

History of Economic Thought

ECO 4343 / 5343 · Fall 2026

Scott Cunningham · Baylor University

Day One

Before we start: introduce yourself



hot-intake.netlify.app

Scan this, pick your name, and take a selfie.

It takes two minutes and I only ask for it once. I use it to learn your names and faces before the second week, which matters in a class that runs on discussion.

Only I can see what you submit.

Poll 1



slido.com
2503490

How many economics courses have you taken before this one?

I want to know what I am working with. Scan, or go to slido.com and enter the code.


Poll 2



slido.com
2503490

What are you hoping to get out of this class this semester?

Type whatever is true. I read all of these.



Part I

The syllabus

One question, fifteen weeks

Does market exchange **harmonize** our interests, or
conceal a conflict between them?

Smith · Malthus · Ricardo · Bentham · Mill · Marx · the marginalists
They do not agree.

Poll 3



slido.com
2503490

**In one word: how would you describe
what economics is the study of?**

Do not look it up. I want the answer you would
have given before you walked in here.

You already believe an answer

You were taught certain things about markets, and by now they feel like facts about the world rather than one answer among several.

This course makes them strange again
by showing you the arguments they beat.

We stop before 1900, on purpose

- ▷ We read the writers before Smith, then Smith, Malthus, Ricardo, Bentham, Say, Senior, Mill, Bastiat, and Marx.
- ▷ We end where the marginalists take over in the 1870s.
- ▷ A course that runs from Aristotle to Keynes gives you a list of names.
- ▷ I would rather you read a small number of difficult writers closely.

What you will be able to do in December

1. Place an argument in the world that produced it.
2. Give an account of a writer in his own words, without slipping in ideas he did not have.
3. Say what a classical argument lines up with today, and what gets lost in the lining up.
4. Account for how economics became mathematical, and what that made harder to see.
5. State a position you disagree with in its strongest form before you argue against it.

The order of the first three matters.

The book

- ▷ Hunt and Lautzenheiser, *History of Economic Thought*, **3rd edition** — the newest one.
- ▷ About \$43 new. The PDF is on Canvas until your copy arrives.
- ▷ You will cite it by page all semester, so get a copy you can mark up.

How you are graded

Component	Weight	When
Tuesday quizzes	20%	Every Tuesday, 11:02 am
Builds	30%	Six, every other Thursday
Exam 1	15%	Thursday, October 1
Exam 2	15%	Tuesday, November 3
Exam 3	15%	Tuesday, December 8
Participation	5%	Attendance

Exam 3 is on the last day of class, not finals week. This course is finished before your finals week starts.

The Tuesday quiz starts at 11:02

- ▷ Ten minutes, on paper, scantron. Handed out at **11:02**, collected at **11:12**.
- ▷ Arrive late and you get whatever time is left. Arrive after 11:12 and you get a zero.
- ▷ There are no makeups, so I drop your two lowest and you do not explain why.
- ▷ It covers the reading, not my lecture, and it comes before I have said anything.

Bring the book.

Three exams, all in class

Exam	When	Weight
Exam 1	Thursday, October 1	15%
Exam 2	Tuesday, November 3	15%
Exam 3	Tuesday, December 8	15%

- ▷ On paper, in this room, no notes and no AI.
- ▷ Miss one and the average of your other two replaces it, automatically.
- ▷ Exam 3 is the last day of class, so this course is over before finals week starts.

The builds

Six times this semester you will build a website that explains one writer's ideas, in the context that produced them.

- ▷ You will use an AI agent to build it.
- ▷ It publishes to a real website that anyone can visit.
- ▷ You present it to each other in class.
- ▷ I will give you the specific assignment before each one.
- ▷ I will show you an example of what I am looking for, so you are not guessing.

I grade one thing: did you do it.

You have to be there

The build is not finished when you post it.
It is finished when you explain it to someone.

- ▷ If you are not in class on presentation day, that build is a zero.
- ▷ If you genuinely cannot be there, send it to me and I will grade it myself.

Attendance has rising marginal cost

Absences	Attendance grade	Marginal cost
0–3	100	—
4	90	–10
5	75	–15
6	50	–25
7	F in the course	—

The last line is not mine. Baylor requires 75% attendance — for us, 21 of 27 meetings.

Poll 4



slido.com
2503490

What do you use chatbots for?

This is anonymous and I am not policing anything. Answer honestly — the whole point of the next hour depends on it.

You also need an AI agent

- ▷ An account with **Claude** or **Codex**, paid tier, about \$20 a month.
- ▷ Treat it like the book. It is a required material and you should budget for it.
- ▷ Prefer a different agent? Ask me and I will probably say yes.
- ▷ If the cost is a real problem, come see me early and privately.

A chatbot is not an agent

A chatbot

You ask a question.

It answers.

One exchange.

You read the reply.

An agent

You give it a task.

It works through steps on its own.

It writes and runs code.

It hands you something built.

Most of you have done the first. Very few have done the second.

Two rules that are really one rule

I require you to build things with an agent.

I also quiz you every Tuesday where no agent can help.

The quiz is what makes the agent legitimate.

If you can do the work, delegating it is reasonable.

If you cannot, you are approving output you cannot judge.



Part II

Thinking about AI agents

Trust Nothing and Verify Everything

In April 2026, Sullivan & Cromwell apologized to a federal judge after a filing contained AI-generated errors, including fabricated legal citations.

The striking part is that the firm's AI training already included the rule:

**Trust Nothing
Verify Everything**

The rule was right. The workflow still failed.

Production and verification are different tasks

Imagine research as two tasks:

Production

Generating drafts, code, tables, simulations, literature summaries, slides, and alternative versions.


Verification

Knowing whether the output is true, faithful, correctly estimated, and worth believing.

Historically these were bundled. AI agents are pulling them apart.

Agents may have the larger comparative advantage in production. But production only becomes research when directed and verified by skilled humans.

AI makes machine time a substitute for human attention



Pre-AI: human attention and machine help were complements. Post-AI: machine time substitutes for human time in production.

$$Q = f(H, M)$$


- ▷ H = human time and attention
- ▷ M = machine time

The question is not whether to use M . It is how much H you intentionally keep.

Deskilling is real, but not deterministic


The gains are real. People will use them.

The risk is also real: if agents let us bypass attention, we may stop building the human capital needed to judge the output.



This is a design problem, not destiny. Build verification workflows that keep judgment in the loop.

Case I: keep effort constant, take the output gain




Keep H^* the same. Do more and better work in the same hours.

- Biggest gains in output.
- Human capital preserved.

This requires us not to take all the time savings.

Case II: reclaim time, but stay safe




Cut hours, but stay in the green. As long as $H > \bar{H}$, output remains above baseline.

- Moderate time savings.
- Still a net gain in output.

Some slack is fine. The curve shifted up enough to absorb it.

Case III: the depreciation trap



The red zone. Skill depreciation without the effort to maintain $H > \bar{H}$ means $Q_{\text{post}} < Q_{\text{pre}}$.

- Worse off than before AI.

The point is not to avoid time savings. The point is to be intentional.




Part III

What are agents?

Agents are different because they act

A spectrum of AI use, from passive to autonomous:



An agent reads files, writes code, runs it, reads the output, and decides what to do next across many steps.

What the agent actually does

An agent is not just a chatbot with better prose. It acts on your project.

<code>ls data/</code>	list files
<code>rg "minimum wage" .</code>	search text and code
<code>Rscript estimate.R</code>	run an R script
<code>python audit.py</code>	run a Python script
<code>pdflatex paper.tex</code>	compile a paper or deck
<code>git diff</code>	inspect changes

It reads files, writes files, runs commands, sees errors, revises, and tries again.

Reasoning models and agents are merging

Reasoning LLMs

o1, o3, DeepSeek R1, Claude extended thinking

- ▷ Trained to deliberate before answering
- ▷ Pure inference: no action in the world
- ▷ Strong at math, code, logic puzzles

AI agents

ReAct, tool use, Claude Code, Cursor, Codex

- ▷ LLM + tools: read, write, run, search
- ▷ Loop: act, observe, decide, repeat
- ▷ Strong at long, multi-step tasks

2025–26: the merger. Frontier systems increasingly *think* and *act*: plan a task, execute steps, check output, and revise.

Claude Code began as a terminal agent

Boris Cherny joined Anthropic in 2024. Late that year he gave Claude direct access to his terminal and asked it to fix his Spotify playlists.

Claude wrote API calls, read its own output, handled errors, and finished the job autonomously. *That was the proof-of-concept.*

- ▷ **Late 2024.** Internal prototype.
- ▷ **Feb. 24, 2025.** Research preview with Claude 3.7 Sonnet.
- ▷ **May 2025.** General availability.
- ▷ **Nov. 25, 2025.** Claude Code lands inside the desktop app.

The capability mattered. The product wrapper made it usable.

Claude is not one interface

The terminal is where Claude Code began. It is still the cleanest place to work with project folders, code, data, and reproducible files.

CLI / terminal

- ▷ Project directories
- ▷ File edits and code runs
- ▷ Repeatable workflows

Desktop / Cowork

- ▷ Conversational work
- ▷ Shared context
- ▷ Faster collaboration loops

Read Chris Cornwell's Claude academic guide:

<https://cornwl.github.io/files/claude-academic-guide.html>

Explain which parts apply to my project.

Help me choose CLI, Desktop/Cowork, or both.

What is the terminal?

The terminal is a text doorway into your computer: instead of clicking folders and menus, you type instructions.

Underneath macOS and Linux is an old Unix idea: files, folders, programs, and permissions.

Bash or zsh is just the language of that doorway.

Claude Code lives there because research work lives there:

code data logs manuscripts outputs

You do not need to become a Unix expert today. The CLI lets the agent operate where the project actually lives.

Why the CLI still matters

The desktop app is excellent for collaboration. The CLI is different: it lets the agent act directly inside a folder.

- ▷ It can read and edit local files.
- ▷ It can run code, see errors, and try again.
- ▷ It can write logs and leave durable artifacts.
- ▷ In a sandbox, it can move quickly without constant permission prompts.

```
claude --dangerously-skip-permissions
```

That flag means: stop asking before tool actions. Use it only in folders where you are comfortable giving the agent that much freedom.

Homework: ask Claude how to install Claude Code from the CLI on your machine, then do the first run in a disposable folder.

Use agents when the work is valuable, slow, and checkable

My view: delegate to agents when a task satisfies all four conditions.

- 1. The output is highly valuable.**

The stakes justify the setup cost.

- 2. The task is time-intensive for humans.**

Hours or days of skilled effort, not minutes.

- 3. The probability humans do it well is low.**


Difficult to execute at high quality consistently.

- 4. The output can be verified.**


You can inspect, compare, rerun, or source-ground it.

Coding is the canonical example, but this also describes literature review, replication, deck creation, and referee-style review.


Where the work happens




What gets delegated



What is left for you



Three things, and only one of them is yours



Production. An agent can now *do* almost anything. It reaches your computer, your software, and outside services through connectors called MCP. Producing output no longer costs you time.

Verification. But you did not do it, so you cannot be sure it is right. When you produced your own work, verification came along for free. Now it is a separate act you have to perform on purpose.

Understanding. And if the output cost you no time at all, you almost certainly do not understand it. If you do not understand it, we have to ask what school is for.

Poll 5



slido.com
2503490

**AI agents are different from chatbots.
Have you ever used one?**

Now that you know what an agent actually is,
answer again with that in mind.

For Thursday

- ▷ Read Hunt & Lautzenheiser, **chapters 1 and 2**.
- ▷ Chapter 1 is feudalism turning into capitalism. Chapter 2 is what came before Smith.
- ▷ Pay attention to the chartered companies with monopolies over trade routes.
- ▷ Get the book, or the PDF off Canvas, today. Set up your agent this week.

First quiz is Tuesday, September 1, at 11:02.