MATH 285J, Spring 2026

# Agentic AI for Autonomous Research in Quantitative Finance

**Lecture**: MS 5233 - Wednesday 3:00 PM–3:50 PM; Friday 2:00 PM–3:50 PM

**Instructor**: Mihai Cucuringu
mihai@math.ucla.edu Office: MS 7620B
https://www.math.ucla.edu/~mihai/fin.htm

**Catalog Description**:
Lecture, 3 hours per week (4 units). This project-based course focuses on the use, development, and extension of AI agents and multi-agent systems that assist with and partially automate the research process in mathematical and statistical finance. You will work with real financial data spanning multiple time scales, from low-frequency asset pricing to high-frequency market microstructure, and learn how to use and improve existing agentic platforms to perform tasks such as literature review, hypothesis generation, data analysis, modeling, experimentation, and research report generation. The course is highly project-based and emphasizes building end-to-end research workflows that generate and test ideas, with particular focus on out-of-sample validation, alpha generation, and economically meaningful performance metrics. A natural direction is to extend prior projects toward publishable-quality work using AI-assisted workflows.

## Course Description

- This project-based course focuses on the design and use of *agentic artificial intelligence systems* for conducting research in quantitative finance, combining theoretical, methodological, computational, and empirical components to build agents that assist with and partially automate the research process from problem formulation to empirical validation and communication.
- A central theme is that strong research systems must not only generate plausible outputs, but withstand *rigorous empirical validation* under realistic assumptions, with emphasis on reproducibility, robustness, and economically meaningful evaluation.
- The course adopts a systems-level perspective by decomposing research workflows into modular components, including literature review, hypothesis generation, data acquisition and preprocessing, modeling, evaluation, and reporting, and studies how to design and coordinate agents that interact with data, code, and external tools.
- Methodologically, the course introduces large language models, tool-using agents, memory and planning mechanisms, and multi-agent architectures, with hands-on development using modular platforms that coordinate specialized agents for tasks such as code generation, experimentation, and report writing.
- A distinguishing feature is the use of financial datasets spanning multiple time scales, from low-frequency asset pricing to intraday and high-frequency limit order book data, enabling exploration of problems across asset pricing, portfolio construction, and market microstructure.
- Students learn to extend, customize, and critically evaluate agentic systems within rigorous empirical workflows, using financial performance metrics such as out-of-sample profit and loss, risk-adjusted returns, drawdowns, turnover, and robustness under realistic conditions, with explicit focus on *alpha generation* and economically meaningful trading performance.
- The course culminates in a reproducible research project in which students design agentic systems to replicate existing work, generate hypotheses, conduct experiments, and produce technical reports, while critically assessing failure modes such as overfitting, data leakage, and non-stationarity.

**Prerequisites**:
Linear algebra (MATH 115A), elementary probability & statistics, and programming experience (eg, PIC 16A). Familiarity with basic machine learning is recommended. No prior finance background required.