

I work at the intersection of economics and machine learning. My research focuses on flexible models for learning causal relationships, to help understand the effects business decisions have on consumer choices. I apply a diverse skillset to drive business impact—designing experiments, implementing ML models, analyzing data, creating dashboards, and presenting results to internal and external stakeholders.

Work Experience

Arena-ai, Machine Learning Scientist **2022–Present**

- Helped build generative model of customer behavior to automate pricing, sales force, and marketing decisions at large enterprise companies. Increased profitability in randomized controlled trials by more than \$20m.
- Set up and designed experiments to measure impact (geo tests, synthetic controls, A/B testing).
- Built pipelines, analyzed data, and constructed dashboards to track model health and business goals.

Instacart, Senior Machine Learning Engineer II (L6) **2020–2022**

- Designed system for exploration and optimization of price markups to support retailer partners.
- Used heterogeneous treatment effect estimation to target promotions to customers and shoppers to balance market and drive cost effective growth.
- Collaborated with search and ranking team to incorporate prices, promotions and predicted inventory in rankings to balance revenue and customer engagement. Generated a 0.7% increase in profit per order.

Facebook, Research Scientist (L5) **2019–2020**

- Worked on a wide variety of empirical, strategic, and financial problems related to the design and operation of the Libra payments network. Analyzed pricing and competition in remittance market. Combined pricing and transaction data from existing stablecoins to calibrate simulations of coin liquidity on Libra network.

Microsoft Research, Research Assistant / Summer Intern **2012–2014, 2017**

- Estimated empirical model of advertiser behavior in online search ad auctions to infer advertiser preferences and predict their responses to changes in the platform. Worked on automated ML-based causal inference.

Cornerstone Research, Senior Analyst **2009–2012**

- Conducted economic and financial analyses to support expert witness in complex corporate litigation.

Technologies and Languages

Languages:	Python, SQL, R, Julia, C#
Technologies:	PyTorch, Presto, Pandas, Polars, data.table, AWS, Git
Statistics:	Causal Inference, heterogeneous treatment effects, econometrics, double ML

Education

PhD Business Economics, Stanford Graduate School of Business **2014–2019**

Coursework: Artificial Intelligence, Machine Learning, Machine Learning for Causal Inference, Convex Optimization, Modern Applied Statistics, Bayesian Statistics, Econometrics, CS106X

BA Mathematics, Pomona College. **2005–2009**

Research

Welfare effects of personalized ranking

with Ilya Morozov and Ayush Kanodia

Published in Marketing Science

Estimated a latent-factorization-based model in PyTorch that learns customer preferences from search and purchase data at a large online retailer. Applied model to measure the impact of personalized rankings on customer welfare and retailer profitability.

Counterfactual Inference for consumer choice

with Francisco Ruiz, David Blei and Susan Athey

Published in Quantitative Marketing and Economics

Compared the ability of econometric and machine learning based models to accurately predict out-of-sample responses to price changes and stockouts in a grocery store. Demonstrated the effectiveness, speed, and interpretability of latent-factorization-based models. Evaluated the potential gains from using flexible models to target personalized marketing and pricing campaigns.

Estimating heterogeneous preferences for restaurants and travel time

with Susan Athey, David Blei, Francisco Ruiz, and Tobias Schmidt

Published in AEA: Papers and Proceedings

Estimated individual-specific consumer preferences for restaurants using Yelp.com data combined with detailed geospatial tracking data. Predicted patterns of spatial competition in a dense urban region with thousands of consumers and restaurants. Used variational Bayesian approach for scalable inference.

Value of data for price targeting

with Susan Athey, Ayush Kanodia, Aaron Kaye, and Mitchell Linegar

Working paper presented at ASSA 2022

An investigation of the “scaling laws” for personalized pricing. Measured changes in incremental profit generated by targeted discounts as the number of customers, products, and length of data history in the training data increases.

Teaching

Applied Causal Inference

2022–2023

Developed and taught a 12-hour online course on causal inference for data scientists. The curriculum covered topics such as propensity scores, synthetic controls, heterogeneous treatment effects, and double ML.

Big Data, Strategic Decisions

2018–2019

Assisted professor in creating 2 new case studies for executive MBA course on data science and machine learning.

Hobbies

Swing dancing. Running. Finding edge cases and loopholes in complex systems.