Digital Economics and AI Tutorial Application
Applicant: Kehang Zhu

Contact Information
Harvard University
Cambridge, MA, USA
Email: kehang_zhu@g.harvard.edu
Phone:617-909-9678
Website: https://kehang-zhu.github.io/

---

1. Graduate Study and Enrollment

I am a fifth-year Ph.D. candidate in EconCS at Harvard University, where I work at the intersection of economics, computer science, and artificial intelligence.

Education & Course of Study (professors in parentheses)

- Ph.D. in Physical Science & Computer Science, Harvard (2027, Expected)
    - Committee: John J. Horton (Chair), David C. Parkes
- S.M. in Computer Science, Harvard (2024)
- B.A. in Physical Science (College Honors), University of Science and Technology of China (2021)

Coursework (professors in parentheses)

- Machine Learning (Weiwei Pan)
- Artificial Intelligence and Language Model (Kung)
- Statistics and Algorithmic Fairness (Cynthia Dwork)
- Game Theory and Mechanism Design (David Parkes)
- EconCS Seminar (Yiling Chen)
- Economics of Information and IT (John J. Horton)
- Negotiation (Keslley Hong)

---

2. Research Interests and Current Projects

My research studies how large language models can be used as subjects and tools in economic experiments, and how such “digital mice” can inform mechanism design and digital platform policy. One line of work builds computational systems that use LLM-based agent simulations to help social scientists explore mechanisms and hypotheses in silico before running costly human experiments. A second line develops open-source software for large-scale online economic experiments to understand the role of AI agents in digital market settings such as multi-party negotiations.

Broadly, my research interests include:

- LLM agents as "digital mice" and automated social science
- Bargaining and negotiation in platforms and marketplaces
- Behavioral and experimental economics with human–AI collaborations
- Digital platforms, online marketplaces, and market design

Below are my main current projects:

(1) Automated Social Science: Language Models as Scientist and Subjects (With Benjamin S. Manning and John J. Horton), *The Quarterly Journal of Economics* - Reject & Resubmit.

There is much work on efficiently estimating econometric models of human behavior but comparatively little work on efficiently generating and testing those models to estimate. There is increasing evidence that Large Language Models (LLM) can simulate humans as experimental subjects with surprising degrees of realism. A natural question is how to efficiently leverage AI simulations to generate hypotheses for testing in the lab with human subjects. *Automated Social Science: Large Language Models as Scientist and Subject* provides a proof of concept that LLMs can autonomously design, run, and analyze LLM-based economic experiments. I have developed a computational system to take a social-scientific scenario of interest and automatically generate hypotheses in the form of structural causal models. The system then designs an experiment, instantiates independent LLM-powered agents that vary along the exogenous parameters of the causal model, runs thousands of simulations, and collects and analyzes the resulting data. Across economic domains—from bargaining to auctions—the system successfully tests hypotheses and produces plausible results that can be validated in experimental settings with people. Crucially, we also find that LLMs are unable to predict the outcomes of LLM based experiments in advance—highlighting the value of simulation. Like human scientists, they must often run experiments to discover how they will actually behave.

(2) Strategic Tradeoffs Between Humans and AI in Multi-Agent Bargaining (With Benjamin S. Manning, John J. Horton, Crystal Qian, Nithum Thain, James Wexler), In preparation for *Management Science*

In ongoing work, I study how individuals use AI systems as tools or as autonomous agents in a classical economic setting, that of multi-party negotiation. In collaboration with Google DeepMind, my work in this area is focused on a new class of economic actors: LLM-powered software agents that make decisions on our behalf. In An Experimental Study of AI Agent Delegation in Multi-Player Negotiation (in preparation for Management Science), my coauthors and I report the results of a large-scale lab experiment (~500 participants) in a multi-player, multi-turn negotiation game, evaluating how people adopt different AI-assistance modes and how well they perform together with AI. Participants were randomly assigned to one of four conditions: (1) no AI assistance, (2) AI-generated recommendations (Advisor), (3) AI-generated feedback (Coach), or (4) full AI delegation (Delegate), with the other players in the same mode. In treatments (2) through (4), participants can choose whether or not to adopt the AI assistance. We find that participants in the Delegate treatment achieve significantly higher gains than those in the Advisor/ Coach and no-AI treatments. Among all AI assistance modes, participants report a

dominant preference for the Advisor mode (44%). Moreover, Delegate generates positive externalities; even non-adopting human players in Delegate-enabled games benefit by receiving higher-quality offers.

---

3. Dissertation Title, Description, and Committee

Working Dissertation Title
Language Models as Economic Agents: Synthetic Laboratories for Optimizing Mechanisms for Digital Markets.

Short Description
My dissertation investigates whether and how large language models can function as **proxies for human decision-makers** in economic environments, and how they can be used to explore and optimize mechanisms for digital markets.

The dissertation is organized around three core questions:

1. **Human-Like Bounded Rationality:**
   To what extent can LLM agents reproduce key qualitative patterns of human behavior in classic experimental settings (e.g., auctions, public goods, bargaining)? When they fail, what dimensions of bounded rationality or social preferences are they missing?
2. **Theoretically Sophisticated Agents for Mechanism Evaluation:**
   Can LLM agents be steered toward behavior that approximates theoretically optimal or equilibrium strategies, and used to stress-test mechanisms (e.g., auction formats, reserve price policies, rules in digital platforms) before running large-scale human or field experiments?
3. **Hybrid Human–AI Experimental Design:**
   How can we combine human and LLM data—through techniques like prediction-powered inference or model-assisted design—to allocate experimental resources more efficiently and to explore large design spaces (e.g., over mechanisms, assistance modes, or communication protocols) that would be prohibitively expensive with humans alone?

The overarching objective is to build a conceptual and empirical foundation for using LLM agents as a **complement** to human subjects in digital economics, with an emphasis on mechanisms relevant to online platforms and AI-enabled marketplaces.

Committee (in formation)
My dissertation is co-advised by:

- **John J. Horton** – Professor, MIT Sloan School of Management (online labor markets, digital platforms, experimental economics)
- **David C. Parkes** – Area Dean and George F. Colony Professor, Harvard John A. Paulson School of Engineering and Applied Sciences (market design, algorithmic game theory, AI and economics.