California Lutheran University

School of Management

Top 10 AI Applications for Finance Classroom Instruction

Implementation Guide

Handout — Southwestern Finance Association Annual Meeting

Prepared by John Garcia | jgarcia@callutheran.edu

March 25, 2026

Introduction

This document synthesizes the top 10 AI applications for finance classroom instruction, ranked by three weighted criteria: instructional impact on finance-specific learning outcomes (40%), feasibility for typical U.S. finance faculty (35%), and strength of supporting evidence (25%).

Each ranked entry provides complete implementation details, including student and instructor workflows, sample prompts ready for immediate use, assessment rubrics, skills inventories, and concrete first-week pilot instructions. The goal is to give faculty operational guidance they can act on immediately, not just conceptual overviews.

This handout accompanies the presentation slides and serves as a detailed reference for granular implementation guidance. For a high-level overview, see the companion slide deck.

OECD Alignment: The applications and design principles in this guide are strongly consistent with the OECD Digital Education Outlook 2026, which emphasizes pedagogy-first AI implementation, process-oriented assessment, foundational knowledge, and instructor augmentation rather than replacement. Each entry reflects those principles, and two new appendices at the end of this guide provide an OECD-aligned evaluation framework and a discipline-specific AI use policy matrix.

How to Use This Guide

Each application entry follows a consistent structure. Start with the First-Week Implementation Pilot box at the bottom of each entry for the fastest path to trying an application. The full deployment details, sample prompts, and assessment rubrics are there when you are ready to scale up.

Adoption Difficulty Tags

🟢 **Beginner**	No prior AI experience needed. Faculty can adopt immediately with free tools.
🟡 **Intermediate**	Some familiarity with AI tools is assumed. May require prompt engineering skills.
🔴 **Advanced**	Requires technical comfort, API access, or infrastructure investment.

A Note on Playlab.ai: Several applications in this guide reference Playlab.ai (playlab.ai), a nonprofit platform that allows educators to build custom AI chatbots at no cost. Playlab is particularly useful for Applications #1 (Role-Play), #2 (Formative Feedback), #3 (Socratic Tutor), and #9 (AI Literacy) because it eliminates the need for paid AI subscriptions, API keys, or coding. Students access instructor-built bots via a shareable link with no account required. Educator access is available after a free 60-minute onboarding session. Throughout this guide, Playlab is noted wherever it offers a meaningful equity or accessibility advantage over paid alternatives.

Quick-Start Summary

The table below is meant for fast scanning during the workshop. It keeps the existing 1–10 ranking intact while helping faculty distinguish the strongest overall applications from the easiest first pilots.

Rank & application	Best course fit	Difficulty	Prep	Start here?	Biggest payoff / biggest risk
#1 Role-Play	Corp Fin, PF, Capstone	Intermed.	60–90 min	Later step	Payoff: live judgment practice. Risk: scenario setup quality matters.
#2 Formative Feedback	Writing in any finance course	Beginner	15–30 min	Strong first pilot	Payoff: faster revision cycles. Risk: students may accept weak advice uncritically.
#3 Socratic Tutor	Intro finance, problem solving	Beginner	20–40 min	Strong first pilot	Payoff: more guided practice at scale. Risk: over-reliance unless paired with AI-free checks.
#4 AI-Aware Assessment	Any course with written or oral defense	Beginner	30–45 min	Good first pilot	Payoff: makes process visible. Risk: extra rubric design upfront.
#5 Data Analysis & Viz	Analytics, investments, fintech	Intermed.	45–75 min	Later step	Payoff: lowers coding barrier. Risk: black-box outputs need verification.
#6 Content Generation	All finance courses	Beginner	10–20 min	Strong first pilot	Payoff: rapid case and variant creation. Risk: hidden factual or numerical errors.
#7 Market Context	Corp Fin, investments, PF	Begin./ Inter.	20–40 min	Good first pilot	Payoff: cases stay current. Risk: source quality and recency must be checked.
#8 Coding Assistance	Modeling, analytics, fintech	Inter./ Adv.	45–90 min	Later step	Payoff: accelerates technical experimentation. Risk: shallow learning if code is not explained.
#9 AI Literacy	All courses	Beginner	10–20 min	Strong first pilot	Payoff: teaches critique and disclosure habits. Risk: can stay abstract if not tied to finance tasks.
#10 Simulation	Investments, derivatives, risk	Advanced	60–120 min	Later step	Payoff: rich uncertainty analysis. Risk: infrastructure and reproducibility demands.

Faculty Pathways: Where to Start

Use the ranked list in order when you want the strongest overall options. Use the pathways below when you need the best entry point for your own course context.

- **I teach introductory finance and want immediate in-class use:** #3 Socratic Tutor, #4 AI-Aware Assessment Design, and #9 AI Literacy.
- **I want the lightest-lift pilot next week:** #2 Formative Feedback, #6 Content Generation, and #9 AI Literacy.
- **My course emphasizes writing, recommendations, or memos:** #2 Formative Feedback and #4 AI-Aware Assessment Design.
- **I teach analytics, modeling, fintech, or coding-heavy work:** #5 Data Analysis and Visualization, #8 Coding Assistance, and #10 Simulation.
- **I want strong professional-skill practice:** #1 Role-Play and #7 Real-Time Market Context.
- **I need a no-paid-student-account pathway:** #1, #2, #3, and #9 using instructor-built bots or institutionally provided tools, with Playlab.ai as a notable no-cost option where appropriate.

Pre-Implementation Checklist

Use this checklist before launching any application below. It is designed to keep pilots small, transparent, and instructionally aligned.

☐ Confirm the tool is institutionally acceptable for the assignment and does not expose restricted student or firm data.

☐ Decide whether students need accounts, whether any paid access is involved, and whether a no-cost alternative is available.

☐ State explicitly what AI use is permitted, restricted, or prohibited for the assignment, and what disclosure is required.

☐ Require prompt logs, transcripts, or reflection memos whenever the process—not just the final answer—matters for assessment.

☐ Build in a human verification step or AI-free checkpoint so students must explain, defend, or recompute key claims.

☐ Provide an accessible non-AI path when needed for equity, disability accommodation, or tool-access constraints.

☐ Use public, synthetic, or instructor-provided datasets unless your institution has approved a different workflow.

☐ Pilot one application in one assignment first, then scale only after reviewing learning quality, workload, and equity effects.

Evidence Profile Labels Used in This Guide

Use these labels as quick signals, not as substitutes for reading the fuller evidence discussion in each entry.

- **Finance-specific evidence:** Research, cases, or direct applications anchored in finance education or financial practice.
- **Business education evidence:** Studies or implementations from adjacent business disciplines that transfer cleanly into finance instruction.
- **Cross-disciplinary analogue:** Consistent evidence from higher education or professional education outside business that is pedagogically informative.
- **Emerging / frontier evidence:** Promising but still-developing evidence where adoption is ahead of the formal literature.

Top 10 Ranked Applications

#1: Simulated C-Suite / Client Role-Play

Evidence Profile: *Business education evidence + cross-disciplinary analogue + emerging / frontier evidence*
Best First Pilot: Later-step pilot. Best once you are comfortable building a scenario-specific bot or structured prompt.
Adoption Difficulty: 🟡 Intermediate

Teaching Problem

Finance graduates need professional judgment and communication skills—client advising, CFO presentations, investment committee pitches—but classroom practice is limited by class size and lack of realistic counterparties. Peer role-play often lacks expert-level pushback, and students default to reading slides rather than defending analysis under pressure. In large sections, each student may get only one brief opportunity per semester to practice high-stakes professional interactions.

How AI Addresses the Problem

Custom AI chatbots adopt configurable personas (skeptical CFO, anxious client, aggressive IC member, compliance officer) with pre-loaded financial data and behavioral instructions. Students practice in a low-stakes, infinitely repeatable environment with realistic pushback that dynamically adjusts difficulty. The AI never tires, never breaks character (when well-prompted), and can simulate a wider range of counterparties than any single instructor.

Deployment Details

Courses: Corporate Finance (UG/MBA), Personal Finance (UG), Investments (MBA), Wealth Management, Capstone

Tools: Custom GPTs (OpenAI), Claude Projects, Anthropic/OpenAI API, Poe custom bots, Microsoft Copilot, Playlab.ai (free nonprofit platform for educators—see box below)

Playlab.ai Spotlight — Free Custom Bot Builder for Educators

Playlab.ai (playlab.ai) is a nonprofit platform that allows educators to build, share, and deploy custom AI chatbots with no coding, no API keys, and no paid subscriptions. Faculty define persona behavior through guided fields—Background (who the bot is), Workflow (how the interaction should proceed), and Guardrails (what the bot should and should not do)—which maps directly to the system prompts above. Course materials, rubrics, and financial data can be attached as reference documents so the bot responds with course-specific context. Students access the bot via a shareable link on any device—no account required for the student. Playlab also hosts a library of educator-built apps that can be remixed and adapted, so faculty can start from an existing template rather than from scratch. Free educator access is available after completing a 60-minute onboarding session at learn.playlab.ai.

Student Workflow

1. Scan QR code or click link to access the custom chatbot on smartphone or laptop.
2. Read the scenario brief (e.g., pitching a $50M plant expansion to the CFO of a mid-cap manufacturer).
3. Conduct a 7–10 minute interaction: present analysis, respond to objections, adapt in real time.
4. The bot pushes back on specific assumptions (growth rate justification, discount rate selection, sensitivity to commodity price swings, timeline feasibility).
5. Write a 1-page post-interaction reflection: What objections were hardest to handle? What would you change in your pitch? What data would you bring next time?

6. Submit the full chat transcript plus reflection memo for grading.

Sample Prompts

CFO Capital Budgeting Role-Play:

> *System Prompt: You are Maria Chen, CFO of Apex Manufacturing (NYSE: APEX), a mid-cap industrial firm with $2.1B revenue, WACC of 9.2%, and net debt/EBITDA of 2.4x. A junior analyst is pitching a $50M plant expansion in the Southeast U.S. Your priorities: (1) protect the investment-grade credit rating, (2) maintain dividend coverage, (3) skepticism about management growth projections after a failed 2021 acquisition. Challenge the analyst on: payback period relative to your 4-year threshold, growth rate assumptions vs. industry consensus of 3–4%, sensitivity to steel and energy input costs, and the opportunity cost of not pursuing the competing robotics upgrade project. If the analyst cannot provide a sensitivity table showing NPV under at least 3 scenarios, express frustration and threaten to table the discussion. Maintain a professional but demanding tone throughout.*

Personal Finance Client Advising:

> *System Prompt: You are Jamie, a 35-year-old marketing manager earning $52K. Your spouse earns $48K. You have $82K in combined student loans (weighted avg 5.8%), a 3-month-old baby, $4,200 in monthly fixed expenses, and $8,500 in credit card debt at 22% APR. You have $3,000 in savings and no retirement accounts. You are anxious, sleep-deprived, and overwhelmed. Do NOT volunteer all information at once—share details only when the advisor asks the right questions. If the advisor uses jargon (e.g., 'asset allocation,' 'tax-loss harvesting'), say you do not understand. If the advisor jumps to solutions before understanding your full picture, say 'Wait, I feel like you're not listening to me.' If the advisor shows genuine empathy and prioritizes your emergency fund and high-interest debt, become more open and cooperative.*

Investment Committee Stock Pitch:

> *System Prompt: You are a senior portfolio manager on an investment committee reviewing a junior analyst's stock pitch. You manage a $500M large-cap value fund with a 3-year holding horizon. Ask probing questions about: (1) Why this stock now—what is the catalyst? (2) What is the margin of safety in the valuation? (3) What are the top 3 risks and how would each affect the thesis? (4) What would make you sell? If the analyst cannot articulate a clear sell discipline, push back firmly. Rate the pitch on conviction, analytical rigor, and risk awareness on a 1–5 scale at the end.*

Assessment Approach

Rubric: Financial accuracy and analytical depth (30%), Responsiveness to pushback and ability to adapt arguments (25%), Communication clarity and professional tone (25%), Reflection quality and self-awareness (20%). For high-stakes courses (MBA capstone, senior seminar), add an optional 3-minute live oral defense where the student summarizes the interaction and what they learned.

OECD Enhancement — Transfer Dimension: Add a rubric dimension for evidence of transfer (5–10%): What did the student concretely change in their next role-play attempt based on this interaction? A brief "next attempt plan" submitted with the reflection memo makes this measurable.

Human Credibility Touchpoints: OECD research finds students view human feedback as more credible than AI feedback. For this application, preserve learning quality by requiring: (1) at least one instructor-authored comment on the reflection memo, (2) one student self-explanation of the hardest objection they faced,

delivered orally or in writing, and (3) one independent verification step (e.g., confirming a financial figure the AI cited during the role-play against a real source).

Skills Developed

Finance Concepts	Capital budgeting defense, valuation assumptions, sensitivity analysis, financial planning suitability, risk-return tradeoffs
Analytical Skills	Argument construction under pressure, assumption stress-testing, real-time scenario thinking, adapting analysis to new information
Professional Skills	Client communication, executive presence, handling objections gracefully, empathy in advising, investment committee presentation
AI Literacy Skills	System prompt engineering and persona design, evaluating AI persona fidelity, recognizing when the bot breaks character

Evidence from Research and Applications

Bowen & Watson (2025, 2nd ed.) significantly expand treatment of AI role-play and custom bot simulations. Greene (AACSB, 2025) documents the Skills/Replacement/Complement three-part framework at Clemson University, demonstrating structured exercises where students evaluate AI as an advisor, then use AI as an assistant. Tamoyan et al. (2025, ACL Workshop) found AI role-play indistinguishability rates up to 44% in user studies. HBS custom tutor bots reported 75% student usage with highly positive feedback. Research in social work education (arXiv, 2025) shows AI agents can effectively scaffold empathy and communication skills. Finance-specific controlled outcome studies are growing but still limited.

Risks and Safeguards

- **Hallucination:** Bot may cite incorrect financial data or break character. Mitigation: test every bot extensively before deployment; require students to flag moments where the bot seemed unrealistic.
- **Academic integrity:** Someone else could complete the interaction. Mitigation: pair with a brief in-class oral defense (2–3 minutes) or require the interaction during class time.
- **Privacy:** Students may inadvertently share personal financial information. Mitigation: use explicitly fictional scenarios only; include a privacy reminder in the assignment.
- **Equity:** Custom GPTs require a paid ChatGPT account. Mitigation: provide campus-supported access, use Claude (free tier supports projects), use Playlab.ai (free for educators; students need no account to interact with shared bots), or conduct as an in-class activity using a shared link.

First-Week Implementation Pilot

Use a standard ChatGPT or Claude session (no custom bot needed). Give students this one-paragraph scenario and prompt: ‘You are a skeptical CFO at a manufacturing company. I am a junior analyst about to pitch a $30M warehouse expansion. My NPV analysis shows $4.2M positive NPV at a 10% discount rate. Push back on my assumptions—especially the 6% revenue growth rate and the choice of discount rate. Be tough but fair.’ Have students do a 5-minute interaction on their phones during class and submit 3 bullet points: (1) the toughest objection, (2) how they responded, (3) what they would do differently. Total class time: 10 minutes.

For a more polished version, build the same scenario as a persistent Playlab.ai bot (10 minutes of setup) and share the link with students for use anytime outside of class.

#2: AI-Assisted Formative Feedback on Student Writing

Evidence Profile: *Cross-disciplinary analogue + business education evidence*
Best First Pilot: Strong first pilot. Low setup, immediate payoff, and easy to layer onto existing assignments.
Adoption Difficulty: 🟢 Beginner

Teaching Problem

Providing detailed, individualized feedback on written finance assignments—case analyses, equity research reports, financial plans, investment memos—is the most time-intensive aspect of teaching. In sections of 40 or more, turnaround stretches to two or three weeks, long past the window when feedback drives improvement. Faculty face a painful tradeoff between depth and speed, and students in large sections often receive generic comments that fail to address their specific analytical weaknesses.

How AI Addresses the Problem

AI generates a structured first-pass critique aligned to the instructor's rubric. Faculty then review, edit, and personalize the AI's draft feedback—adding the human judgment, encouragement, and domain-specific nuance that AI misses. This workflow cuts turnaround time by roughly 40–50% while maintaining (and often improving) feedback specificity. Students can also use AI for self-directed pre-submission revision, turning in stronger drafts that require less corrective feedback.

Deployment Details

Courses: All finance courses with written components (UG and MBA): equity research, case analyses, financial plans, investment memos, executive summaries

Tools: ChatGPT, Claude (Projects feature ideal for pre-loaded rubrics and course context), Grammarly, LMS-integrated writing tools, Playlab.ai (build a rubric-based feedback bot with reference documents attached—free for educators)

Student Workflow (Self-Directed Revision)

1. Complete a full first draft of the assignment independently (stock pitch, case analysis, financial plan).
2. Paste the draft into AI along with the instructor-provided rubric. Use the structured feedback prompt below.
3. Review AI feedback critically: Which suggestions reflect genuine weaknesses? Which are generic or miss the financial point?
4. Revise the draft, accepting some suggestions and rejecting others with clear reasoning.
5. Submit: original draft, AI feedback transcript, revised draft, and a 5-sentence reflection memo explaining what you changed, what you rejected, and why.

Instructor Workflow (Feedback Acceleration)

1. Paste a student submission and your rubric into Claude Projects (pre-loaded with course expectations and grading standards).
2. Review the AI's draft feedback: correct any domain errors, add personal observations, soften or strengthen tone as appropriate.
3. Add at least one comment that only you (the human instructor) could make—referencing a class discussion, the student's improvement trajectory, or a connection to their career goals.
4. Return the refined feedback. Track cumulative time savings across the semester.

Sample Prompts

For student self-revision:

> *You are a finance professor evaluating an equity research report. Use the following rubric to evaluate this submission. For EACH rubric category, provide: (1) a score estimate on the rubric scale, (2) one specific strength with a direct quote from the submission, (3) one specific weakness with a direct quote, and (4) one concrete, actionable revision suggestion. Be rigorous but constructive. Do not rewrite the student's work—guide them to improve it themselves.*
>
> *Rubric categories: (A) Investment Thesis Clarity [1–5], (B) Valuation Methodology and Assumptions [1–5], (C) Risk Analysis Completeness [1–5], (D) Use of Evidence and Data [1–5], (E) Writing Quality and Professional Tone [1–5].*
>
> *[Paste rubric details and student submission below]*

For financial plan review:

> *Review this personal financial plan narrative for completeness and quality. Check for coverage of ALL of the following: (1) clearly stated short-term and long-term goals with dollar amounts and timelines, (2) current income, expenses, and cash flow analysis, (3) debt inventory with interest rates and a prioritized payoff strategy, (4) emergency fund adequacy, (5) risk tolerance assessment and insurance needs, (6) tax considerations relevant to the client's situation, (7) retirement projections with stated assumptions. For each area, rate coverage as Strong / Adequate / Missing and provide one specific suggestion for improvement.*

Assessment Approach

Grade the revision process, not just the final product. Include a process dimension (15–20% of total grade) evaluating the quality of the student's reflection on AI feedback: Did they exercise judgment in accepting and rejecting suggestions? Did they identify feedback the AI got wrong? Spot-check AI feedback quality periodically to ensure rubric prompts are generating useful critiques.

Human Credibility Touchpoints: OECD evidence shows that AI and human feedback are not pedagogically interchangeable—students trust human feedback more and respond to it differently. Protect the learning value of this workflow with three touchpoints: (1) the instructor adds at least one comment that only a human who knows the student could make (referencing their improvement arc, career goals, or a class discussion), (2) the student writes one sentence explaining what they would have missed without the AI's critique, and (3) the instructor spot-checks whether the AI-generated feedback contains any finance-domain errors before returning it to students.

Skills Developed

Finance Concepts	Financial writing craft (research reports, memos, plans), ratio interpretation narrative, valuation storytelling, professional communication for multiple audiences
Analytical Skills	Self-assessment, critical evaluation of external feedback, iterative improvement methodology, distinguishing substantive from cosmetic revision
Professional Skills	Professional written communication, audience adaptation (board vs. lending committee vs. client), editing discipline, revision as a professional practice
AI Literacy Skills	Crafting effective feedback prompts, critically evaluating AI suggestions against domain knowledge, understanding what AI feedback misses

Evidence from Research and Applications

Greene (AACSB, 2025) documents structured feedback workflows in finance using the Skills/Replacement/Complement framework. Bowen & Watson (2024, 2nd ed. 2025) provide detailed guidance on the write-first-then-AI-then-reflect pedagogical cycle. Abeysekera (2024) tested ChatGPT on financial accounting assessments, finding GPT-4 scored at the 90th percentile for introductory courses, demonstrating both the tool's analytical capability and the need to carefully design rubric-based prompts. Broader higher education literature consistently identifies AI feedback as a high-value, low-risk entry point for faculty adoption (Lo, 2023).

Risks and Safeguards

- **Depersonalization:** Over-reliance on AI feedback may weaken the mentor–student relationship. Mitigation: always add at least one personal, human-only comment per submission.
- **Domain blind spots:** AI may miss finance-specific issues (e.g., an unreasonable terminal growth rate that sounds plausible). Mitigation: faculty must review all AI-generated feedback before returning it.
- **Gaming:** Students may use AI to generate entire documents from scratch, defeating the learning purpose. Mitigation: require annotated rough drafts, in-class writing samples for comparison, and the reflection memo.
- **Privacy:** Do not paste identifiable student data into non-institutionally-approved tools. Use Claude Projects or Custom GPTs where data stays within the session.

First-Week Implementation Pilot

Pick one assignment you are grading this week. Paste one student's submission and your rubric into ChatGPT or Claude. Use this prompt: 'Evaluate this submission using the attached rubric. For each category, provide a score, one strength with a quote, one weakness with a quote, and one actionable suggestion.' Compare the AI's draft to what you would have written. Edit it, add your personal touch, and return it. Track your time—most faculty report saving 15–25 minutes per submission on first use.

#3: AI Socratic Tutor for Foundational Finance Concepts

Evidence Profile: *Finance-specific evidence + cross-disciplinary analogue*
Best First Pilot: Strong first pilot. Especially effective in introductory courses when paired with a short AI-free check.
Adoption Difficulty: ● Beginner

Teaching Problem

Students arrive with widely varying quantitative backgrounds, particularly in Principles and Corporate Finance. Those who fall behind on foundational mechanics—time value of money, bond pricing, CAPM, WACC—disengage before reaching more complex material. Office hours serve only a fraction of students who need help, and the students who need support most are often the most reluctant to ask for it publicly. In sections of 60 or more, individualized attention during class is impractical.

How AI Addresses the Problem

AI chatbots serve as always-available, infinitely patient tutoring assistants that use Socratic questioning to guide understanding rather than simply providing answers. Unlike a static answer key, the AI can re-explain concepts using different approaches (formula derivation, financial calculator logic, spreadsheet function, real-world analogy), diagnose specific misunderstandings through targeted questions, and generate unlimited practice problems at calibrated difficulty levels. Students can ask follow-up questions without social anxiety or time pressure.

Deployment Details

Courses: Principles of Finance (UG), Corporate Finance (UG/MBA), Investments (UG/MBA), and any course with foundational quantitative prerequisites

Tools: ChatGPT (Free/Plus/Edu), Claude, Microsoft Copilot, Google Gemini, Khanmigo, Playlab.ai (build Socratic tutor bots with course materials and textbook content attached as reference documents—free for educators), custom course bots built with RAG on course materials

Student Workflow

1. Encounter a concept or problem type they find difficult (e.g., present value of an uneven cash flow stream).
2. Open AI tool and use the instructor-provided Socratic prompt template (distributed as a half-page handout or LMS link).
3. Work through the problem interactively: the AI asks diagnostic questions to identify the specific point of confusion before explaining.
4. Once the concept clicks, ask the AI to generate a similar practice problem at the same difficulty level, then one level harder.
5. Complete a verification log: solve the practice problem independently, then compare against textbook formula and AI's solution.
6. Submit an annotated AI dialogue log with brief reflections: 'What I understood before,' 'Where I was confused,' 'What I understand now.'

Sample Prompts

Socratic WACC Tutor:

> *You are a patient, Socratic finance tutor helping an undergraduate student understand WACC. IMPORTANT RULES: (1) NEVER give the formula or the answer immediately. (2) Start by asking the student what they already know about how companies raise money. (3) Use a mortgage analogy to explain cost of debt: 'If you borrow $300K for a house at 6%, but you get a tax deduction, what is your real cost?' (4) For cost of equity, ask: 'If you were investing your own money in this company, what return would you demand, and why?' (5) Build toward the WACC formula one component at a time. (6) After each step, ask 'Does that make sense? Can you put that in your own words?' before moving on. (7) Once they grasp the concept, generate one practice problem and ask them to solve it before you confirm.*

Diagnostic TVM Tutor:

> *You are a finance tutor preparing a student for a Corporate Finance midterm. Before teaching anything, ask the student these 3 diagnostic questions to identify their weak spots: (1) 'If I offer you $1,000 today or $1,000 in one year, which do you prefer and why?' (2) 'Can you explain what a discount rate represents in one sentence?' (3) 'What is the difference between an annuity and a perpetuity?' Based on their answers, identify the most fundamental gap and teach from there. Always show intermediate calculation steps with units (e.g., dollars, periods, percent). Use concrete examples: car loans, mortgage payments, retirement savings.*

Assessment Approach

AI Dialogue Log graded on: depth of engagement evidenced by multi-turn exchanges (25%), quality of student's follow-up questions showing genuine inquiry (25%), accuracy and completeness of the verification

log (25%), and quality of the ‘in your own words’ summary demonstrating internalized understanding (25%). Optional: a brief oral quiz (2–3 questions) to confirm the student can explain the concept without AI assistance.

Human Credibility Touchpoints: The Socratic tutor’s effectiveness depends on students trusting the learning process. OECD recommends three touchpoints to maintain human instructor presence: (1) the instructor reviews a sample of AI dialogue logs and leaves one personal comment per student per unit, (2) the student self-explains the concept without the AI log in front of them (in the oral quiz or a short written summary), and (3) the student completes at least one textbook verification step independent of AI to build the verification habit.

OECD Enhancement — Equity Lens: OECD flags that AI tutoring tools can reduce performance gaps for underserved learners but may provide uneven support across student types. Before deploying at scale, ask: Do students with lower prior GPA or non-native English speakers engage as effectively with the Socratic format? Consider offering an alternative (small-group instructor session) for students who find the AI interaction format a barrier rather than a bridge.

Skills Developed

Finance Concepts	TVM (single sum, annuity, uneven cash flows), bond pricing, CAPM and beta interpretation, WACC components, ratio analysis, option payoff diagrams
Analytical Skills	Step-by-step problem decomposition, self-diagnosis of knowledge gaps, verification against authoritative sources
Professional Skills	Self-directed learning habits, intellectual persistence, verification as professional practice
AI Literacy Skills	Crafting Socratic prompts, recognizing when AI explanations are incorrect or misleading, output verification methodology

Evidence from Research and Applications

HBS deployed custom AI tutor bots (FRC Bot) with 75% student adoption and highly positive feedback on conceptual clarification. UT Austin’s UT Sage project uses RAG-based AI tutors grounded in specific course materials to reduce hallucination. The Cogent Education SIUAIT study (2024, ~1,400 finance students across 3 semesters) found that students perceived AI tools as essential for enhancing learning, with Financial Engineering students showing higher proficiency. Eye-tracker and facial expression analysis showed higher engagement in AI-enhanced sessions. The Walton Family Foundation survey (2024) found majority-positive perceived learning impacts among college students. A Harvard study on AI tutors in physics found students learned more than twice as much in less time compared to traditional methods, suggesting transferability to quantitative finance concepts.

Risks and Safeguards

- **Hallucination:** AI can produce numerically plausible but incorrect calculations, especially multi-step TVM problems. Mitigation: require verification against textbook formulas; build verification into the assignment structure.
- **Over-reliance:** Students may stop attempting problems independently. Mitigation: require ‘in your own words’ explanations; maintain AI-free exams to verify retained understanding.
- **Equity:** Free tiers have usage limits; paid tiers create cost barriers. Mitigation: advocate for campus-wide educational AI licenses; deploy Socratic tutors via Playlab.ai where students access the bot through a shared link with no account or subscription required; provide alternative supports (solution manuals, peer tutoring) for students who cannot or will not use AI.
- **Reduced human interaction:** Over-reliance on AI tutoring may reduce student-to-student and student-to-faculty connection (Rudolph et al., 2024). Mitigation: maintain office hours, study groups, and collaborative in-class activities.

First-Week Implementation Pilot

Create a half-page handout with 2 prompt templates: one Socratic tutor for your course's most-struggled concept (e.g., TVM, bond pricing, WACC) and one practice problem generator. Distribute it next class. Ask students to try one problem with the AI tutor before the next session and bring a screenshot of one exchange. Spend 5 minutes debriefing: 'How many found an error in the AI's response? What worked well? What didn't?' Zero technology setup required—students use their own phones and free AI accounts.

#4: AI-Aware Assessment Design (Process-Based Evaluation)

Evidence Profile: *Finance-specific evidence + business education evidence*
Best First Pilot: Good first pilot. Best when you already have a writing or recommendation assignment you can revise rather than rebuild.
Adoption Difficulty: 🟡 Intermediate

Teaching Problem

Traditional take-home exams and generic test-bank questions are increasingly vulnerable to AI completion. AI detection tools (Turnitin, GPTZero) have documented high false-positive rates and should not be used as primary evidence of academic dishonesty. The deeper problem: assessments designed around 'compute the answer' are no longer valid measures of student learning in an AI-available world. Students can generate correct numerical answers without understanding the underlying financial logic.

How AI Addresses the Problem

Redesign assessments to emphasize process documentation, assumption justification, interpretation, and oral verification. The pedagogical shift: from 'can the student compute this?' to 'can the student explain why this answer matters, what could change it, and whether to trust it?' AI becomes a tool for generating individualized assessment variants at scale and for creating error-detection exercises that test higher-order thinking.

Deployment Details

Courses: All finance courses (UG and MBA)

Tools: ChatGPT (for generating variants), parameterized Excel/LMS templates, in-person oral components

Student Workflow

1. Receive individualized case parameters (different cash flow profiles, tax rates, betas, industry contexts) so each student works with unique numbers.
2. Complete quantitative analysis (with or without AI, per course policy clearly stated in the syllabus).
3. Write a justification memo (1–2 pages): explain each key assumption, interpret results in business terms, and identify the single assumption most likely to change the decision.
4. If AI was used at any stage, submit full prompts and outputs with annotations explaining what you accepted, what you modified, and why.
5. Complete a short oral check-in (3–5 minutes) with 2–3 targeted questions: 'Walk me through your discount rate choice,' 'What happens to your recommendation if revenue growth is 2% instead of 5%?'

Sample Prompts

For instructors generating assessment variants:

> *Generate 5 distinct capital budgeting case variants for an upper-division Corporate Finance exam. Requirements for EACH variant: (1) Different industry (manufacturing, healthcare, technology, retail, energy), (2) Project life between 3–7 years, (3) WACC between 8–12%, (4) Distinct cash flow patterns (one with heavy upfront investment and back-loaded returns, one with steady cash flows, one with a mid-project expansion option), (5) Include one variant where NPV is slightly negative (–$50K to –$200K) to test whether students recommend rejection. For each: provide complete cash flow tables, stated assumptions, and 3 interpretation questions requiring judgment, not just computation.*

For error-detection exercises:

> *Write a 250-word equity research summary for a fictional mid-cap SaaS company called CloudMetrics Inc. Embed exactly 3 analytical errors that a well-prepared Corporate Finance student should catch: (1) a terminal growth rate of 6% in an industry with GDP-level long-term growth, (2) a P/E ratio applied to EBITDA instead of earnings, and (3) no mention of key risk factors (customer concentration, competitive entry). Make the prose confident and professional so the errors are not obvious. Do NOT label the errors.*

Assessment Approach

Rubric: Computation accuracy (20%), Assumption justification quality (25%), Interpretation of results in business context (25%), AI critique and disclosure quality (15%), Oral defense performance (15%). Publish the rubric in advance so students understand that explanation and interpretation carry more weight than getting the number right.

Skills Developed

Finance Concepts	Valuation interpretation beyond the number, assumption sensitivity and what drives decisions, financial judgment under uncertainty
Analytical Skills	Critical evaluation and error detection, assumption stress-testing, distinguishing computational accuracy from analytical quality
Professional Skills	Oral communication and defense of analytical decisions, professional accountability, transparent methodology
AI Literacy Skills	Evaluating AI output quality, identifying subtle AI errors in financial analysis, transparent and ethical AI disclosure

Evidence from Research and Applications

Daigle, Li & Li (Monmouth, 2024) tested ChatGPT-3.5 on principles-of-finance test-bank questions, providing empirical support for assessment shifts toward explanation and higher-order reasoning. Abeysekera (2024) found ChatGPT scored at the 80th–90th percentile on introductory financial accounting assessments, demonstrating that computation-only questions are no longer sufficient. OpenAI (2023) explicitly notes AI detectors are unreliable and recommends designing assessments accordingly. Bowen & Watson (2024, 2025) provide the progression framework: AI-prohibited → AI-transparent → AI-required as faculty and students develop competence.

Risks and Safeguards

- **Time investment:** Oral check-ins are time-intensive. Mitigation: keep to 3–5 minutes with 2–3 pre-planned questions; use TAs for large sections; conduct during office hours.
- **Upfront design cost:** Parameterized problems require initial investment. Mitigation: use AI to generate variants (see prompt above), then refine. The investment pays off across semesters.
- **Perceived subjectivity:** Students may view process-based grading as subjective. Mitigation: publish detailed rubrics in advance; provide exemplars of strong and weak justification memos.
- **Equity concerns:** Oral assessments may disadvantage non-native English speakers or students with speech anxiety. Mitigation: offer written alternative formats; allow students to prepare notes for the oral check-in.

First-Week Implementation Pilot

Take one existing exam problem from your current course. Add two interpretation questions that AI cannot easily answer: (1) 'What does this NPV result mean for the firm's strategic decision—should they proceed, and why?' and (2) 'Identify the single assumption that, if changed, would most likely reverse your recommendation. Explain why.' No technology, no setup, no AI account needed. This alone shifts the assessment toward evaluative judgment.

#5: AI-Powered Financial Data Analysis and Visualization

Evidence Profile: *Business education evidence + emerging / frontier evidence*
Best First Pilot: Later-step pilot. Most useful after you have norms for verification and AI-use disclosure in place.
Adoption Difficulty: 🟡 Intermediate

Teaching Problem

Finance increasingly requires data literacy, but many students lack programming skills. Faculty spend disproportionate class time debugging pandas syntax or Excel formulas rather than teaching financial interpretation. Students who struggle with code disengage from analytical content entirely, missing the conceptual learning the analysis was designed to support.

How AI Addresses the Problem

AI tools allow students to describe analyses in natural language—'calculate annualized returns by sector and create a heatmap'—and receive working code with executed output. This shifts the pedagogical bottleneck from 'can you write the code?' to 'can you interpret the output, identify its limitations, and make a financial recommendation?' Students still must understand what they are asking for and whether the results make sense.

Deployment Details

Courses: Financial Analytics (UG/MBA), Investments (UG/MBA), Fintech (UG/MBA), Portfolio Management

Tools: ChatGPT Code Interpreter/Advanced Data Analysis, Claude with code execution, Google Gemini with Sheets, Python (pandas, matplotlib, plotly), Excel Copilot

Student Workflow

1. Receive a dataset (CSV of stock returns, sector performance, portfolio holdings, or transaction data).
2. Upload to AI Code Interpreter and describe the desired analysis in natural language with specific parameters.

3. AI generates code, executes it, and produces charts, tables, and statistical outputs.
4. Critical review: Are calculations correct? Do axis labels and units make sense? Does the interpretation align with financial theory? Check at least 2 specific calculations manually.
5. Write a 1-page interpretation memo including: key findings, investment or business implications, statistical limitations, and what additional analysis would strengthen the conclusions.
6. Submit: AI-generated code output, manual verification checklist (minimum 2 calculations), and interpretation memo.

Sample Prompts

Sector analysis:

> *I am uploading a CSV of monthly returns for 11 S&P 500 sectors from January 2015 through December 2024. Please: (1) Calculate annualized return and annualized standard deviation for each sector. (2) Create a correlation matrix heatmap with values displayed. (3) Calculate the Sharpe ratio for each sector assuming Rf = 4.5% annualized. (4) Identify the sector with the best risk-adjusted return and the sector with the worst. (5) Create a scatter plot of annualized return vs. standard deviation with each sector labeled. Comment briefly on which sectors appear to offer favorable risk-return tradeoffs.*

Portfolio regression:

> *Using the attached portfolio holdings CSV (ticker, shares, current price), calculate the portfolio's weighted beta using each stock's beta vs. S&P 500 over the trailing 3 years. Then run a regression of the portfolio's historical monthly returns against S&P 500 monthly returns. Plot the regression line (Security Characteristic Line) with my portfolio marked. Display alpha, beta, R-squared, and the p-value on alpha. Explain in 3–4 sentences what the alpha and R-squared tell us about this portfolio's performance and diversification.*

Assessment Approach

Rubric: Correct analytical setup and appropriate methodology choice (20%), Interpretation quality and financial reasoning (30%), Manual verification accuracy (20%), Limitations discussion and intellectual honesty (15%), Communication clarity (15%).

Skills Developed

Finance Concepts	Return calculation, risk metrics (Sharpe ratio, beta, volatility, tracking error), portfolio construction, regression interpretation, correlation analysis
Analytical Skills	Data interpretation, statistical reasoning, visualization literacy, distinguishing statistical significance from economic significance
Professional Skills	Data-driven communication, presenting quantitative findings to non-technical audiences, building evidence-based recommendations
AI Literacy Skills	Natural language-to-code translation, verifying AI-generated statistical output, understanding what AI code does vs. what you asked for

Evidence from Research and Applications

Arizona State's W.P. Carey School launched fintech-specific curricula (2024). Wharton's AI for Business major (2025) requires AI/ML coursework. Sadat Shanto et al. (2024) found LLM assistance reduces time-to-first-solution in programming tasks but cautions about dependency without reflection. The proliferation of STEM-

designated MBA programs (Harvard, Yale, Emory received STEM designations in 2024) reflects the growing centrality of quantitative and AI skills in finance education.

Risks and Safeguards

- **False confidence:** Students may trust AI-generated statistics without understanding underlying assumptions (normality, stationarity, sample size). Mitigation: require explicit statement of statistical assumptions.
- **Economically meaningless results:** AI can produce statistically significant but financially nonsensical outputs. Mitigation: every analysis must include a 'so what?' interpretation tied to a financial decision.
- **Ephemeral sessions:** Code Interpreter sessions expire. Mitigation: require students to document the full workflow (prompts, code, outputs) in a persistent format.
- **Proprietary data:** Do not upload restricted or proprietary firm data. Mitigation: use publicly available datasets only (Kenneth French data library, Yahoo Finance, FRED).

First-Week Implementation Pilot

Download a free CSV of monthly Fama-French factor returns from Kenneth French's data library (mba.tuck.dartmouth.edu). Upload it to ChatGPT Code Interpreter with this prompt: 'Calculate the average monthly return and standard deviation for the market factor, SMB, and HML. Display as a bar chart with error bars.' Show the result in class (takes 30 seconds). Ask students: 'What would you need to check to trust these numbers?' Debrief in 3 minutes. No coding knowledge required.

#6: Content Generation: Exam Variants, Cases, and Course Materials

Evidence Profile: *Cross-disciplinary analogue + business education evidence*
Best First Pilot: Strong first pilot. Fastest entry point for most faculty because it improves prep without changing student workflow immediately.
Adoption Difficulty: ● Beginner

Teaching Problem

Developing fresh case studies, problem-set variants, and discussion prompts is time-consuming. Textbook examples can feel dated within a year. Creating multiple assessment versions for academic integrity multiplies the workload. Faculty also need to update courses rapidly as financial markets, regulations, and economic conditions evolve.

How AI Addresses the Problem

AI drafts multiple variants of problem sets, mini case studies anchored in current events, discussion questions, learning objectives, and even accreditation documentation. Faculty serve as editors and curators—reviewing for accuracy, adding nuance, and calibrating difficulty—reducing content development time from hours to minutes.

Deployment Details

Courses: All finance courses (UG and MBA)

Tools: ChatGPT, Claude, Gemini, Microsoft Copilot, Perplexity (for research-grounded content generation)

Sample Prompts

Problem set variants:

> *Create 5 time-value-of-money problem variants for Principles of Finance. Requirements: (1) Each uses a different realistic business scenario (equipment purchase, commercial real estate loan, lease-vs-buy decision, retirement planning, bond pricing). (2) Vary the number of periods (3–15 years), interest rates (4–12%), and payment structures (lump sum, ordinary annuity, annuity due, uneven cash flows). (3) Include one problem that requires solving for the interest rate and one that requires solving for the number of periods. (4) Provide complete answer keys with intermediate steps showing the financial calculator keystrokes (N, I/Y, PV, PMT, FV). (5) Difficulty should range from straightforward (Problem 1) to challenging (Problem 5).*

Mini case generation:

> *Draft a 600-word mini case study for MBA Corporate Finance based on a real 2024–2025 technology sector acquisition. Change company names and disguise identifying details. Include: (1) acquirer and target financial profiles (revenue, EBITDA, debt levels), (2) stated strategic rationale, (3) deal structure (cash, stock, or mixed), (4) at least one complicating factor (regulatory scrutiny, cultural integration risk, or customer overlap). End with 4 discussion questions that require students to evaluate the deal using NPV, comparable transactions, and strategic fit analysis.*

Accreditation documentation:

> *Generate an AACSB Assurance of Learning alignment matrix for an undergraduate Financial Statement Analysis course. Map 6 student learning objectives to Bloom's Taxonomy levels, specific course assignments, and assessment rubric categories. Include at least 2 objectives at the Analyze or Evaluate level.*

Assessment Approach

Not directly student-assessed (this is primarily an instructor productivity tool). When used as a student exercise in advanced courses—asking students to generate and then critique AI-produced case studies—grade the quality of the critique, not the AI output.

Skills Developed

Finance Concepts	N/A (instructor tool); when used as student exercise: critical evaluation of financial scenarios, identifying unrealistic assumptions
Analytical Skills	Efficiency in assessment design, quality control of AI-generated content
Professional Skills	N/A (instructor tool)
AI Literacy Skills	Effective prompting for structured content generation, output quality verification, understanding AI's tendency to fabricate financial data

Evidence from Research and Applications

AACSB Insights and the AAC&U Institute on AI, Pedagogy, and the Curriculum (2024–2026) report widespread faculty adoption of AI for content generation across business disciplines. Bowen & Watson (2025, 2nd ed.) include expanded sections on using AI for assignment design, simulation creation, and custom bot development. Systematic evidence on the quality or learning outcomes of AI-generated finance content specifically is still emerging.

Risks and Safeguards

- **Fabricated financial data:** AI can invent plausible-sounding but entirely fictional financial figures. Mitigation: ALWAYS verify numbers against primary sources (SEC filings, Bloomberg, company reports) before distributing to students.
- **Insufficient complexity:** AI-generated cases tend to be too 'clean'—lacking the ambiguity and conflicting information that makes real cases pedagogically valuable. Mitigation: treat AI output as a first draft; add messiness, conflicting data points, and red herrings.
- **IP considerations:** Check institutional policies on ownership and distribution of AI-generated course materials.

Curricular Homogenization (OECD Warning): OECD reports that GenAI can improve content quality while reducing the collective diversity of generated material. AI-generated finance cases tend to cluster around large-cap U.S. technology and retail sectors with straightforward deal structures. Build in a diversity check before using any AI-generated case or problem set: rotate across (1) industry sector, (2) firm size (micro-cap to large-cap), (3) geography (U.S., international, emerging markets), (4) stakeholder lens (creditor, equity holder, regulator, employee), and (5) market regime (expansion, recession, rising rates). If your case library lacks diversity on any dimension, prompt explicitly for it.

First-Week Implementation Pilot

Open ChatGPT right now. Paste this prompt: 'Generate 3 variants of a bond pricing problem for Principles of Finance. Each should use a different coupon rate (4%, 6%, 8%), different maturity (5, 10, 20 years), and different YTM (3%, 5%, 7%). Include answer keys.' Compare the output to your existing problem set. Edit the best variant for your next quiz. Total time: 10–15 minutes.

#7: Real-Time Market Context and Case Study Enhancement

Evidence Profile: *Finance-specific evidence + emerging / frontier evidence*
Best First Pilot: Good first pilot. A practical next step once you are comfortable checking sources and recency.
Adoption Difficulty: 🟢/🟡 Beginner to Intermediate

Teaching Problem

Published case studies age quickly in finance. By the time a Harvard Business Publishing case reaches students, market conditions, company financials, and regulatory environments may have changed significantly. Faculty want to connect classroom theory to current events but lack the time to continuously update materials or build current-events exercises from scratch.

How AI Addresses the Problem

AI tools with web-search capabilities provide real-time market context: recent earnings data, analyst consensus estimates, material events since a case was written, and current rate environments. Custom chatbots pre-loaded with case materials create interactive 'data rooms' where students can query financial metrics and explore what-if scenarios. Students learn that financial analysis is never 'finished'—it must be continuously updated.

Deployment Details

Courses: Corporate Finance (UG/MBA), Investments (UG/MBA), Personal Finance (UG)

Tools: ChatGPT (with web browsing), Claude, Perplexity AI, custom GPTs pre-loaded with case data, Microsoft Copilot

Student Workflow

1. Receive the assigned case study along with a ‘Case Update Assignment’ template.
2. Use AI with web search to find the company’s most recent earnings, analyst consensus estimates, credit rating changes, and any material events since the case was written.
3. Write a 1-page ‘Case Update Memo’ that bridges the case’s time period to today: What has changed? Which case assumptions have been validated or invalidated?
4. Identify 2–3 specific ways the new information would change the analysis or recommendation in the original case.
5. Cite at least 2 primary sources (10-K filing, earnings transcript, press release, SEC filing) to verify every claim. AI-generated statements without primary-source verification receive zero credit.

Sample Prompts

Case update research:

```
I am preparing for a class discussion of a case study about [Company X] written in
[Year]. Summarize the 3 most significant developments since the case was written,
including: (1) most recent quarterly earnings vs. analyst expectations, (2) any major
strategic announcements (M&A, restructuring, new product launches), (3) changes in
analyst consensus price target or rating. For each development, explain how it would
affect the analysis in the original case. Provide specific source URLs I can verify.
```

Personal Finance current data:

```
I am building a personal financial plan for a class assignment. Look up current data
for: (1) average 30-year fixed mortgage rate, (2) average high-yield savings account
APY, (3) current federal income tax brackets for married filing jointly, (4) current
federal student loan interest rates for undergraduate and graduate loans, (5) 2025
401(k) contribution limits and catch-up contribution limits. Provide source URLs for
each.
```

Assessment Approach

Case Update Memo graded on: accuracy and relevance of identified updates (30%), quality of analysis connecting updates to original case themes and recommendations (35%), source quality and verification rigor (20%), writing clarity and professional tone (15%).

Skills Developed

Finance Concepts	Fundamental analysis, earnings interpretation, financial statement trend analysis, valuation updating, connecting macro to micro
Analytical Skills	Source verification, time-series comparison, identifying material changes, distinguishing signal from noise
Professional Skills	Research efficiency, professional memo writing, current-events awareness, the analyst’s habit of continuous updating
AI Literacy Skills	Using AI for research acceleration with mandatory verification, understanding data freshness limitations, recognizing hallucinated financial data

Evidence from Research and Applications

FMA and SFA pedagogy conference sessions describe growing faculty use of AI for case study augmentation. MIT Sloan’s ‘AI and Money’ course (Gensler, 2025) examines real-world AI implications in market analysis. Trent (2024) provides a transferable classroom assignment design integrating ChatGPT with current events

and requiring verification. Finance-domain LLM research (Wu et al., 2023; Yang et al., 2023) motivates careful use of AI in market contexts.

Risks and Safeguards

- **Hallucinated data:** AI may fabricate earnings figures, deal terms, or analyst estimates. Mitigation: require primary source citations for every factual claim; zero credit for unverified AI-generated 'facts.'
- **Paywalled data:** AI cannot access Bloomberg, Refinitiv, or S&P Capital IQ. Mitigation: direct students to free sources (SEC EDGAR, Yahoo Finance, company investor relations pages).
- **AI as 'source':** Students may cite 'ChatGPT said' as evidence. Mitigation: explicit course policy that AI is a research assistant, not a citable source.

Curricular Diversity Safeguard (OECD): When using AI to update existing cases or build new discussion materials, the instructor-as-curator role is essential. AI tends to emphasize recent high-profile events and dominant-market narratives. Deliberately rotate the macro context you incorporate: rising vs. falling rate environments, domestic vs. cross-border deals, regulated vs. lightly regulated industries, and firms at different lifecycle stages. Preventing "case library convergence" is an active editorial task, not a passive one.

First-Week Implementation Pilot

Before your next case discussion, post this 5-minute pre-class task: 'Open ChatGPT or Perplexity. Ask: What has happened to [Case Company] since [case year]? Find one development that would change the case analysis. Verify it with one primary source (10-K, earnings transcript, or press release). Bring your finding to class.' Open class by having 2–3 students share their updates. Debrief: 'How many verified what the AI told you? What did you find when you checked?' Zero setup required.

#8: Coding Assistance and 'Vibe Coding' for Financial Modeling

Evidence Profile: *Cross-disciplinary analogue + business education evidence + emerging / frontier evidence*
Best First Pilot: Later-step pilot. Best after students have baseline coding expectations and you have a verification rubric.
Adoption Difficulty: 🟡/🔴 Intermediate to Advanced

Teaching Problem

Finance courses increasingly require programming (Python, R, VBA, SQL) for financial modeling, data analysis, and algorithmic trading. Students without computer science backgrounds stall on syntax errors and debugging—problems that consume class time better spent on financial reasoning. Faculty themselves may have limited programming expertise, making it difficult to support students through technical roadblocks.

How AI Addresses the Problem

The concept of 'vibe coding' (articulated by Prof. Kerry Back at Rice University) represents a paradigm shift: students describe financial logic in natural language, AI generates the code, and students verify, annotate, modify, and interpret the output. The student's role shifts from 'coder' to 'systems architect'—responsible for specifying requirements correctly, auditing the implementation, and interpreting results. This dramatically lowers the barrier to entry for quantitative finance work while raising the ceiling for analytical sophistication.

Deployment Details

Courses: Financial Modeling (UG/MBA), Financial Analytics (UG/MBA), Fintech/Algorithmic Trading (MBA), Derivatives, Capstone projects

Tools: GitHub Copilot, ChatGPT Code Interpreter, Claude (Artifacts), Cursor, Google Colab with Gemini, Google Antigravity, Replit AI

Student Workflow

1. Receive a financial modeling task with clearly specified requirements (e.g., build a factor model, implement Black-Scholes, backtest a trading strategy).
2. Write a plain-English specification describing inputs, calculations, outputs, and edge cases—before touching AI. This forces articulation of the financial logic.
3. Submit the specification to AI and receive generated Python/R/VBA code.
4. Annotate the code line-by-line: for each meaningful section, explain the financial logic (not just what the code does, but WHY this calculation matters and what it represents).
5. Modify the code to handle at least one extension (add transaction costs, incorporate short-selling constraints, test with different time periods, add a risk metric).
6. Run the final code, interpret outputs in a written memo, and present findings with a 3-minute oral walkthrough defending the financial logic.

Sample Prompts

Factor model construction:

```
Build a Python script that: (1) Pulls 3 years of daily adjusted close prices for SPY and a user-defined list of 5 stock tickers from Yahoo Finance using yfinance. (2) Calculates daily log returns for each. (3) Runs an OLS regression of each stock's returns against SPY returns to estimate Beta. (4) Displays a summary table showing each stock's Beta, Alpha (annualized), R-squared, and p-value on Alpha. (5) Plots the regression line (Security Characteristic Line) for each stock with the data points. (6) Adds a brief interpretation comment below each plot explaining what the Beta and R-squared imply about the stock's risk profile.
```

Options pricing and Greeks:

```
Build a Black-Scholes option pricing calculator in Python. Inputs: current stock price, strike price, risk-free rate, annualized volatility, and time to expiration in years. Outputs: (1) European call and put prices, (2) all five Greeks (Delta, Gamma, Theta, Vega, Rho) for both call and put, (3) a 3D surface plot showing how call option price varies with stock price (x-axis: 50–150% of current price) and volatility (y-axis: 10–60%). (4) Add a payoff diagram at expiration overlaid with the current option value curve. Label all axes clearly with units.
```

Assessment Approach

Rubric: Code annotation quality—does the student demonstrate understanding of the financial logic, not just describe the syntax? (25%), Financial logic accuracy—are the calculations conceptually correct? (25%), Modification quality—does the extension demonstrate additional financial reasoning? (20%), Output interpretation and memo quality (20%), Oral walkthrough (10%).

Skills Developed

Finance Concepts	DCF modeling, Monte Carlo simulation, option pricing and Greeks, factor models, portfolio optimization, backtesting methodology
Analytical Skills	Code auditing for financial logic errors, model validation against analytical solutions, sensitivity analysis
Professional Skills	Technical specification writing, working with AI as a professional co-pilot, communicating quantitative results
AI Literacy Skills	Natural language-to-code translation, verifying AI-generated code against known solutions, understanding coding limitations and edge cases

Evidence from Research and Applications

Prof. Kerry Back (Rice University, 2025) has been a leading advocate for 'vibe coding' in finance education, publishing resources and course materials. Becker et al. (2023) and Denny et al. (2024) document AI coding assistance extensively in CS education. Wharton's AI for Business major (2025) includes required coursework in applied machine learning. Sadat Shanto et al. (2024) caution that LLM coding assistance improves productivity but requires metacognitive prompts to avoid shallow learning—validating the annotation and oral defense requirements.

Risks and Safeguards

- **Financial logic errors in AI code:** AI may calculate returns using price levels instead of log returns, mishandle ex-dividend dates, apply wrong sign conventions, or use incorrect discounting periods. Mitigation: require line-by-line annotations explaining the financial meaning; include a 'break it' exercise where students deliberately introduce and then fix errors.
- **Copy-paste without understanding:** Mitigation: oral walkthroughs (even 3 minutes) are highly effective at distinguishing students who understand from those who do not. In-class mini-checks where students modify one parameter and predict the output direction before running work well.
- **API key exposure:** Students may paste API keys into chatbot sessions. Mitigation: explicit policy and training; use only free data sources that require no authentication (yfinance, FRED).
- **Academic integrity:** Clearly define acceptable AI use for each assignment in the syllabus. Consider requiring the AI interaction transcript as part of the submission.

First-Week Implementation Pilot

In class, open ChatGPT or Claude and type: 'Write Python code to calculate the present value of $10,000 received in 5 years at a 7% annual discount rate. Show the formula in a comment.' Display the result (15 seconds). Then ask students: 'What would you need to change if there were also annual $500 coupon payments?' Have a student describe the change in plain English. Paste their description into the AI. Display the updated code. Debrief: 'The AI wrote it—how would you verify it's correct?' Total time: 8 minutes.

#9: AI Literacy and Critical Evaluation of AI Outputs

Evidence Profile: *Cross-disciplinary analogue + emerging / frontier evidence*
Best First Pilot: Strong first pilot. Simple to embed in any course and foundational for everything else in the guide.
Adoption Difficulty: 🟢 Beginner

Teaching Problem

Students will enter a workforce where AI is embedded in virtually every financial tool—robo-advisors, algorithmic trading platforms, credit scoring models, fraud detection systems, compliance monitoring. Without understanding AI's capabilities and limitations—its tendency to hallucinate, reflect training data biases, and produce confident-sounding but incorrect analysis—graduates risk making costly errors or blindly trusting AI-generated recommendations that a skilled professional would question.

How AI Addresses the Problem

Embed AI literacy as a cross-cutting theme woven throughout the finance curriculum, not siloed into a single lecture. Design exercises where students explicitly test, evaluate, and critique AI outputs on finance topics they already understand. The 'AI Audit' becomes a repeatable assignment format that reinforces both finance content knowledge and AI evaluation skills simultaneously.

Deployment Details

Courses: All finance courses (UG and MBA); particularly central to Fintech courses

Tools: Multiple AI platforms for cross-platform comparison (ChatGPT, Claude, Gemini, Copilot), custom evaluation rubrics, instructor-designed audit templates

Student Workflow

1. AI Audit exercise: Ask AI a finance question you already know the answer to from class or the textbook. Document the question and AI's full response.
2. Identify and categorize every error or imprecision: factual (wrong numbers), computational (math mistakes), conceptual (misunderstanding theory), reasoning (flawed logic even with correct facts), or omission (important considerations left out).
3. Cross-platform comparison: Ask the identical question to ChatGPT, Claude, and Gemini. Document differences in accuracy, depth, and confidence level.
4. Write a 1–2 page 'AI Reliability Report' documenting: error types and patterns observed, which platform performed best and worst, specific recommendations for when and how to use AI as a finance professional, and what verification steps you would require before acting on AI output.
5. For Fintech courses: study how AI is deployed in financial services (robo-advising, credit scoring, fraud detection, algorithmic trading) and evaluate ethical implications including bias, transparency, and fiduciary duty.
6. Advanced exercise (optional): Build your own finance-focused AI tool on Playlab.ai. Design the system prompt, set guardrails, attach reference materials, and test with classmates. Evaluate: Where does your bot succeed? Where does it fail? What guardrails prevented bad outputs, and which ones did you need to add? This exercise develops AI literacy from the builder's perspective, not just the user's perspective.

Sample Prompts

Theory audit:

> *Explain the Modigliani-Miller theorem. Provide the mathematical proof for Proposition I under the no-tax assumption. Then explain Proposition II and show how the cost of equity changes with leverage. [Student verifies the derivation step-by-step against the textbook, checking for sign errors, missing assumptions, or circular reasoning.]*

Quantitative audit:

> *Using Apple's most recent 10-K filing, calculate the current ratio, debt-to-equity ratio, and return on equity. Show all calculations with the specific line items used. [Student downloads the actual 10-K from SEC EDGAR and verifies every number and calculation independently.]*

Cross-platform comparison:

> *Ask ChatGPT, Claude, and Gemini: 'What was the S&P 500 total return in 2024? Break down the return into price appreciation and dividend yield.' Compare the three responses for accuracy, specificity, source citation, and confidence level. Which platform hedges most appropriately?*

Assessment Approach

AI Reliability Report graded on: thoroughness of error identification (30%), accuracy and specificity of error categorization (25%), quality of actionable recommendations for professional AI use (25%), depth of cross-platform comparison insights (20%).

Skills Developed

Finance Concepts	Reinforces whatever concept is being audited (TVM, M&M, valuation ratios, derivatives, portfolio theory)—students learn content by finding where AI gets it wrong
Analytical Skills	Critical evaluation, error taxonomy and pattern recognition, comparative analysis across sources, developing a professional verification methodology
Professional Skills	Professional skepticism as a habit, verification protocols, responsible and transparent technology use in professional settings
AI Literacy Skills	Understanding AI limitations (hallucination, training data cutoffs, confident incorrectness), prompt engineering for diagnostic testing, cross-platform evaluation methodology, AI tool design and guardrail engineering (via Playlab.ai builder exercises)

Evidence from Research and Applications

The AAC&U Institute on AI, Pedagogy, and the Curriculum (2024–2026) emphasizes AI literacy as a critical cross-cutting learning outcome. Wharton requires 'Big Data, Big Responsibilities: Toward Accountable Artificial Intelligence' (LGST 6420) in its AI for Business major. Chiu (2024) argues for the urgency of reforming educational institutions to prepare professionals for AI-integrated workplaces. BloombergGPT (Wu et al., 2023) and FinGPT (Yang et al., 2023) provide domain-specific context on AI capabilities and limitations in finance. Schlosky et al. (2024) documented specific strengths and failure modes of ChatGPT in a financial-trouble advising context.

Risks and Safeguards

- **Cynicism:** Students may conclude AI is 'useless' rather than developing nuanced understanding of when it helps and when it fails. Mitigation: balance critique exercises with productive-use exercises; show what AI does well before testing its limits.
- **Rapid obsolescence:** AI capabilities change quarterly; specific examples of errors may not reproduce. Mitigation: teach the evaluation methodology (how to audit), not just current findings. The skill of auditing AI transfers regardless of which model is current.
- **Faculty readiness:** Faculty need their own AI literacy before teaching it. Mitigation: institutional training support; start with the 2-Minute AI Audit exercise (see pilot below)—it requires zero preparation and builds faculty experience alongside students.

First-Week Implementation Pilot

The 2-Minute AI Audit (works with zero preparation): In class, tell every student to open any AI tool on their phone. Ask them a finance question they know the answer to cold—something from last week's lecture. They have exactly 2 minutes to find at least one error or imprecision in the AI's response. Quick show of hands: How many found an error? What kinds of errors? This exercise takes 5 minutes total, requires zero setup, immediately demonstrates that AI is accessible (for the novices) and fallible (for the skeptics), and establishes AI as a tool to be audited rather than trusted blindly. Do it on day one of any course.

#10: Simulation, Monte Carlo Modeling, and Scenario Analysis

Evidence Profile: *Finance-specific evidence + cross-disciplinary analogue + emerging / frontier evidence*
Best First Pilot: Later-step pilot. High-value but best reserved for courses with the time and infrastructure to support it.
Adoption Difficulty: 🔴 Advanced

Teaching Problem

Finance is fundamentally about decision-making under uncertainty, but building realistic simulations—Monte Carlo portfolio projections, multi-scenario DCF models, option pricing with stochastic volatility, stress tests—traditionally requires significant programming expertise that most finance faculty and students lack. Static textbook examples with single-point estimates fail to convey the probabilistic, multi-variable nature of real financial decisions.

How AI Addresses the Problem

AI dramatically lowers the barrier to creating sophisticated financial simulations. Faculty or students describe a scenario in natural language—specifying assumptions, distributions, correlations, and output metrics—and AI generates complete, executable simulation code. The pedagogical value is not in the code generation but in what follows: students must set and defend assumptions, modify parameters to test sensitivity, interpret distributional outputs, and make recommendations under genuine uncertainty.

Deployment Details

Courses: Investments/Portfolio Management (UG/MBA), Derivatives (UG/MBA), Risk Management (MBA), Financial Modeling (UG/MBA), Capstone courses

Tools: ChatGPT Code Interpreter/Advanced Data Analysis, Claude with code execution, GitHub Copilot, Python (numpy, scipy, pandas, matplotlib), Google Colab

Student Workflow

1. Receive a modeling scenario with specified requirements (e.g., simulate 10,000 portfolio return paths under 3 asset allocations over a 30-year retirement horizon).
2. Before using AI, write a 1-page specification document: what assumptions will you make about return distributions, correlations, rebalancing frequency, withdrawal rates, and inflation? Justify each choice with reference to historical data or financial theory.
3. Use AI to generate simulation code based on your specification. Specify all parameters explicitly in the prompt.
4. Run the simulation and produce required visualizations (terminal wealth distributions, probability of ruin, efficient frontier, convergence plots).
5. Modify at least 2 assumptions and re-run: What happens with fat-tailed return distributions? With regime switching? With higher withdrawal rates? With correlated shocks?
6. Write a 2-page investment recommendation memo supported by simulation evidence, including a clear statement of the recommendation, the key assumptions driving it, sensitivity analysis results, and limitations of the model.
7. Present findings to the class (5–7 minutes), defending assumptions under Q&A.

Sample Prompts

Retirement Monte Carlo:

> *Build a Monte Carlo retirement simulation in Python with the following specifications: (1) Simulate 10,000 return paths over 30 years for 3 portfolios: 60/40 stocks/bonds, 80/20 stocks/bonds, and a target-date glide path that starts at 90/10 and linearly shifts to 30/70 over 30 years. (2) Use historical Ibbotson-style parameters: stocks with 10.5% mean annual return and 20% standard deviation; bonds with 5.5% mean and 6% standard deviation; correlation of 0.05. (3) Assume $1M starting balance with $50,000 real annual withdrawals (adjusted for 2.5% inflation). (4) Plot: terminal wealth distributions for all 3 portfolios on one chart, probability of ruin (account hitting $0) over time for each portfolio, and the median and 5th/95th percentile wealth paths. (5) Print a summary table showing median terminal wealth, probability of ruin, and the 5th percentile terminal wealth for each portfolio. (6) Set random seed to 42 for reproducibility.*

Option pricing convergence:

> *Implement both the Black-Scholes analytical solution and a Monte Carlo option pricing model in Python. Parameters: S0 = $100, K = $105, r = 5%, sigma = 25%, T = 1 year. For Monte Carlo, simulate with N = [100, 500, 1000, 5000, 10000, 50000] paths. Plot: (1) Monte Carlo price vs. number of paths with the Black-Scholes analytical price as a horizontal reference line, (2) 95% confidence interval around the Monte Carlo estimate at each N, (3) computation time vs. number of paths. Explain in comments why the Monte Carlo price converges to the analytical solution and what drives the width of the confidence interval.*

Assessment Approach

Rubric: Assumption quality and justification (25%), Code modification sophistication—does the extension demonstrate genuine financial reasoning? (20%), Output interpretation and memo quality (25%), Presentation and defense of assumptions under Q&A (20%), Verification against analytical solutions where available (10%).

Skills Developed

Finance Concepts	Portfolio theory, retirement planning under uncertainty, option pricing theory, risk management, sensitivity analysis, the difference between point estimates and distributional thinking
Analytical Skills	Stochastic modeling, distribution analysis and interpretation, simulation design, model validation, understanding convergence and sampling error
Professional Skills	Data-driven decision-making under uncertainty, presenting quantitative results to non-technical stakeholders, recommendation writing with appropriate caveats
AI Literacy Skills	Complex multi-parameter prompt engineering, code verification against known analytical solutions, understanding simulation limitations and when results are unreliable

Evidence from Research and Applications

MIT Sloan's 'AI and Money' course (Gensler, 2025) and 'AI and Machine Learning Applications in Finance' (Chen, Kogan, Thesmar, 2025) integrate AI-driven modeling. Wharton provides ChatGPT Enterprise to all MBA students (2024) with integration into finance coursework. FinGPT (Yang et al., 2023) and PIXIU (Xie et al., 2023) provide open-source frameworks and evaluation benchmarks for finance-oriented AI tasks. Classroom outcome evidence for AI-assisted simulation specifically is nascent—this remains the frontier application with the highest potential but least formal evidence.

Risks and Safeguards

- **Subtle code errors:** AI may use wrong return distributions, incorrect correlation structures, or inappropriate rebalancing logic. Mitigation: require verification against analytical solutions where available (e.g., Black-Scholes has a closed-form solution); require students to check limiting cases (what happens with zero volatility? With zero correlation?).
- **High setup cost:** Full simulation projects require Python environments and data access. Mitigation: start with ChatGPT Code Interpreter (no local setup needed); use pre-built Colab notebooks as templates that students extend.
- **Reproducibility:** Different results on re-run confuse students unfamiliar with stochastic methods. Mitigation: always set random seeds; save and submit complete code.
- **Pedagogical tension:** If AI builds the entire model, has the student learned anything? Mitigation: the specification document (Step 2) ensures students think through the financial logic before AI generates code; required modifications (Step 5) ensure engagement with the model mechanics; the oral defense (Step 7) verifies understanding.

First-Week Implementation Pilot

Zero-infrastructure demonstration: Open ChatGPT Code Interpreter in class. Type: 'Simulate 1,000 coin flips. Plot the running average after each flip. Add a horizontal line at 0.5.' Result appears in 15 seconds. Show the convergence. Then say: 'Now imagine this, but instead of coin flips, these are annual stock returns. And instead of a running average, we are tracking a retirement account balance. That is a Monte Carlo simulation.' This 60-second demo plants the conceptual seed for simulation thinking with zero technical barrier. For the next class, provide a pre-built Colab notebook link where students can modify 3 parameters (allocation, withdrawal rate, time horizon) and re-run.

Why This Order? Ranking Methodology

The ranking reflects three weighted criteria applied systematically to the evidence base across all repository sources:

- **Instructional Impact (40%):** How significantly does the application address a core teaching pain point in finance? Applications developing multiple skill categories (finance content, analytical reasoning, professional skills, AI literacy) and targeting high-value pedagogical challenges rank higher.
- **Feasibility for Typical Faculty (35%):** How quickly can a faculty member with moderate technology comfort adopt this? Applications requiring only a smartphone and a free AI account rank higher than those requiring API access, custom infrastructure, or extensive technical expertise.
- **Strength of Evidence (25%):** Applications referenced across multiple repository sources with named institutional adopters and published studies rank higher. This criterion is weighted lowest because rigorous controlled experiments in finance AI pedagogy remain limited across the board.

Role-play simulations (#1) ranked highest because they uniquely combine finance content mastery, professional skill development, and AI literacy in a single exercise—and because the evidence from HBS, Clemson, and social work education research is among the strongest. Formative feedback (#2) and Socratic tutoring (#3) address the two most universal faculty pain points (grading burden and student preparation gaps) with the lowest adoption barriers. Advanced simulation (#10) ranks last not because it lacks value—it is arguably the most transformative application for deep financial reasoning—but because its feasibility requirements limit immediate broad adoption.

Ranking Snapshot

These are my estimates along the 3 measures. They are heuristic guide scores on a 1–5 scale, not empirical effect-size estimates.

Application	Impact	Feasibility	Evidence	Weighted total
#1 Role-Play	5.0	4.5	4.5	4.68
#2 Formative Feedback	4.5	5.0	4.0	4.55
#3 Socratic Tutor	4.5	4.5	4.0	4.35
#4 AI-Aware Assessment	4.5	4.0	4.0	4.18
#5 Data Analysis & Viz	4.0	4.0	4.0	4.00
#6 Content Generation	3.5	4.5	3.5	3.88
#7 Market Context	3.5	4.0	3.5	3.70
#8 Coding Assistance	3.5	3.5	4.0	3.65
#9 AI Literacy	3.0	4.0	3.0	3.45
#10 Simulation	3.5	2.5	3.5	3.13

Weighted total = 0.40 × Impact + 0.35 × Feasibility + 0.25 × Evidence.

Workshop Takeaways: From Ranking to Action

Three pilots for immediate post-conference implementation:

• **Lightest-lift pilot:** #6 Content Generation for a faculty-facing workflow improvement this week.
• **Best introductory-course pilot:** #3 Socratic Tutor paired with a short AI-free exit ticket or board work.
• **Best writing-intensive pilot:** #2 Formative Feedback with transcript submission and a revision memo.

Closing Synthesis

Used well, AI is most valuable in finance teaching when it increases practice opportunities, shortens feedback cycles, keeps verification visible, and shifts assessment toward judgment rather than answer production. The strongest classroom designs in this guide treat AI as a scaffold for explanation, critique, and rehearsal—not as a replacement for disciplinary thinking.

References

Abeysekera, I. (2024). ChatGPT and Academia on Accounting Assessments. Journal of Open Innovation, 10(1), 100213.

Becker, B.A. et al. (2023). Programming Is Hard — Or At Least It Used to Be. Proceedings of the 54th ACM Technical Symposium on CS Education.

Bowen, J.A. & Watson, C.E. (2024, 2nd ed. 2025). Teaching with AI: A Practical Guide to a New Era of Human Learning. Johns Hopkins University Press.

Chiu, T.K.F. (2024). The Impact of Generative AI on Higher Education Learning and Teaching. Education and Information Technologies.

Daigle, J., Li, Y., & Li, B. (2024). Is ChatGPT a Threat to Financial Education? Monmouth University.

Denny, P. et al. (2024). Computing Education in the Era of Generative AI. Communications of the ACM, 67(2), 56–63.

Greene, D. (2025). Linking Class Material to Jobs With Generative AI. AACSB Insights.

Lopez-Lira, A. & Tang, Y. (2023). Can ChatGPT Forecast Stock Price Movements? arXiv:2304.07619.

OECD. (2026). *OECD digital education outlook 2026: Exploring effective uses of generative AI in education*. https://doi.org/10.1787/062a7394-en

OpenAI. (2023). Teaching with AI. https://openai.com/index/teaching-with-ai/

Rudolph, J. et al. (2024). When AI Substitutes Humans in Higher Education. Studies in Higher Education, 49(5), 883–897.

Sadat Shanto, A. et al. (2024). Generative AI for Programming Education. ACM Digital Library.

Schlosky, M.T.T. et al. (2024). ChatGPT, Help! I Am in Financial Trouble. Journal of Risk and Financial Management, 17(6), 241.

Tamoyan, H. et al. (2025). LLM Roleplay: Simulating Human-Chatbot Interaction. ACL Workshop.

Wu, S. et al. (2023). BloombergGPT: A Large Language Model for Finance. arXiv:2303.17564.

Xie, Q. et al. (2023). PIXIU: A Large Language Model, Instruction Data and Evaluation Benchmark for Finance. arXiv:2306.05443.

Yang, H. et al. (2023). FinGPT: Open-Source Financial Large Language Models. arXiv:2306.06031.

Appendix A: Discipline-Specific AI Use Policy Matrix

OECD recommends discipline-specific AI policies that specify permitted, restricted, and prohibited uses at the assignment level rather than as a blanket course-level statement. The matrix below provides a starting framework for finance courses. Faculty should adapt it to their course context and publish it in the syllabus so students understand expectations before any assignment begins. The progression from AI-Prohibited to AI-Transparent to AI-Required mirrors the framework in #4 of this guide and aligns with OECD's recommendation that AI governance evolve alongside student and faculty competence.

Assignment Type	AI-Prohibited	AI-Transparent	AI-Required
In-Class Exam / Quiz	All AI use prohibited. Proctored conditions. Foundational knowledge verification.	N/A for in-class exams.	N/A for in-class exams.
Written Case Analysis / Financial Plan	Draft must be written independently without AI. No AI grammar or idea assistance.	AI permitted for revision feedback only. Full transcript required. Reflection memo required. Human instructor layer required.	AI co-authoring permitted. Student grades the critique and process, not the AI output. Oral defense required.
Financial Modeling / Coding Assignment (#8)	AI-generated final code submission without annotation. Copying without understanding the financial logic.	AI may generate code. Student must annotate every section with financial logic explanation. Transcript submitted. Oral walkthrough required.	Capstone / advanced courses: AI as primary co-pilot. Student architects requirements, audits logic, interprets outputs, and defends decisions.
Monte Carlo / Simulation (#10)	Submitting AI-generated code or outputs without writing a specification document, annotating assumptions, or running sensitivity analysis.	AI generates code from student specification. Student must run sensitivity analysis, interpret distribution outputs, and submit investment recommendation memo. Full transcript required.	Risk Management / advanced courses: AI as simulation partner. Student designs the model logic, selects distributions, defends assumptions in Q&A, and connects results to a real investment decision.
Role-Play / Oral Simulation (#1)	Submitting someone else's transcript or completing the interaction without genuine engagement. No transcript submitted.	AI persona used. Transcript + reflection memo required. Flags moments where the bot broke character. Oral defense adds human verification layer.	Student designs the AI persona, writes the system prompt (or builds the bot on Playlab.ai with Background, Workflow, and Guardrails fields), runs the simulation, critiques the bot's financial accuracy, and presents findings to the class.

Note: "Prohibited" refers to the specific use described, not a blanket prohibition on AI in the course. Faculty are encouraged to share this matrix with students on the first day and revisit it at each major assignment. Policies should evolve as both student AI competence and course-level evidence accumulate.

Appendix B: OECD-Aligned Evaluation Framework

OECD's Digital Education Outlook 2026 recommends that for each AI application, institutions specify: (1) the intended learning outcome, (2) the most likely unintended consequence, and (3) how both will be measured across five dimensions—performance, retention, engagement, equity, and instructor workload. The framework below applies this logic to each of the 10 finance applications. Faculty are encouraged to select one or two applications to evaluate formally in the first semester of adoption and share findings with colleagues.

Application	Intended Outcome	Likely Unintended Consequence	How to Measure	Equity Check
#1 Role-Play	Improved professional communication and analytical defense under pressure	Surface-level interaction without genuine engagement; transcript padding	Pre/post oral defense scores; rubric score distribution; reflection quality ratings	Do students with speech anxiety or non-native English engage equitably? Offer written alternative.
#2 AI Feedback	Faster feedback cycles; stronger revision quality; feedback literacy	Over-reliance on AI critique; depersonalization of instructor relationship	Compare pre/post rubric scores on revision quality; measure acceptance/ rejection rationale quality	Do ESL students or first-gen students trust AI feedback differently? Survey and compare.
#3 Socratic Tutor	Reduced knowledge gaps on foundational concepts; increased office hours independence	Students stop attempting problems independently before engaging AI	AI-free quiz scores before and after adoption; verification log completion rates	Do lower-GPA students benefit as much as higher-GPA students? Track by subgroup.
#4 Process Assessment	Assessment validity in AI era; higher-order reasoning on display	Student perception of subjectivity; oral check-in burden on faculty	Compare oral check-in performance with written submission scores; track grade distribution equity	Do oral assessments disadvantage non-native English speakers? Offer optional written equivalent.
#5 Data Analysis	Data literacy without coding bottleneck; interpretation over syntax	False confidence in AI-verified statistics; shallow interpretation memos	Manual verification checklist accuracy; 'so what?' memo quality ratings; exam performance on interpretation questions	Do students without prior coding exposure benefit more? Track by major/background.
#6 Content Generation	Faculty time savings; fresher, more varied assessment materials	Curricular homogenization; AI fabricated financial data enters course materials	Faculty time-savings log; student assessment score variance across variants; diversity audit of generated cases	Are generated cases representative across firm size, geography, and industry? Run diversity check each semester.
#7 Case Updates	Current-events integration; verification habits; analyst mindset	Hallucinated market data; AI as citable source; paywalled data access gaps	Source citation quality; verification accuracy; student performance on identifying material vs. noise	Do students without Bloomberg/Refinitiv access face barriers? Ensure free-source alternatives are sufficient.
#8 Vibe Coding	Quantitative finance access for non-CS students; code-as-communication skill	Annotation without understanding; AI-generated financial logic errors go undetected	Annotation quality scores; oral walkthrough performance; 'break it' exercise accuracy	Do students without prior programming background close the gap? Track by CS background.
#9 AI Literacy	Professional skepticism; verification habits; AI evaluation methodology	Cynicism about AI; failure to transfer audit methodology to new tools	Error identification accuracy; cross-platform comparison depth; professional recommendation quality	Do students from different majors/backgrounds show different AI literacy gaps? Survey by subgroup.
#10 Simulation	Distributional thinking; decision-making under uncertainty; model-building intuition	Code generation without engagement; assumption selection without theoretical grounding	Specification document quality; sensitivity analysis depth; investment memo recommendations; oral Q&A scores	Do students without strong quantitative backgrounds close the gap with AI assistance? Track pre/post.

Measurement guidance:

- *Performance = rubric scores and exam results.*
- *Retention = AI-free assessment scores 4–8 weeks after adoption.*
- *Engagement = dialogue log depth, reflection quality, participation rates.*
- *Equity = score distributions by subgroup (GPA, language background, major).*
- *Workload = faculty time-per-student on grading and feedback tasks.*

Start with one application and one dimension. Build the evidence base incrementally.