AI for Economic Research: Dynamic Models, Language, and Agents
Lecture 1.1: AI and the Economics Profession — Cognitive Capital

Zhigang Feng

2026-07-07 10:16:36

Lecture 1 in One Map

One lecture, three decks, one claim: AI is new *cognitive capital* — cheap, fast, non-rival — and this course trains you to deploy it on real research.

you are here ▼

1.1 FRAME
cognitive capital · growth

what kind of economic object is AI?

PRICE OF COGNITION ↓

1.2 PROOF
DeepRamsey · GRAM · meas.

can AI actually do frontier research?

CONTEXT, NOT CAPABILITY

1.3 PRACTICE
roadmap · tooling · capstone

how does this course make it mine?

REPLICATION-FIRST

front door: no prerequisites — bring a research question and skepticism

downstream: Lec02 opens the toolbox — the predictive and generative stacks

Act 1.1 (here) *frames AI as an economic object: why now (capability up, the price of cognition down, record adoption — as of early 2026), what kind of capital it is (non-rival, self-replicating), and where it enters growth — with the measurement gap as the standing caveat.*

Why This Course, Why Now

Capability up, cost down, adoption at record speed — a moving snapshot, early 2026

The Frontier Is Rising Fast

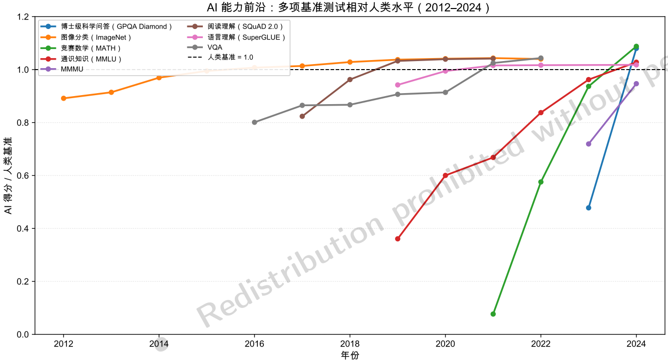

AI benchmark performance, 2012–2024: AI now *exceeds* human baselines on most standard tasks. Source: Stanford AI Index; 苗建军、袁哲、冯志钢，《人工智能经济学导论》（2027）.

- For economists, the headline is not "AI beats humans on benchmarks" but that a *general-purpose cognitive input* is now cheap, fast, and improving — the kind of shock that belongs in a growth or labor model.

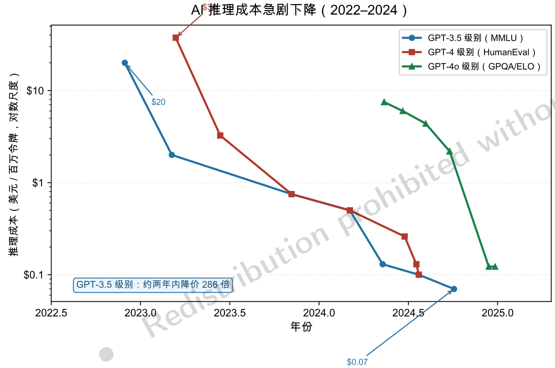

Inference cost has fallen ∼280× in three years; capabilities have compounded simultaneously. Source: 苗建军、袁哲、冯志钢,《人工智能经济学导论》(2027).

- The economically relevant variable is the *price of cognition*. When the marginal cost of a competent draft, proof sketch, or block of code falls by orders of magnitude, the optimal mix of human and machine effort in every task shifts.

Adoption at Historic Speed

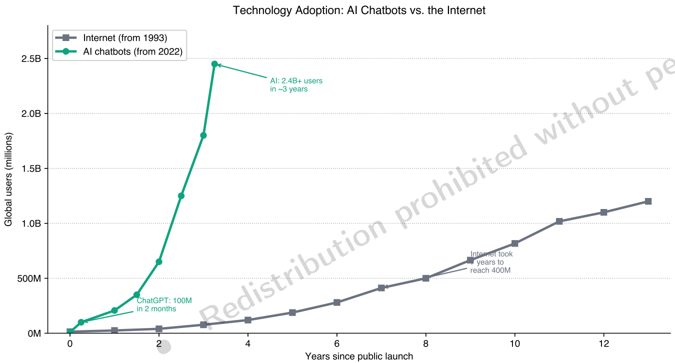

ChatGPT reached 100M users in two months — the fastest consumer technology adoption in history. Source: 苗建军、袁哲、冯志钢,《人工智能经济学导论》(2027).

- Diffusion this fast is rare: most general-purpose technologies took decades to reach households. The speed compresses the window in which firms, workers, and policy adjust.

Diffusion into Firms, and a Multipolar Frontier

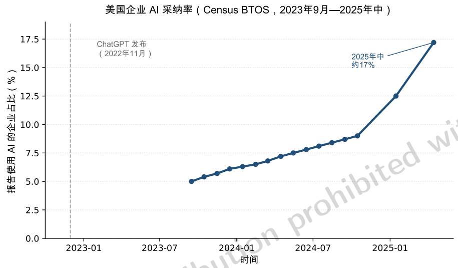

U.S. business AI adoption (Census Bureau BTOS), 2023–2024: rapid diffusion across sectors and firm sizes. Source: 苗建军、袁哲、冯志钢，《人工智能经济学导论》（2027）.

- **Multipolar frontier (as of early 2026):** U.S. labs (OpenAI, Anthropic, Google DeepMind) sit alongside strong open-weight and Chinese models (DeepSeek, Qwen) that have closed much of the capability gap and cut the cost of academic experimentation.
- **Investment, hedged:** capital flooding in is the macro story — and *reported, as of early 2026*, several frontier labs (e.g. Anthropic, OpenAI) and Chinese players such as MiniMax are discussed as potential IPO candidates. Treat all such figures as a moving target.

Why Economists Should Care: AI as a General-Purpose Technology

- **The GPT thesis.** Like the steam engine, electricity, and information technology, AI looks like a *general-purpose technology*: it is
  - *pervasive* — usable across essentially every sector and task;
  - *improving* — getting cheaper and more capable on a steep curve;
  - *innovation-spawning* — it complements and triggers new methods, products, and business models.
- That profile is exactly what growth theory studies. So AI is not only a computer-science topic — it is an object for **growth, labor, and productivity** analysis.

Framing for the course. We treat AI as an economic input — a new kind of capital — and ask the economist's questions: what does it substitute for, what does it complement, how is it priced, and how does it show up in output and the distribution of income?

What Makes a Technology "General-Purpose"?

- **The test (Bresnahan & Trajtenberg, 1995).** A general-purpose technology (GPT) is (i) *pervasive* across sectors, (ii) *improves* steadily over time, and (iii) spawns *complementary innovations* that compound its effect.
- **The ones that delivered** — reshaping growth over *decades*:
  - the *steam engine*; *electricity* (the dynamo); the *internal-combustion engine*;
  - the *semiconductor* / *computer*; the *internet*.

Complementarities take time. David (1990), "The Dynamo and the Computer": factories needed ∼40 years to redesign around electricity before the productivity gains showed up. GPT payoffs *lag* the invention — the same J-curve we should expect for AI.

...But Not Everything Hyped Is a GPT

Many technologies were announced as transformational — even "general-purpose" — and never delivered economy-wide:

- **Nuclear fission power** — "too cheap to meter" (1954); stalled on cost, safety, and regulation.
- **Supersonic passenger flight** (Concorde) — retired; a niche, not a transport revolution.
- **3D printing** — valuable for prototyping; not the predicted manufacturing revolution.
- **The Segway** — hyped to "redesign cities"; a commercial flop.
- **Web3 / blockchain** — promised to rewire finance and the web; limited diffusion so far.
- *...and the perennial "a decade away" list: the hydrogen economy, VR / the metaverse, hyperloop.*

GPT status is an *ex-post* verdict, not a press release. The honest question for AI is *which bucket* it lands in — a genuinely open, empirical question, and one reason this course treats it carefully rather than triumphantly.

AI as Cognitive Capital

Non-rival, self-replicating, scalable: a new factor of production

AI as Cognitive Capital: The Economic Framing

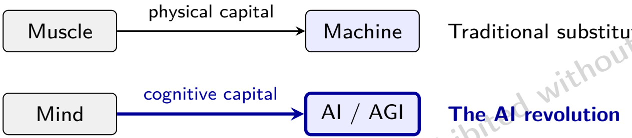

- **What Makes AI Different from Other Capital?**
  1. *Scalability:* Copy model weights instantly; serve 10^6 users in parallel
  2. *Replication:* Training a new human takes 20 years; training a new AI takes days
  3. *Non-rivalry:* Unlike physical capital, AI can be used simultaneously
  4. *Recursion:* AI can generate data and code to improve itself

- **Implication:** Cognitive labor transforms from a biological constraint into a scalable, manufacturable capital good.

The Economics of AI Adoption: Task-Level Production

- Think of output as produced task by task. Within task i, human cognitive labor h_i and AI inference m_i combine with elasticity of substitution σ_i:

$$H_i = \left[\varphi_i \, h_i^{\frac{\sigma_i - 1}{\sigma_i}} + (1 - \varphi_i)\,(\theta_i m_i)^{\frac{\sigma_i - 1}{\sigma_i}}\right]^{\frac{\sigma_i}{\sigma_i - 1}}$$

- $\theta_i =$ AI *capability* on task i; $\varphi_i =$ weight on human labor; $\sigma_i =$ how easily AI substitutes for the human.

- **Reading it.** As θ_i rises and AI's price falls, the cost-minimizing mix tilts toward m_i — quickly where σ_i is high (close substitute), slowly where it is low (complement).

The two unobservables here — the capability index θ_i and the adoption margin (next slide) — are exactly what Topic 1.2's measurement example tries to recover from data.

When Is a Task Automated? The Adoption Margin

- Task i is automated when the value of automating exceeds the upfront cost:

$$E_i(t) = 1 \iff V_i(t) \geq C_i^{\text{upfront}}(t)$$

- **Value** scales with the wage saved, the price of inference, and how often the task runs:

$$V_i(t) = \frac{(w_i - p_t) \cdot \text{Scale}_i}{\rho}$$

  where w_i = human wage, p_t = AI inference price, ρ = discount rate.

- **Cost** is convex in required precision: low-tolerance tasks (surgery, formal proofs, audited accounts) stay expensive to automate even when raw capability is high.

Adoption Follows an S-Curve

- Putting the margin in motion: as p_t falls and θ_i rises, tasks cross the threshold *in order* — easy, high-scale tasks first; hard or low-scale tasks much later.
- The result is the familiar *S-curve* of diffusion, and a long tail of tasks that resist automation.

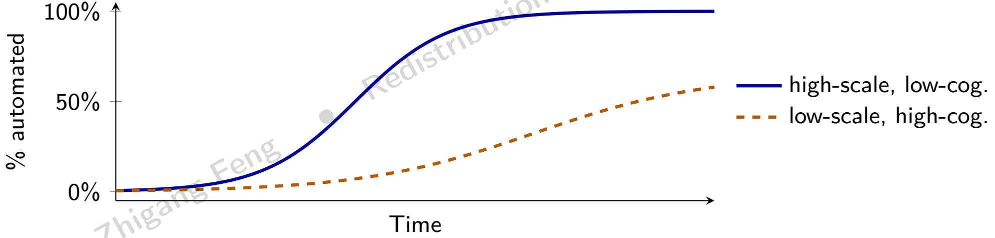

Intangible, Non-Rival Capital — and a Measurement Gap

- AI capital is **intangible** (model weights, prompts, pipelines, skills) and **non-rival** (the same model serves many users at once) — unlike a rival machine that one plant uses at a time.
- National accounts were built to measure *tangible, rival* capital. They struggle with software, data, and now AI: much of the value shows up as better decisions, faster code, and improved service quality that GDP captures poorly.
- This is the modern face of the *Solow productivity paradox*: "you can see the AI age everywhere but in the productivity statistics" — partly a measurement problem, partly a diffusion-lag problem.

Why it matters. If AI is non-rival cognitive capital that the accounts mismeasure, then standard growth and productivity numbers are a *lower bound* on its effect — and measuring it correctly is itself a research agenda (we build toward one in Topic 1.2).

Where AI Enters the Economy

Three channels into growth — and why the statistics may understate all of them

Three Channels: How AI Enters Growth

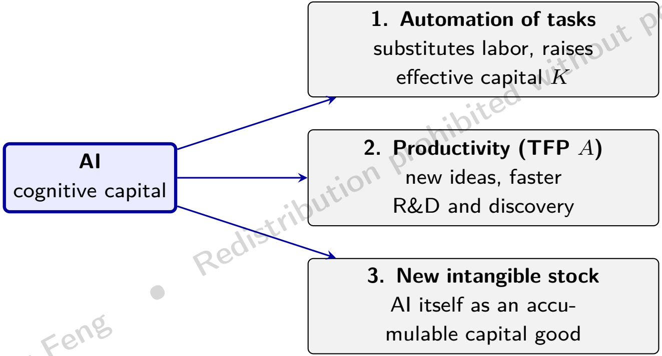

- The three channels map onto the three terms of a growth-accounting identity — which we make precise next.

A Light Growth-Accounting Lens

- Standard Solow decomposition of output growth:

$$\Delta \ln Y \;=\; \alpha\,\Delta \ln K \;+\; (1-\alpha)\,\Delta \ln L \;+\; \Delta \ln A$$

  where α is capital's share, and $\Delta \ln A$ is the Solow residual (TFP).

- **Where AI shows up:**
  - in $\Delta \ln K$ — automation and capital deepening (Channel 1);
  - in $\Delta \ln A$ — AI as an ideas / TFP shifter (Channel 2);
  - as a *new intangible capital input* that the standard K measure omits (Channel 3).

- This is deliberately a *lens*, not a model. The deep dynamic-programming and growth foundations — Bellman equations, optimal growth, heterogeneous agents — arrive in Lec. 4 and Lec. 6.

Task Gains Don't Auto-Sum to Macro TFP

- It is tempting to add up task-level productivity gains into an aggregate number. But aggregation is weighted and bottlenecked:

$$\Delta A_{\text{macro}} \approx \sum_i \omega_i \, \Delta a_i$$

  where ω_i is task i's share and Δa_i its productivity gain.

- **Bottlenecks dominate.** If a few hard-to-automate tasks ($\Delta a_i \approx 0$) carry large weight, they cap the aggregate (a Baumol / weakest-link logic): automating the easy 80% barely moves macro TFP if the binding 20% does not budge.

A sober benchmark. Acemoglu (2024), *The Simple Macroeconomics of AI*, estimates no more than ∼0.66% TFP gain over 10 years (possibly ∼0.53%), because only ∼4.6% of tasks are currently exposed to AI automation. Big task-level wins, modest measured-macro effect — so far.

Bridge: From Framing to Practice

- **Where we are.** AI is cheap, fast, diffusing, and best understood as *non-rival cognitive capital* — a general-purpose technology that reshapes growth and labor through three channels, with effects that today's statistics likely understate.

Read this lecture as a moving snapshot. The AI frontier is changing on a timescale of *months* — capabilities, costs, market structure, and policy are all in flux. The economics of AI is an **active, unsettled research area**: our understanding (and many of the numbers here) will keep evolving quickly. Treat results as current best estimates, not settled facts.

Next (1.2): leave the framing behind and watch AI *actually do* frontier economic research — two live case studies and a measurement example. *Lecture 2* ("What is AI") then opens the toolbox: the technique families (ML / DL / RL / LLM / agentic) and where AI is powerful *but bounded*.

Arc: ***1.1 FRAME*** *(cognitive capital) → 1.2 PROOF (AI doing research) → 1.3 PRACTICE (course + capstone).*