AI at Your Side: Thinking, Learning, and Working with Artificial Intelligence

Kyiv School of Economics

October 6 and 7, 2026

Based on the book *AI at Your Side: The Student's Guide to Smarter Learning*, Sebastian Galiani and Raul A. Sosa, Oxford University Press, 2026.

Course Description

Artificial intelligence is rapidly changing how we learn, analyze problems, conduct research, write code, and make decisions. But using AI well requires more than learning how to prompt a chatbot. The central question of this course is: How can we use AI to expand our capabilities without outsourcing the thinking that ultimately makes our work valuable?

The course develops a practical framework built around three principles: augmentation rather than replacement, thoughtful automation, and responsible and ethical use of AI. The premise is that AI dramatically lowers the cost of many intellectual tasks—searching, coding, computation, simulation, drafting, and exploring alternatives—but this makes human judgment more, rather than less, important. Choosing the right question, understanding assumptions, evaluating evidence, interpreting results, and recognizing when an answer should not be trusted remain human responsibilities. This reflects the central philosophy of *AI at Your Side*: AI should augment human intelligence rather than replace it.

Rather than treating AI as a separate technical subject, the course uses it to tackle increasingly demanding intellectual tasks: generating and developing ideas, conducting research, learning difficult material, understanding and extending formal models, programming, analyzing data, and making decisions under uncertainty.

The course is organized around live demonstrations of AI-assisted workflows. Students see the entire process unfold: the initial question, the interaction with AI, the resulting analysis or code, the mistakes and unexpected results, and the process of questioning, verifying, and revising what the AI produces. The examples and prompts are provided so that students can reproduce and extend the applications independently after class.

The final lecture turns the perspective around. Instead of asking how we can use AI, we examine how economists can study AI itself. Drawing on ongoing research, we compare simple chatbots, self-correcting systems, and autonomous AI agents, and consider what these technologies may imply for productivity and the organization of knowledge work.

Course Objectives

By the end of the course, students should be able to:

- Use AI as a partner in reasoning rather than simply as an answer-generating machine.
- Distinguish tasks that can profitably be automated from those that require human judgment.
- Use AI to generate, challenge, and develop ideas.
- Conduct AI-assisted research while verifying sources, evidence, and claims.
- Use AI to understand, modify, and stress-test quantitative models.

- Build, understand, and audit code with AI without treating the resulting program as a black box.
- Use AI for data analysis, simulation, and decision-making under uncertainty.
- Understand the emerging distinction between chatbots, self-correcting AI workflows, and autonomous agents.

Course Outline

Lecture 1 — AI at Your Side: How to Work with Artificial Intelligence

What can frontier AI systems actually do, and where do they fail? We begin with an intuitive introduction to large language models and the rapidly changing AI landscape. We then develop the principles that organize the course: augmentation rather than replacement, thoughtful automation, and responsible and ethical use.

Particular attention will be paid to the uneven or "jagged" frontier of AI capabilities: these systems can perform remarkably sophisticated tasks while failing unexpectedly at apparently simple ones. The implication is that effective AI use requires both experimentation and verification.

The lecture introduces a basic framework for interacting with AI, but the emphasis is not on memorizing prompting techniques. Instead, we ask how the division of labor between human and machine should change as the cost of intellectual execution falls.

Central question: What should I delegate to AI, and what should I continue to do myself?

Lecture 2 — Thinking with AI: Ideas, Research, and Discovery

AI is particularly powerful before we know exactly what we are looking for. This lecture examines how AI can expand the space of ideas we consider, challenge our initial framing of a problem, and help turn a vague intuition into a precise and researchable question.

We then move from ideation to research. AI can dramatically reduce the cost of exploring a literature, identifying competing explanations, understanding unfamiliar papers, and mapping what is known. But it can also invent references, mischaracterize papers, and create a false impression of consensus.

We will work from an initial question through ideation, literature discovery, verification, and synthesis, ending with a structured research map. The distinction is important: ideation asks *What could I study?* while literature review asks *What do I need to read, and how should I organize what I find?*

Central question: Can AI help us think of better questions rather than merely give us faster answers?

Lecture 3 — Learning Difficult Things with AI: Mathematics and Economics

This lecture uses mathematics and formal economic models to explore a deeper use of AI: learning by dialogue.

Rather than asking AI to solve a problem, we use it as a tutor, critic, graphical assistant, and model stress-tester. We will work through a formal economic model of time allocation, derive its implications, visualize the equilibrium, and then deliberately change its assumptions.

What happens when work hours cannot be freely chosen? When overtime is paid at a different wage? When constraints become nonlinear? Which conclusions survive, and which depend critically on the assumptions of the original model?

The objective is to show how AI can turn a static model in a textbook into an interactive intellectual laboratory. The emphasis is on interpretation, assumptions, comparative statics, and understanding why a result changes—not simply obtaining the solution.

Central question: Can AI help us understand a model deeply enough to know what happens when we break it?

Lecture 4 — Programming with AI: From an Idea to a Working Model

Generative AI is radically reducing the cost of programming. But producing code and understanding what the code is doing are very different things.

This lecture develops an AI-assisted programming workflow:

Model → Pseudocode → Code → Execute → Inspect → Test → Revise

We will build an economic choice model from scratch. We begin by specifying the economic environment and then use AI to translate the model into code, simulate individual choices, construct a likelihood function, and estimate the model.

Rather than treating AI-generated code as a black box, we will examine the economically and computationally important lines: what each is doing, why it is needed, what could go wrong, and how we can test whether the program is doing what we think it is doing.

The point is not to turn students into software engineers. It is to show how AI makes programming accessible as an analytical instrument while leaving responsibility for the model, its assumptions, and its interpretation with the user. This follows the book's distinction between using AI to generate code and understanding enough to inspect how that code executes.

Central question: If AI can write the code, what does the human still need to understand?

Lecture 5 — AI for Data, Finance, and Decision-Making under Uncertainty

Business and financial decisions rarely depend on a single number. They depend on assumptions, alternative scenarios, correlations, uncertainty, and an understanding of what could go wrong.

This lecture applies AI to valuation, portfolio analysis, data visualization, sensitivity analysis, and simulation. We move from point estimates to distributions of possible outcomes and use AI-assisted programming to conduct Monte Carlo analysis.

The emphasis throughout is on auditing rather than merely producing results. A model can run perfectly and still answer the wrong question. An AI-generated simulation can look convincing while embodying incorrect assumptions or correlation structures. We therefore ask not simply whether the code runs, but whether the economic assumptions, data, and outputs make sense.

We will also use AI to identify hidden assumptions, construct alternative scenarios, and conduct pre-mortem analysis: if our model or investment thesis turns out to be badly wrong, what are the most plausible reasons why?

Central question: How can AI improve decisions without creating false confidence in quantitative outputs?

Lecture 6 — AI at the Research Frontier: From Chatbots to Autonomous Agents

The final lecture moves from using AI to studying AI.

Recent AI systems increasingly operate not merely as chatbots but as agents: systems capable of planning, executing tasks, inspecting their own outputs, using tools, and revising their work. What does this change?

Drawing on ongoing experimental research, we compare three forms of AI assistance—standard chatbots, systems allowed to inspect and repair their own work, and autonomous agents—across demanding analytical tasks. We examine the research question, experimental design, measurement, results, and broader implications.

The lecture uses this research to ask a larger question: what happens to the organization of knowledge work when AI moves beyond providing advice and begins executing complex, multi-step intellectual tasks?

Central question: What happens when AI moves from answering questions to executing complex intellectual tasks?

Format

The course combines 6 lectures of 75 minutes each with live demonstrations of AI-assisted workflows. Rather than asking students to configure tools and execute exercises during class, the instructor will work through the applications live, showing the interaction with the AI, the resulting code or analysis, and the process of inspecting, questioning, and revising its output.

Students will receive the slides, prompts, and examples used in class so that they can reproduce and extend the exercises independently afterward. This format allows class time to focus on the substantive questions—how to formulate the problem, what to delegate, what to inspect, and when to trust the result—rather than on software setup and troubleshooting.

No prior background in artificial intelligence or computer science is required. Familiarity with basic economics, statistics, or finance will be useful for some of the applied examples.

Interested in this course? Email me at sgaliani@tulane.edu