BUSINESS SCHOOL

IT ALL STARTS HERE

BERLIN I LONDON I MADRID I PARIS I TURIN I WARSAW

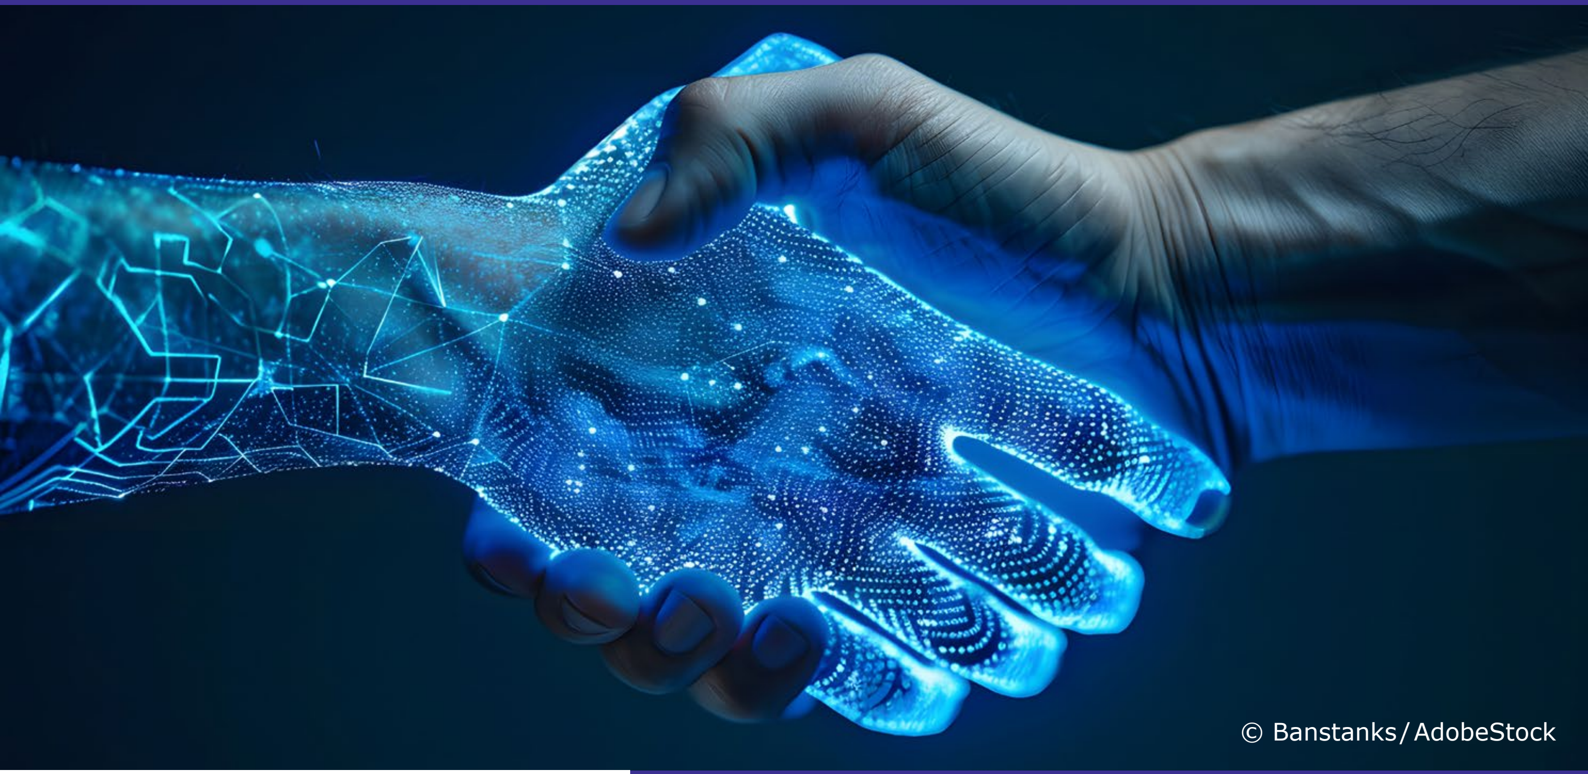

LIGHTS - Technology

Talk machine to me

ESCP Impact Paper No.2026-35-EN

Alara TASCIOGLU
ESCP Business School

ESCP RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ESCP Impact Paper No.2026-35-EN

Talk machine to me

Alara Tascioglu*

ESCP Business School

Abstract

Every year, roughly 900 to 1,000 second-year ESCP Bachelor students build, stress-test, and defend real AI systems on platforms they have never used before. In Year 1 they built a custom GPT for business analysis, and in Year 2 an agentic customer-support system. More than fifty of these student-built systems are now public. Getting there meant solving a curriculum problem business schools have not faced before: generative AI content changes faster than any planning cycle, and every deployment decision encodes values about fairness, energy, governance, and human oversight. My response is a just-in-time curriculum framework delivered as a gamified, GenAI-native build experience: a stable pedagogical backbone paired with annually rebuilt frontier content, and a build-stress-recommend assessment cycle that forces value-laden judgement into the work students produce. Students met these trade-offs informally inside their graded builds; in Year 2 one team chose accuracy and absorbed the higher energy cost, then defended that choice. A five-question values protocol, scheduled for first explicit use in Year 3, formalises this judgement rather than introducing it. The evidence shows the architecture can be rebuilt annually and delivered at three-campus scale, with recommendations for the next five years of AI education in business schools. Whether the framework improves student judgement is the subject of the next study.

Keywords: generative AI, business education, curriculum design, sustainability, AI literacy

*Assistant Professor, ESCP Business School

Talk machine to me

Introduction

In Year 2's third case study (2025-26), a second-year ESCP student group built a customer support agent on an unfamiliar platform. The case study was meant to teach them how to evaluate an AI system. What they took away was the cost of one.

When I finalised my syllabus in June 2024, GPT-4 was the frontier. By September, the terrain had shifted. By December, an assignment built around a prompting technique felt like a relic. Generative AI moves faster than any academic planning cycle, and the speed is not only technical but normative: every choice about what to teach and which tools to deploy encodes values about fairness, environmental responsibility, and human agency.

Business schools can treat AI as a technical module and hope the fundamentals transfer. I argue they also need a curriculum rebuilt annually around judgement, teaching not just how AI works but how to decide whether to use it, for whom, and at what cost.

I began designing "AI for Big Data Management" in early 2024 and delivered its first iteration that September. Since then, I have rebuilt it annually for second-year Bachelor students at the Paris, Madrid, and Turin campuses.

I design the shared curriculum and assessment system centrally, and campus instructors deliver it. The central challenge is not what to teach, but how to build a course architecture that survives a technology reinventing itself every semester while keeping human judgement inside the build, where students defend the choices in their systems rather than treat ethics as a closing lecture.

The domain pairing is the distinctive design contribution. My students arrive with an introduction to generative AI but only half of their coding sequence, so this course is their first real encounter with business analytics and data science, a field they often believe is dry. That first impression matters. What makes the design original is not using AI tools in a classroom, now common, but pairing that first encounter with a gamified, GenAI-native delivery. Modern AI grew out of data science and amplifies it, so I teach them together, and I built the delivery so that gaining real analytics and big-data competence feels like discovery rather than a slog.

In the framework, the pedagogical backbone stays stable across years, while content, tools, and cases are rebuilt annually. Inside it, an assessment cycle asks students to build, stress-test, and recommend. A third element, a five-question values protocol I developed across the first two cycles, is scheduled for Year 3 and offered as a design proposition, not a delivered instrument. The framework itself is the primary contribution, with the assessment cycle and protocol as components.

Locked content may be dated before the course begins, so the course should model the adaptability it asks of students, who see the syllabus change between cohorts. If generative AI makes packaged knowledge easy to access, the competitive question shifts from who has the best content to who can renew it fastest.

I designed the framework to work across knowledge levels and contexts. What is new is not the individual principles, each with prior work in pedagogy or AI literacy, but how they bind into one architecture that forces annual renewal by design. Modular course design lets an instructor swap a single unit and call the course updated; this configuration instead makes renewal unavoidable, tying the assessment to tools students have not seen, so the course cannot quietly settle around last year's cases. Because its parts are tightly coupled rather than modular, the architecture needs central coordination to stay coherent as its content is renewed every year, much as a tightly coupled design needs a single point of integration when a change in one part affects the others (Ulrich, 1995). That is why, for now, I extend the framework myself rather than releasing it as an open template. It builds on AI literacy work (Long & Magerko, 2020) and the call to reconceive critical thinking under AI (Larson et al., 2024).

Three design principles

Three principles make annual renewal feasible without heroic effort. Together they define what stays fixed, what gets rebuilt, and what the course aims to produce. Figure 1 shows how the parts fit together.

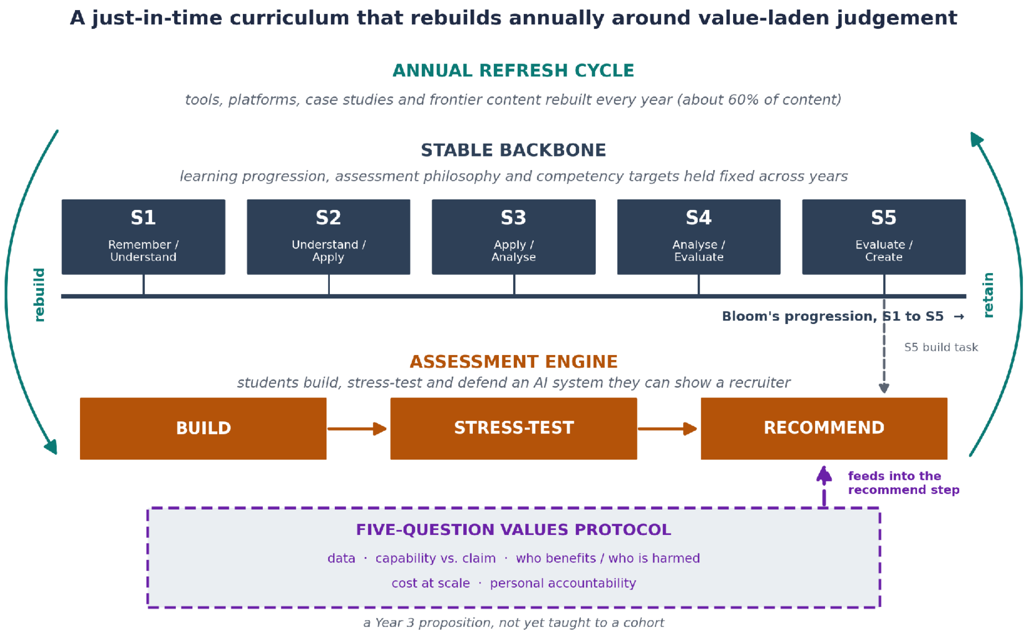

Figure 1. The just-in-time curriculum framework.

Fix the backbone, refresh the frontier.

The first principle separates skeleton from musculature. From my Year 1 syllabus (September 2024), every piece of content carries a tag: backbone (learning progression, assessment philosophy, competency targets) or refresh (tools, platforms, case data, regulatory examples). Bloom's revised taxonomy (Anderson & Krathwohl, 2001) and Biggs's constructive alignment (Biggs, 1996) keep objectives, activities, and assessment aligned even as the content changes. Session 1 always orients students in the field. Session 5 always requires them to build something and evaluate it. What they explore and build changes by design.

Teach tools as instances, not destinations.

Within that structure, I teach the frontier rather than only the fundamentals. Data analysis, classification, forecasting, and workflow design are taught through GenAI-mediated build tasks, not as separate technical modules. Students encounter the actual tools and platforms they will use in their first jobs, accepting the risk that a tool taught in October may not exist in March of the following year. Every encounter frames a tool as an instance of a broader category, so students learn both the specific skill and the transferable pattern. The aim is adaptive expertise (Hatano & Inagaki, 1986): the capacity to solve novel problems on platforms they have never seen, not only to execute familiar techniques.

Assess judgement under unfamiliar conditions.

Tool-specific AI skills have a half-life of months. What I want students to leave with is meant to endure: the disposition to approach unfamiliar technology critically. Every session embeds a "try it, measure it, compare it, question it" cycle, and students choose their platform, approach, and deliverable form, treating that choice as a pedagogical move. The artefact they build becomes the site where that judgement is rehearsed. If last year's slides answer this year's exam, the test measures memory, not judgement.

These principles operate across a five-session backbone. Each session is designed against a Bloom's level, so cognitive demand rises across the term while the topic surface refreshes annually. Table 1 shows the mapping.

Table 1. Five-session backbone mapped to Bloom's revised taxonomy (Anderson & Krathwohl, 2001)

Session	Focus	Bloom's level	Backbone activity
S1	AI and big data fundamentals, ethics, regulation, sustainability	Remember / Understand	Vocabulary alignment, framework recall
S2	Data mining, machine learning, classification metrics	Understand / Apply	Run algorithms on a labelled dataset, interpret outputs
S3	Prompt and context engineering, time-series forecasting, failure modes	Apply / Analyse	Compare strategies and diagnose where each fails
S4	AI investment debate, system building, and the five-question protocol (introduced in Year 3)	Analyse / Evaluate	Structure debate, then introduce the five critical questions (Year 3 deployment)
S5	Workflow design, embeddings, agent evaluation	Evaluate / Create	Build, stress-test, recommend

The framework in action: year-over-year evidence

The Year 1-to-Year 2 rebuild is the clearest evidence that the cadence I designed is feasible. One cycle does not prove durability, but it shows the architecture holds.

Each cohort reaches 900 to 1,000 second-year students across three campuses. A shared end-of-course evaluation is mandatory under ESCP policy. By 1 May 2026, the build-stress-recommend cycle had produced more than fifty distinct student-built spaces on ESCP's Hugging Face organisation page (huggingface.co/ESCP), after excluding instructor templates and duplicate re-uploads. Those spaces are public under the students' names, so each team ships a working system rather than only handing in a file. One team's delivery-route optimization app runs the team's analysis notebooks at runtime and wires a language model to an automation agent, all from a public repository. One student named what the course is for: *"it's good we learn how to use AI responsibly instead of banning it"* (anonymous Year 1 evaluation). Each team also submitted a short memo explaining the trade-off it accepted. One team recorded its choice: *"we shipped [a teammate's] build because its design coherence won the trade-off, and we accepted that the [email] integration and other [model-specific] features would not be in the deployed product"* (anonymous Year 2 team report). I read the responses now finalised across the three campuses to surface design tensions and recurring student judgements, not cross-campus satisfaction or learning effects, which the forthcoming study is designed to test. Where I quote a student, the line illustrates a tension the shared curriculum surfaces, not an outcome attributable to any campus or section. Table 2 documents which elements held as backbone and which were rebuilt across the two cycles.

Table 2. Year-over-year content evolution across the two rebuilds I led: backbone vs. refresh (CS = case study, Backbone = stable across years, Refresh = rebuilt annually, Growth = structural expansion of hours and credits)

Element	Year 1 (2024–25)	Year 2 (2025–26)	Type
Bloom's progression (5 sessions)	Stable	Stable	Backbone
Python + Google Colab	Stable	Stable	Backbone
Sustainability/ethics (recurring)	Stable	Stable	Backbone
Build-stress-recommend assessment	Stable	Stable	Backbone
CS1: KPMG client analysis	Retained	Retained	Backbone
CS2: societal analysis case	Pfizer/EY analytics	Future Jobs with AI	Refresh
CS3	Prompt engineering	AI Agent evaluation	Refresh
Sessions 4–5 content	Prompt engineering focus	AI bubble + agentic systems	Refresh
Course weight	10 hrs / 1 ECTS	15 hrs / 2 ECTS	Growth

Note. Cohort sizes: 914 enrolled in Year 1, approximately 988 enrolled in Year 2 across the three campuses. Year 3 (2026–27) is in design at the programme level, with further growth in scheduled hours and credit weight anticipated.

In Year 2, I rewrote about 9 of the 15 scheduled hours, rebuilding Sessions 3–5 with new slides, activities, platforms, and case data while keeping Sessions 1–2 and the learning objectives stable. The routine is simple: audit each content tag against the frontier, replace the stale cases, test a platform, rewrite the affected sessions, and leave the backbone untouched. Between the two years, agentic AI moved from research curiosity to production deployment and the EU AI Act (European Parliament, 2024) entered its first enforcement phase. A curriculum frozen at Year 1 would have prepared students for a world that no longer existed.

Year 1's CS3 deliverable was a custom GPT for automated business analysis. Year 2's is a working customer support agent assembled on a platform new to the cohort. The deliverable shape stays consistent: a built, stress-tested AI system the student can present to a recruiter, defending the trade-offs.

Recent evidence strengthens the case for this structure. Bastani et al. (2025) show that students given GPT-4 access without scaffolding perform substantially worse after access is removed than students who never had it, while designed safeguards largely mitigate this erosion. The build-stress-recommend cycle is a curriculum-level instance of that scaffolding logic in the spirit of productive failure (Kapur, 2008): the stress phase forces confrontation with failure modes, and the recommend phase requires judgement that survives the AI's absence.

Five critical questions as a decision protocol

This protocol is a design proposition scheduled for Year 3; it has not yet been taught to a cohort. The judgement it formalises, however, already appears in the delivered work described below.

Session 4 is where the course pivots from skill-building to judgement. Students debate the AI investment bubble, surfacing the claims AI products make about themselves. I began developing these questions in September 2024 and refined the sequence by observing the first two cohorts. Some distil tensions I watched recur: the gap between what a tool demonstrated and what its marketing promised, stakeholder-harm blind spots in societal-analysis cases, and the cost that surfaced once teams put an agent into a prototype. The Year 3 rule will be simple: before recommending any AI system, a student answers all five in order and shows where the evidence for each answer comes from:

1. What data was this trained on, and who curated it?

2. What does the system actually do versus what the marketing says?

3. Who benefits if this works, and who is harmed if it fails?

4. What is the total cost of running this at scale, including energy and compute?

5. Would I trust this system with a decision that affects someone's livelihood?

The order matters: the questions trace an AI system's lifecycle from training data (Q1) to accountability (Q5). The Year 2 cohort met these questions informally before the protocol was named. A Year 2 team's KPMG client-risk report made Q3 explicit, choosing to *"favor higher recall … at the cost of lower precision"* on the reasoning that *"missing a risky client can be more costly than over-flagging"* (anonymous Year 2 team report). Another team's human-handoff fallback for ambiguous queries was Q5 in practice, declining to let the system settle the higher-stakes calls. The Year 3 protocol formalises this kind of move, asking the questions in a defined order rather than waiting for them to surface.

CS3 already places sustainability inside the recommendation criteria (Schwartz et al., 2020). The Year 3 deployment will bind Q3–Q5 to the recommend step of the build-stress-recommend cycle, requiring students to name the sustainability trade-off and defend it, making this judgement a routine part of the build.

Limitations

Four caveats bound the claims that follow. The framework is, so far, the work of a single designer. I created it and rebuild it each year, which keeps the design coherent but concentrates it in one person, even as I run pre-session meetings with the instructors, to brief them on each rebuild and see how the materials land. One Year 1 respondent named the delivery load early: *"I believe the course is in its early stages of development, so much will get changed in the following years... I think the course is relevant and heavy enough to warrant a second teacher working alongside Alara to help her manage the students."* (anonymous Year 1 evaluation). My route to scaling is not to hand the design over but to make it teachable: for Year 3 I am building a structured playbook and an instructor-training model, so the curriculum can be delivered faithfully by others.

A second tension sits inside the design. I designed the course to be demanding but fun, with walkthrough videos for catch-up and case studies that run as small quests with concrete deliverables. In the evaluations I read, the combination registers more often as energising than as overload, though for some students the demand tips into overload, a gap the catch-up videos do not always close. Minimal scaffolding creates space for discovery but also for unproductive struggle, a line the design manages deliberately. A Year 1 respondent captured the failure mode: *"The class activities were difficult to follow. We jumped directly to difficult content without learning basics which made the learning process even more difficult."* (anonymous Year 1 evaluation). Among recurring positive themes in the free-text responses, students wrote that *"the examples and explanations made complex topics more accessible and relevant"* and that the course *"was very enriching intellectually... it helped me develop my critical thinking skills"* (anonymous Year 1 evaluation). The two cohorts are different second-year groups, so the contrast marks the two poles of one unresolved split, not a year-over-year recovery. The framework is not a single-campus pilot but runs across all three campuses for close to 1,000 students a year, and the cohort keeps growing. A structured comparison of how the rebuild lands across campuses is work for the forthcoming study. Year 2 added explicit "abandon this path" checkpoints, and Year 3 will pilot tiered scaffolding tracks.

A fast-moving, centrally designed curriculum also requires something of the instructors: because the materials change each year, teaching it sometimes means learning alongside the students. I support that deliberately, sharing session materials ahead of time, running optional pre-session walkthroughs, and helping with unfamiliar build platforms. One instructor reflected afterward: *"I learned a lot through the course... the experience for me was definitely enriching."* (anonymous instructor, Year 2). The point is not a shortcoming on anyone's part; it is that a curriculum changing this fast carries tacit knowledge a syllabus alone does not transfer.

What this paper documents is design and its reasoning, not measured learning outcomes. A forthcoming study will analyse pre/post AI self-efficacy scales, task-transfer tests, and campus comparison. Until then, the claims rest on design logic, behavioural observation, and one rebuild cycle. The five-question protocol is part of that design rather than delivered evidence: the two prior cohorts met the assessment cycle but not the protocol as a named diagnostic.

Design propositions for the next five years

Two years and three campuses produce four recommendations for business schools navigating AI education. These are design propositions tested against one rebuild cycle, offered for further validation, not conclusions. Each addresses a different actor (Table 3).

Table 3. Four design propositions, by stakeholder

For	Recommendation	Why it matters
Accreditation bodies	Recognise values embedded in assessment criteria, hands-on activities, case studies, and final deliverables as substantive evidence of integration, rather than treating a stand-alone ethics module as the primary indicator	Standalone ethics modules get crowded out by the next refresh, while values built into assessment criteria persist across content rebuilds
Programme directors	Budget protected design time (roughly three weeks per year) and platform access costs (a few hundred euros) for the annual refresh	The hidden cost is institutional willingness to ring-fence faculty time for curriculum renewal rather than additional teaching
Faculty designers	Invest in instructor training and structured playbooks so a single coherent design can be delivered faithfully by others	Teaching at the edge of one's own familiarity is hard; trained delivery and a clear playbook transfer the capability without fragmenting the design across a committee
Assessment owners	Demand assessments built around judgement under unfamiliar conditions: build, stress-test, defend on platforms students have not seen	An assessment a student could pass by recalling last year's content fails to test the disposition this course aims to build

A values-embedded curriculum shows its worth when students make more explicit and defensible recommendations, weighing fairness against speed. When the support-agent team chose accuracy and absorbed the energy cost, they were doing what a business education still has to teach: reading a trade-off no slide deck had pre-decided for them (Larson et al., 2024). When content itself is cheap, the durable advantage may be the institution's capacity to rebuild its curriculum every year rather than any single year's version of it, a proposition this two-year demonstration motivates but cannot yet confirm. That capacity lives in faculty time, infrastructure access, and assessment design, and the willingness to fund it is the institutional asset that AI content-volatility makes urgent.

Next September, I will rebuild the course for the third time, and pilot the cycle as a parallel executive-education track. I already know I will throw out the agentic platform CS3 used in Year 2. I do not yet know whether the five critical questions will land as a sequence or fragment under classroom pressure. I would revise the design if any of three things happens: if annual rebuilding weakens cross-campus delivery, if the Year 3 scaffolding fails to reduce the overload the evaluations show, or if the five-question protocol does not make students' trade-off reasoning more explicit. Adaptability is what the course is really teaching. Students must adjust to tools that did not exist when they enrolled, and a course earns the right to ask that only by adapting itself, year after year. The capacity to rebuild is the point.

References

Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives.* Longman.

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences,* 122(26), e2422633122. https://doi.org/10.1073/pnas.2422633122

Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education*, 32(3), 347–364. https://doi.org/10.1007/BF00138871

European Parliament. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act).* Official Journal of the European Union. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

Hatano, G., & Inagaki, K. (1986). Two courses of expertise. In H. Stevenson, H. Azuma, & K. Hakuta (Eds.), *Child development and education in Japan* (pp. 262–272). W.H. Freeman.

Kapur, M. (2008). Productive failure. *Cognition and Instruction,* 26(3), 379–424. https://doi.org/10.1080/07370000802212669

Larson, B. Z., Moser, C., Caza, A., Muehlfeld, K., & Colombo, L. A. (2024). Critical thinking in the age of generative AI. *Academy of Management Learning & Education*, 23(3), 373–378. https://doi.org/10.5465/amle.2024.0338

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). ACM. https://doi.org/10.1145/3313831.3376727

Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM,* 63(12), 54–63. https://doi.org/10.1145/3381831

Ulrich, K. (1995). The role of product architecture in the manufacturing firm. *Research Policy,* 24(3), 419–440. https://doi.org/10.1016/0048-7333(94)00775-3

I thank the instructors across Paris, Madrid, and Turin who deliver the curriculum and assessments I designed, supervise in-class hands-on activities, and manage the logistical complexity of multi-campus delivery.