Liying Qiu

Ph.D. Candidate

412-616-4397 · liyingq@andrew.cmu.edu · https://liyingqiu.github.io

Education

Carnegie Mellon University 2027 (expected)
Ph.D. in Business Technologies & Marketing (Minor in Machine Learning)
M.S. in Business Technologies

Duke University 2020
M.A. in Economics

Peking University 2018
B.S. in Finance

Research Interests

Substantive: consumer search, personalization, pricing, behavioral economics

Methodological: generative models, reinforcement learning, structural models

Publications

Qiu, L., Huang, Y., Singh, P. V., & Srinivasan, K. (2025). Personalization, consumer search, and algorithmic pricing. ***Marketing Science***, 44(6), 1278-1298.

Working Papers

Qiu, L., Singh, P. V., Mehta, N., & Srinivasan, K. A structural model of consumer search leveraging generative AI. [Job Market Paper]

- *Best Presentation Winner in LLM and Text Mining Flash Session, ISMS Marketing Science Conference 2025*

Qiu, L., Singh, P. V., & Srinivasan, K. Testing theoretical alignment in large language models: Evidence from insurance choices.

- *Major Revision at* ***Information Systems Research***

Work in Progress

Qiu, L., Zhang, S., & Singh, P. V. Are simpler machine learning models fairer? Evidence from a large-scale gamification experiment in banking.

Fellowships, Awards, and Honors

Prabuddha De and Family Fellowship (CMU)	2026
47th Annual ISMS Marketing Science Doctoral Consortium	2025
Dipankar and Sharmila Chakravarti Doctoral Fellowship Award (CMU)	2025
46th Annual ISMS Marketing Science Doctoral Consortium	2024
Henry J. Gailliot Presidential Fellowship (CMU)	2024
Graduate Student Assembly/Provost Conference Funds (CMU)	2024
Henry J. Gailliot Presidential Fellowship (CMU)	2023
PNC Research Assistantship Award (CMU)	2023
Litzenberger Family Fund (CMU)	2022
Tepper Dean's Fund (CMU)	2021
William Larimer Mellon Fellowship (CMU)	2020
Merit-based Tuition Award (Duke)	2018
Highest Distinction Undergraduate (PKU)	2018
Tan-Siu-Lin Scholarship (PKU)	2017
National Scholarship (PKU)	2016
CF40 & RK Scholarship (PKU)	2015
Kwanghua Scholarship (PKU)	2015
First-Class Freshman Scholarship (PKU)	2014

Conference Presentations

INFORMS Annual Meeting	2025
47th Annual ISMS Marketing Science Conference	2025
46th Annual ISMS Marketing Science Conference	2024
Workshop on Information Systems and Economics	2023
Conference on Information Systems and Technology	2023
Air Liquide Data Science Team	2023
Wharton Business & Generative AI Conference	2023
45th Annual ISMS Marketing Science Conference	2023
Production and Operations Management Society Annual Conference	2023

Workshop on Information Technology and Systems	2022
Conference on Information Systems and Technology	2022
44th Annual ISMS Marketing Science Conference	2022

Teaching Experience

Instructor
Math Skills Workshop (*MBA*)

Teaching Assistant
Data Exploration and Visualization (*MBA, Undergraduate*)
Digital Marketing and Social Media Strategy (*MBA, MSBA, Undergraduate*)
Modern Data Management (*MBA*)
Economics of AI in Business Research (*PhD*)

Guest Lecturer
Digital Marketing and Social Media Strategy (*Undergraduate*)
Structural Models & Quantitative Methods (*PhD*)

PhD Coursework

Economics
Behavioral Economics, Causal Inference, Computational Methods in Economics, Dynamic Discrete Choice, Econometrics, Game Theory, Industrial Organization, Information Economics, Mechanism Design, Microeconomics, Structural Econometrics

Machine Learning
Advanced Deep Learning, Advanced Natural Language Processing, Computer Vision, Convex Optimization, Deep Reinforcement Learning, Introduction to Machine Learning, Machine Learning Ethics and Society, Machine Learning and Game Theory, Probabilistic Graphical Models, Prompt Engineering

Statistics & Business
Economining, Estimating Dynamic and Structural Models, Bayesian Inference & Decision, Bayesian Statistics in Marketing, Human-AI Interactions in Business Decision Making, Human and Algorithmic Bias, Proseminar in Quantitative Marketing, Structural Models & Quantitative Methods

Media Article

Qiu, L., Malik, N. 6 Reasons Why Prices Change in Real Time. ***Tepperspectives***, February 2026

Service

Reviewer: Management Science, WITS 2022, CIST 2024, CIST 2025, CIST 2026, ICIS 2026

Student Coordinator: Conference on AI and Human Decisions 2026

Skills

Programming: ArcGIS, C, MATLAB, Python, R, SQL, Stata

Languages: Chinese, English

Referees

Kannan Srinivasan (co-chair)
H.J. Heinz II Professor of Management, Marketing and Business Technology
Carnegie Mellon University
kannans@andrew.cmu.edu

Param Vir Singh (co-chair)
Carnegie Bosch Professor of Business Technologies and Marketing
Carnegie Mellon University
psidhu@cmu.edu

Nitin Mehta
Ellison Professor of Marketing
University of Toronto
Nitin.mehta@rotman.utoronto.ca

Yan Huang
Associate Professor of Business Technologies
Carnegie Mellon University
yanhuang@cmu.edu

Selected Abstracts

A structural model of consumer search leveraging generative AI

Liying Qiu, Param Vir Singh, Nitin Mehta, Kannan Srinivasan (Job Market Paper)

Abstract: Independent sellers on digital platforms face a fundamental information disadvantage in making marketing decisions: while rich individual-level consumer search journeys, including queries, clicks, comparisons, and purchases, exist in click-stream data, platforms typically provide only aggregated metrics to sellers, such as query-level click shares and purchase shares. These aggregate measures are useful for monitoring performance but are insufficient for directly recovering the consumer preferences and search costs that drive search and purchase behavior. To bridge this information gap, we propose a lightweight and scalable agentic framework that transforms seller-visible aggregates into the specific underlying consumer preferences and search costs governing search and purchase decisions. The framework is designed for seller-side implementation: it relies only on information commonly available from platform dashboards and does not require sellers to observe individual-level search trajectories, estimate complex structural models, or maintain large-scale proprietary consumer databases. The framework first uses a feature-extractor agent based on general-purpose large language models (LLMs) to identify utility-relevant variables. These extracted variables, together with seller-visible aggregate measures, are then passed to a fine-tuned preference-generator agent that outputs consumer preferences and search costs. We fine-tune the preference-generator using detailed consumer search trajectories spanning 696 categories, 27,064 products, and 106,632 consumers. Our results show that the framework can recover consumer preferences and search costs in both future-period market holdouts and new product categories excluded from training, suggesting that the model learns a transferable mapping from aggregate dashboard information to consumer search primitives. Counterfactual experiments further demonstrate that sellers can make better marketing decisions using our framework than using either their own purchase data or dashboard data alone. Overall, this paper shows that generative AI can bridge the gap between rich platform-side consumer data and limited seller-side information, providing independent sellers with lightweight, scalable, interpretable, and decision-relevant insights for marketing decisions.

Personalization, consumer search, and algorithmic pricing

Liying Qiu, Yan Huang, Param Vir Singh, Kannan Srinivasan (Published in **Marketing Science**)

Abstract: Our study investigates the impact of product ranking systems on artificial intelligence (AI)-powered pricing algorithms. Specifically, we examine the effects of "personalized" and "unpersonalized" ranking systems on algorithmic pricing outcomes and consumer welfare. Our analysis reveals that personalized ranking systems, which rank products in decreas-

ing order of consumer's utilities, may encourage higher prices charged by pricing algorithms, especially when consumers search for products sequentially on a third-party platform. This is because personalized ranking significantly reduces the ranking-mediated price elasticity of demand and thus incentives to lower prices. Conversely, unpersonalized ranking systems lead to significantly lower prices and greater consumer welfare. These findings suggest that even in the absence of price discrimination, personalization may not necessarily benefit consumers because pricing algorithms can undermine consumer welfare through higher prices. Thus, our study highlights the crucial role of ranking systems in shaping algorithmic pricing behaviors and consumer welfare.

Testing theoretical alignment in large language models: Evidence from insurance choices

Liying Qiu, Param Vir Singh, Kannan Srinivasan (Major Revision at **Information Systems Research**)

Abstract: Large Language Models (LLMs) are increasingly proposed as tools for simulating consumer decision-making in business research. While their outputs often resemble human choices, it remains unclear whether these decisions reflect coherent applications of formal decision theories or arise from alternative statistical regularities. This paper introduces the Theory-Based Chain of Verification (TBCV), a diagnostic framework to assess whether LLM-generated decisions are internally consistent with specified theoretical models. TBCV proceeds in two steps: inferring decision-theoretic parameters or heuristics (e.g., loss aversion) from model choices, and verifying whether those parameters predict consistent behavior under systematically perturbed scenarios. To illustrate, we apply TBCV to GPT-4-turbo using a dataset of real-world home insurance choices that are well explained by Prospect Theory, providing a strong benchmark for theoretical alignment. While GPT-4-turbo's choices approximate aggregate human patterns, 85% of its follow-up decisions violate the Prospect Theory parameters implied by its prior selections. In contrast, its behavior is highly consistent with Extremeness Aversion, systematically favoring compromise options across choice sets. TBCV offers a principled method to distinguish between different forms of alignment, enabling more rigorous use of LLMs in business research and economic modeling.

Are simpler machine learning models fairer? Evidence from a large-scale gamification experiment in banking

Liying Qiu, Shunyuan Zhang, Param Vir Singh (Work in Progress)

Abstract: The increasing complexity of machine learning (ML) models presents challenges in terms of interpretability, thereby fueling the demand for inherently interpretable, simpler

models. In this paper, we utilize data from a large-scale gamification randomized control trial (RCT) within the banking sector to examine the trade-off between ML model simplification and fairness objectives. Through empirical analysis, we demonstrate that simpler models not only exhibit reduced accuracy but also demonstrate increased bias. Our findings reveal that when simpler models are utilized, valuable information that specifically benefits protected groups is disregarded. Furthermore, our results underscore that simpler models intensify the "impossibility of fairness" dilemma. As simpler models are made to satisfy one fairness criteria through the enforcement of constraints, they result in a heightened violation of other fairness criterion. Specifically, when a simple model is designed to achieve fairness criterion of equal opportunity (statistical parity), it significantly compromises alternate fairness criterion of statistical parity (equal opportunity) when compared to complex models. Our findings emphasize the critical need to carefully consider the trade-offs between model simplification and fairness goals in order to prevent the exacerbation of disparities across different fairness criteria.