Generative AI in Markets and Science: Bias, Search, and Expertise

Yafei (Debbie) Liang | Application to the NBER Digital Economics and AI Tutorial

Haas School of Business, University of California, Berkeley
debbieliang-123@berkeley.edu | Berkeley, California

Academic background and research interests

I am a third-year Ph.D. student in Quantitative Marketing at the University of California, Berkeley, Haas School of Business (2024-present), advised by J. Miguel Villas-Boas. I have completed doctoral coursework in economic theory, econometrics, industrial organization, behavioral economics, and quantitative marketing. I hold an M.S. in Electrical Engineering and Computer Sciences (May 2022) and a B.A. in Computer Science and Data Science (December 2020), both from UC Berkeley.

My prior experience includes platform competition and ratings research as a predoctoral associate at Harvard Business School's Platform Lab, and machine learning systems research at Berkeley's Sky Lab. I also worked on advertising models as a machine learning engineer at Meta and contributed to clinical prediction and climate-related social media research.

My current research asks how generative AI changes the information people receive, the expertise they seek from others, and the way consumers search and choose merchants. I study these questions through digital advertising, consumer valuation, scientific collaboration, and generative search, combining laboratory and field experiments with panel data analysis and machine learning. Projects 1 and 2 described below have been accepted for presentation at the Symposium on Statistical Challenges in Electronic Commerce Research (SCECR 2025), the Conference on Artificial Intelligence, Machine Learning, and Business Analytics (AIML 2025 and 2026), the INFORMS Workshop on Data Science (2026), the Conference on Information Systems and Technology (CIST 2026), and the Hawaii International Conference on System Sciences (HICSS 2027).

1. From AI Implicit Associations to Market Outcomes: Evidence from Lab and Field Experiments

With Mingduo Zhao

Research on algorithmic bias shows that cost-efficient ad delivery can generate gender disparities even under gender-neutral targeting (Lambrecht and Tucker, 2019), while models trained on human language can reproduce historical racial and gender associations (Caliskan et al., 2017). Generative AI extends these concerns from algorithmic allocation to the decentralized production of market information. Many firms and individuals can cheaply generate plausible content in which subtle choices of wording or imagery reproduce demographic stereotypes, making their effects difficult to detect and audit. We ask how these associations enter market signals, affect consumer demand, and respond to model fine-tuning and disclosure. We study two applications: AI-generated product-review summaries and video advertisements.

AI summaries and consumer valuation

We first adapt the Implicit Association Test (IAT) to measure how models associate Black and White faces with positive and negative descriptors. We find negative racial associations across multiple closed-source and open-source models. Our focal model, Gemma, exhibits a larger association gap than the human IAT benchmark used in the study, raising concern that generated content could amplify existing human stereotypes.

We next examine whether these associations enter product information. For 839 products, we generate paired summaries of identical customer reviews using matched Black and White seller images. Holding the product and customer reviews fixed, we find that summaries generated with Black seller images receive lower sentiment scores than those generated with White seller images.

To test the downstream consequences for demand, we recruited 976 U.S. participants on Prolific for an incentive-compatible laboratory experiment involving 16 household products. Independently randomizing summary text and the displayed seller photograph separates responses to AI-generated information from direct responses to seller race. The Black-conditioned summaries lower willingness to pay for the same product. Direct Preference Optimization (DPO) fine-tuning nearly closes both the sentiment and willingness-to-pay gaps, connecting an intervention in content generation to reduced disparities in consumer valuation. Holding the summary text fixed, we find that the effect of an AI versus human source label on willingness to pay depends on whether the summary is described as representative of the underlying reviews or potentially selective.

AI video advertising and realized demand

Working with the marketing team of our field partner, a wine company, we reproduce how marketing professionals use AI to generate advertising prompts and then video ads. Despite gender-neutral instructions to maximize conversion, the pipeline generates male characters in 44% of basic-wine ads and 88% of premium-wine ads. We ask whether this gender–price association reflects professional judgment and helps the merchant maximize sales.

We benchmark AI's choices against an incentive-compatible experiment in which business owners and marketing practitioners select ads to maximize conversion. We use AI to reverse the character's gender in each video ad while holding other creative elements fixed. Randomly assigning the resulting male- and female-character versions across respondents identifies how character gender affects professional selection. Professionals' male-ad selection rate is 4.7 percentage points higher for premium than basic wine, compared with a 44-percentage-point gap in AI's male-character generation rate. AI thus exhibits a substantially stronger gender–price association than the human benchmark.

We then compare store visits and purchases generated by matched male- and female-character ads on TikTok to test whether AI's default gender choices maximize conversion. Our campaigns saturate the advertising cells, with matched gender-flipped ads randomly assigned across cells under equal budgets and reach optimization. For premium wine, the female-character version generates 49% more store visits and 42% more purchases per two-second video view than the male-character version. This reversal shows that AI's default gender choice is not conversion-optimal in our setting: its strong premium–male association directs creative production toward a less effective ad. A visible AI disclosure increases conversion, while adding information about AI's demographic associations reduces conversion relative to AI disclosure alone, linking disclosure design to the sales performance of AI-generated ads.

2. Generative AI and the Boundaries of Scientific Collaboration

With Ke Rong, Tianxi Yang, and Feng Zhu

Scientific production increasingly relies on teams to combine specialized expertise (Jones, 2009). Generative AI may substitute for collaborators' language support while helping researchers understand unfamiliar methods and engage with expertise outside their own fields. We distinguish substitution for language services from

complementarity with research expertise, asking how AI changes the human contributions assembled in scientific teams.

Using SSRN author–quarter panels from 2020 through 2026, we compare changes around ChatGPT 3.5's release across language and specialization groups. Non-native English-speaking researchers' relative decline in native English-speaking collaborator share predates the release, leaving the language-substitution prediction unsupported. Specialists, however, show larger postrelease increases than generalists in collaborator research-background distance and team knowledge breadth. These changes involve more distant new collaborators and greater participation by generalists, consistent with AI helping specialists integrate unfamiliar expertise.

Linked OpenAlex records associate broader team expertise with journal publication within 12 months of posting and with citations. Authors with stronger prior publication records also obtain larger relative publication and citation gains after ChatGPT's release. The findings suggest that AI can complement human expertise and broaden knowledge combinations while advantages associated with prior research output persist.

3. Generative Search, Consumer Choice, and Merchant Reach

With Mingduo Zhao and Yeyu Zhang

Search costs shape which merchants consumers consider and choose (Stigler, 1961; Weitzman, 1979). We study whether agentic AI reduces these costs by comparing alternatives before consumers visit merchant pages, reshaping the purchase funnel and merchant reach. Meituan is a Chinese platform that combines Yelp-style listings and reviews with Amazon-style purchasing and local delivery. Its agentic AI assistant accesses internal merchant and review databases to clarify needs, compare options, and assist with ordering. Linked platform records capture the full search–click–purchase funnel.

Using the agentic AI assistant's rollout, we compare users with and without access. Preliminary estimates associate access with less manual browsing, higher click-to-order conversion, and click share shifting from the largest merchants toward mid-sized alternatives, consistent with earlier screening and broader merchant consideration. Conversation-level analyses examine how initial intent, clarification, and recommendations relate to return visits and purchases across sessions, including whether early recommendations predict eventual merchant choice and whether nontransactional exchanges predict later shopping.

Interest in the tutorial

The tutorial would help me sharpen three questions: whether mitigating demographic associations in AI improves consumer and firm outcomes; how AI substitutes for or complements human expertise in scientific collaboration; and how agentic search changes consumer search costs and merchant competition. I would value feedback on disclosure experiments and consumer responses, identifying changes in the value of human expertise, and separating improved consumer search from platform steering. These questions closely match the organizers' expertise in digital advertising, the economics of AI, and platform design. I would contribute evidence from laboratory and field experiments, scientific collaboration records, and platform search data, alongside experience building machine learning systems for digital advertising.