Application to the NBER Digital Economics and AI Tutorial

January 20-21, 2027 | Stanford University

Haosen He

Department of Economics, University of California, Berkeley
660 Evans Hall #3880, Berkeley, CA 94720
haosen_he@berkeley.edu

Graduate study

I am a third-year PhD student in Economics at the University of California, Berkeley. My co-advisors are Ben Handel and Nano Barahona. I have completed two years of doctoral coursework, including the core sequences in economic theory, macroeconomics, and econometrics, followed by field courses in industrial organization, mathematical economics, mechanism design and agency theory, game theory, and political economy. My research combines economic theory with empirical and experimental methods to study digital platforms and artificial intelligence.

Research interests and ongoing projects

My research focuses on the regulation of digital platforms and the economic behavior of AI agents. My social media projects examine how participation externalities affect monetization and regulation. My AI projects ask whether agents' resource choices serve delegated objectives and how their participation changes market outcomes.

Social media regulation

My paper *Optimal Social Media Regulation* studies a collective participation problem: as more people join a platform, staying off becomes increasingly costly, so even people who derive little intrinsic value from the platform may choose to participate. The central regulatory problem is that participation choices depend on other users' decisions while the platform controls how participation is monetized. Using mechanism design and extreme-point methods, I characterize how the platform optimally combines subscription pricing and advertising when consumers differ in their valuations. I then examine optimal regulation when advertising intensity is non-contractible, showing how the appropriate policy depends on the platform's ability to monetize advertising. Within the model, weak advertising monetization calls for subsidies, intermediate monetization for taxes, and sufficiently strong monetization for a ban. The paper is currently under review.

A related empirical project, *The Equilibrium Welfare Effects of Social Media Age Restrictions*, examines how Australia's social media age restrictions affect platform monetization and consumer welfare among adult users. By allowing age groups to differ in willingness to pay and fear of missing out, the model links changes in the composition of users to the platform's choice of monetization. The model predicts that restricting minors' participation can shift monetization from advertising toward user payments, potentially benefiting adults. I find strong empirical evidence supporting the theory's predictions. This approach highlights why a policy aimed at minors can also affect adult consumers through the platform's response.

Economics of AI

In *Thinking on a Budget: Revealed Preferences over Computational Resources in Large Language Models*, joint with Shachar Kariv and Mingduo Zhao, we study whether AI agents allocate computational resources in ways that serve their delegated objectives. Delegating a task also delegates decisions about what information to acquire, how long to deliberate, and whether to verify an answer. Revealed-preference tests assess whether choices can be rationalized by a stable utility function. For a delegated agent, a further question is whether that utility aligns with the principal's objective.

Our experiment observes both resource demand and task performance. Five open-weight language models allocate budgets across worked examples, deliberation, and verification as relative prices vary. Separately, we assign every feasible resource bundle to each model and measure its performance on the same tasks. Because the same model chooses and uses these inputs, the design allows us to compare its revealed demand with the returns that different allocations actually produce. We use these measurements to refine the Critical Cost Efficiency Index (CCEI) by requiring rationalizing utilities to respect the measured performance ordering. This distinguishes consistency with some utility from consistency with a utility aligned with task performance.

We find a substantial gap between choice consistency and alignment with the delegated objective. Models make relatively coherent resource choices, but score poorly on the stricter test requiring rationalizing utilities to respect measured task performance. The results show why choice consistency alone is insufficient for evaluating delegated agents: systematic resource choices need not serve the objective the agent was assigned.

In joint work with Debbie Liang, we study how AI employers change search, reputation, and competition in online labor markets. We ask how lower search and evaluation costs affect matching, and whether agents' use of richer information reduces the hiring disadvantage faced by workers with little prior experience. We also examine how improved search and bargaining interact with competition for scarce workers to determine workers' earnings and the division of surplus.

Interest in the tutorial

The tutorial and the accompanying Digital Economics and AI meetings would broaden my exposure to current research and help me identify new questions for my dissertation. I would also value the opportunity to exchange ideas with research leaders and fellow graduate students and get to know others working on related questions.