

Marketing: Selected Doctoral Theses

“Essays on MarTech: Learning to Design, Deliver, and Diffuse Interventions”

Author: Jeremy Yang (2021)

Committee: Juanjuan Zhang (co-chair), Sinan Aral (co-chair), Dean Eckles

Abstract:

Chapter one develops an algorithm to predict the causal effect of influencer video advertising on product sales. A summary statistic, motion-score, or m-score, is proposed to capture the extent to which a product is advertised in the most engaging parts of a video. Pixel-level product placement is located with an object detection algorithm and pixel-level engagement is estimated as a saliency map by fine-tuning a deep 3D convolutional neural network on video-level engagement data. M-score is then defined as pixel-level engagement-weighted advertising intensity of a video. The algorithm is constructed and evaluated with influencer video ads on TikTok. Causal effects of video ads on product sales are identified by exploiting variation in video posting time. Videos of higher m-score indeed lift more sales. This effect is sizable, robust, and more pronounced among impulsive, hedonic, or inexpensive products. The mechanism can be partially traced to influencers’ incentives to promote themselves rather than the product. How various stakeholders in entertainment commerce can use m-score in a scalable way to optimize content, align incentives, and improve efficiency are discussed.

Chapter two proposes a method to optimize a targeting policy that maximizes an outcome observed only in the long term. Traditionally, this typically requires delaying decisions until the outcome is observed or relying on simple short-term proxies for the long-term outcome. The method builds on the statistical surrogacy and off-policy learning literature to first impute the missing long-term outcomes and then approximate the optimal targeting policy on the imputed outcomes via a doubly robust approach. It is applied in large-scale proactive churn management experiments at The Boston Globe by targeting optimal discounts to its digital subscribers to maximize their long-term revenue. It is shown that conditions for the validity of average treatment effect estimation with imputed outcomes are also sufficient for valid policy evaluation and optimization; furthermore, these conditions can be somewhat relaxed for policy optimization. The method is also validated empirically by comparing it with a policy learned on the ground truth long-term outcomes, they are shown to be statistically indistinguishable. It also outperforms a policy learned on short-term proxies for the long-term outcome.

Chapter three investigates how network embeddings can be applied to the study of product diffusion. Three sets of results are documented using a combination of real and simulated datasets: First, node embeddings can predict adoption decisions and timing better than standard centrality-based summary statistics. Second, node embeddings as control variables reduce the bias in the estimation of peer effect, especially when tie formation depends on unobservables. Third, graph embeddings based on the diffusion process as a whole reveal meaningful similarities between different diffusion processes such as simple vs. complex contagion.

“Machine Learning Methods for Targeting and New Product Development”

Author: Artem Timoshenko (2019)

Committee: John Hauser (co-chair), Duncan Simester (co-chair), Dean Eckles

ABSTRACT:

The dissertation consists of four essays on the applications of machine learning methods to targeting and product development. The first essay addresses the problem of identifying customer needs from user generated content. Traditionally, market research relies on interviews and focus groups to identify customer needs. User-generated content (UGC), such as online reviews, social media, and call-center data, provides an opportunity to identify customer needs more efficiently. Established methods are not well-suited for large UGC datasets because much of the content is uninformative or repetitive. We propose a machine learning approach

for identifying customer needs from UGC and evaluate the method using a custom dataset matching Amazon reviews for oral care products to the customer needs identified from the interviews by professional analysts.

The second essay addresses the problem of training robust targeting policies. Targeting policies are used in marketing to match different firm actions to different customers. For example, retailers want to send different promotions to different customers, real estate agents want to show different homes, and car dealers want to propose different prices. We conduct two large-scale field experiments to evaluate seven methods widely used to design targeting policies. The findings compare the performance of the targeting methods and demonstrate how well the methods address common data challenges. The challenges we study are covariate shift, concept shift, information loss through aggregation, and imbalanced data. We show that more complex methods perform better when the training data is ideal, but they also deteriorate faster in the presence of the challenges that affect the quality of the training data.

The third essay addresses the question of efficiently evaluating targeting policies. Firms typically compare the performance of different targeting policies by implementing the champion versus challenger experimental design. These experiments randomly assign customers to receive marketing actions recommended by either the existing (champion) policy or the new (challenger) policy, and then compare the aggregate outcomes. We discuss an alternative experimental design and propose an estimation approach to improve the evaluation of targeting policies. The alternative experimental design randomly assigns customers to marketing actions. This allows evaluation of any targeting policy without requiring an additional experiment, including policies designed after the experiment is implemented. The proposed estimation approach identifies customers for whom different policies recommend the same action and recognizes that for these customers there is no difference in performance. This allows for a more precise comparison of the policies. We illustrate the advantages of the alternative experimental design and the proposed estimation approach using data from an actual field experiment.

“Consumer Inattention, Uncertainty, and Marketing Strategy”

Author: Xinyu Cao (2018)

Committee: Juanjuan Zhang (chair), Tony Ke, John Hauser

Abstract:

This dissertation investigates the implications of consumer inattention and uncertainty for firms' advertising and pricing decisions. The first chapter is an overview of the problems addressed in the dissertation and the main findings.

The second chapter develops a theory-based, cost-effective method to estimate the demand for new products using choice experiments. The premise is that consumers are uncertain about their valuation of a new product and need to spend costly effort to learn their valuation. The effort consumers spend is affected by the probability of their choice being realized, and as a result will change the manifested demand curve derived from choice experiments. We run a large-scale choice experiment on a mobile game platform, where we randomize the price and realization probability of a new product. Data support our theoretical hypothesis. We then estimate a structural model of consumer decisions. The structural estimates allow us to accurately infer actual demand based on choice experiments of small to moderate realization probabilities.

The third chapter examines firms' advertising strategy on social media under consumers' limited attention. Advertising on social media faces a new challenge as consumers can actively select which advertisers to follow. A Bayesian learning model suggests that consumers with limited attention may rationally choose to unfollow a firm. This happens if consumers already know about the firm's value well and if the firm advertises too intensely. However, we find that intensive advertising may still be the optimal strategy for firms. If a firm is perceived as providing low value, it will want to advertise aggressively to change consumers' mind; if a firm is perceived as providing higher value, it will also want to advertise intensively, but in an effort to crowd-out advertising messages from its competitors. Tracking company accounts of 49 TV shows on the most popular tweeting website in China provides empirical evidence that both popular and non-popular firms advertise intensively, although the number of followers does go down when a firm advertises too intensively.

The fourth chapter investigates channel coordination in search advertising. Given that consumers have limited attention, there are only a limited number of advertising slots on search engine platforms that can attract positive number of clicks. A manufacturer can sponsor retailers to advertise its products while at the same time compete with them in a position auction with limited number of slots. We prescribe the optimal cooperative search advertising strategies for the manufacturer. We find that it may not be optimal for a manufacturer to cooperate with all of its retailers, even when these retailers are ex ante the same. This finding reflects the manufacturer's tradeoff between higher demand and higher bidding cost caused by more intensified competition. With two asymmetric retailers, the manufacturer should support the retailer with the higher channel profit per click to get a higher position than the other retailer. The manufacturer should take a higher position than a retailer when its profit per click via direct sales exceeds the channel profit per click of the retailer. We also investigate how a manufacturer uses both wholesale and advertising contracts to coordinate channels with endogenous retail prices.

“The Effect of Message Framing on Initial Choices, Satisfaction, and Ongoing Engagement”

Author: Eleanor Putnam-Farr (2015)

Committee: John Hauser (chair), Drazen Prelec, Renee Gosline, Jason Riis

Abstract:

Essay 1: "Yes/No/Not right now." Yes/No response formats can increase response rates even in non-forced-choice settings. Although yes/no response formats have been used to increase enrollment rates in several different types of programs, their use has been largely limited to forced choice settings. Across two field experiments, we demonstrate a substantial advantage in click-through rates for a yes/no response format over traditional opt-in response formats in an email context where choice is not forced. The increase in click-through rate does, under certain conditions, also persist through downstream program enrollment and participation. Finally, though noting that the yes/no format advantage is probably multi-determined, we discuss several potential psychological mechanisms, which are particularly relevant in non-forced choice settings.

Essay 2: The Effect of Benefit Quantification on Goal Setting and Persistence We look at how language used to describe rewards, specifically the quantification of an expected reward, might lead participants to create specific targets for their own performance based on that language. Through a combination of field and lab experiments, we demonstrate that the use of a high number to describe rewards leads to higher interest and enrollment, but also higher expectations of performance and a higher drop-out rate from the program when the reward is difficult to achieve. Marketers should be aware of this issue when describing benefits to potential customers, particularly if they wish to motivate persistent behavior.

Essay 3: A Picture is Worth a Thousand Words: Photographs as Anchors for Expectations Marketers often use images to increase the vividness of their communications to customers. These images should make the message easier to process, thereby increasing liking for the product and certainty of expectations. In a series of experiments, I demonstrate that images do indeed increase certainty of preference estimates, both within and between respondents, but may have more mixed effects on the valence of estimates of preference. I also begin to examine how these estimates might impact evaluations of actual products and propose some areas for future exploration.

“Consumer Heterogeneity, Uncertainty, and Product Policies”

Author: Song Lin (2015)

Committee: Birger Wernerfelt (chair), Duncan Simester, Juanjuan Zhang

Abstract:

This dissertation consists of three essays on the implications of consumer heterogeneity and uncertainty for firms' strategies. The first essay analyzes how firms should develop add-on policies when consumers have heterogeneous tastes and firms are vertically differentiated. The theory provides an explanation for the seemingly counter-intuitive phenomenon that higher-end hotels are more likely than lower-

end hotels to charge for Internet service, and predicts that selling an add-on as optional intensifies competition, in sharp contrast to standard conclusions found in the literature.

The second essay examines how firms should develop product and pricing policies when customer reviews provide informative feedback about improving product or service quality. The analysis provides an alternative view of customer reviews such that they not only can help consumers learn about product quality, but also can help firms learn about problems with their products or services.

The third essay studies the implications of cognitive simplicity for consumer learning problems. We explore one viable decision heuristic - index strategies, and demonstrate that they are intuitive, tractable, and plausible. Index strategies are much simpler for consumers to use but provide close-to-optimal utility. They also avoid exponential growth in computational complexity, enabling researchers to study learning models in more-complex situations.