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 individuallevel, 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 1 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 rootmean-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.