Prof. Dr. Joerg Osterrieder

# Content Reference: AI Agents in Finance

Digital Finance – BSc Course – Standalone Lecture

## Learning Objectives

After engaging with this reference sheet and the companion lecture, you will be able to:

- **Distinguish** an AI agent from a classical LLM chatbot (tool use, memory, goal-seeking) [Understand]
- **Identify** four live 2026 deployment patterns in finance (research, trading, back-office, risk) [Apply]
- **Analyse** the governance gaps that arise when agents trade capital or bind contracts [Analyze]
- **Evaluate** whether a given finance task is suitable for agentic automation [Evaluate]

## 1. Key Definitions

- **AI agent** — an LLM given a goal, a memory, and a set of tools; it plans, acts, observes, and replans until the goal is met or a budget is exhausted.
- **Tool use** — the agent calls named functions (web search, database query, trade order, email send) through a structured interface.
- **Memory** — short-term (context window) and long-term (vector store, scratchpad) state the agent uses to avoid re-solving the same sub-problems.
- **Orchestrator** — a planner agent that decomposes a goal and routes sub-goals to specialist agents.
- **Agentic workflow** — a multi-step, multi-turn process with non-trivial branching that a single LLM call cannot handle.
- **Human-in-the-loop (HITL)** — human approval required before irreversible actions (send payment, execute trade, sign contract).
- **Autonomy budget** — the dollar, time, or risk cap under which an agent may act without further approval.
- **Grounding** — anchoring the agent's reasoning in verifiable sources (real data, authenticated APIs) rather than model memory alone.

## 2. Core Concept: The Four Deployment Patterns (2026)

**Where agents already earn their keep in finance**

1. **Research agents.** Analyst-grade research: earnings-call summaries, peer comps, draft investment memos. Ubiquitous at bulge-bracket banks by Q1 2026 (BofA, Goldman Sachs, Morgan Stanley have each announced internal deployments).
2. **Execution agents (low-stakes).** Trade routing within tight price bands, portfolio rebalancing under approved rules, liquidity monitoring; always with HITL for outlier trades.
3. **Back-office agents.** KYC file assembly, AML case triage, reconciliation break investigation; 50–70% of cycle time removed in published case studies.
4. **Risk-ops agents.** Alert triage (which of 10,000 flags a day actually matter), scenario generation for stress tests, model-validation assistants.

## 3. Key Figures & Data

- **Anthropic Claude Code** and **Cognition Devin** represent the first commercially available software-engineering agents (public 2024–2025); both are now used in financial firms' IT departments.

- **Morgan Stanley** reports its GenAI assistant is used by 98% of its ∼16,000 advisors for research summarisation (company earnings call, 2024).
- **JPMorgan IndexGPT** launched as a GenAI-backed research tool in 2024; internal reports cite 20–40% productivity gain on research tasks.
- **Brynjolfsson et al.** (NBER 2023) found a 14% productivity gain at a Fortune 500 contact centre from GenAI assistants, concentrated in the bottom quintile of performers.
- **SEC Chair Gensler** (speech, Oct 2024) warned that AI-driven hyper-correlation between portfolios is the next systemic-risk frontier.
- **EU AI Act** (force August 2024; phased enforcement through 2027) classifies credit scoring, insurance underwriting, and certain fraud detection as *high-risk* AI, requiring conformity assessment.

## 4. Worked Example

**Cost-benefit of a back-office KYC agent**

**Scenario.** A mid-size bank processes 400 new corporate-client KYC files per month. A specialist analyst takes 4 hours per file; fully loaded cost \$80/hour. An agent assembles the file, flags gaps, and produces a draft; the analyst reviews in 1 hour.

| | |
|---|---|
| Current cost | $400 \times 4 \times \$80 = \$128{,}000$/mo |
| Agent-assisted | $400 \times 1 \times \$80 = \$32{,}000$/mo |
| Agent compute | \$0.50 per file $\times 400 = \$200$/mo |
| Saving | $\$128{,}000 - \$32{,}200 =$ **\$95,800**/mo |
| Annual saving | **\$1.15M** |

Caveats: (i) accuracy must be $\geq$ current baseline; otherwise the regulatory tail risk dwarfs the saving. (ii) Drift over time means a model-risk-management programme (SR 11-7 style) is part of the cost. (iii) The saving compresses as competitors roll out the same tool.

## 5. Self-Check Questions

1. Name two finance tasks where HITL is non-negotiable and explain why in one sentence each.
2. Identify the single most important difference between an AI agent and a rules-based RPA bot.
3. The agent in the worked example starts hallucinating supporting evidence. What is the first governance control you would tighten?

*[Answers hidden in student version.]*

## 6. Further Reading


- Brynjolfsson, E., Li, D. & Raymond, L. (2023). "Generative AI at Work." NBER Working Paper 31161.
- Bommasani, R. et al. (2021). "On the Opportunities and Risks of Foundation Models." Stanford CRFM.
- European Commission (2024). *The EU Artificial Intelligence Act.* Regulation (EU) 2024/1689.
- Weidinger, L. et al. (2022). "Taxonomy of Risks posed by Language Models." FAccT 2022.
- BoE/PRA (2024). SS1/23: *Model Risk Management Principles for Banks.*


---