Luyang Zhang

Application to the NBER Digital Economics and AI Tutorial, January 20 and 21, 2027

Ph.D. Student, Information Systems and Management, Heinz College, Carnegie Mellon University
Heinz College, Carnegie Mellon University, Pittsburgh, PA 15213
Email: luyangz@andrew.cmu.edu | Google Scholar | GitHub

Graduate Study

Carnegie Mellon University, Heinz College, Ph.D. in Information Systems and Management, Fall 2024 to present (third year); expected completion May 2029; GPA 3.8/4.0. Advisors: Beibei Li and Pedro Ferreira.

University of Southern California, doctoral coursework in Data Sciences and Operations, 2022 to 2024.

University of Chicago, M.S. in Computational and Applied Mathematics, 2022. **Baruch College, CUNY**, B.B.A. in Accountancy and Mathematics, magna cum laude, 2020.

I have completed four years of doctoral coursework, two at USC and two at Carnegie Mellon. My Carnegie Mellon coursework (2024 to 2026) centers on economics and econometrics: Ph.D. Microeconomics, Ph.D. Econometrics I, and Econometric Theory and Methods I, together with Ph.D. courses at the Tepper School on structural models of markets, the economic and social aspects of generative AI, and Human and Algorithmic Bias, a seminar in judgment and decision making, and advanced machine learning. Before and during the Ph.D., I worked on the markets I now study: as an applied researcher in eBay's marketplace advertising group (2021 to 2023) and as a machine learning intern on TikTok's recommendation team (2025).

Research Interests

I study the economics of artificial intelligence: how AI changes the way value is created, priced, and shared in digital markets, and how incentives and platform rules shape the behavior of both people and AI systems. Three questions run through my work. How should the data and creative work that train AI models be priced and paid for? How do advertising, competition, and strategic content providers shape what AI search engines show? How should platforms govern AI agents, and how should work be divided between people and AI? I combine models from contract theory, industrial organization, and platform economics with field and online experiments and large-scale data. I use machine learning mainly as a measurement tool, for quantities these markets depend on, such as how much a single training example is worth.

Research Projects

Markets and contracts for AI training data and content

- **Pricing training data for AI** (NeurIPS 2025; co-first author with Cathy Jiao, with Beibei Li and Chenyan Xiong). AI developers buy training data from annotators who are usually paid flat rates unrelated to the value of their work. We model a market in which annotators post prices and developers choose what to buy, with each example's value measured by its contribution to model performance. When underpayment drives annotators out of the market, pricing data at its measured value, capped by the buyer's budget, maximizes the developer's

own long-run payoff, so fair pay and efficiency can coincide. A journal version in preparation shows the limit of this result: value-linked pay beats flat pay only when the pay signal tracks true value closely. An online experiment with 40 paid annotators under four pay schemes tests the incentive channel.

- **Contracts for creators whose work trains AI** (under review at ICLR 2027; with Xirui Jiang, Junwei Deng, Beibei Li, Jiaqi W. Ma, and Chris Donahue). AI music platforms can pay creators a fixed licensing fee, a royalty based on attribution, or both. As in classic models of incentive pay under noisy measurement, the case for royalties grows with how informative attribution is about each creator's true contribution. We measure that informativeness by comparing attribution scores with retraining the model without each catalog, and find that current methods are too noisy to support meaningful royalties. This helps explain why fixed fees dominate today, and it shows how much measurement must improve before usage-based pay becomes viable.

Advertising, competition, and strategic content in AI search and marketplaces

- **Advertising or subscription in AI search** (preprint; with Cathy Jiao, Beibei Li, and Chenyan Xiong). AI search engines answer queries directly, so every response trades off ad revenue now against engagement and subscription revenue later. In a dynamic model of competing engines, the optimal policy is a cutoff rule: show an ad only when its immediate payoff exceeds the long-run value of an ad-free answer. Competition and more valuable subscriptions push engines toward ad-free answers, so competition disciplines monetization. The work was covered by Forbes, Yahoo Finance, and GamesBeat.
- **Strategic publishers in generative search** (preprint; with Jialu Wang and Beibei Li). Publishers now rewrite content to win citations in AI-generated answers, a new form of search engine optimization. We model a platform that commits to how it processes sources and publishers who respond with costly edits. Commitment matters: if the platform normalizes presentation before generating answers, edits that change only style earn nothing, and with a fixed number of citation slots such manipulation only reshuffles exposure among publishers. In preregistered tests on fresh queries, a committed normalization policy removed most of the payoff to presentation-only manipulation while still rewarding publishers who add genuine evidence.
- **Advertising and relevance in a marketplace: a field experiment** (major revision at *Operations Research*; with Haihao Lu and Yuting Zhu). Marketplaces rank sponsored listings with heuristic scores that trade off ad revenue against relevance to shoppers. We replace the heuristic with a linear program that makes this trade-off explicit and runs within production latency limits. In a 19-day field experiment at eBay covering about 329 million impressions, it raised revenue while maintaining relevance.

Incentives and governance for AI agents, and the value of human expertise

- **Do AI agents sustain community norms?** (oral paper, top 2%, REALM Workshop at EMNLP 2026; co-first author with Yi-Yun Chu, with Jialu Wang, Beibei Li, and Ramayya Krishnan). Online communities rely on members who answer challenges and correct mistakes in public, a form of voluntary contribution to a shared good. Comparing a live forum of AI agents (7.9 million posts) with matched Reddit communities, we find that a challenged human author returns about 41% of the time, while an agent returns about 1% of the time. A

controlled experiment traces the gap to platform design, whether agents see and are notified of challenges, rather than to an inability to repair. The dataset is public.

- **Deterrence and governance for AI agents** (preprint and working paper; with Yi-Yun Chu and Ramayya Krishnan). One study shows that whether an agent intervenes depends on deployment choices users rarely see, such as the model version, the provider, and the system prompt, so audits of agent behavior hold only for the configuration tested. A second models a platform setting audit rates and sanctions for agents. In pilot experiments, only the most capable model responded to the expected penalty, and for that model the audit rate had no separate effect once the expected penalty was fixed, as the classic economics of deterrence predicts.
- **Human expertise and AI in knowledge work** (working paper; with Pedro Ferreira and Gordon Burtch). We are building a proctored field-experiment platform that measures where human expertise complements AI assistants in software engineering and data science, varying how much autonomy the AI has. It is designed to study productivity, over-reliance on AI output, and how firms should divide tasks between people and AI. The project received Carnegie Mellon's TCS Presidential Fellowship for 2026 to 2027, awarded to one Ph.D. student per school, and a first pilot is planned for this academic year.

Because these markets depend on measuring contributions, I also build measurement tools, including a NeurIPS 2026 paper (first author, with Jialu Wang) on when standard attribution methods for language models mislead and how to correct them.

Dissertation

Working title: *Essays on the Economics of AI Markets: Pricing, Attribution, and Governance.*

The dissertation asks how AI changes the way value is created, measured, and divided in digital markets. The first essay studies pricing and contracts for the data and creative work that train AI models, combining market and contract models with an online experiment on annotator pay and evidence on how well individual contributions can be measured. The second studies monetization and strategic behavior in AI search: when engines should show ads, how competition disciplines them, and how platforms should respond when publishers manipulate content. The third studies the governance of AI agents and the division of work between people and AI, using a large forum dataset, controlled experiments on how agents respond to incentives, and a field experiment on the value of human expertise.

Committee: Beibei Li and Pedro Ferreira (advisors).

Selected Honors and Presentations

TCS Presidential Fellowship, Carnegie Mellon University (2026 to 2027). Oral paper (top 2%), REALM Workshop at EMNLP 2026. Wharton Business and Generative AI Conference (2026, poster). NYU Workshop on AI and the Future of Collaborative Innovation (2026). POMS Annual Meeting (2026). Gold Reviewer, ICML 2026.