

Information Technology: Doctoral Theses

“Essays on Online Platforms and Human-Algorithm Interactions ”

Author: Alex Moehring (2024)

Committee: Catherine Tucker (chair), Nikhil Agarwal, Dean Eckles

Abstract:

This dissertation contains three chapters that analyze how algorithms on social media platforms influence the content that users engage with and how individuals incorporate algorithmic predictions in their decision-making. In Chapter 1, I study how engagement maximizing news feed algorithms on social media affect the credibility of news content with which users engage. This allows me to estimate the extent to which engagement-maximizing algorithms promote and incentivize low-quality content. In addition, I evaluate how the ranking algorithm itself can be designed to promote and encourage engagement with high quality content. In Chapter 2, I analyze how the introduction of a new non-personalized news feed impacts user engagement quantity, quality, and diversity on the Reddit platform. I find that this auxiliary feed increases the share of users that engage with news-related content and the diversity of engagement within news categories and within articles from publishers across the political spectrum increases as a result of the feed. In Chapter 3, in collaboration with Nikhil Agarwal, Tobias Salz, and Pranav Rajpurkar, we study human-AI collaboration using an information experiment with professional radiologists. Results show that providing (i) AI predictions does not always improve performance, whereas (ii) contextual information does. Radiologists do not realize the gains from AI assistance because of errors in belief updating – they underweight AI predictions and treat their own information and AI predictions as statistically independent.

“Essays on Online Platforms and Human-Algorithm Interactions ”

Author: Zanele Munyikwa (2024)

Committee: John Horton (chair), Erik Brynjolfsson, Daniel Rock

Abstract:

This dissertation comprises three chapters that dissect the evolving interface between artificial intelligence (AI) and knowledge work, with a particular focus on language and writing-based AI technologies. Chapter 1 describes the scope, scale, and economic value of writing skills within the labor market, presenting descriptive statistics and employing hedonic wage analysis to estimate the salary premiums associated with writing skills. By analyzing administrative data and job postings, it underscores the economic significance of writing proficiency across various professional domains, emphasizing its crucial role in the contemporary workforce.

In Chapter 2, attention shifts to the direct influence of generative AI on knowledge work through a randomized field experiment in the copywriting industry. This analysis explores how AI-driven tools not only enhance content creation but also impact productivity, creativity, and workers’ subjective feelings of ownership over their work. The findings illustrate the transformative potential of AI in reshaping creative professions and significantly altering traditional workflows.

Chapter 3 investigates the effects of algorithmic resume writing assistance on hiring outcomes for knowledge workers in an online labor market, employing causal text analysis and mediation analysis to uncover the mechanisms through which AI influences hiring decisions. This chapter focuses on text as a mediator, examining how AI-induced adjustments to linguistic properties such as formality and error correction mediate the relationship between AI tool use and hiring outcomes. This examination reveals how subtle changes to linguistic properties can significantly affect job seekers’ success, underscoring the practical benefits and complexities of integrating AI into hiring practices.

Together, these chapters offer a comprehensive exploration of how AI technologies, particularly those focused on language and writing, are redefining the landscape of knowledge

work. By highlighting the interactions between writing skills, technological innovation, and employment, this dissertation sheds light on the critical role of writing in the contemporary job market and the significant impact of AI advancements on professional content creation.

“Artificial Intelligence in Labor Market Matching ”

Author: Emma Wiles (2024)

Committee: John Horton (chair), Catherine Tucker, Dean Eckles

Abstract:

In my dissertation I study three applications of AI in labor market matching. In my first chapter I show that AI-improved but not entirely written resumes make workers more likely to be hired with no negative downstream implications to employers or to match quality. However, in my second chapter I show that when employers are given entirely AI written drafts of a job post, the jobs posted are more generic and less likely to make a hire. Lastly, I provide evidence that non-technical workers can use AI to upskill into data science, however those skills do not persist in absence of AI assistance.

My first chapter investigates the association between writing quality in resumes for new labor market entrants and whether they are ultimately hired. I show this relationship is, at least partially, causal: in a field experiment in an online labor market with nearly half a million jobseekers, treated jobseekers received algorithmic writing assistance on their resumes. I find that the writing on treated jobseekers' resumes had fewer errors and was easier to read. Treated jobseekers were hired 8% more often, at 10% higher wages. Contrary to concerns that the assistance takes away a valuable signal, I find no evidence that employers were less satisfied with the quality of work done, using star ratings, the sentiment of their reviews, and their probability of rehiring a worker. The analysis suggests digital platforms and their users could benefit from incorporating algorithmic writing assistance into text-based descriptions of labor services or products without downstream negative consequences.

In my second chapter, I study a randomized experiment conducted on an online labor market that encouraged employers to use a Large Language Model (LLM) to generate a first draft of their job post. Treated employers are 20% more likely to post the job and decrease time spent writing their job post by 40%. Among the posted jobs, treated employers receive 5% more applications. Despite this, they are 18% less likely to hire. I find no evidence that this is driven by treated employers receiving lower quality applicants. Moreover, despite the large increase in the number of jobs posted, there is no difference in the overall number of hires between treatment and control employers. These results imply that the treatment lowered the probability of hiring among at least some jobs which would have otherwise made a hire. I rationalize these results with a model in which employers with heterogeneous values of hiring can attract better matches by exerting effort to precisely detail required skills. I show how a technology that lowers the cost of writing and imperfectly substitutes for effort causes more posts, but lowers the average hiring probability through both marginal posts (as these are less valuable) and inframarginal posts (as the technology crowds out effort and makes the job posts more generic). I provide evidence for these mechanisms using employer screening behavior and the embeddings of the job posts' texts.

In my third chapter, we investigate if LLMs can be used to help non-technical workers adapt to technology induced, rapidly changing skill demands by “upskilling” into a more technical skillset. With coauthors at Boston Consulting Group, we run a randomized control trial on knowledge workers, who have no data science experience, to test whether workers paired with LLMs are able to perform data science tasks to the level of real data scientists. We give consultants at BCG data science problems, representative of what the data scientist role at the company demands, but which GPT-4 cannot solve on its own. We find that treated workers given access to and training in using ChatGPT are more likely to correctly solve all three tasks, and can perform at the level of real data scientists without GPT-4 on the coding task. These results suggest that LLMs can be used to help workers gain new skills to meet the evolving, more technical demands of the labor market, but that for some types of tasks the work of non-technical workers is not interchangeable with data scientists'.

“Essays on Information Technology, Human Capital, and the Future of Work ”

Author: Sebastian Steffen (2022)

Committee: Erik Brynjolfsson (chair), John Horton, Sinan Aral, Prasanna Tambe (The Wharton School, University of Pennsylvania)

Abstract:

This dissertation contains three essays concerning the economics of information technology, human capital, and the future of work. In the first essay, 'Occupational Change: Automation and Reskilling Risks', I develop a methodology to study occupational skill demands and estimate the returns to skills, by leveraging novel data from over 200 million online job postings from 2010 until 2020. I find large heterogeneity in skill returns across industries and identify potential (re)skill investment opportunities for workers.

In the second essay, 'Digital Resilience: How Work-From-Home Feasibility Affects Firm Performance', I build on the methodology and data from the previous chapter to measure how feasible it is for firms to shift their workforce to remote work. Using these data, I then causally identify how much remote work practices aided firms' resilience against the Covid-19 pandemic, as measures by sales, net income, stock market returns, and volatility. The findings highlight that firms need to strategically manage the labor composition and digitization of their organizations, and consider that work-from-home practices, besides their many other advantages, are an effective way to hedge against operational risks.

In the final essay, 'Treating the Symptoms or the Cause? Substantive and Symbolic Talent Acquisition in Response to Data Breaches', I use the data from the first chapter to study firms' hiring responses to data breaches. Advancing the theory of substantive and symbolic IT adoption to complementary human capital acquisitions, I find that firms significantly increase their hiring for cybersecurity as well as public relations and legal workers after suffering breach. I also find that public scrutiny can serve as an effective mechanism to shift firms' hiring investments toward substantive, rather than symbolic measures. Given the increase in the volume and severity of cyberattacks, these results provide important and timely insights into firms' responses and incentives to more substantively safeguard their data.

“Essays on the Design of Online Marketplaces and Platforms”

Author: David Holtz (2021)

Committee: Sinan Aral (chair), Dean Eckles, John Horton

Abstract:

In Chapter 1, I estimate the impact of increasing the extent to which content recommendations are personalized. By analyzing the results of a randomized experiment on approximately 900,000 Spotify users across seventeen countries, I find that increasing recommendation personalization increased the number of podcasts that Spotify users streamed, but also decreased the individual-level diversity of Spotify users' podcast consumption and increased the dissimilarity between the podcast consumption patterns of different users across the population. Additional analysis shows that exposure to more personalized recommendations affected not only algorithmically-driven content consumption, but also the content that users consumed organically. The shifts in consumption diversity I observe can affect user retention and lifetime value, and impact the optimal strategy for content producers. These results indicate that personalized recommendations have the potential to create an "engagement-diversity trade-off" when firms optimize solely for consumption.

In Chapter 2, I propose methods for obtaining unbiased estimates of the total average treatment effect (TATE) when conducting experiments in online marketplaces, and test the viability of said methods using a simulation built on top of scraped data from Airbnb. The baseline approach to experimentation -- an individual-level, Bernoulli-randomized experiment analyzed using the difference-in-means treatment effect estimator -- is likely to yield biased TATE estimates when used in online marketplaces, due to, e.g., competition between sellers in the marketplace. The methods proposed in this chapter, such as graph cluster randomization and exposure modeling, draw on the existing literature on experimentation in networks, and depend on modeling the market as a network, in which an edge exists between two items if they might complement or substitute for one another. I find that blocked graph cluster randomization can reduce the bias of TATE estimates in online

marketplaces by as much as 64.5%, however, this reduction in bias comes with a substantial increase in root-mean-square error (RMSE). I also find that fractional neighborhood treatment response (FNTR) exposure models and inverse probability-weighted estimators have the potential to further reduce bias, depending on the choice of FNTR threshold. These results are robust across different treatment interventions, outcomes, levels of network mis-specification, clustering approaches, market structures, levels of demand, and data generating processes.

In Chapter 3, I conduct two large-scale meta-experiments on Airbnb in an attempt to estimate the actual magnitude of bias in TATE estimates from marketplace interference. In both meta-experiments, some Airbnb listings are assigned to experiment conditions at the individual-level, whereas others are assigned to experiment conditions at the level of clusters of listings that are likely to substitute for one another. The two meta-experiments measure the impact of two different pricing-related interventions on Airbnb: a change to Airbnb's fee policy, and a change to the pricing algorithm that Airbnb uses to recommend prices to sellers. Analysis of the fee policy meta-experiment reveals that at least 32.60% of the treatment effect estimate in the Bernoulli-randomized meta-experiment arm is due to interference bias. I also find weak evidence that the magnitude and/or direction of interference bias in online marketplaces depends on the extent to which a market is supply- or demand-constrained. Analysis of the pricing algorithm meta-experiment does not produce a statistically significant estimate of the magnitude of TATE estimate bias due to marketplace interference, but does highlight the difficulty of detecting interference bias when treatment interventions require intention-to-treat analysis.

“How Should We Measure the Digital Economy?”

Author: Avinash Collis (2020)

Committee: Erik Brynjolfsson (chair), Sinan Aral, Catherine Tucker

Abstract:

Gross domestic product (GDP) measures production and is not meant to measure well-being. While many people nonetheless use GDP as a proxy for well-being, consumer surplus is a better measure of consumer well-being. This is increasingly true in the digital economy where many digital goods have zero price and as a result, the welfare gains from these goods are not reflected in GDP or productivity statistics. Chapter 1 proposes a way of directly measuring consumer's economic well-being using massive online choice experiments. It finds that digital goods generate a large amount of consumer surplus that is currently not captured in GDP. For example, the median Facebook user needed a compensation of around \$48 to give it up for a month. Building up on these results, Chapter 2 extends the GDP framework to include welfare gains from new and free goods and construct a new metric called GDP-B, where B stands for benefits. It finds that including the welfare gains from Facebook would have added between 0.05 and 0.11 percentage points to GDP-B growth per year in the US. Chapter 3 proposes a way of measuring network effects on multi-sided platforms using choice experiments. It also models digital platforms allowing for heterogeneity in demand elasticity and network effects across users of different types. It then calibrates the model using an empirical application to Facebook and simulates six different taxation and regulatory policies. Chapter 4 looks at the impact of social media on subjective well-being and academic performance through a randomized controlled trial of University students. Chapter 5 summarizes the research agenda moving forward and concludes with a framework for measuring different aspects of well-being in the digital economy.