Anran Liu

Application to the NBER Digital Economics and AI Tutorial, January 2027

Contact information

anrliu@ucdavis.edu · (+1) 530-460-8175 · Department of Economics, University of California, Davis, 1 Shields Avenue, Davis, CA 95616

Graduate study

I am a fourth-year Ph.D. student in economics at the University of California, Davis. I enrolled in September 2023, passed the qualifying exams after the first year, and received an M.A. in 2024. My primary fields are development and labor economics, and my secondary fields are urban economics and machine learning in economics. My dissertation committee is Giovanni Peri (chair), Diana Moreira, and Arman Rezaee. From January to May 2027, I will be a Visiting Student Researcher at Stanford's Center on China's Economy and Institutions (SCCEI).

I hold an M.M. in agricultural economics and management from Renmin University of China (2023) and a B.A. in economics from Capital University of Economics and Business (2020).

I taught myself machine learning outside my economics coursework. From October 2025 to September 2026, I was a research intern at the Shanghai Artificial Intelligence Laboratory. There I used reinforcement learning to train LLM agents to search for information and use tools, and to train vision-language models to answer questions about documents. I am first author of papers at EMNLP 2026 and ACM Multimedia 2026 and co-first author of papers at ICML 2026 and AAAI 2026.

Research interests

I study the frictions that keep households and firms out of markets, and how policy and technology change them. In my work on rural China, the friction is distance. I evaluate policies that move people to jobs and policies that bring markets to villages. In my work on AI, the frictions are trust and information. What an AI agent remembers affects whether others can trust its promises, and when applicants use the same AI tools, screeners may no longer tell them apart.

I combine household panel data from the staggered rollout of national programs in China with experiments on LLM agents, in which memory, identity, and information can be set directly. My machine learning background lets me build and run these experiments myself.

Next, I plan to study which identity and memory rules trading agents should follow, and how platforms can verify what applicants report when AI writes the application. I will already be at Stanford during the tutorial. It would let me get feedback on these AI projects from researchers who study platforms and AI.

Research in digital economics and AI

Designing the Shadow of the Future: Memory, Identity, and Relational Contracts among AI Agents (working paper)

As AI agents begin to trade on behalf of firms, whether relational contracts work depends on what each agent remembers and whom it recognizes, and developers and platforms choose both. I show that for a broad class of memory designs, whether trust can replace paid enforcement such as escrow depends on two numbers, how strongly the record deters cheating and how reliably it identifies who cheated. If persistent identity makes records more reliable, the best design pairs it with memory that forgets once cheating is deterred. Holding fixed which relationships survive, a platform that sells enforcement prefers weak identity and long memory, because wrongful punishment sends trade to paid enforcement.

In experiments with LLM agents from three model families, AI buyers keep promises only to counterparts whose records make cheating costly. After being cheated, AI suppliers return to trust at a speed set by their memory, which lets me price each design in escrow fees. Agents that summarize their own history recall cheating that never happened.

Who Holds the Pen? Shared AI Advisors and the Defense of Screening Signals (working paper)

Lenders, employers, and insurers screen applicants on what applicants report about themselves, and more applicants now write these reports with the same few AI tools. In public income-screening tasks, an AI advisor that only helps applicants fill in their forms, without seeing the model, lowers the accuracy of the retrained screener more than an optimizer that is given the model weights. Screening suffers because AI writes the reports, not because AI can see the model. A platform has to verify self-reported information or stop relying on it.

Bringing E-commerce to the Countryside: Impacts on Rural Household Welfare and Economic Inequality (with Mingzhe Wang, Jia Tang, and Dongyue Chen, revise and resubmit at *World Development*)

This paper evaluates China's national program to bring e-commerce to rural counties, which paired online platforms with rural logistics and training. The program raised household income, with larger gains for lower-income families.

Generative AI and the Minimum Scale of Knowledge Work (work in progress)

This project asks whether generative AI lets one person do knowledge work that once needed a team in a large city. Building on knowledge hierarchies (Garicano 2000) and entry with fixed costs (Hopenhayn 1992, Melitz 2003), I predict more employment in small firms, more solo work, and a shift of codifiable knowledge work away from knowledge-hub cities. Preliminary estimates from Occupational Employment and Wage Statistics data show that occupations with more codifiable specialization lost employment in knowledge-hub cities after 2022, and remote work explains little of this decline.