AI Agents in Finance: In-Class Exercises

Standalone Lecture — 3 Short Exercises

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

How to use these exercises

- Three short in-class exercises, ˜8–10 minutes each.
- Attempt on paper first, then discuss with a partner.
- Solutions are hidden by default; instructor build reveals them via `\solutionstrue`.

Companion to `lecture_ai_agents_finance`.

Exercise 1: Agent-Automatable Tasks

Task. List three tasks in financial services where an AI agent can realistically operate in 2026. For each, name the human role it substitutes or augments, and state whether it operates *autonomously* or *with a human-in-the-loop*.
Difficulty: Introductory — tests recall and current-state awareness.

Exercise 2: Error Scale in an Agent Workflow

Scenario. An AI agent executes 5,000 trades per day with a 0.05% critical-error rate (e.g. wrong ticker, wrong side, order size off by $10\times$). Average trade notional: \$100,000. Expected damage per critical error: 0.1% of notional. **Task.**

Compute:

a. Expected critical errors per day.

b. Expected daily loss in dollars.

c. Expected annual loss (250 trading days).

Difficulty: Introductory — arithmetic, but forces you to internalise "rare errors at scale."

Exercise 3: Who Is Accountable?

Scenario. An AI agent hosted by your firm executes an unauthorised trade that loses $2M for a retail client.

Task. In 3–4 sentences, identify the most likely accountable party among (a) the client, (b) your firm (the deploying broker), (c) the LLM / model vendor, (d) the developer who wrote the agent, (e) no one. Justify using a 2026 regulatory principle. *Difficulty: Intermediate — conceptual + current regulation.*