# NBER Digital Economics and AI Statement

Calvin Jahnke, Princeton University
4th Year, PhD in Economics
Industrial Organization
jahnke@princeton.edu
+1(920)574-5674

## 1 Research Interests

In my current work, I study two digital economies: open source software and the attention economy. Artificial intelligence factors into both of these projects, in allowing us to measure text in ways not previously possible and in being the focus of structural models that allow us to probe opaque recommendation algorithms. Machine learning is also forcing us to ask new questions, like how to measure market power in an attention economy that lacks traditional pricing structures and is increasingly powered by machine learning models.

The program will directly benefit my current research projects. For our feedback style paper, we study the role of toxicity in feedback from maintainers on the Linux Kernel; a setting unique from traditional workplaces, understanding how feedback functions in this project is key to understanding how open source software is created and maintained, especially as it grows in importance in the digital economy. In our second project, we develop a novel method for measuring platform market power in a setting where there are no traditional prices over which to estimate mark-ups. Having developed this novel methodology, we are working towards producing policy-relevant results. Understanding the current frontier of approaches to understanding digital economies and AI will help us position this work to make meaningful contributions and generate further inquiries into these systems.

The tutorial will also lay the foundation for my future research. I plan to extend my work bringing structural modeling to algorithmic systems, applying these techniques to study platform design, the substance of recommended content, and digital market competition. I am also interested in understanding prediction markets, how they work, and what we can learn about the expectations of market participants through their betting behavior. The program will also help me forge connections with researchers with similar interests, and I look forward to conversations and collaboration beyond the Institute with the other participants.

My research exhibits the bidirectional relationship between machine learning and economics: using ML to measure and understand economic phenomena while applying economic theory to make algorithmic systems more transparent and accountable. This approach addresses critical policy challenges at the intersection of AI and digital markets, from platform regulation to algorithmic fairness. The NBER tutorial will provide essential exposure to methods and connections with researchers tackling similar questions, positioning me to make meaningful contributions to both fields.

# 2 Current Research

**Recommendation Algorithms and Welfare in the Market for Attention**
*with Vartan Shadarevian*

**Abstract**

We study the welfare implications of recommendation algorithms in attention markets, documenting how these machine learning systems strategically manipulate search costs to increase platform engagement. Combining randomized audits with structural estimation, we develop a novel methodology that infers users' underlying preference parameters without requiring sensitive user data. In over 2,000 systematic audits of a video-sharing platform, we find that recommendations systematically respond to multiple dimensions of user behavior in ways that prioritize engagement over user value. Simulations from our calibrated model demonstrate that engagement-maximizing recommendation strategies favor significantly longer videos than would maximize user welfare, and the engagement-value gap shrinks as the value of the outside option grows. By comparing engagement-maximizing and value-maximizing strategies, our approach provides a measure of platform market power in settings without traditional prices. Our structural audit methodology also offers regulators a template for monitoring algorithmic behavior without requiring platform cooperation or access to proprietary systems.

**Status**

This work exists as a draft submitted to fulfill the third-year paper requirement for Princeton University's Economics PhD program. The paper includes results from our random audit of the platform's recommendation algorithm, demonstrating sophisticated platform responses to randomly-generated user histories. By modeling user demand and the platform's response to user choices, we show how we can back out the platform's priors over parameters governing user demand, extracting meaningful data about users from the algorithm. Our simulations illustrate the incentive divergence between platform and users, with the platform systematically pushing users toward longer-duration content. We are improving the project with an updated model of user content demand, a stratified sampling procedure for the platform audits, and an improved estimation procedure. I presented this work at the 2025 Machine Learning in Economics Summer Conference at the University of Chicago.

While machine learning is often seen as a tool to advance the study of economics, this paper introduces a method for using economics to advance understanding of machine learning models. In addition to developing a method to use trained machine learning algorithms for demand estimation, we show how structural models can help researchers understand the behavior of black-box algorithms.

**Manager Feedback Style and Worker Productivity**
*with Gàbor Nyèki and Carolyn Tsao*

**Abstract**

This paper studies how a manager's tone when giving feedback to workers affects individual productivity and output quality. We construct a novel panel dataset that links software engineers and managers to their email communications and code contributions on the largest open source software project, the Linux kernel. We identify tones used in the emails (e.g., toxic, polite, encouraging) using natural language processing and machine learning techniques. We find a strong negative relationship between manager toxicity and engineer productivity. Using an instrumental variables design to address endogeneity in a manager's choice of tone, we find that receiving toxic feedback from a manager reduces the likelihood that an engineer completes a programming task, increases the amount of time to task completion, and decreases the likelihood that an engineer completes more tasks in the next 30 days.

**Status**

This work exists as a completed extended abstract including results from our main specification where we implement a judge-IV design to predict message toxicity using a reviewer's other messages as an instrument. We plan to add robustness to the results by leveraging engineers who move between reviewers and to expand the scope of the results by estimating effects on engineers' career outcomes. I presented this paper at the *IZA PhD Workshop in Labor and Behavioral Economics* in June 2025.

This paper employs a range of Natural Language Processing techniques to quantify the toxicity of code review messages, including BERT-based models developed in the computer science literature to measure the toxicity of millions of messages posted to the Linux kernel. To validate these measures and capture additional tone dimensions, we develop n-gram based measures.

For our ongoing work, I fine-tune a GPT-2 language model on the code review corpus. I then leverage the next-token probability distribution to extract information-theoretic metrics. Specifically, I measure entropy over the token distribution and surprisal (negative log-likelihood) of the selected tokens. These metrics provide both controls to isolate toxicity effects and serve as interpretable measures of feedback informativeness, demonstrating a novel application of language models beyond their typical predictive use cases.