Alex Moehring
919-606-3475
moehring@mit.edu
alexmoehring.com

100 Main St (E62-365)
Cambridge, MA 02142

Dear Organizers,

I am extremely excited to apply for the 2023 Digitization tutorial. As a brief introduction, I am a fifth year PhD student in Information Technologies at MIT Sloan advised by Prof. Catherine Tucker. My current research examines questions related to social media platforms and how humans and algorithms interact. During the first two years of my program, I focused on coursework that emphasized applied microeconomics, industrial organization, and applied econometrics. In addition, I have taken courses in statistics and machine learning.

To better describe my research interests, I will briefly discuss three ongoing projects that focus on digital platforms and AI. I am confident attending the tutorial would greatly improve my current projects in addition to providing an opportunity to interact with graduate students from other institutions interested in similar topics.

An Empirical Model of the Reddit News Feed

In this project, which I am developing for my job market paper, I seek to understand the tradeoffs faced by social media platforms when optimizing for user engagement. To be more specific, I use data from Reddit to estimate a discrete choice model of engagement on the platform and estimate the counterfactual where Reddit implements a personalized engagement maximizing ranking of posts. A key identification challenge in this project is credibly modeling the impact a post's newsfeed rank has on future engagement, as we are worried that potential outcomes (or latent post quality) are correlated with a post's rank. To overcome this issue, I leverage Reddit's public mirror of its codebase they maintained until 2017. This allows me to inspect the ranking code, reconstruct the score used to rank posts, and identify the effect of rank on engagement using a regression discontinuity design.

With these RD based instruments for a post's rank, I then estimate a discrete choice model of engagement and simulate the counterfactual ranking where each user is shown their personalized engagement maximizing ranking on a large subreddit specializing in political news. Preliminary results suggest that this counterfactual reduces the diversity of news across the political spectrum that users are exposed to and engage with. I also intend to study the sentiment of the comments and the credibility of the publishers that users are engaging with, though these results are still in development.

News Feeds and User Engagement: Evidence from the Reddit News Tab

In this project I study how the introduction of a new non-personalized news feed impacts user engagement quantity, quality, and diversity on Reddit. In June 2018, Reddit introduced the News tab on iOS devices that surfaces popular content from a curated list of news-related communities. We leverage this natural experiment to identify the causal effects of the News tab on iOS user engagement in a

difference-in-differences design. We find that the News tab increases the share of iOS devices that engage with news-related content and the new content induced by the News tab is not meaningfully lower in quality relative to the existing comments. We also find that the diversity of engagement within news categories and within articles from publishers across the political spectrum increases as a result of the News tab. These results suggest that non-personalized feeds can be an important tool to mitigate algorithmic filter bubbles and need not come at the expense of reduced user engagement.

Human Experts and Artificial Intelligence: The Value of Human Input in Diagnostic Imaging
Joint with Nikhil Agarwal, Tobias Salz, and Pranav Rajpurkar

This is an experimental project that aims to study how radiologists incorporate AI assistance in their decision making. Over the past year we have been running and evaluating an experiment to test (1) how important contextual information is in evaluating the relative performance of humans and AI, (2) how well humans incorporate the information contained in the AI assistance (i.e. do they overweight or underweight the AI signals), and (3) how to design the collaboration between humans and AI. We have constructed an experimental interface that elicits radiologist assessments and diagnostic decisions under randomly assigned information environments to test such hypotheses. Relative to the existing literature, we also have estimated an economic model that will allow us to identify how human experts are incorporating AI predictions and understand the extent to which they over- or underweight the AI signals in their assessments. We find that, relative to the Bayesian benchmark, radiologists underweight AI signals in their assessments. In addition, radiologists exhibit another bias by treating the AI signal as conditionally independent of their own. These errors combined lead result in radiologists exhibiting worse decisions when given the AI assistance under certain signal realizations and a small improvement on average.

I would like to thank you for organizing the tutorial and your consideration.

Sincerely,

Alex Moehring