Yue Li

1. Contact Information

Yue Li, Ph.D. Student (third year, expected graduation in 2029), Information Technology Management, Scheller College of Business, Georgia Institute of Technology, Atlanta, GA, USA

Email: yue.li@scheller.gatech.edu

2. Graduate Study

I am a third-year Ph.D. student in Information Technology Management at Georgia Tech's Scheller College of Business, enrolled since Fall 2024, with a doctoral minor in Artificial Intelligence from the College of Computing. My training combines doctoral economics coursework (Mathematics for Economists, Microeconomic Theory II, Econometrics I & II, Advanced Game Theory) with doctoral seminars on the economics of AI, analytical modeling, structural modeling, and causal inference, and graduate CS courses in machine learning, artificial intelligence, and computer vision.

3. Research Interests

I study the economics of AI in firms and markets. AI's value depends on what people put into it and what they do with its output, and both involve asymmetric information. Upstream, workers produce the data AI learns from, but their effort is hidden and the AI they train may replace them (moral hazard). Downstream, users decide whether to follow AI advice, but they hold information the AI lacks (private information). In both settings, what looks like resistance to AI (under-training it, ignoring its advice) can be a rational response to incentives and information. The remedy is then better contract and system design. I study the upstream problem with contract theory (Project 1) and the downstream problem with field data and a randomized experiment (Project 2).

4. Research Projects in Progress

Project 1. Training the Technology That Replaces You: Moral Hazard, Strategic Resistance, and Contract Remedy under AI Learning

With Lizhen Xu and D.J. Wu

Question. *Upstream: who trains the AI?* Firms increasingly build AI from their own employees' work: labeled cases, demonstrated procedures, and corrections of model output. This training data is valuable, but collecting it creates an incentive problem: training effort is hard to observe or contract on, and the better the AI a worker trains, the more likely it can take over the worker's job. Will workers put effort into training an AI that may replace them, and which contracts make them do it?

Model. In a two-period principal–agent model under a linear contract, the worker's unobservable effort in period 1 trains an AI. In period 2, the firm observes the AI's capability and chooses whether to retain the worker to work alongside the AI or replace him with it and keep all AI output. The firm cannot commit to this choice in advance. Training therefore cuts both ways for the worker: if he is retained, a more capable AI raises the output on which his pay is based (augmentation), but the same capability makes the AI a better substitute for him and raises the firm's temptation to replace him (automation).

Findings.

1. **Strategic training cap.** Because the firm cannot commit, AI capability itself becomes the dismissal trigger. The worker trains the AI only up to the highest capability at which he keeps his job, and the firm responds by weakening his incentives. Beyond the standard agency cost, this creates two AI-specific losses: the AI is under-trained, and, when the AI is strong, the firm retains workers it would be efficient to replace. *The fear of layoff prevents the layoff.* These losses grow without bound as AI becomes easier to train and more productive on its own.
2. **Flexibility makes it worse; commitment fixes it.** Our main model sits between two extremes: the firm can fire the worker but cannot rewrite his contract. A promise never to fire removes the strategic training cap. Allowing the firm to rewrite the contract after seeing the AI lowers both welfare and profit even below the main model, since each improvement the worker makes to the AI reduces his future pay. The first-period contract loses its grip on training: he trains only for a small renewal rent.
3. **Contracts that pay at dismissal restore efficiency, with gains for both sides.** We propose two contract remedies paying the worker if he is replaced: severance is a fixed lump sum, and a royalty is a share of the AI's later output. When the AI is weak, a royalty works best: the firm fires only workers who under-train, so workers train beyond the commitment level, and moderate royalty rates can make both sides strictly better off. When the AI is strong, replacement is efficient and severance works best: a fixed payment for reaching the efficient AI capability keeps the firm willing to fire and turns the employment contract into a contract to buy training. A royalty falls short here because it is paid from the very output that replacement transfers to the firm.

Contribution. The paper contributes to the economics of AI and to research on human–AI collaboration by modeling AI capability as the endogenous product of workers' hidden effort. This turns the firm's inability to commit into a new source of inefficiency in training. We show it is a contracting problem with a contracting solution: at every level of AI strength, some contract makes both sides better off.

Project 2. Adoption and Reliance: When Shoppers Reject AI Sizing Advice

With Yi Gan, Pan Li, and D.J. Wu

Question. *Downstream: who follows it?* Recommender systems create value only if users consult them, share the information the AI needs, and act on its advice. We ask when consumers decline AI advice, how declining relates to decision quality, and whether it reflects private information rather than algorithm aversion. The tension is a two-sided information asymmetry: the tool knows product and historical fit data the shopper cannot see, and the shopper knows their body, fit preference, and experience, which the tool observes only partly through disclosure.

Data and design. We link an apparel brand's AI size-recommender logs to purchases and returns (1.03 million order lines, July–December 2024) and run a randomized test of tool availability at two other retailers (about 9 million visitors, April–May 2026). We treat AI use as three sequential decisions: whether to consult the tool, whether to disclose body measurements, and whether to buy the recommended size. Returns measure decision quality.

Findings.

1. **Access does not create value.** Randomly offering the tool slightly lowers conversion (about 1.2%) and net-of-returns value (about 1.7%), with no significant change in returns.

2. **Use is selective, and advice is often ignored.** Shoppers consult the tool for only 8.3% of eligible items, disproportionately those hardest to size. Only 36.1% of advised items are bought in the recommended size, and disclosing measurements barely changes this.
3. **Overriding can be informed.** Among shoppers who disclose measurements, items bought against the AI's advice are returned 2.44 percentage points *less* often than accepted items (3.98 within the same product). There is no such gap for shoppers who give only basic inputs.

Contribution. Disagreement with AI need not signal algorithm aversion. It can reflect consumers combining AI advice with private information the AI does not observe. Whether disagreement is valuable depends on what the user has disclosed, so AI reliance is better understood as a sequence of choices than as a single accept-or-reject decision.

5. Publications

Li, Y., Deng, H., & Zhang, C. (2026). Alice in land of games: Investigating behavior spillover effects of users' engagement and social connections across games. *Decision Support Systems, 202,* Article 114611. https://doi.org/10.1016/j.dss.2025.114611

6. Dissertation

I will propose my dissertation in Fall 2028. It will study how asymmetric information between people and AI shapes AI's economic value, building on the projects above.

Advisors: Dr. D.J. Wu (Georgia Tech), Dr. Lizhen Xu (Georgia Tech).