

Information Technology: Doctoral Theses

“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.

“Essays on Information Technology, Intangible Capital, and the Economics of Artificial Intelligence”

Author: Daniel Rock (2019)

Committee: Erik Brynjolfsson (chair), Andrew Lo, Chad Syverson, Prasanna Tambe

Abstract:

This dissertation discusses the role of intangible and technological investments in the dynamics of productivity growth, the valuation of firms, and the employment of human capital. The first essay describes the Productivity J-Curve. General purpose technologies (GPTs) such as AI enable and require significant complementary investments, including business process redesign, co-invention of new products and business models, and investments in human capital. These complementary investments are often intangible and poorly measured in the national accounts, even if they create valuable assets for the firm. We develop a model that shows how this leads to an underestimation of output and productivity in the early years of a new GPT, and how later, when the benefits of intangible investments are harvested, productivity will be overestimated. Our model generates a Productivity J-Curve that can explain the productivity slowdowns often accompanying the advent of GPTs, as well as the follow-on increase in productivity later. We use our model to assess how AI-related intangible capital is currently affecting measured total factor productivity (TFP) and output. We also conduct a historical analysis of the roles of intangibles tied to R&D, software, and computer hardware, finding substantial and ongoing effects of software in particular and hardware to a lesser extent.

The second essay investigates the role of engineering talent in predicting the market value of publicly traded firms, then looks in greater detail at AI Talent. Engineers, as implementers of technology, are highly complementary to the intangible knowledge assets that firms accumulate. This paper seeks to address whether technical talent is a source of rents for corporate employers, both in general and in the specific case of the surprising open-source launch of TensorFlow, a deep learning software package, by Google. First, I present a simple model of how employers can use job design as a tool to exercise monopsony power by partially allocating employee time to firm-specific tasks. Then, using over 180 million position records and over 52 million skill records from LinkedIn, I build a panel of firm-level investment in technological human capital (information technology, research, and engineering talent quantities) to measure the market value of technological talent. I find that on average, an additional engineer at a firm is correlated with approximately \$855,000 more market value. Consistent with that finding, AI-intensive companies rapidly gained market value

following the launch of TensorFlow. Using a difference-in-differences approach, I show that the launch of TensorFlow is associated with an approximate increase of 4-7% in firm market value for firms employing workers with AI skills compared to firms without AI talent exposure.

The third essay is about which tasks done by workers in the U.S. economy are Suitable for Machine Learning (SML). Advances in machine learning (ML) are poised to transform numerous occupations and industries. This raises the question of which tasks will be most affected by ML. We present a model of labor demand in the presence of new technology and labor constraints following Autor, Levy, and Murnane (2003). We then apply the rubric evaluating task potential for ML in Brynjolfsson and Mitchell (2017) and extended in Brynjolfsson, Mitchell, and Rock (2018) to build measures of “Suitability for Machine Learning” (SML) and apply it to 18,112 tasks in O*NET. We find that ML has the potential to affect many occupations in the economy, though few (if any) jobs can be completely automated by ML. We discuss the distribution of sensitivity to ML technologies across regions and industries, finding that the effects of ML will follow different patterns than earlier waves of automation.

The fourth essay discusses shifts in the employment of routine labor. A large literature has documented occupational shifts in the US away from routine intensive tasks. Theories of skill-biased technological change differ in whether they predict changes in occupational mix within firms, or merely across different firms or industries. Using LinkedIn resume records, BLS OES data, and Compustat employee counts, we estimate occupational employment for publicly traded US firms from 2000 through 2016. We find that faster employment growth among firms that disproportionately employ non-routine workers is the most important cause of SBTC, followed by within firm occupational mix rebalancing. The entry of new firms also plays a role, although firm exit is slightly routine-worker biased. R&D leads firms to have a larger share of routine workers. These results are most consistent with a theory of routine task demand reduction caused by the diffusion of infra-marginally implemented new technologies. We also introduce a new measure of business labor dynamism, capturing the frequency with which firms change their occupational mix. Consistent with trends in productivity and other measures of business and labor market dynamism, this measure has decreased steadily since 2000

“Three Essays on Social Influence in Social Advertising Using a Large-Scale Randomized Field Experiment”

Author: Shan Huang (2018)

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

Abstract:

In this thesis, I examine the causal relationships among products, social influence and network-embedded human behaviors, in the context of social advertising. Social advertising places social cues (e.g., likes) in ads, utilizing the power of social influence (the effects of social cues in ads) to encourage ad engagement. I collaborate with a social networking app for a large-scale randomized field experiment on its social ads. In the experiment, the presence and the number of social cues were randomly assigned among 57 million ad-user pairs (more than 37 million subjects and across 71 products in 25 product categories). Integrating the experimental evidence and the data of individuals, products, ads and network structures, my studies address the incentives, magnitude, contagion patterns and viral factors (i.e., characteristics of products, behaviors and individuals) of social influence in social advertising and product adoptions.

“Truce Structures: Examining Cross-Professional Coordination in the Wake of Technical and Institutional Change”

Author: Arvind Karunakaran (2018)

Committee: Wanda Orlikowski (chair), Katherine Kellogg, John Van Maanen, JoAnne Yates

Abstract:

This research examines the structures, processes, and mechanisms that facilitate cross-professional coordination during periods of technological and institutional change. My study draws on a 24-month ethnographic study, combined with historical data and quantitative analysis, of 911 emergency management organizations in the United States.

In Chapter 2, I focus on the mechanisms to facilitate cross-professional coordination in conditions that are marked by protracted jurisdictional conflicts. My findings articulate the importance of truce structures - an ensemble of truce roles and organizational forms - that are intended to address protracted jurisdictional conflicts among symmetrical professions such as police officers and firefighters. I further find that the coevolution of truce roles and organizational forms resulted in the emergence of a specific truce profession - in this case, that of 911 Public Safety Telecommunicators. The truce profession serves to triage, direct, and channel contested tasks among the conflicting professions without bringing those professions into direct contact with each during the initial stages of coordination when the "definition of the situation" is getting worked out.

In Chapter 3, I turn to examining how the truce professionals navigate what I call status-authority asymmetry in order to effectively coordinate with the focal professionals. Conducting within-shift comparisons of coordination encounters between 911 dispatchers and police officers, I identify that the bounded publicization tactic performed via the open radio channel allows dispatchers to generate peer knowledge about individual non-compliance. Through this process, dispatchers navigate the status-authority asymmetry and orchestrate effective cross-professional coordination.

My focus in Chapter 4 shifts to examining how truce professionals respond to the public's increased digital scrutiny, and consider the consequences for organizational accountability. My findings suggest that the public's increased use of mobile phones and social media to monitor and report on organizations and their workers can, under some conditions, end up worsening accountability. I unpack the processes that generate this paradox of public accountability, showing how these processes reshape the work of truce professionals and produce a vicious cycle of coordination that worsens organizational accountability. I end with a concluding chapter that discusses the implications of my dissertation for research on cross-professional coordination, accountability, and technological change.

“Operating in the Shadows: The Productive Deviance Needed to Make Robotic Surgery Work”

Author: Matthew Beane (2017)

Committee: Wanda Orlikowski (chair), Katherine Kellogg, John Van Maanen

Abstract:

Though a 2.5-year mixed-method study comparing robotic surgical practice to traditional surgical practice, I explore how crucial outcomes require productive deviance: norm- and policy-challenging practices that are tolerated because they produce superior outcomes in the work processes governed by those norms and policies. My empirical focus was fortunate - I show that productive deviance is likely especially important in the first ten to twenty years of significant technical reconfiguration of surgical work. I open my dissertation through a comparative empirical introduction to my context and a review of the literature on deviance in organizations.

The second chapter of my thesis is a history of how the surgical profession has relied on productive deviance for integrating new technologies since the early 1800s, ending with a deeper treatment on robotic surgery.

My third chapter focuses on how only a very few surgical residents managed to gain confidence and competence with robotic surgical methods given significant barriers to such learning. In contrast to what the

standing literature on learning would predict and in tension with the norms for learning within the surgical profession, these residents engaged in a suite of practices I call "shadow learning" - involving premature specialization, abstract rehearsal and undersupervised struggle. I explore how each of these practices both allow progress and create unintended negative consequences for the profession.

My fourth chapter explores a case in which surgical teams routinely used new, well-maintained robotic surgical devices and occasionally faced the stressful and practically difficult task of using an under-maintained, unreliable surgical robot. In this chapter, I show quantitatively that patients did just as well on the degraded robot, and I outline the often invisible, undervalued "compensatory work" that professionals did to ensure such outcomes. The main contribution here is to explicitly treat affect as integral to coordinated work that is grounded in suboptimal material arrangements. Through these studies, I solidify and enrich our conception of productive deviance and show how it is critical for a range of professional and organizational outcomes.