Mohammad Mosaffa

Ph.D. Student in Quantitative Marketing, Cornell University
Cornell Tech and SC Johnson College of Business

New York, NY
mm3322@cornell.edu

Application: NBER Digital Economics and AI Tutorial: January 2027

GRADUATE STUDY AND ENROLLMENT

I am a fourth-year Ph.D. student in Quantitative Marketing at Cornell University. I began the program in August 2023 at the SC Johnson College of Business and have been based at Cornell Tech since January 2025. During my first two years, I completed graduate coursework in both economics and computer science. On the economics side, my training includes microeconomic theory, Econometrics I and II, applied game theory, structural models, and choice models, complemented by subsequent coursework in dynamic choice and counterfactual reasoning at the University of Washington. On the computer science side, I completed courses in causal inference with machine learning, deep generative models, computer vision, natural language processing, and the structure of information networks. Previously, I earned an M.Sc. in Mathematics, specializing in statistical machine learning, from the University of British Columbia.

RESEARCH INTERESTS AND DOCTORAL DIRECTION

I study how information and incentives shape economic decisions in digital markets. My research examines how firms value consumer data, how news outlets influence political perceptions through their content choices, how audiences respond to political discussions, and how AI intermediaries transform consumer search and purchasing decisions. I combine economic theory, causal inference, and machine learning to investigate these questions using structured behavioral data and unstructured data, including images and text.

My working dissertation title is *Machine Learning and the Economics of Information in Digital Markets*. My doctoral committee consists of Omid Rafieian, Emaad Manzoor, and Sachin Gupta (Cornell University) and Hema Yoganarasimhan (University of Washington).

RESEARCH PROJECTS

The Privacy–Utility Trade-Off of Location Tracking in Ad Personalization

(with Omid Rafieian; extended abstract at *ACM Economics and Computation 2026*; full manuscript, revise and resubmit at *Management Science*)

Digital advertising relies on the premise that more consumer data improve personalization. Yet consumers leave both digital footprints through behavioral data and physical footprints through geographical data, raising the question of how these two sources of information complement or substitute for each other. Using more than 10 million advertising impressions, we compare counterfactual targeting policies that use geographical data, behavioral data, both, or neither. We find that geographical data and behavioral data are complementary when behavioral data are sparse, with almost a 20% improvement in targeting performance during the cold-start stage. As behavioral data accumulate, they increasingly substitute for geographical data, leaving little incremental value from continued collection. Thus, the privacy–utility trade-off is *dynamic*: collecting geographical data early may improve targeting, while continued collection yields diminishing benefits. These findings also have implications for competition, as newer platforms with limited behavioral data may depend more heavily on geographical data than established platforms with richer behavioral data.

Visual Polarization Measurement Using Counterfactual Image Generation

(with Omid Rafieian and Hema Yoganarasimhan; extended abstract at *ACM Economics and Computation 2025*; accepted at the *NeurIPS 2026 ML×OR Workshop*)

News outlets convey political slant not only through words but also through the images they select to portray politicians. While existing research has extensively studied verbal polarization, measuring visual polarization remains challenging due to the high dimensionality of images. We develop a unified framework combining economic theory, generative models for counterfactual image generation, and multi-modal deep learning. We model editors' image choices through a utility framework to define visual slant and polarization, generate counterfactual images that vary politicians' facial expressions while preserving other characteristics, and link our measures to a news outlet prediction problem. Applying our framework to images of 30 politicians across 20 major news outlets, we document significant ideological slant and substantial heterogeneity in visual polarization. We validate our measures against non-image-based measures of media slant and document an increasing trend in visual polarization, particularly since 2016.

ONGOING RESEARCH

Building on my research on personalization and political polarization, my ongoing work focuses on the economics of large language models (LLMs) and AI. I study how these technologies can support causal inference in digital settings and how AI systems use information, respond to incentives, and make economic decisions.

In *Identifying Dynamic Causal Effects with High-Dimensional Unstructured Data*, I study causal inference methods that rely on outcome models, such as synthetic control, treatment-assignment models, such as propensity-weighted DiD, or either model, as in doubly robust approaches. In settings involving high-dimensional unstructured data, I show why identification may require information from both models simultaneously. I propose *Propensity-Anchored Synthetic Difference-in-Differences (P-SDiD)*, which combines propensity-based weighting with pretreatment outcome matching to address these complementary identification challenges. I apply the estimator to study how political discussions causally affect audience attention in YouTube podcasts.

In *From Engagement to Causal Rewards: Aligning Large Language Models with Audience Attention*, we study how conversational content affects audience attention in long-form YouTube podcasts. Since creators choose both what to discuss and when, observed engagement may reflect content selection rather than its causal effect on attention. We investigate how causal measures of audience response can inform the design of reward functions for large language models, moving beyond correlational engagement metrics. Our broader goal is to understand how attention-based incentives shape content production and how causal evidence can guide AI-generated content.

In a separate project, *Revealed Preferences of AI Agents*, we study how AI agents learn consumer preferences and make personalized recommendations across domains, from movies to consumer products. Unlike traditional personalization, which relies on collecting users' behavioral histories, AI agents can directly elicit preferences through questions and voluntary information sharing. We investigate whether this approach can substitute for extensive data collection while maintaining recommendation quality. More broadly, we examine the economic value of actively elicited information and its implications for personalization, consumer privacy, and data collection.

INTEREST IN THE TUTORIAL

I hope to attend the NBER Digital Economics and AI Tutorial to exchange ideas with researchers studying how digital technologies and AI shape economic behavior and market outcomes. I would particularly value discussions on how to connect evidence from high-dimensional images and text data, as well as AI agents' choices, to economic theory and causal mechanisms. These conversations would help me sharpen the economic questions in my ongoing work on causal inference with unstructured data and the behavior of AI agents.