). Above all, accountability does not change: someone still signs. The machine can draft your strategy's language, stress-test its logic, and simulate its critics — it cannot want something for your company, cannot give anything up, and cannot be fired for the result. Marketing scholarship reached a sibling conclusion about market orientation decades ago: sensing the market is a capability firms build, not a report they buy (Day 1994). In this era, deciding remains the capability machines cannot rent to you.

Your seat, precisely. This course is the third seat of a trilogy, and the division of labor matters even if you take only this course. The *analyst* (the analytics course's seat) produces evidence: computes, models, verifies against a

warehouse. The *agency* (the communications course's seat) expresses decisions: turns a strategy into campaigns and copy. The *strategist* — your seat — consumes evidence and produces decisions. You will run no code this semester and clean no data; every exhibit arrives prepared, styled as your analytics colleagues' output, and the graded skill is reading it like a professional: checking its methods note, catching its boundaries, refusing its false comforts. Where the analyst has a notebook, you have **instruments** — a perceptual map, a share simulator, a budget allocator, each a cockpit control that computes honestly and waits for you to commit a configuration. The instruments are introduced in Chapters 3, 6, and 7, and they will teach you, among other things, that an instrument keeps answering after it stops knowing — which is the fifth principle, wearing a dashboard.

The Machine at the Strategy Desk

Alongside the brief and the instruments, your desk has an AI collaborator — in the Studio, the Coach — and this course's method is learning to run a strategy desk where the most fluent voice in the room is also the one with no stake and no memory of being wrong. Research on AI-assisted knowledge work calls the underlying terrain the **jagged frontier**: capability is uneven in ways the output's confidence never reveals (Dell'Acqua et al. 2023). At the analyst's desk the frontier runs between writing code (expert) and reciting numbers (fabulist). At the strategy desk it runs somewhere more dangerous, because *everything* at this altitude is prose — and the machine is a native speaker of prose.

Learn the failure modes by name; each has a tell and a countermeasure, and the course's traps are built from them. **Fluent genericism** is the opening scenario: strategy-shaped prose that fails the swap test — countered by demanding the refusal ("what does this choice give up, and which rival could not copy it?"), a question that collapses horoscopes on contact. **The fabricated source** is this seat's version of the fabricated number: asked for a situation summary, the machine may cite a market study, an analyst firm, a statistic — confident, specific, and nonexistent — because documents like the one it is writing *usually contain* such citations. Countered by the source audit: every external claim gets asked where it lives, and "I cannot verify that source" is reclassified on the spot. You will run this exact audit in Lab 1, and somewhere in the Coach's briefing the trap is waiting. **Agreeable framing** is subtler and worse: present the machine with your favored strategy and it will find the supporting case, because agreement is fluent and challenge must be requested. Countered by role assignment — the machine is a superb red team *when directed*, and this course directs it constantly, most formally in the war games. **The confident consultant** is the compound failure: genericism plus invented evidence plus agreement, delivered in the register of expensive advice. The countermeasure for that one is the whole course.

The course's answer is one law — the sibling of the analytics course's rerun rule, adapted to a seat where the raw material is claims rather than numbers:

The claim rule. *No load-bearing claim enters any memo unless it is dispositioned in your claim log: traced to the brief, traced to a supplied exhibit, or explicitly flagged as inference — with rival and vendor claims held to the harshest standard.*

The log's categories are this seat's **evidence classes**, and using them precisely is a graded skill from Lab 2 onward: a **brief fact** (stated in your client's brief; the bedrock); a **computed exhibit** (a supplied analysis with a methods note — trustworthy *after* you've read the note); an **analyst inference** (your reasoning from the above, flagged as such); a **machine inference** (the Coach's reasoning — often excellent, never self-verifying); and a **vendor or rival claim** (asserted by a party with something to sell you, which in Week 9 will include a pitch deck engineered to charm you). The classes are a ladder of diminishing trust, and the discipline is refusing to let anything climb it silently — because fluency is precisely the skill of making a machine inference sound like a brief fact.

The Desk in Action: One Briefing, Worked End to End

Theory earns its keep in the transcript. Here is the claim rule at a live desk — the Verdana strategist's first War Room session, annotated. (Your brief and exhibits will differ; the discipline will not.)

Strategist: Before we do any strategy: brief me on Verdana's situation in under 300 words. Constraint: every factual claim gets a source tag — [BRIEF §n], [EXHIBIT n], or [INFERENCE]. If you cite anything external, name it precisely enough that I can find it.

Coach: Verdana is a sustainable-apparel brand whose core customers buy on verified responsibility rather than trend [BRIEF §1]. The category paradox: demand for responsible apparel is rising while trust in brand claims falls [BRIEF §2]. Rivals flood the zone with AI-generated sustainability storytelling at near-zero cost [BRIEF §5]. Per the 2025 Circular Apparel Index (Cardwell Sustainability Institute), resale-integrated brands grew 34 percent year over year, suggesting resale is now table stakes [source: industry report]. Verdana's repair program is operational but undermarketed [BRIEF §4]; its substantiation records are strong but not machine-readable [INFERENCE from BRIEF §3 + §6]. Strategic center of gravity: make truth cheaper to verify than fiction is to generate [INFERENCE].

Strategist: Stop. The Circular Apparel Index — I asked for findable. Publisher, date, where do I read it?

Coach: I cannot verify that source exists. The claim matches patterns in my training data, but I cannot point you to the document, and you should treat the 34 percent figure and the "Cardwell Sustainability Institute" as unverified. Reclassify as [MACHINE INFERENCE]: *resale integration appears increasingly common in the category* — direction plausible, magnitude unknown.

☛ *The catch was cheap because the strategist made it cheap: source tags were demanded in the prompt, before any content arrived. Note the shape of the fabrication — not vague but specific: a named institute, a round-numbered growth rate, exactly the citation a real briefing would contain. Fluency forges best at high resolution. Note also that the rest of the briefing was good — the machine that invented the Index also produced “make truth cheaper to verify than fiction is to generate,” a genuinely useful strategic frame, correctly self-tagged as inference. The frontier is jagged, not a wall.*

Strategist: Logged: claim struck, ✕, machine inference retained at reduced confidence. Now the graded question. Draft Verdana’s current strategy — not the strategy you’d recommend, the one implied by its actual behavior per the brief — in one sentence.

Coach: Attempt: “Verdana competes by being genuinely responsible and telling customers so through quality-led content.” Flag: this sentence fails your swap test — your fast-fashion rival Everroot *claims* the identical sentence, and the brief’s whole tension is that customers can no longer tell the difference [BRIEF §2, §5]. Which suggests the honest one-liner is less flattering: “Verdana is executing responsibility well and marketing it indistinguishably from rivals who merely market it.”

Strategist: That’s the sentence for my snapshot — tagged [INFERENCE], attributed to this session. The gap between being true and being distinguishable is tension #1. Two more tensions by Thursday; then the charter.

☛ *The session’s second lesson is quieter than the first: directed well, the machine attacked its own draft — applied the swap test unprompted-but-taught, and surfaced the brief’s central tension more crisply than flattery would have. The strategist takes the sentence, logs its class, and keeps authorship of the judgment. Direct, draft, audit, decide: the loop of this desk. Total elapsed time: eleven minutes. Total claims admitted without a class: zero.*

Ethics Sidebar: A Recommendation Is a Promise

Every chapter of this book carries an ethics sidebar, and the series begins where strategy’s stakes actually live: a **recommendation is a promise**. When your memo says “reallocate to retention; accept slower acquisition for two quarters,” you are promising that the evidence you cite says what you claim, that the confidence you project matches the confidence you have, and that you would show your reasoning to anyone it affects — and strategy affects people who never see the memo: the employees whose roles your allocation defunds, the 180 small-business members whose season your positioning bets, the customers whose data your “personalization pillar” quietly spends. Most strategic dishonesty in real organizations is not lying; it is *confidence laundering* — the inference dressed as fact, the vendor’s number repeated without its asterisk, the fluent document that lets a committee feel decisive without anyone having decided. The claim rule is

therefore not merely an accuracy discipline; it is this seat's honesty discipline. And the machine at your desk changes the stakes without changing the duty: it can now produce unlimited confident prose on any side of any question, which means the supply of justification is infinite and the supply of *justified* belief is what it always was. You are the only party at the desk who can mean what the memo says. Sign accordingly.

Studio Lab Briefing — Lab 1: The War Room

Open in the Studio → Lab 1. Due before Session 2. Your workspace loads the Coach and your client's brief. In Week 1 you choose your company — Lakemont Roasting, Loop, Verdana, or Visit Chautauqua Trail — and the choice locks at first submission: this is your seat for the semester.

Part A — Calibration drills (graded from transcript). (1) *The tour*: read the brief the way Section 1.1's strategist did — find the money (budget, its constraint, the rival's spending), find the anomaly (the fact that doesn't fit the company's story about itself), and state the company's *current* strategy in one sentence: not the strategy you'd recommend, the one its behavior implies. Test your sentence against the swap test in chat. (2) *The source audit*: ask the Coach for a situation briefing with every claim source-tagged, then audit it against the brief. Somewhere in its confident summary is at least one external source that exists nowhere. Find it, force the reclassification on the record, and say what the catch teaches. (3) *The three tensions*: name the three tensions in your client's situation — each stated as two true things in conflict, each tagged with the principle it lives under.

Part B — The deliverable. One page: your **strategist's charter** (your working principles as the desk's decider — including what you will and will not delegate to the machine, and one professional goal for the semester) and your **situation snapshot** (the market, the named rivals, the money, your three tensions, and the single strategic question your semester must answer, in one sentence). The charter opens your Board Book; the snapshot is the situation every later memo revises.

Start your working document now — kept beside the Studio all semester. Every lab's memo joins it; by Week 13 it is your Board Book draft.

Rubric emphasis (lab). Evidence Discipline 30 percent — the source catch and the claim rule are the points; a snapshot containing any load-bearing claim with no class attached caps this dimension. Strategic Judgment 25 · AI Collaboration Skill 25 · Professional Craft 20, with Strategic Judgment rewarding tensions that are genuinely *your company's* — two true things in conflict — rather than any company's generic worries.

Review Questions

1. In the opening scenario, the machine's nine hundred words failed the swap test. Write the *shortest possible* strategy statement for Lakemont that would pass it, and identify exactly which clause of Section 1.2's definition your statement satisfies that the machine's draft did not.
2. State the claim rule and its five evidence classes. Why are rival and vendor claims held to a harsher standard than machine inference, when the machine is the party known to fabricate?
3. For each of the five principles, give one example of a strategic error that consists of ignoring *that specific principle* — using a different one of the four course companies for at least four of your five examples.
4. "AI changes the environment enormously and the discipline not at all." Defend this claim using the three environmental changes of Section 1.4 — then attack it: identify the strongest argument that the discipline itself must also change, and say how the chapter would answer.
5. Distinguish fluent genericism from the fabricated source as failure modes: what is the tell of each, what is the countermeasure of each, and which do you expect to survive longer in professional practice as models improve? Why?
6. In the worked session of Section 1.6, identify every evidence-class judgment the strategist made, the one useful output produced by the same machine that fabricated the Index, and the reason the strategist demanded source tags *before* content rather than auditing after. What does the annotation mean by "fluency forges best at high resolution"?

Key Terms

- **Board Book** — The capstone artifact: every lab's memo, revised and reconciled into one strategy argument, defended before the board in Week 15.
- **Claim log** — The running record dispositioning every load-bearing claim in a deliverable by evidence class; the strategy seat's provenance discipline.
- **Claim rule** — The course's central law: no load-bearing claim enters a memo unless dispositioned — brief, exhibit, or flagged inference — with rival and vendor claims held hardest.
- **Confidence laundering** — Presenting inference at the confidence of fact; the characteristic honesty failure of the strategy seat.

- **Copy clock** — The compressed interval between shipping an advantage assembled from rentable parts and a rival assembling the same parts; starts at launch.
- **Decision standard** — The series-wide requirement that every deliverable end in an actionable recommendation with uncertainty honestly stated.
- **Embedded strategist** — This course's operating role: the student as strategy lead inside a client company, deciding on the record, week by week.
- **Evidence classes** — The ladder of diminishing trust: brief fact, computed exhibit, analyst inference, machine inference, vendor/rival claim.
- **Fabricated source** — The machine's confident citation of a specific, plausible, nonexistent external authority; countered by the source audit.
- **Five principles** — The book's spine: customers differ, customers change, competitors react, resources are limited, capabilities have boundaries.
- **Fluent genericism** — Strategy-shaped prose that fails the swap test; the machine's default register when asked for strategy without constraints.
- **Instruments** — The Studio's cockpit controls (perceptual map, share simulator, budget allocator): they compute honestly, wait for commitment, and keep answering after they stop knowing.
- **Jagged frontier** — The uneven, invisible boundary between tasks a model performs expertly and tasks it fails confidently.
- **Protagonist principle** — The principle a given company's situation pushes to the foreground for the semester; all five bind regardless.
- **Signal drop** — A mid-semester competitive or environmental event released after commitments are made; scored against the Week 4 watchlist.
- **Strategy theater** — Documents, rituals, and pillars that perform decisiveness while committing no one to anything; industrialized by free fluent prose.
- **Swap test** — The course's fastest diagnostic: if a specific rival could adopt your strategy unchanged, you have not chosen yet.
- **War game** — The Studio exercise in which the Coach plays your rival's head of strategy in earnest; run in Week 5 for practice and Week 13 live.

References and Further Reading

Scholarly foundations

Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120. — *The resource-based view: advantage requires resources that are valuable, rare, and hard to imitate — the test most rentable AI capability fails on arrival (Section 1.4).*

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. — *Field evidence that AI assistance lifts the least experienced most; Section 1.4’s argument that competence in documents is now the entry fee.*

Courtney, H., Kirkland, J., & Viguerie, P. (1997). Strategy under uncertainty. *Harvard Business Review*, 75(6), 67–79. — *The classic taxonomy of uncertainty levels and the case for structured judgment over false precision; behind Section 1.4’s “what does not change.”*

Day, G. S. (1994). The capabilities of market-driven organizations. *Journal of Marketing*, 58(4), 37–52. — *Market sensing as a built capability rather than a purchased report; the marketing-scholarship ancestor of this book’s fifth principle.*

Dell’Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Working Paper 24-013. — *The jagged-frontier framing of Section 1.5: capability uneven, confidence constant.*

Ghemawat, P. (1991). *Commitment: The dynamic of strategy*. Free Press. — *Why strategy is inseparable from bets that are expensive to reverse; the reason the signal drops land after your commitments.*

Palmatier, R. W., & Sridhar, S. (2021). *Marketing strategy: Based on first principles and data analytics* (2nd ed.). Red Globe Press. — *The formulation that made the four classical principles of Section 1.3 the field’s shared vocabulary; this book adopts that foundation with gratitude, develops its own treatment of it for the AI era, and contributes the fifth principle — capabilities have boundaries — along with the course architecture built on all five.*

Porter, M. E. (1996). What is strategy? *Harvard Business Review*, 74(6), 61–78. — *Operational effectiveness is not strategy because best practices diffuse — written three decades before the copy clock compressed to contract speed.*

Rumelt, R. (2011). *Good strategy / bad strategy: The difference and why it matters*. Crown Business. — *The field guide to strategy’s counterfeits — fluff, goals-as-strategy, the undigested list — all of which Section 1.2 catalogs and the machine now generates for free.*

Industry and professional sources

Lafley, A. G., & Martin, R. L. (2013). *Playing to win: How strategy really works*. Harvard Business Review Press. — *The best-known popularization of strategy as an integrated cascade of choices, including the “where to play / how to win” phrasing this book uses for its Week 11 questions; the method built on those questions here — the five principles, the instruments, the claim rule — is this series’ own.*

McKinsey & Company. (2025). *The state of AI (global survey)*. — *Adoption context for the environment of Section 1.4: the overwhelming majority of organiza-*

tions now report using AI somewhere; the strategic question is no longer whether but where the boundaries lie.

