Essays on Personalization and Market Design in Mobile Advertising

Omid Rafieiankoopaei

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2020

Reading Committee:

Hema Yoganarasimhan, Chair

Arvind Krishnamurthy

Simha Mummalaneni

Program Authorized to Offer Degree:
Foster School of Business

University of Washington

Abstract

Essays on Personalization and Market Design in Mobile Advertising

Omid Rafieiankoopaei

Chair of the Supervisory Committee:
Professor Hema Yoganarasimhan
Marketing and International Business

Personalization and market design are two pillars of mobile advertising that resulted in a dramatic growth in this industry over the past few years. In this dissertation, I focus on these two topics and study the new challenges and opportunities that arise in mobile advertising industry. In particular, I study the consequences of improved user trackability and sequential ad delivery in mobile in-app advertising. Chapter 2 focuses on the former and examine the privacy implications of it in a competitive environment, whereas Chapter 3 studies the latter and quantifies the effects of variety of previous ads on users' engagement with the next ad. In Chapters 4 and 5, I combine the findings from the first two chapters and develop an Adaptive Ad Sequencing framework that determines how a publisher can personalize a sequence of ads for a user to maximize engagement, as well as how it can monetize these sequences through auctions to optimize revenues.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to take this opportunity to express my gratitude to everyone who helped me in the past five years and made these years a wonderful part of my life. I would like to thank my dissertation chair, Hema Yoganarasimhan, who has been involved in every step of my development as a researcher and provided me with unlimited support. Hema, words cannot describe how much I am indebted to you for being a mentor, a teacher, and a friend. Thanks for taking a chance on me and guiding me through every step along the way. I look up to you and I hope, some day, I will be able to impact someone else's life the way you impacted mine.

I would like to thank other members of my supervisory committee, Amin Sayedi, Arvind Krishnamurthy, Simha Mummalaneni, Oliver Rutz, and Jacques Lawaree, whose constant support was an invaluable resource for me. Amin, you are the reason I am here today. Thanks for trusting me and showing me the way during my Ph.D. I will be always grateful to have your support and friendship. Arvind, it is a privilege to know you. I was lucky to have access to your judgment and advice. I learned from you in every interaction we had. Simha, thank you for your kind support throughout my Ph.D. life. My work benefited greatly from your thoughtful and detailed input. Oliver, thank you for your encouragement and support in this entire process. I would also like to thank Jacques whose classes on microeconomics and contract theory made me appreciate the beauty of the field even more.

I would like to express my appreciation to Jaime Banaag, Beau Kirkeby, and other staff members at the Foster Business School. My warm gratitude goes to my friends in Seattle who made my social life brighter with their presence: Behnaz, Shahryar, Pegah, Sareh, Eugene, Aravinda, Amir, Nadia, Emisa, Koosha, Majid, Ebrahim, Hessam, Maria, Keehea, Melissa, Anupama, Michelle, Chethana, and Maggie. I am and will always be grateful to my dear friends whom I have not seen in a while but been able to talk to and share experiences and memories with. To name a few: Hamid, Leila, Navid, Ali, Navid, Maryam, Yasaman, Saba, Arash, Sepehr, Sarah, Saba, Rashid, Ahmad, Hadi, Arman, and Mehdi.

Finally, I want to thank my family. I would like to thank my parents, Akbar and Nafiseh, who supported me unconditionally in every step of my life. Everything I have achieved in my life, I owe it you two. I would like to thank my brother Hamid, and my sisters Hoori and Homa. Being the youngest child in the family gave me the privilege to have access to the experiences of my siblings. They introduced me to new adventures and experiences I would have had to discover with trial and error if it was not for them. I would also like to thank my in-laws, Niusha, Anubhav, and Mohammad, whom I love and hold as close to my heart as my own siblings.

DEDICATION

to my beloved parents, Nafiseh and Akbar.

&

to my dear siblings, Hamid, Hoori, and Homa.

Chapter 1

INTRODUCTION

The mobile advertising industry has seen an unprecedented growth over the past decade. In 2018, mobile advertising became the dominant advertising channel in the U.S., surpassing TV and desktop advertising in share of the total media ad spend [34]. The growth in mobile advertising is particularly attributed to mobile in-app advertising, that is, ads shown inside mobile apps [33]. As consumers, advertisers, and app developers engage more with mobile in-app advertising, it is increasingly important to understand different elements of this market.

Like other digital advertising formats, two key factors contribute to the growth of the mobile in-app advertising industry – personalization and market design. Personalization in the advertising context is the delivery of the right ad to a user based on the user's characteristics and past history. As such, personalized delivery of ads can create value in the market by enhancing user engagement with ads. While personalization only focuses on users as the main player, market design brings advertisers to the equation to guarantee certain outcomes in a multi-sided market. In particular, market design allows the advertising industry to optimally monetize its service. Together, personalization and market design facilitate value creation and extraction in the advertising market.

In this dissertation, I study personalization and market design in the mobile in-app advertising industry. Although these two elements are present in the digital environment, there are features unique to in-app advertising that give rise to new challenges and opportunities in the mobile advertising ecosystem. My goal in this dissertation is to explore some of these new directions by combining theoretical and empirical frameworks to better understand the mobile advertising domain.

The outcome has been four separate papers in this domain that I present in the following four chapters. I provide a brief overview of each chapter and how each chapter relates to another below:

1. **Chapter 2 – Personalization and Privacy.** In this chapter, I focus on a unique feature of mobile in-app advertising – user trackability. While the user-tracking properties of mobile in-app advertising facilitate behavioral targeting, they have raised concerns among privacy advocates. This has resulted in an ongoing debate on the value of different types of targeting information, the incentives of ad-networks to engage in behavioral targeting, and the role of regulation. To understand the gains from targeting, I first build a machine learning targeting framework to personalize ad delivery in a mobile in-app advertising platform and document the gains from doing so. I then take this framework to a competitive auction environment and study the platform's decision to allow targeting through lenses of revenue-efficiency trade-off. I find that the platforms may have economic incentives to limit extensive personalization and behavioral targeting as it thins out the market and hurts their revenues. An implication of this finding is that self-regulation can be achieved by design. This chapter is based on the paper *Targeting and Privacy in Mobile Advertising* [107].

2. **Chapter 3 – Dynamic Effects of Advertising.** In this chapter, I focus on a unique aspect of in-app advertising – sequential ad placement. In this form of advertising, users are exposed to a sequence of potentially different ads within a session. This gives rise to a series of questions related to the effects of ad sequences. In particular, I study the effects of variety of previous ads on user's clicking behavior on the next ad. I empirically document that when exposed to a higher variety of previous ads, users are more likely to click on the next ad. I then explore the sources for the effects of variety and identify the sequential organization of exposures as a major source. This chapter is based on the paper *How Does Variety of Previous Ads Influence Consumer's Ad Response?* [106].

3. **Chapter 4 – Adaptive Ad Sequencing: Value Creation.** The findings of Chapter 3 highlights an inter-temporal trade-off in allocation of ads: while increasing variety can help generate more clicks, it can come at the cost of showing more irrelevant ads. In this chapter, I combine the insights from Chapter 2 and 3, and propose a framework called "Adaptive Ad Sequencing" that determines how a publisher can personalize a sequence of ads for a user to maximize engagement. I document the gains from adopting this framework compared to a series of benchmarks that are often used in practice. This finding suggests that the publisher can create value by adaptive sequencing of ads. This chapter is based on the paper *Optimizing User Engagement through Adaptive Ad Sequencing* [104].

4. **Chapter 5 – Adaptive Ad Sequencing: Value Extraction:** While Chapter 4 shows a path to create value in the market through Adaptive Ad Sequencing, it is not clear whether the publisher can extract optimal revenue from this created value. In this chapter, I address the problem of value extraction and propose a unified framework with two components – (1) a theoretical framework to derive the revenue-optimal dynamic auction in a setting where the publisher uses a forward-looking allocation rule, and (2) an empirical framework that involves the structural estimation of advertisers' private valuations. I document significant revenue gains from using the revenue-optimal dynamic auction compared to the revenue-optimal static auction. These gains stem from the improvement in the match between users and ads in the dynamic auction. The revenue-optimal dynamic auction also improves all key market outcomes, such as the total surplus, average advertisers' surplus, and market concentration. This chapter is based on the paper *Revenue-Optimal Dynamic Auctions for Adaptive Ad Sequencing* [105].

While all the chapters work with the same setting and data source, there are differences in the samples I used given the purpose of each chapter. As such, I present the setting and data in each chapter separately. Some information may be a bit redundant, but the aim is to keep the chapters fully contained. Further, the contributions and implications of each chapter are discussed separately.

Chapter 2

PERSONALIZATION AND PRIVACY

2.1 *Introduction*

2.1.1 Mobile Advertising and Targeting

Mobile advertising now constitutes the largest share of total digital ad spend [35]. The popularity of mobile advertising stems from an ad format unique to the mobile environment: in-app ads or ads shown inside apps. These ads have excellent user-tracking properties, and allow ad-networks to stitch together user data across sessions, apps, and advertisers.[1] Thus one of the main attractions of in-app advertising is its ability to facilitate behavioral targeting [31].

While the advertising industry has lauded the trackability of in-app ads, consumers and privacy advocates have derided them, citing privacy concerns. Advertisers argue that tracking allows consumers to enjoy free apps and content, and see relevant ads, whereas users demand higher privacy and limits on behavioral tracking and targeting [32]. Responding to consumer concerns, regulatory bodies have started taking action. For example, the European Union's General Data Protection Regulation agreement requires users to opt into, rather than opt out of, behavioral targeting [77].

Even as consumers, businesses, and regulators are trying to find the right balance between consumer protection and business interests, we do not have a good understanding of the key issues at the core of targeting and privacy. For example, to what extent does targeting improve the efficiency of the advertising ecosystem, what is the value of different

[1] Advertisers and ad-networks have access to a unique device ID associated with the mobile device referred to as IDFA (ID For Advertisers) in iOS devices, and AAID (Android Advertiser ID) in Android devices. This device ID is highly persistent and remains the same unless actively re-set by the user.

types of targeting information, and what are the incentives of different players in the advertising industry to engage in user-tracking and behavioral targeting? The lack of a cohesive framework to analyze these issues hampers our ability to have an informed discussion and to form policy on them.

2.1.2 Research Agenda and Challenges

In this chapter, we seek to address this gap by providing a unifying framework to answer the following sets of questions related to targeting and privacy in the advertising ecosystem.

The first set of questions relates to targeting and efficiency. How can ad-networks use the data available to them to develop targeting policies? How can we evaluate the performance of these policies in both factual and counterfactual settings? In particular, what are the gains in CTR from adopting an efficient (CTR-maximizing) targeting policy? The second set of questions relates to the value of targeting information. We are particularly interested in the relative value of contextual vs. behavioral information. The former captures the context (when and where) of an impression, and the latter summarizes an individual user's past app usage, ad exposure, and ad response. Contextual information is privacy-preserving, whereas behavioral information is based on user-tracking and therefore impinges on users' privacy. Third, we are interested in quantifying the revenue-efficiency trade-off and ad-network's incentives to enable different forms of targeting. What is the empirical relationship between efficiency and ad-network revenues? What is the optimal level of targeting from the perspective of different players in the market? Finally, to what extent are the ad-network's and advertisers' incentives aligned?

There are three main challenges that we need to overcome to satisfactorily answer these questions. First, to develop efficient targeting policies, we need to obtain accurate estimates of CTR for all ads that could have been shown in an impression (i.e., counterfactual ads), and not just the ad that was actually shown in that impression. Thus, we need exogenous variation in the ad allocation mechanism to evaluate counterfactual targeting policies. Second, to quantify the value of different pieces of targeting information, we need a model that can accurately predict whether a targeted impression will lead to a click or not. Models with poor predictive ability will lead to downward bias in the estimates of the value of information. Third, we need an underlying model of strategic interactions that can quantify market outcomes (e.g., ad-network and advertiser revenues) under different targeting regimes. Without an economic model that puts some structure on the ad-network's and advertisers' utilities, we cannot make any statements on their incentives to target and/or the extent to these incentives are aligned.

2.1.3 Our Approach

We present a unified and scalable framework that coherently combines predictive machine learning models with prescriptive economic models to overcome the challenges listed above. Our framework consists of two main components. The first, a machine learning framework for targeting, addresses the first two challenges of obtaining counterfactual CTR estimates and achieving high predictive accuracy in this task. The second is an analytical model that incorporates competition and characterizes the ad-network's and advertisers' profits under different targeting regimes. This addresses the third challenge of linking targeting regimes to ad-network and advertiser revenues.

The main goal of the first component is to estimate the match value between an impression and an ad, where match value can be interpreted as the CTR of an impression-ad combination. Once we have match values for all impression-ad combinations, we can use them to define and evaluate any counterfactual targeting strategy. Match values are thus the key primitives of interest here, and we infer them by combining ideas from causal inference with predictive machine learning models. Our approach consists of three parts – (a) a filtering procedure, (b) a feature generation framework, and (c) a learning algorithm. The goal of the filtering procedure is to identify the set of ads for which we can generate accurate counterfactual estimates of CTR for each impression. If the platform uses a deterministic ad allocation mechanism (as is the common practice in the industry), then this set is null, by definition. However, in our setting, there is exogenous variation in the ad allocation process, which gives us a non-empty set of counterfactual ads for each impression. Our filtering procedure determines this set by identifying the ads which have a non-zero propensity of being shown in a given impression. Next, our feature generation framework relies on a set of functions to generate a large number of features that capture the contextual and behavioral information associated with an impression. Using these functions, we generate a total of 160 features that serve as input variables into a CTR prediction model. Finally, we use XGBoost, proposed by [22], a fast and scalable version of boosted regression trees, as our learning algorithm.

In the second component, we focus on the ad-network's incentives to allow targeting. In an influential paper, [82] conjecture that while high levels of targeting can increase efficiency in the market, it can reduce the ad-network's revenue by softening the competition between advertisers. We propose a theoretical framework that allows us characterize this revenue-efficiency trade-off under counterfactual targeting regimes. To take this framework to data, we need an estimate of each advertisers' valuation for a given impression. This valuation can be decomposed into two sets of primitives: (a) match valuations or CTRs for all impression-ad combinations, and (b) advertisers' click valuations for each impression. While match valuations are already available from the machine learning targeting framework, we need to infer click valuations by inverting advertisers' observed equilibrium bids [53]. The product of these two entities gives us each advertiser's value of a given impression, which allows us to quantify the ad-network's revenue, advertisers' surplus, and total surplus under different targeting regimes.

We apply our framework to one month of data from the leading mobile ad-network from a large Asian country. The scale and scope of the data are large enough to provide realistic substantive and counterfactual results. For our analysis, we sample over 27 million impressions for training, validation, and testing, and use another 146 million impressions for feature generation. A notable feature of our setting is the use of a quasi-proportional auction allocates impressions to ads using a probabilistic rule: an advertiser's probability of winning an impression is proportional to her bid. This induces randomization or exogenous variation in ad allocation, which in turn, allows us to estimate match valuations for counterfactual ad-impression combinations. At the same time, the auction mechanism preserves the strategic linkage between bids and advertisers' click valuations, which allows us to estimate click valuations from the bid data. Our setting thus facilitates the separate identification of both match and click valuations.

2.1.4 Findings and Contribution

We first discuss the results from the machine learning model for targeting. We present both factual and counterfactual evaluations of our model. In the factual evaluation, we use goodness-of-fit measures to evaluate how well our model can predict the observed outcome. We find that our model predicts the outcome on a hold-out test set with substantial accuracy: it achieves a Relative Information Gain (RIG) of 17.95% over a baseline model that simply predicts the average CTR for all impressions. Next, we find that behavioral information contributes more to the predictive accuracy of the model compared to contextual information. In the second part of our evaluation, we consider the efficient targeting policy, wherein each impression is allocated to the ad with the highest estimated CTR in that impression. We show that this efficient targeting policy can increase the average CTR in the ad-network by 66.80% over the current system.

Next, we link advertisers' targeting strategies to ad-network's incentives and revenues. First, we theoretically prove that in an efficient auction mechanism (e.g., second-price auction) – (a) the total surplus in the system monotonically increases as the extent of targeting increases, but (b) the ad-network's revenues are not monotonic; it may or may not increase with more granular targeting. So we take our theoretical framework to data and perform empirical counterfactuals to compare ad-network revenues under different targeting regimes.

In particular, we consider four targeting regimes that relate to our research agenda – full (impression-level targeting), behavioral (user-level targeting), contextual (app-time-level targeting), and no targeting. We find that the ad-network's revenue is maximized when it restricts targeting to the contextual level even though doing so lowers total surplus, i.e., allowing behavioral targeting thins out the market, which in turn reduces ad-network revenues. Therefore, the ad-network has economic incentives to adopt a privacy-preserving targeting regime, especially if it cannot extract additional surplus from advertisers through other mechanisms. On the advertisers' side, we find that although a majority of them prefer a regime where the ad-network allows behavioral targeting, not all do. An important implication of our findings is that it may not be necessary for an external entity such as EU/FCC to impose privacy regulations, in light of ad-networks' economic incentives.

This chapter makes several contributions to the literature. First, from a methodological perspective, we propose a novel machine learning framework for targeting that is compatible with counterfactual analysis in a competitive environment. A key contribution of our targeting framework is in combining existing ideas from causal inference literature with recent machine learning literature to generate counterfactual estimates of user behavior under alternative targeting regimes. Further, we present an efficient auction framework with targeting that characterizes advertisers' utility function under any targeting regime and provides a direct link to the estimation of market outcome such as efficiency and revenue. Second, from a substantive perspective, we provide a comprehensive comparison of contextual and behavioral targeting, with and without the presence of competition. To our knowledge, this is the first study to compare the role of behavioral and contextual targeting on market outcomes. Third, from a managerial perspective, our results demonstrate a non-monotonic relationship between targeting granularity and revenues. While our findings may depend on the context of our study, our framework is generalizable and can be applied to most standard advertising platforms that use deterministic auctions as long as the platform randomizes ad allocation over a small portion of the traffic (which would satisfy the unconfoundedness assumption). Finally, from a policy perspective, we identify the misalignment of the ad-network's and advertisers' incentives regarding behavioral and contextual targeting and information disclosure. We expect our findings to be of relevance to policy-makers interested in regulating user-tracking and behavioral targeting in the advertising space.

The rest of this chapter is organized as follows. In §2.2, we discuss the related literature. We introduce the setting and data in §2.3. In §2.4, we present our machine learning framework for targeting, and in §3.6 we presents a series of results on efficiency gains from targeting. Next, in §2.6, we develop a theoretical framework for analyzing the revenue-efficiency trade-off and a corresponding empirical analysis of auctions with targeting. In §2.7, we present the results on market outcomes under counterfactuals targeting regimes. Finally, in §2.8, we conclude with a discussion on the generalizability of our framework and our main contributions.

2.2 Related literature

First, this chapter relates to the computer science literature on CTR prediction [93, 61, 21]. These papers make prescriptive suggestions on feature generation, model selection, learning rates, and scalability. Our work differs from these papers in two main ways. First, we develop a filtering procedure that allows us to obtain accurate CTR estimates for both the ad shown during an impression as well as counterfactual ads not shown in the impression. Thus, unlike the previous papers, our framework can be used to develop and evaluate different targeting policies. Second, we quantify the value of different types of information in the targeting of mobile ads, whereas the previous papers were mainly concerned with click prediction.

This chapter also relates to the literature on ad-network's incentives to allow targeting. [82] were one of the first to conjecture the trade-off between value creation and market thickness. They argue that too much targeting can thin out markets, which in turn can soften competition and make the ad-network worse off. This is particularly the case when there is significant heterogeneity in the distribution of advertisers' valuation of impressions [20]. Building on this insight, a growing stream of analytical papers show that there is a non-monotonic pattern between the extent of targeting and ad-network revenues [13, 2, 65, 26, 121]. A key difference between these papers and ours is that we do not make any distributional assumptions on the match values in our analytical model.

In spite of the increasing interest from the theoretical side, there has been limited empirical work on this topic with mixed findings. In an early paper, [135] present a structural model to estimate advertisers' valuations and show that targeting benefits both advertisers and the ad-network. In a similar context, however, [7] present a case study of two keywords and show that coarsening CTR predictions (worse targeting) can help a search advertising ad-network generate more revenue. However, unlike this chapter, neither of these papers can effectively design or evaluate counterfactual targeting regimes because their data come from highly targeted eco-systems without any randomization in ad-allocation. More broadly, ours is the first empirical paper to view revenue-efficiency trade-off through the lens of privacy and quantify the ad-network's incentives to preserve users' privacy.

Next, our work relates to the literature on the interplay between privacy and targeting. [49] use data from a series of regime changes in advertising regulations and show that restricting targeting reduces response rates and thereby advertisers' revenues. Similarly, [48] and [129] highlight the perils of excessive targeting as users perceive increased targeting as a threat to their privacy. Please see [47] for an excellent review of targeting in online advertising and [1] for a detailed discussion of consumer privacy issues. This chapter contributes to this literature by providing the first empirical evidence in support of the possibility of self-regulation in this market.

Finally, this chapter adds to the growing literature on applications of machine learning in marketing, which focus on prediction problems; see [128] and [29] for excellent summaries. This chapter contributes to this stream by demonstrating how a combination of theory-driven frameworks and machine-learning methods can be used to go beyond prediction and help answer important substantive and prescriptive questions.

2.3 Setting and Data

2.3.1 Setting

Our data come from the leading mobile in-app advertising network of a large Asian country, which had over 85% market-share in the category in 2015. The ad-network works with over 10,000 apps and 250 advertisers and it serves over 50 million ads per day (about 600 auctions per second). This ad-network specializes in the Android Operating System. At the time of our study, smartphone penetration was reasonably high in the country with over 60% of the population having access to smartphones. The share of Android OS was over 85% of the market in this country in 2015, which is consistent with its share worldwide [115].

Players

There are four key players in this marketplace.

Users: Individuals who use apps. They see the ads shown within the apps that they use and may choose to click on the ads.

Advertisers: Firms that show ads through the ad-network. They design banner ads and specify their bid as the amount they are willing to pay per click, and can include a maximum budget if they want to. Advertisers can target their ads based on the following variables: app category, province, connectivity type, time of the day, mobile operators, and mobile brand of the impression. The ad-network does not support more detailed targeting (e.g., behavioral targeting) at this point in time.

Publishers: App developers who have joined the ad network. They accrue revenues based on the clicks generated within their app. Publishers earn 70% of the cost of each click in their app (paid by the advertiser), and the remaining 30% is the ad-network's commission.

Ad-network or Platform: It functions as the matchmaker between users, advertisers, and publishers. It runs a real-time auction for each impression generated by the participating apps and shows the winning ad during the impression. The platform uses a CPC pricing mechanism, and therefore generates revenues only when clicks occur. [2]

[2] An impression lasts one minute. If a user continues using the app beyond one minute, it is treated as a new impression and the platform runs a new auction to determine the next ad to show the user.

Auction Mechanism

The platform uses a *quasi-proportional* auction mechanism [94]. Unlike other commonly-used auctions (e.g., second price or Vickrey), this auction uses a probabilistic allocation rule:

$$\pi_{ia} = \frac{b_a q_a}{\sum_{j \in \mathcal{A}_i} b_j q_j} \tag{2.1}$$

where π_{ia} is the probability that advertiser a with bid b_a and quality score q_a wins impression i, and $\mathcal{A}_i$ denotes the set of advertisers participating in the auction for impression i. The quality score is an aggregate measure that reflects the advertiser's potential profitability for the platform. Currently, the platform does not use impression-specific quality scores; rather it uses an advertiser-specific quality score that remained constant during our observation period.

Because of the probabilistic nature of the auction, the ad that generates the highest expected revenue for the platform is not guaranteed to win. Rather, advertiser a's probability of winning is proportional to $b_a q_a$.[3] Further, advertisers are only charged when a user clicks on their ad. The cost-per-click for an impression is determined using a next-price mechanism similar to that of Google's sponsored search auctions. In this case, the amount that the winning ad is charged per click is the minimum amount that guarantees its rank among the set of bidders. For example, suppose that there are three advertisers with bids 1, 2, and 3, and quality scores 0.1, 0.2, and 0.3, bidding on an impression. Then, the product of bid and quality score for the three advertisers are 0.1, 0.4, and 0.9, respectively. In this case, if the second-ranked bidder wins the auction, he only needs to pay $\frac{1 \times 0.1}{0.2} = 0.5$, since it is the minimum bid amount that guarantees that he will be ranked higher than the third-ranked bidder. Formally, we can write the cost-per-click for ad a in impression i as:

$$CPC_{ia} = \inf \left\{ b' \Big| \sum_{j \in \mathcal{A}_i, j \neq a} \mathbb{1}(b' q_a \leq b_j q_j) = \sum_{j \in \mathcal{A}_i, j \neq a} \mathbb{1}(b_a q_a \leq b_j q_j) \right\}, \tag{2.2}$$

[3]From a practical perspective, probabilistic auctions ensure that individual users are not exposed to the same ad repeatedly within the same app-session (which can be irritating). In contrast, in a deterministic auction, the same advertiser would win all the impressions until his budget runs out.

where $\sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b_a q_a \leq b_j q_j)$ is essentially the number of ads whose product of bid and quality score is lower than ad a, and the infimum over this set finds the minimum bid (b') that guarantees ad a's rank. Finally, note that the platform uses a fixed reserve price, r_0, for all impressions. It is the minimum bid that is accepted by the platform. So, if an advertiser is not willing to pay at least r_0 per click, he is automatically out of competition.

2.3.2 Data

We have data on all the impressions and corresponding clicks (if any) in the platform for a 30-day period from 30 September, 2015, to 30 October, 2015. For each impression, we have data on:

- Time and date: The time-stamp of the impression.
- AAID: Android Advertising ID is a user re-settable, unique, device ID that is provided by the Android operating system. It is accessible to advertisers and ad networks for tracking and targeting purposes. We use it as the user-identifier in our main analysis.
- App ID: A unique identifier for apps that advertise through the platform.
- Ad ID: Identifier for ads shown in the platform.
- Bid: The bid that the advertiser has submitted for her ad. Advertisers bids do not change across impressions in our sample.
- Cost-per-Click (CPC): The price that the winning advertiser has to pay, if she wins the impression and a click occurs. This is calculated by the ad-network based on Equation (2.2).
- Location: This includes the province as well as the exact location of a user, based on latitude and longitude.
- Connectivity type: It refers to the user's type of connectivity (e.g., Wi-Fi or cellular data).
- Smartphone brand: The brand of user's smartphone (e.g., Samsung, Huawei, etc.).
- MSP: The user's mobile-phone service provider.

- ISP: The user's Internet service provider.
- Click indicator: This variable indicates whether the user has clicked on the ad or not.

The total data we see in this one month interval is quite large. Overall, we observe a total of 1,594,831,699 impressions and 14,373,293 clicks in this time-frame, implying a 0.90% CTR.

2.3.3 Data Splits and Sampling

We use penultimate two days of our sample period (October 28 and 29) for training and validation, and the last day for testing (October 30). We also use the preceding history from September 30 to Oct 27 (referred to as global data) to generate the features associated with these impressions. The splits of data are shown in Figure 2.1. Note that we do not fit our model on the global data because we do not have sufficient history to generate features for these impressions. Further, constraining all the three data-sets – training, validation, and testing – to a three-day window has advantages because recent research has shown that data freshness plays an important role in CTR prediction, i.e., using older history for prediction can lead to poor predictive performance [61].

We draw a sample of 728,340 unique users (out of around 5 million) seen on October 28, 29, and 30 to form our training, validation, and test data-sets.[4] In Appendix §A.5.4, we show that this sample size is sufficient and that larger samples do not significantly improve model performance.

Figure 2.1 presents a visual depiction of the sampling procedure. Rows represent users. The impressions by users in our sample are shown using black points. There are 17,856,610 impressions in the training and validation data, and 9,625,835 impressions in the test data. We have an additional 146,825,916 impressions by these users in the time preceding October 28, which form global data. These impressions will be solely used for feature generation (and not for model fitting). Note that both our user-based sampling procedure and feature

[4] Another approach would be to randomly sample impressions in each split of the data. However, this would not give us the complete user-history for each impression in the training, validation, and test data-sets. This in turn would lead to significant loss of accuracy in user-level features, especially since user history is sparse. In contrast, our user-based sampling approach gives us unbroken user-history.

Figure 2.1: Schema for data generation.

generation approach (see Appendix §A.2) require us to be able to identify and track users. For this purpose, we use the AAID variable as our user identifier.

2.3.4 Summary Statistics

We now present some summary statistics on our training, validation, and test data, which constitutes a total of 27, 482, 444 impressions.

Table 2.1 shows the summary statistics of the categorical variables in the data. For each variable, we present the number of unique values, the share of top three values that the categorical variable can take, and the number of non-missing data. While we always have information on the app, ad, and time-stamp of the impression, the other variables are sometimes missing. The shares are shown after excluding the missing variables in the respective category.

We observe a total of 263 unique ads and 9709 unique apps in the data. The top three sub-categories in each have large shares and there is a long tail of smaller apps and ads. Moreover, as shown in Figure 2.2, we find that the top 37 ads account for over 80% of the

Variable	**Number of categories**	**Share of top categories** 1^{st}	2^{nd}	3^{rd}	**Number of impressions**
App	9709	37.12%	13.56%	3.05%	27,482,444
Ad	263	18.89%	6.71%	6.31%	27,482,444
Hour of the Day	24	7.39%	7.32%	6.90%	27,482,444
Province	31	25.25%	6.65%	6.51%	21,567,898
Smartphone Brand	8	46.94%	32.30%	9.53%	25,270,463
Connectivity Type	2	54.64%	45.36%		27,482,444
ISP	9	68.03%	14.02%	7.09%	10,701,303
MSP	3	48.57%	43.67%	7.76%	26,051,042

Table 2.1: Summary statistics for the categorical variables.

impressions, and similarly, the top 50 apps account for 80% of impressions.

Next, we present some descriptive analysis that examines the role of contextual and behavioral information in predicting CTR. A context is characterized by the 'when' and 'where' of an impression. As such, we define a unique context as a combination of an app and a specific hour of the day. Figure 2.3 shows the histogram of CTR for different contexts. As we can see, there is a significant amount of variation in CTR across contexts, which suggests that contextual information can be informative for predicting clicks. Next, to understand the role of behavioral information, we focus on the length of history available for a user. Figure 2.4 shows the CDF of the length of history for all the impressions and clicks. It suggests that users with longer histories are less sensitive to ads. Most of the clicks come from users with shorter histories, while most impressions come from users with longer histories. Thus, user-history or behavioral information also seem helpful in explaining the clicking behavior observed in data.

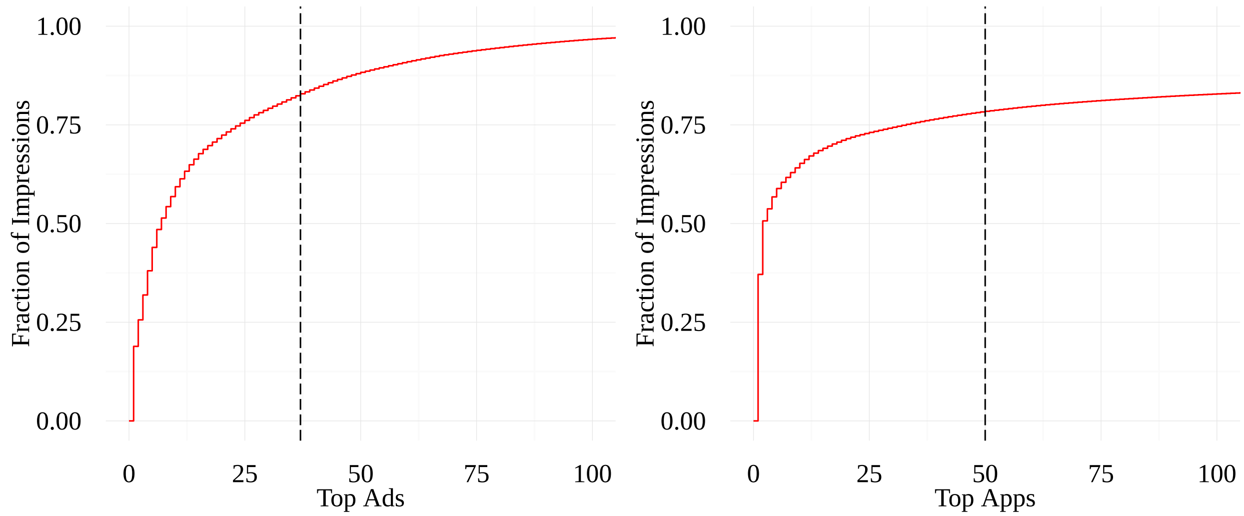

(a) Ads (b) Apps

Figure 2.2: Cumulative fraction of impressions associated with the top 100 ads and top 100 apps.

2.4 Machine Learning Framework for Targeting

In this module, our goal is to develop a framework that can accurately estimate the gains in efficiency or the CTR for any targeting policy. To do that, we first need to specify and train a machine learning model that accurately predict the match between an impression and an ad, i.e., predict whether an impression will generate a click or not, for both factual and counterfactual ads.

This section is organized as follows. We first define our problem in §2.4.1. Next, in §2.4.2, we discuss our empirical strategy. Here, we explain the need for, and the extent of, randomization in our data generating process, and propose a filtering approach that establishes the scope of our framework in estimating both factual and counterfactual targeting policies. In §2.4.3, we present the details of our feature generation framework. Finally, in §2.4.4, we discuss our estimation procedure which consists of the learning algorithm, the loss function,

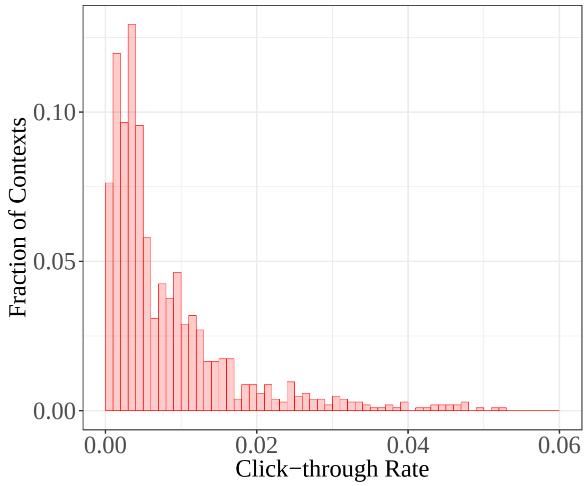

Figure 2.3: Histogram of click-through rates (CTR) for different contexts. Context is defined as a unique combination of an app and an hour of the day (where and when of an impression)

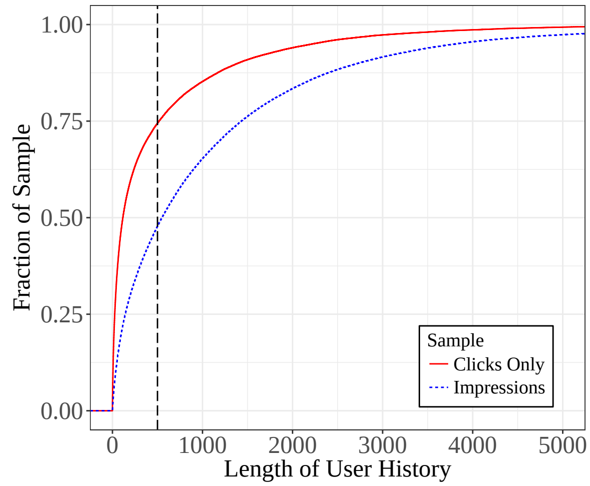

Figure 2.4: Empirical CDF of the length of user history for impressions and clicks (truncated at 5000). History is defined as the number of previous impressions from September 30 till the current impression.

and the validation method.

2.4.1 Problem Definition

Consider a setting with N impressions and A ads. We begin with a formal definition of a targeting policy.

Definition 1 *A targeting policy τ is defined as a mapping between impressions to ads such that each impression is allocated one ad. For example, $\tau(i) = a$ means that targeting policy τ selects ad a to be shown in impression i.*

In order to evaluate the effectiveness of a targeting policy, we first need an accurate prediction of CTR for each ad for a given impression in our data. That is, for each impression i and ad a, we need to estimate $\Pr(y_{i,a} = 1)$, where $y_{i,a}$ is the indicator that ad a receives a click when it is shown in impression i. This brings us to the formal definition of match value matrix:

Definition 2 *Let $m_{i,a} = \Pr(y_{i,a} = 1)$. The $N \times A$ match value matrix M is defined as:*

$$M = \begin{bmatrix} m_{1,1} & m_{1,2} & \dots & m_{1,A} \\ m_{2,1} & m_{2,2} & \dots & m_{2,A} \\ \vdots & \vdots & \ddots & \vdots \\ m_{N,1} & m_{N,2} & \dots & m_{N,A} \end{bmatrix}, \tag{2.3}$$

where N denotes the total number of impressions in our data and A denotes the total number of ads competing for these impressions. There is a corresponding $N \times A$ matrix of outcomes Y, which consists of elements $y_{i,a}$. Note that we only observe the realized outcome for one element in each row or impression i for Y, which corresponds to the ad $a = a_i$ which was actually shown in that impression in our data. The rest of the elements are treated as potential or unrealized outcomes.

In this section, our goal is to develop a machine learning framework to estimate this match value matrix. We can use our estimated match value matrix, $\hat{M}$, to perform the following analyses:

1. **Evaluate model performance:** We can evaluate the predictive performance of our model using the observed outcome. Let τ_0 denote the *current targeting policy*, such that:

$$\tau_0(i) = a_i, \tag{2.4}$$

where a_i is the ad that is actually shown in impression i. Since we observe the actual outcomes for y_{i,a_i}, we can evaluate how well our $\hat{m}_{i,a_i}$'s estimate these outcomes.

2. **Evaluate the gains from efficient targeting policy:** Using the match value matrix, we can evaluate the expected CTR of any counterfactual targeting policy τ as follows:

$$\hat{m}^{\tau} = \frac{1}{N} \sum_{i=1}^{N} \hat{m}_{i,\tau(i)} \tag{2.5}$$

In particular, we are interested in the *efficient targeting policy*, τ^*, determined by our model that allocates each impression to the ad with the highest CTR for that impression:

$$\tau^*(i) = \underset{a}{\operatorname{argmax}}\, \hat{m}_{i,a} \tag{2.6}$$

In §2.5.2, we quantify the gains in average CTR from efficient targeting over the current system.

2.4.2 Empirical Strategy

We now present our empirical strategy to estimate matrix M. At a high-level, our goal is to build a model to predict whether an impression i showing ad a will receive a click or not, based on the joint distribution of impressions and clicks in our data. That is, we seek to estimate a function $f(X_{i,a})$ such that:

$$m_{i,a} = Pr(y_{i,a} = 1) = f(X_{i,a}), \tag{2.7}$$

where $X_{i,a}$ is a set of features that are informative of whether impression i showing ad a will receive a click. Since this problem can be interpreted as function evaluation, we turn to machine learning algorithms that can capture complex relationships between the covariates and the outcome without imposing strong parametric restrictions on $f(\cdot)$.

While machine learning methods can flexibly learn the function f from the data, their prediction power is bounded by the joint distribution of co-variates and outcome (click) in the data. That is, these methods can accurately predict the outcome for an observation only if that observation could have been observed in the data. This requirement gives rise to two main challenges in evaluation of counterfactual targeting policies:

Challenge 1 *Function f cannot learn $m_{i,a}$ from the data if ad a could have never been shown in impression i, i.e., ad a has zero propensity of being shown in impression i.*

The reason is simple: if ad a could have never been shown in impression i, then the set of features $X_{i,a}$ is not within the joint distribution of the observed data. For example, if the ad for a fashion clothing brand was never shown in a sports app, then it is not possible to recover the fashion ad's click probability in the sports app.

It is worth noting that if the platform runs a deterministic auction (e.g., second price auction), the set of ads that could have won the auction (and hence been shown during an impression) is a singleton. Similarly, the set of ads that can be shown in an impression in highly targeted environments would be very small. Therefore, data-sets generated without any randomization in the ad allocation mechanism will not allow researchers to push the scope of their analysis beyond the set of actual outcomes observed in the data. Randomization in ad allocation is thus necessary if we want to use our framework to evaluate the effectiveness of counterfactual targeting policies. This brings us to our first remark, that addresses Challenge 1.

Remark 1 *Any ad participating in the auction for impression i ($\forall\ a \in \mathcal{A}_i$) has a non-zero propensity of being shown in impression i.*

This is a direct result of the quasi-proportional auction run by the platform. As shown in Equation (3.1), each ad that participates in an auction has a non-zero probability of winning. This claim is the equivalent of the *positivity* or *overlap* assumption in the causal inference literature [114].

While any kind of randomization can help overcome Challenge 1, we need to know the distribution of randomization to be able to correctly infer the click probability of counterfactual ads in any given impression i, i.e., infer m_{ia} for ads $a \neq a_i$. If ads are randomized according to an unobserved rule, we may run into selection issues and obtain biased estimates of m_{ia}. We can characterize this challenge as follows.

Challenge 2 *Function f cannot correctly infer match values (m_{ia}s) for counterfactual ads, if the allocation rule is a function of an unobserved variable that is correlated with the outcome.*

The following example helps illustrate this challenge: suppose that ad a_Y is targeted more towards younger users, whereas ad a_O is targeted more towards older users. Now, if younger users have a higher probability of click, failure to account for users' age will lead us to attribute the better performance of ad a_Y to the ad, rather than to users' age. In the causal inference literature, this is usually known as endogeneity or selection on unobservables [134].

In our setting, we can simulate the allocation rule using the observed covariates. This gives us the unconfoundedness assumption, which we characterize in Remark 2.

Remark 2 *For any impression i, ad allocation is independent of the set of the potential outcomes for participating ads ($a \in \mathcal{A}_i$), after controlling for the observed covariates:*

$$\{y_{i,a}\}_{a \in \mathcal{A}_i} \perp\!\!\!\perp a_i \mid X_{i,a} \tag{2.8}$$

Again, the allocation rule in Equation (3.1) directly satisfies the unconfoundedness assumption because everything in the right-hand side of this equation is known. First, for each i, we can infer the set of ads competing ($\mathcal{A}_i$) from our data because we observe all the targeting variables that can induce variation in $\mathcal{A}_i$. Second, advertisers do not change their bids and the platform does not customize the quality score for each impression. Hence, $b_a q_a$ remains constant throughout our study and we can easily infer propensity scores π_{ia} from the data, controlling for $\mathcal{A}_i$.

Together, in light of Remarks 1 and 2, we can estimate the match values $m_{i,a}$, not only for the ad that is shown in impression i, but for any counterfactual ad that could have been shown with non-zero propensity score. Naturally, estimates for small ads with a very small probabilities of winning will be noisy. However, it is possible to overcome this issue by focusing on the top 37 ads that constitute over 80% of our data. We first discuss our procedure for identifying the set of all participating ads in each impression that have non-zero propensity scores. Next, we discuss how we estimate these propensity scores and assess covariate balance.

Filtering Procedure

As discussed earlier, if ad a could have never been shown in impressions i, we cannot accurately estimate the match value for that impression-ad combination $m_{i,a}$. As such, we need to identify the set of participating ads in each impression and filter those that have zero propensity of being shown. In general, two factors influence whether an ad is available to participate in an auction for an impression.

- Targeting: Targeting by advertisers is the main reason why some ads are unavailable to compete for certain impressions, and therefore have zero probability of being shown in them. For example, if an ad chooses to target only mornings, then it is not considered in the auctions for impressions in evenings. In that case, we should filter out this ad for all impressions in the evening. While limited, targeting is nevertheless present in our setting, and mainly happens on province, time, and app categories. Hence, for each impression i, we filter out all ads that were excluded from the auction for i because of targeting.
- Campaign availability: Second, some ads may be unavailable to compete for a given impression because their ad campaigns may not be running in the system when the impression happens. This could happen either because the advertiser's budget has been exhausted, or because the advertiser has exited the market. Therefore, for each impression i, we filter out ads that were unavailable when it happens. Empirically, we find that campaign availability is not a major factor that leads to ad-filtering since we focus on top ads.[5]

We now construct a filtering matrix $E_{N \times A} = [e_{i,a}]$ that filters out ads for each impression based on the factors discussed above, where each element $e_{i,a}$ takes value 1 if ad a has a non-zero probability of winning impression i, and 0 otherwise. Each row in this matrix shows which ads are competing for an impression.[6] However, our filtering may not be accurate for observations with missing targeting variables. Therefore, for all the analyses that use filtering, we only focus on the *Filtered Sample*, which consists of the impressions in the test data for which all targeting variables are non-missing. Figure 2.5 shows the empirical CDF

[5]Only six ads experience budget exhaustion (at least once) in the training data, four of which are completely out for auctions in the test data.

[6]Note that this information is not directly observed but inferred from advertisers' targeting decisions and campaign availability.

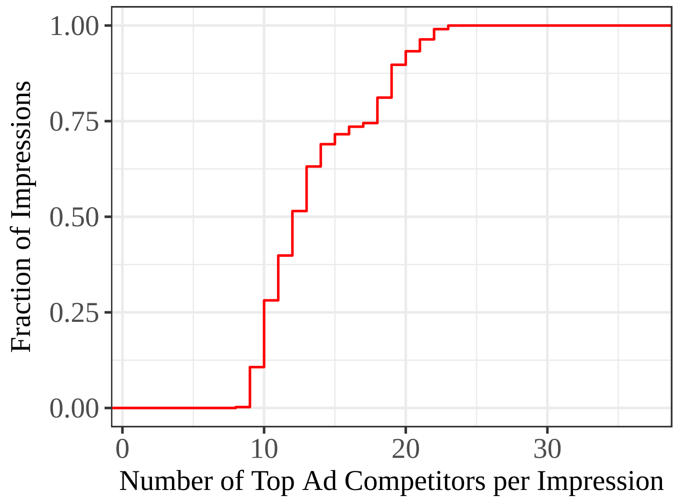

Figure 2.5: Empirical CDF of the number of competitors (of the top 37 advertisers) per impression for the Filtered Sample.

of the number of competing ads for each impression in the Filtered Sample data among top 37 ads per impression. Note that almost all impressions have at least 8 top ads competing for it, and the median impression has 13 top ad competitors.

Propensity Score Estimation and Covariate Balance

As discussed earlier, the accuracy of our counterfactual match value estimates is predicated on the independence of assignment to ads and potential outcomes, given the observed covariates. While we know that this is theoretically true in our setting because of the allocation rule (Remark 2), we nevertheless need to empirically demonstrate the validity of this remark in our setting.

The standard practice in these cases is to assess and show covariate balance, since it is a necessary condition for unconfoundedness assumption. In simple settings, where both treatment and control are randomly assigned with a fixed probability to the entire population, we can easily assess balance by comparing the pre-treatment variables across treatment and control groups. Our case is more complicated because of two reasons: (1) assignment to ads is not fully random, but random given the propensity scores, and (2) since we focus on top 37 ads, we have more than two treatment arms. To assess covariate balance, we therefore need to take the following steps:

1. *Step 1 – Propensity Score Estimation:* The first step is to estimate the propensity score π_{ia} for all a and i. Since we have multiple treatments, the dependent variable is a categorical variable with multiple classes. We use a multi-class XGBoost to estimate propensity scores given the success of machine learning methods in propensity score estimation [92]. Please see Appendix §A.1.1 for more details.

2. *Step 2 – Assessing Covariate Balance:* Once we have the estimates for propensity scores, we can assess the balance for all pre-treatment covariates. In our case, these covariates are the variables that advertisers can target on: province, app, time of the day, smartphone brand, connectivity type, and MSP. To assess balance, we need to show that the inverse propensity weighted distribution of each pre-treatment variable is the same across all the ads. Following the norm in the literature, we use standardized difference of the weighted mean of a covariate when assigned to ad a and the population mean of covariate and rule for balance if this difference is below 0.2 [92]. Please see Appendix §A.1.2 for details on our balance measures and results.

2.4.3 Feature Generation Framework

As discussed in §2.4.2, our goal is to build a model, $m_{ia} = Pr(y_{ia} = 1) = f(X_{ia})$ to accurately predict whether an impression i will receive a click or not. As such, we first need a vector of features $X_{i,a}$ that captures the factors that affect whether the user generating impression i will click on ad a.

It is important to generate an exhaustive and informative set of features since the predictive accuracy of our model will largely depend on the quality of features we use. Given our research agenda, our features should also be able to capture the contextual and behavioral information associated with an impression over different lengths of history preceding the impression (long-term, short-term, and session-level). To achieve these objectives, we adopt the main ideas from the functional feature generation framework proposed by [137]. There are three advantages of doing so. First, her function-based approach allows us to generate a large and varied set of features using a parsimonious set of functions. Second, it allows for a natural mapping between feature inputs and feature classification. Third, the general class of features she suggests have been shown to have good predictive power in this class of problems.

We now present a short overview of our feature functions and feature categorization below, and refer interested readers to Appendix §A.2 for a more detailed description.

Inputs for Feature Functions

To generate a set of features for each impression, we use feature functions that take some inputs at the impression level and output a corresponding feature for that impression. Our feature functions typically need two types of inputs:

- *Impression-specific information:* Each impression in our data can be uniquely characterized by three types of information – (1) contextual information that captures the context (where and when) of the impression, i.e., which app serves this impression and at what time (hour of day) is the impression being shown, (2) behavioral information that denotes the identity of the user generating this impression, and (3) ad-related information, that denotes the identity of the ad which was shown during this impression.
- *History:* This input characterizes the history over which we aggregate to calculate the output of our functions. We define three different levels that capture the long-term (approximately one month), short-term (3 days), and ongoing session-level history. Besides, we characterize the history in such a way that we can update the features in real-time.

To reduce the dimensionality of our feature sets and boost the speed of our feature generation framework, we group the smaller apps (below top 50) into one app-category and all the smaller ads (below top 37) into one ad-category. Thus, our features do not distinguish the context of smaller apps (ads) as separate from each other, though they are able to distinguish them from the top apps (ads). Please see Appendix §A.2.1 for a complete formal definition of the inputs for feature functions.

Feature Functions

One challenge we face is that most of the information characterizing an impression-ad combination is categorical in nature, e.g., the app showing the ad, the user seeing the ad. As a result, approaches that include all these categorical raw inputs and their interactions as covariates are prone to the curse of dimensionality. So we define functions that take these raw inputs as well as their interactions and map them onto a parsimonious set of features that reflect the outcome of interest – CTR.

We present an overview of our feature functions in Table 2.2 along with their functionality (see Appendix §A.2.2 for a detailed description of the feature functions). These functions take different inputs based on the focal impression and return outputs that are integers or real numbers. These inputs are basically interactions of different raw inputs. The following examples give a high-level overview of what these functions do. Let p_i, t_i, u_i, and a_i denote the app, hour, user, and ad associated with impression i. If the function *Impressions* is given p_i, u_i, and a_i and long-term history as inputs, it simply returns the number of times user u_i has seen ad a_i inside app p_i from the start of the data till the time at which impression i occurred. However, if it is only given u_i and short-term history, it returns number of impressions user u_i has seen across all apps and ads over the last three days. Using this logic, we give different sets of inputs to these functions and generate 98 features for each impression i. In addition, we include a few standalone features such as dummies for each of the top ads, the user's mobile and Internet service providers, latitude, longitude, and connectivity type. Overall, we have a total of 160 features for each impression-ad (ia) combination. Together, these features capture the interactive effects of advertising that are documented in the literature such as carryover effects [118], spillover effects [85], and effects of ad variety [106]. Please see

Function	**Functionality**
Impressions	Number of impressions for a given set of inputs over a pre-specified history
Clicks	Number of clicks for a given set of inputs over a pre-specified history
CTR	Click-through rate for a given set of inputs over a pre-specified history
AdCount	Number of distinct ads shown for a given set of inputs over a pre-specified history
Entropy	Dispersion of ads shown for a given set of inputs over a pre-specified history
AppCount	Number of distinct apps used by a given set of inputs over a pre-specified history
TimeVariability	Variance in the user's CTR at different hours of the day over a pre-specified history
AppVariability	Variance in the user's CTR across different apps over a pre-specified history

Table 2.2: Feature functions.

Appendix §A.2.3 for the full list of features.

Feature Categorization

All our features capture one or more type of information – contextual, behavioral, and ad-specific. To aid our analysis, we therefore classify features based on the type of information used to generate them and group them into the following (partially overlapping) categories:

- *Contextual* features (F_C): These are features that contain information on the context of the impression – app and/or hour of the day.
- *Behavioral* features (F_B): These are features that contain information on the behavior of the user who generated the impression.
- *Ad*-specific features (F_A): These are features that contain information on the ad shown during the impression.

The three feature sets form our full set of features $F_F = F_B \cup F_C \cup F_A$. We now present a few examples of features generated using the *Clicks* function to elucidate this classification. The total clicks made by user u_i across all apps, ads, and hours of the day in the past

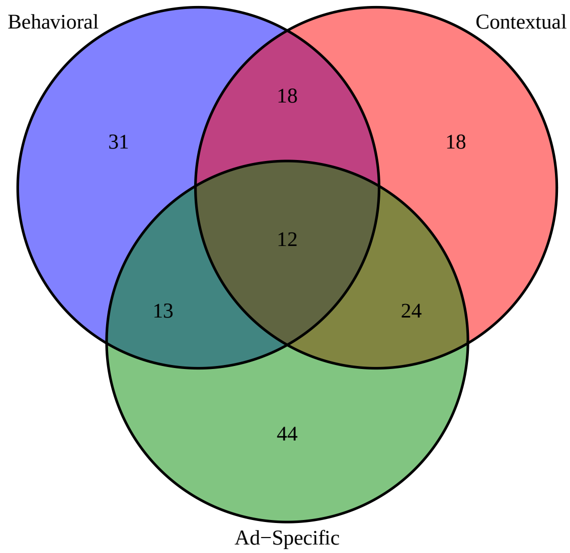

Figure 2.6: Venn diagram of the three feature sets, with the number of features in each region.

month is a purely behavioral feature since it only contains information on the behavior of the user who generated impression i. On the other hand, the total clicks made by user u_i in the app p_i over the last month is both a behavioral and contextual feature since it contains information on both the behavior of u_i as well as the context (app p_i) in which she made these clicks. Finally, the total clicks received by ad a_i over the last one month across all users, apps, and times is a purely ad-specific feature since it only reveals information about the ad's propensity to receive clicks. Thus a feature can contain any combination of behavioral, contextual, or ad-specific information depending on the inputs used to generate it. Please see Table A.1 in the Appendix for a mapping between each feature and the categories that it falls under and Figure 2.6 for a Venn diagram of our classification system.

2.4.4 Learning Algorithm: XGBoost

We now discuss the final step of our machine learning framework – the learning algorithm, which helps us learn the function $f(X_{i,a})$. It provides a mapping between our feature set

$(X_{i,a})$ and the match value or click probability, as: $f(X_{i,a}) = m_{i,a} = Pr(y_{i,a} = 1)$. Given that we want to maximize the predictive accuracy of the model, we do not want to impose parametric assumptions $f(\cdot)$. The problem of function evaluation is fundamentally different and harder than the standard approach used in the marketing literature, wherein we simply evaluate parameters after assuming a functional form. In the latter, the researcher only needs to search over the set of parameters given functional form, whereas in the former we have to search over the space of functions. Therefore, we turn to machine learning algorithms that are designed for this task.

Specifically, we employ the XGBoost algorithm proposed by [22]. XGBoost is a variant of the standard boosted regression trees, and is one of the most successful prediction algorithms developed in the last few years. It has been widely adopted in both academia and industry.[7] At a high level, boosted regression trees can be thought of as performing gradient descent in function space using shallow trees as the underlying weak learners [17, 41]. While boosted trees have been around for over a decade, [22]'s implementation is superior to earlier implementations from both methodological and implementation standpoints.[8] We refer interested readers to Appendix §A.3 for a more detailed description of XGBoost and now focus on two key components of our implementation: loss function and validation procedure.

To train any learning model, we need to specify how the model should penalize model fit, i.e., the difference between the observed outcome y_{i,a_i} and model prediction $\hat{m}_{i,a_i}$ (where a_i refers to the ad shown in impression i). This is done using a loss function, which the machine

[7]Boosted trees in general, and XGBoost in particular, perform exceptionally well in tasks involving prediction of human behavior. Examples include store sales prediction, customer behavior prediction, product categorization, ad CTR prediction, course dropout rate prediction, etc. Indeed, almost all the KDD cup winners have used XGBoost as their learning algorithm (either as a standalone model or in ensembles) since 2015.

[8]First, from a methodological standpoint, XGBoost can be interpreted as performing Newton boosting in the function space (as opposed to gradient descent), and thereby uses information from the Hessian as well. Thus, both the quality of the leaf structure and the leaf weights learned are more accurate in each step. Second, XGBoost uses a trick commonly used in Random Forests – column sub-sampling,which reduces the correlation between subsequent trees. Third, XGBoost employs a sparsity-aware split finding, which makes the algorithm run faster on sparse data. Finally, from an implementation perspective, XGBoost is highly parallelized, which makes it fast and scalable.

learning algorithm minimizes. Since our outcome variable is binary, we use logarithmic loss (log-loss) as our loss function. It is the most commonly used loss function in the CTR prediction literature [**?**] and has some attractive properties, e.g., a faster convergence rate compared to other loss functions such as squared loss [113]. The log-loss for a model with predictions $\hat{\mathbf{M}}$ when the prediction matrix is Y can be written as:

$$\mathcal{L}^{log\ loss}(\hat{M}, Y) = -\frac{1}{N}\sum_{i=1}^{N}\left(y_{i,a_i}\log\left(\hat{m}_{i,a_i}\right) + (1 - y_{i,a_i})\log\left(1 - \hat{m}_{i,a_i}\right)\right) \tag{2.9}$$

Note that while the log-loss function takes as inputs the two matrices $\hat{M}$ and Y, the metric is calculated only over those ad-impression combinations that are actually observed in the data.

Validation is an important part of training any machine learning model. The boosting algorithm is designed to continuously update the prediction rule (or current estimate of $f(\cdot)$) to capture more and more complex relationships between the features $X_{i,a}$, in order to predict $y_{i,a}$. Since we do not impose any assumptions on the parametric form of $f(\cdot)$, this will likely lead to over-fitting, i.e., the model will evolve to fit too closely to the training data and perform poorly out-of-sample. Validation helps us avoid this problem by using parts of the data to validate the model. This ensures that the chosen model, $f(\cdot)$, will have a good out-of-sample performance. Please see Appendix §A.3.2 for a full description of our validation procedure.

2.5 Results from the Machine Learning Targeting Models

Recall that the goal of our machine learning framework is to estimate the matrix M defined in Equation (2.3). As such, our $\hat{M}$ contains CTR estimates for: (1) the ads shown in the data, and (2) counterfactual situations, i.e., ads that could have been shown. In §2.5.1, we focus on the actual data and present results on the predictive performance of our framework on the observed sample. We also document the contribution of behavioral vs. contextual information to our framework in this section. Next, in §2.5.2, we focus on the counterfactual estimates in $\hat{M}$ and evaluate the gains in CTR from an efficiently targeting policy. Finally, in §2.5.3 we discuss robustness and scalability.

2.5.1 Predictive Performance of the Machine Learning Model

Evaluation Metric

To evaluate whether a targeting model improves our ability to predict clicks, we first need to define a measure of predictive accuracy or an evaluation metric. In line with our loss function, we use "Relative Information Gain" or *RIG*, which is defined as the percentage improvement in log-loss over the baseline that simply predicts average CTR for all impressions. Formally:

$$RIG(\hat{M}, Y) = \left(1 - \frac{\mathcal{L}^{log\ loss}(\hat{M}, Y)}{\mathcal{L}^{log\ loss}(\bar{Y}, Y)}\right) \times 100, \tag{2.10}$$

where $\bar{Y}$ is a $N \times A$ matrix, each of whose element is equal to $\frac{1}{N}\sum_{i=1}^{N} y_{i,a_i}$, i.e., the average observed outcome of the sample or the average CTR of the data. Average CTR is the simplest aggregate metric available from any data, and using it as the baseline prediction tells us how well we can do without any model. It is important to control for this baseline because if the average CTR is very high (close to 1) or very low (close to zero, as in most e-commerce settings, including ours), a naive prediction based on the average CTR leads to a pretty good log-loss. Normalizing the log-loss with the average CTR reduces the sensitivity of the metric to the data distribution [61]. Nevertheless, we need to be careful when interpreting *RIG*s computed on different datasets because there is no obvious normalization in those cases [**?**].

In Appendix §A.5.1, we present four other commonly used evaluation metrics – (1) Mean Squared Error, (2) AUC, (3) 0/1 Loss, and (4) Confusion Matrix. We discuss the pros/cons of these metrics and demonstrate the performance of our model on them.

Predictive Accuracy of the Full Targeting Model

We now discuss our framework's ability to predict the actual outcomes in the data. Table 2.3 shows the gains in prediction for: (1) training and validation data, and (2) test data.

Evaluation Metric	Training and Validation	Test
LogLoss for Full Model	0.041927	0.044364
LogLoss for Baseline Model	0.051425	0.054070
RIG of Full Model	18.47%	17.95%

Table 2.3: LogLoss and *RIG* (in percentage) shown for training, validation, and test data.

The first row depicts the log-loss for the Full model (which uses the set of all features and trains the XGBoost model). The second row depicts the log-loss for the baseline model, which simply predicts the average CTR for the dataset for all impressions. The third row is the *RIG* of the Full model compared to the baseline model.

The *RIG* of the Full model over the baseline is 17.95% on the test data, a substantial improvement in CTR prediction problems. This suggests that the data collected by the ad-network is quite valuable and that our machine learning framework has significant predictive power on whether an impression-ad combination will receive a click.[9]

The *RIG* improvement for training and validation data is 18.47%, which is somewhat higher than 17.95% for the test data. There are two potential reasons for this. First, all statistical models estimated on finite data have higher in-sample fit than out-of-sample fit. Indeed, this is the main reason we use the test data to evaluate model performance. Second, the difference could simply reflect the differences in the underlying data distributions for the two data-sets. As discussed in §2.5.1, we cannot compare *RIG* across data-sets because it is co-determined by the model and data. Thus, the difference between the *RIG* values across

[9]One could argue that the significant predictive power of the Full model is due to the weak benchmark, which simply predicts average CTR for all impressions. Therefore, we also evaluate the performance of the Full model against two other baseline models: (1) Ad-specific CTR, and (2) Targeting-area-specific CTR. The first model relates to ad-networks' quality scoring practice; it predicts the average CTR for each ad as the match-value for impressions showing that ad. The second model resembles the current targeting practice in the platform and predicts the average CTR for each targeting area (defined as the intersection of all targeting variables) as the match value for all impressions within that targeting area. With these benchmark models as the denominator in Equation (2.10), we find that the Full model has a RIG of 16.86% over the Ad-specific model and 10.06% over the Targeting-area-specific model.

the data-sets is not necessarily informative.

Value of Information: Behavioral vs. Contextual Features

We now examine the impact of different types of features on the predictive accuracy of our model. This is important for two reasons. First, data storage and processing costs vary across feature types. For example, some user-specific behavioral features require real-time updating, whereas pure-contextual features tend to be more stable, and can be updated less frequently. In order to decide whether to store and update a feature or not, we need to know its incremental value in improving targeting. Second, the privacy and policy implications of targeting depend on the features used. For example, models that use behavioral features are less privacy-preserving than those that use purely contextual features. Before adopting models that are weaker on privacy, we need objective measures of whether such models actually perform better.

Recall that our features can be categorized into three broad overlapping sets – 1) Behavioral, denoted by F_B, 2) Contextual, denoted by F_C, and 3) Ad-specific, denoted by F_A. We now use this categorization to define two models:

- **Behavioral model:** This model is trained using behavioral and ad-specific features, without including any contextual features. Formally, the feature set used is $(F_B \cup F_A) \setminus F_C$.
- **Contextual model:** This model is trained using only contextual and ad-specific features, without including any behavioral features. The feature set for this model is $(F_C \cup F_A) \setminus F_B$.

Both models include ad-specific features that are neither behavioral nor contextual, e.g., the total impressions received by the ad shown in the impression in the past month.[10] They also use the same loss function and training algorithm, and only differ on the set of features used. Hence, it is possible for us to directly compare the *RIG* of one model over another within

[10] We can also specify Behavioral and Contextual models that ignore ad-specific information. The qualitative results on the relative value of behavioral and contextual information for that case are similar to those presented here.

RIG over Baseline	Full Sample	Top Ads and Top Apps	Filtered Sample
Behavioral Model	12.14%	14.82%	14.74%
Contextual Model	5.25%	5.98%	6.77%
Full Model	17.95%	22.85%	22.45%
No. of Impressions	9,625,835	6,108,511	4,454,634
% of Test Data	100%	63.5%	46.28%

Table 2.4: Comparison of Behavioral and Contextual models for different samples of test data.

the same data.[11]

The results from these two models and their comparisons with the Baseline model are presented in Table 2.4. First, consider the results for the full test data (presented in the second column). The Behavioral model has a 12.27% *RIG* over the baseline, which is considerably higher than 5.12%, the *RIG* of the Contextual model over the baseline. Together, these findings suggest that, from a targeting efficiency perspective, behavioral information is more effective compared to contextual information in mobile in-app advertising.

This difference in the effectiveness of the two models directly relates to the extent of variation in the information used by the two models. The variation in behavioral features is much higher than the variation in contextual features because behavioral features are generated from unique behaviors of over 700,000 users, whereas the total number of unique contexts is limited (to 1200). Hence, the level of granularity of contextual features is much lower and the Contextual model can only learn from aggregate outcome estimates across these limited contexts. Its ability to predict positive labels (i.e., clicks) is therefore much weaker compared to that of the Behavioral model.

[11] As discussed in §2.5.1, *RIG*s are not directly comparable across different data-sets. Simply put, in Table 2.4, comparisons within a column are interpretable, but comparisons across a row are not.

One possible critique of the above analysis is that it does not exploit the full capacity of contextual information since we treat all the non-top ads as one advertiser category and all the non-top apps as one app category during feature generation (see §2.4.3). To address this issue, we consider a sub-sample of the test data which only consists of impressions that were shown in a top app and showed a top ad and re-run all the above comparisons. This accounts for 63.5% of our test data. The performance of our Full model on this subset of the data is even better than that on the full sample because there is no information loss on the ads or apps. The findings on the relative value of behavioral vs. contextual features are even stronger in this data-set, which suggests that our results in the full sample were not driven by the lack of good contextual information.

Finally, in the last column, we show the performance of our model on the Filtered Sample (described in §2.4.2), which is the sample that we use for conducting our counterfactual analysis. Our qualitative findings remain the same for this sample too.

2.5.2 Counterfactual Analysis: Efficiency Gains from CTR-maximizing Targeting Policy

We now focus on an important counterfactual question from the platform's perspective – If the platform employs an efficient targeting policy, such that each impression is allocated the ad with the highest predicted CTR in that impression, to what extent can it improve the CTR in the system?

Recall that τ_0 and τ^* denote the current and efficient targeting policy, as defined in Equations (2.4) and (2.6), respectively. We can then use the following equation to calculate the gains in average CTR:

$$\rho(\tau^*, \tau_0; N_F) = \frac{\hat{m}^{\tau^*}}{\hat{m}^{\tau_0}} = \frac{\frac{1}{N_F}\sum_{i=1}^{N_F} \hat{m}_{i,\tau^*(i)}}{\frac{1}{N_F}\sum_{i=1}^{N_F} \hat{m}_{i,\tau_0(i)}}, \tag{2.11}$$

where N_F is the number of impressions in the Filtered sample. It is crucial to conduct this couterfactual on the Filtered Sample (instead of the Full sample) for the reasons discussed in §2.4.2.

We find that an efficient targeting policy based on our machine learning model increases average CTR by 66.80% over the current regime. This is a substantial improvement, and suggests that targeting based on behavioral and contextual features can lead to significant efficiency gains.

Next, we examine how efficiency-gain varies by impression. Specifically, for each impression we calculate the percentage improvement in CTR with efficient targeting as

$$\left(\frac{\hat{m}_{i,\tau^*(i)}}{\hat{m}_{i,\tau_0(i)}} - 1\right) \times 100,$$

and examine the distribution of this metric over impressions. In Figure 2.7, we show the histogram of this percentage improvement in CTR for the impressions in the Filtered Sample. We document considerable heterogeneity in CTR improvements across impressions: the median improvement in CTR is about 105.35%, implying that efficient targeting policy can make over half the impressions twice as clickable as the current system. The peak at the left side of the graph (at one) denotes cases where the $\tau_0(i) = \tau^*(i)$, where the platform happened to randomly selected the right ad that maximizes eCTR.

This overlap between our efficient targeting policy and actual data allows us to evaluate the efficient targeting policy by inversely weighting the propensity scores for the actual outcomes in the overlapping area. This is a model-free approach known as *importance sampling*, which is commonly used in the policy evaluation literature [28]. We present the details of this approach in Appendix §A.4 and show that it establishes 65.53% improvement in average CTR which is similar to our findings based on Equation (2.11).

In sum, we find that an efficient targeting policy leads to significant gains in clicks for the platform, using both model-based and model-free approaches. Nevertheless, a key question that remains unanswered is whether an efficient targeting policy is also revenue maximizing for the platform. Therefore, in §2.6, we incorporate competition and examine the relationship between efficiency and revenues.

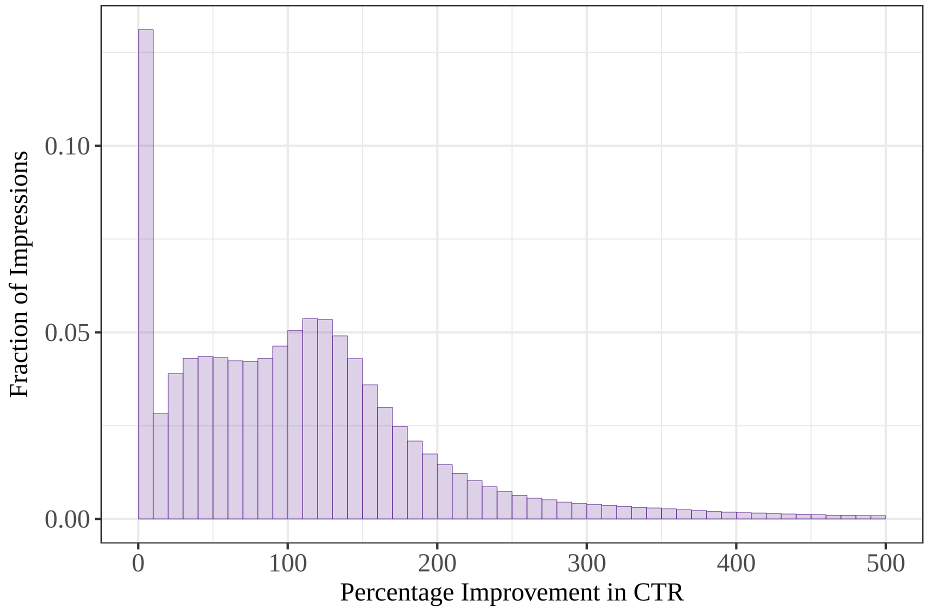

Figure 2.7: Histogram of percentage improvement in CTR over the current system using the efficient targeting policy.

2.5.3 Scalability and Robustness

We perform extensive checks on robustness of all aspects of our machine learning approach and its scalability. We discuss these tests briefly here, and refer readers to Appendix §A.5 for details.

First, in Appendix §A.5.1, we show that our results are robust even if we use other evaluation metrics (AUC, MSE, 0/1 Loss, and confusion matrix). Second, in Appendix §A.5.2, we confirm that XGBoost is the best learning algorithm for our prediction task by comparing its performance to five other commonly used algorithms (Least Squares, LASSO, Logistic Regression, Classification And Regression Tree, and Random Forests). Third, in Appendix §A.5.3, we run a few robustness checks on the feature generation framework by considering alternative ways of aggregating over history as well as app-specific dummies. Again, we find no improvement in the model's predictive performance under these different specifications. Fourth, in Appendix §A.5.4, we present some checks to establish that our data sample is sufficient and large enough to produce reliable results. Specifically, we find that the *RIG* gains start stabilizing with the sample of 100,000 users, and that our sample of 728,340 users is more than sufficient for our purposes. Finally, in Appendix §A.5.5, we show that our results are not sensitive to the validation procedure used to pick the tuning parameters by comparing with other methods, e.g., hold-out validation and k-fold cross validation.

2.6 Analysis of Revenue-Efficiency Trade-off

In §3.6, we showed that the ad-network can substantially increase CTR with efficient targeting. However, that analysis was silent on the ad-network's incentives to target and agnostic to revenues. In this section, we seek to answer two sets of important questions by focusing on competition and incentives. First, to what extent is the ad-network incentivized to allow targeting and is there an optimal level of targeting from its perspective? Second, how does the total surplus accrued by advertisers vary with targeting levels, and is there heterogeneity in advertisers' preferences on the optimal level of targeting?

Incentives are particularly important in this context because if the platform is incentivized to not allow behaviorally targeted bids, then we may naturally converge to a regime with higher consumer privacy protection. In contrast, if the platform is incentivized to allow behavioral targeting, then an external agency (e.g., government) may have to impose privacy regulations that balance consumers' need for privacy with the platform's profitability motives. Similarly, if a substantial portion of advertisers prefer a more restrictive targeting regime, then the mobile ad-industry can self-regulate. So we seek to quantify the platform's and advertisers' profits under different levels of targeting.

We now present an analytical framework to quantify the ad-network's revenue-efficiency trade-off. This section proceeds as follows. In §2.6.1, we present a simple example to fix ideas and highlight the platform's efficiency-revenue trade-off. In §2.6.2, we present a stylized analytical model that characterizes the total surplus and platform revenues under different targeting strategies. In §2.6.3, we take this analytical model to data and present an empirical

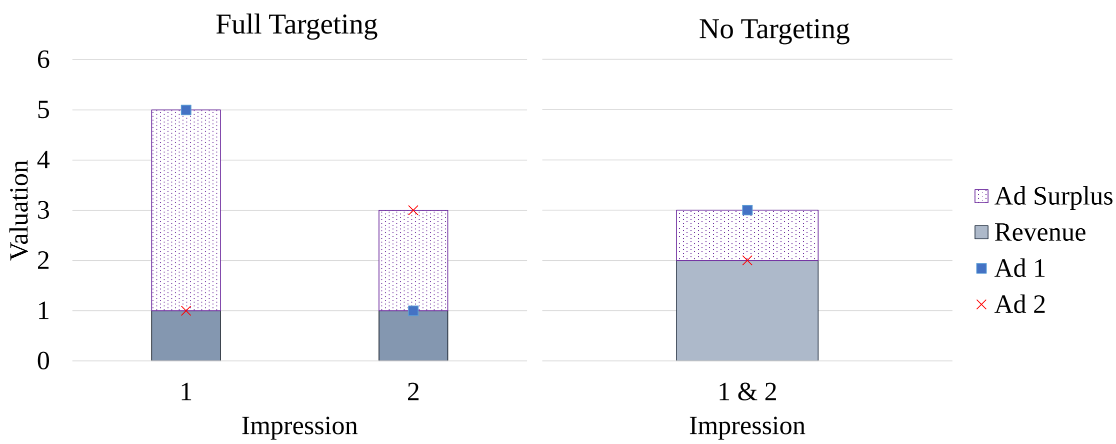

Figure 2.8: Market outcomes under full vs. no targeting. The platform sells two impressions. Ad 1 and Ad 2 have valuations 5 and 1 for impression 1, and valuation 1 and 3 for impression 2, respectively. When bundled together, advertisers cannot distinguish between ads, giving an aggregate value of 3 and 2 to Ad 1 and 2, respectively. The entire shaded area in each case shows the total surplus generated. The area on the top is the share of advertisers and that in the bottom goes to the platform. See Appendix §A.6 for a detailed analyses of this example.

analysis of auctions with targeting.

2.6.1 A Simple Example

In an important paper, [82] argue that micro-level targeting can thin auction markets, which in turn can soften competition and make the platform worse off. In Figure 2.8, we present a simple example to illustrate this idea. In this example, we consider a platform with two impressions and two advertisers whose valuations for these impressions do not align: advertiser 1 has much higher valuation for impression 1 compared to impression 2, whereas the opposite is true for advertiser 2. Assume that the platform uses second price auctions with Cost per Impression (CPI) pricing, where the highest bidder wins the impression and pays the bid of the second-highest bidder. We consider two regimes. In the Full Targeting regime, the platform allows advertisers to submit targeted bids for each impression. In the No Targeting case, advertisers cannot distinguish between the two impressions, and therefore have to submit the same bid for both the impressions (i.e., no targeted bidding). As shown in Figure 2.8, the platform cannot extract sufficient revenue if advertisers can distinguish between impressions (full targeting). However, the platform is able to extract more revenue by not revealing the identity of these impressions since advertisers are forced to rely on their aggregate valuation for both impressions together in this case. This example thus illustrates the platform's trade-off between value creation and value appropriation, and highlights the platform's incentives to limit advertisers' ability to target.

2.6.2 Analytical Model of Auction with Targeting

We now develop a simple analytical model that captures the trade-offs discussed above. To reflect the idea of narrow targeting and thin markets as envisioned by [82], we make two modeling choices. First, the idea of revenue loss in thin markets is due to the use of efficient auctions that guarantee that the highest valuation bidder will win [79]. While efficiency is satisfied in many auction mechanisms, we focus on second-price auctions since it is the most commonly used auction in online advertising. Moreover, it has the truth-telling property that makes our analysis more tractable. Second, the idea of narrow targeting by advertisers requires the pricing mechanism to be per-impression. In a cost-per-click mechanism, advertisers do not care about the match value of impressions since they are charged per click. For these reasons, we consider a setting where the platform uses a second-price auction mechanism with CPI pricing. Nevertheless, it is worth noting that neither of these two assumptions are essential to our analysis. Later in this section, we discuss how our results can be extended to other efficient auction mechanisms and/or cost-per-click pricing.

As before, we consider a platform that receives N impressions and serves A advertisers.

Let v_{ia} denote ad a's private valuation from impression i, and let V denote the value matrix:

$$V = \begin{bmatrix} v_{1,1} & v_{1,2} & \dots & v_{1,A} \\ v_{2,1} & v_{2,2} & \dots & v_{2,A} \\ \vdots & \vdots & \ddots & \vdots \\ v_{N,1} & v_{N,2} & \dots & v_{N,A} \end{bmatrix} \tag{2.12}$$

If an advertiser a can distinguish between all the impressions, s/he will submit targeted bids for each impression i. In a second-price auction, that is equivalent to a's valuation for impression i, $v_{i,a}$.

However, the extent to which advertisers can target depends on the level of targeting allowed by the platform. If certain information is not disclosed, advertisers may not be able to distinguish two impressions i and j. In such cases, a risk-neutral bidder's valuation for both impressions is the same and is equal to the expected value from the bundle of i and j [121]. For example, if the platform does not allow targeting at the app level, then advertisers cannot distinguish between impressions in two different apps, and their optimal strategy would be to submit the same bids for the impressions in both apps. Formally:

Definition 3 *Let I_l denote the set of impressions in bundle l. A targeting regime $\mathcal{I} = \{I_1, I_2, \dots, I_L\}$ denotes the platform's decision to bundle N impressions into L bundles such that advertisers can only bid for bundles and not impressions within the bundle. As such, impressions are only distinguishable across bundles, but not within a single bundle. That is, for bundle I_j, the advertiser a has the valuation $\frac{1}{|I_j|}\sum_{k \in I_j} v_{ka}$.*

This definition characterizes all targeting regimes from impression-level targeting to no targeting. Impression-level targeting occurs when each impression is a bundle ($L = N$), i.e., an advertiser can distinguish between all impressions and place targeted bids for each impression. In contrast, no targeting denotes the case where the platform bundles all impressions into one group ($L = 1$) implying that an advertiser can only have one valuation aggregated over all impressions ($\frac{1}{N}\sum_{i=1}^{N} v_{ia}$ for any a). Any intermediate strategy where $1 < L < N$ can be interpreted as partial targeting. An example of partial targeting is app-level targeting, where each bundle is an app and impressions are distinguishable across apps but not within apps.

We can characterize the relative granularity of two targeting regimes as follows:

Definition 4 *Let $\mathcal{I}^{(1)}$ and $\mathcal{I}^{(2)}$ denote two targeting regimes such that $\mathcal{I}^{(1)} = \{I_1^{(1)}, \ldots, I_{L_1}^{(1)}\}$ and $\mathcal{I}^{(2)} = \{I_1^{(2)}, \ldots, I_{L_2}^{(2)}\}$. Targeting regime $\mathcal{I}^{(1)}$ is at least as granular as $\mathcal{I}^{(2)}$ if for any $I_j^{(1)} \in \mathcal{I}^{(1)}$, there exists a $I_k^{(2)} \in \mathcal{I}^{(2)}$ such that $I_j^{(1)} \subseteq I_k^{(2)}$. In words, if two impressions i and j are distinguishable in $\mathcal{I}^{(2)}$, then they will be distinguishable in $\mathcal{I}^{(1)}$.*

We can use this definition to compare the granularity of two targeting regimes. For example, app-user level targeting is more granular than app level targeting. Now, the main question that the platform faces is at what level of granularity they should disclose information and allow targeting. Since we focus on the second-price auction, the highest-bidding ad in any impression wins that impression, and pays second-highest bid. This auction also guarantees truth-telling property, i.e., for each bundle, advertisers submit their aggregate valuation for that bundle as derived in Definition 3. The following proposition determines the relationship between the granularity level of targeting and market outcomes such as surplus and revenue:

Proposition 1 *Consider two targeting regimes $\mathcal{I}^{(1)}$ and $\mathcal{I}^{(2)}$ such that $\mathcal{I}^{(1)}$ is at least as granular as $\mathcal{I}^{(2)}$. Let $S^{(j)}$ and $R^{(j)}$ denote the total surplus and platform's revenue under targeting regime $j \in \{1, 2\}$. Then, for any distribution of valuations: $S^{(1)} \geq S^{(2)}$, but there is no fixed relationship between $R^{(1)}$ and $R^{(2)}$.*

Proof. *See Appendix §A.7.1.* ■

As the granularity of targeting increases, the total surplus generated increases, but the platform's revenue can go in either direction (unless we impose strong distributional assumptions on match values). Thus, while the matches are more efficient with more granular targeting, the platform may not be able to appropriate these efficiency gains. It is worth emphasizing that our analysis of revenue and surplus holds for any efficient auction because of the Revenue Equivalence Theorem [96, 111]. Further in Appendix §A.8, we show that the same qualitative findings hold for a cost-per-click pricing mechanism. Finally, note that this proposition is not applicable to a quasi-proportional auction since this is not an efficient mechanism.

2.6.3 Empirical Analysis of Auctions with Targeting

We now take this analytical model to data and examine market outcomes under different targeting regimes. Since the examination of revenue-efficiency trade-off requires an efficient auction, our analytical model focuses on a second-price auction with a pay-per-impression payment scheme. However, we must notice that the mechanism in our data is a quasi-proportional auction with a pay-per-click payment scheme. Thus, our empirical analysis involves counterfactual evaluation of settings different from the one in our data.

As illustrated in our analytical model, the primary estimand that we require for our empirical analysis of auctions with different levels of targeting is matrix V defined in Equation (2.12). We can characterize each element in matrix V as follows:

$$v_{i,a} = v_a^{(c)} m_{i,a}, \tag{2.13}$$

where $v_a^{(c)}$ is the private valuation ad a gets from a click and $m_{i,a}$ is the match valuations or expected CTR of ad a if shown in impression i.

Identification Strategy and Counterfactual Validity

To perform counterfactuals, we need to identify the elements of matrix V. While we cannot directly identify $v_{i,a}$ from the data, we can separately identify both elements in the right-hand side of Equation (2.13) – the click valuation of ad a ($v_a^{(c)}$) and the match valuation of ad a when shown in impression i ($m_{i,a}$). To the extent that these two elements are policy invariant primitives, our counterfactual analysis is valid. We now describe the basis on which these two estimands are identified and the conditions under which these are policy invariant primitives:

- *Click valuations* are identified given advertisers' strategic bidding behavior in the current auction environment. The main assumption required is that advertisers select the bid that maximizes their utility. We can then specify advertisers' utility functions under the current auction observed in the data and use a first-order-condition to invert their observed bids to obtain consistent estimates of their click valuations. This is the standard identification strategy employed in the auctions literature [53, 6]. It is worth noting that click valuations are policy invariant, although advertisers' bidding strategy can change under different auction mechanisms.

- *Match valuations* are identified given the unconfoundedness assumption: controlling for observed covariates, ad allocation is random. Please see §2.4.2 for a detailed discussion. Intuitively, match value estimates are policy invariant as long as users' underlying utility model for clicking on ads does not change under a different policy or auction.[12]

In sum, the identification of both click and match valuations is possible in settings that satisfy the unconfoundedness assumption while preserving the linkage between bids and click valuations. Figure 2.9 presents a Venn diagram of settings where each component is identified. It also highlights how common settings such as a second-price auction or a fully randomized ad allocation fail in this dual identification task. In standard auction mechanisms (e.g., second-price auctions) the identification problem stems from the deterministic allocation rule, which makes identification of match valuations impossible. In contrast, in a fully randomized experiment, there is no relationship between an advertiser's private click valuation and her observed bid, which makes the identification of click valuations impossible. To our knowledge, our setting (i.e., a quasi-proportional auction) is the *only* one in the literature that allows for the identification of both these components.

Estimation of Advertisers' Click Valuations

We now discuss the estimation of advertisers' click valuations, $v_a^{(c)}$s, based on the identification strategy discussed in §2.6.3. The standard approach in the structural auctions literature

[12]While we learn users' utility model flexibly using XGBoost without imposing restrictive functional form on the utility function, we still require the underlying utility model to be policy invariant. This is equivalent to treating potential outcomes as structural parameters in potential outcomes framework [66].

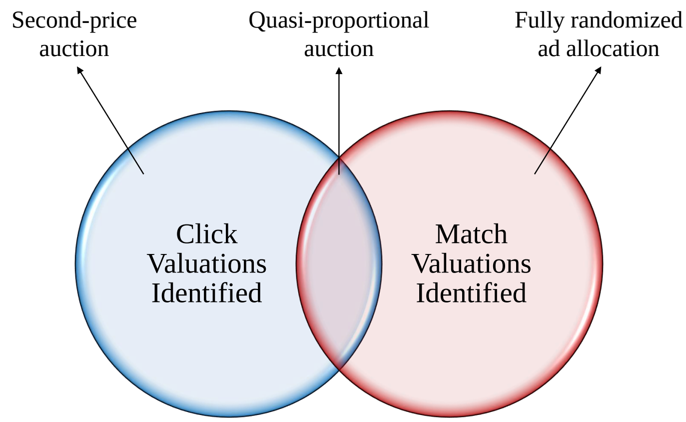

Figure 2.9: Venn diagram depicting settings where click valuations and match valuations are identified.

is to assume that agents (advertisers in this case) are utility-maximizing and derive the click valuations by inverting the equilibrium bidding function [53, 6].

In our empirical setting, we observe that advertisers only submit one bid and do not change it (across impressions). Thus, we model advertiser a's bidding decision as a single-shot optimization, where he selects a bid b_a to maximize his own expected utility across all the impressions that he bids on. Let $\mathcal{G}_a$ denote advertiser a's beliefs about the joint distribution of the click valuations and quality scores of other advertisers bidding on the impressions that a is competing for.

Next, we define advertiser a's cost function as the expected payment that he has to make for each click that he receives, given bid b_a, i.e., $c_a(b_a) = \mathbb{E}_{\mathcal{G}_a}[CPC_{ia}]$. Similarly, let $\pi_a(b_a)$ denote advertiser a's expected probability of winning an impression given bid b_a, i.e., $\pi_a(b_a) = \mathbb{E}_{\mathcal{G}_a}[\pi_{ia}]$. Since the allocation function is proportional, we assume that $\pi_a(b_a) = \frac{b_a q_a}{b_a q_a + Q_{-a}}$, where Q_{-a} is a constant reflecting the competitors' bids and quality scores.[13]

[13]While there is no guarantee that $\pi_a(b_a)$ has the quasi-proportional form, it is easy to show by simulated experiments that it is a very accurate approximation. Further, advertisers know that the platform runs a quasi-proportional auction, so it is reasonable to assume that they rely on this functional form.

We can then characterize advertisers' equilibrium bidding strategy in our platform by taking the First Order Condition (FOC) of their expected utility. This FOC can then be inverted to obtain the click valuations as shown in Proposition 2.

Proposition 2 *Consider a platform which runs quasi-proportional auctions where the allocation rule and cost-per-click are given in Equations* (3.1) *and* (2.2) *respectively. Suppose that the cost function* $c_a(b_a)$ *is twice differentiable,* $\{b_a^*\}_{a\in\mathcal{A}}$ *is the set of observed bids, and* $\frac{b_a^* c_a''(b_a^*)}{c_a'(b_a^*)} + 2 \geq 0$ *for all ads. Then, we can write the click valuation as:*

$$v_a^{(c)} = c_a(b_a^*) + \frac{b_a^* c_a'(b_a^*)}{1 - \pi_a(b_a^*)}. \tag{2.14}$$

Proof. *See Appendix §A.7.2.* ■

We can obtain consistent estimates of click valuations from Equation (2.14) as long as we can observe/infer costs and bids from our data. We make three simplifications that make this task straightforward in our setting: (1) advertisers' probability of winning is close zero, i.e., $\pi_a \approx 0$, (2) advertisers' cost-per-click is approximately their bid, i.e., $c_a(b_a) \approx b_a$, and (3) the first-order condition in Equation (2.14) is satisfied for all advertisers, including reserve price bidders.[14] These three simplifications are reasonable in our empirical setting. First, as shown in Figure A.1 in the Appendix, even top ads that won most impressions have a small probability of winning, justifying the first simplification. Second, on average, we find that an advertiser's cost-per-click is over 92% of their bid, which provides support for the second simplification, i.e., $c_a(b_a) \approx b_a$. Finally, the third simplification is also reasonable: only 11 out of 37 ads are reserve price bidders.

With the three simplifications outlined above, Equation (2.14) can be approximated as

$$\hat{v}_a^{(c)} \approx 2b_a^*. \tag{2.15}$$

[14]The equality in Equation (2.14) may be invalid for reserve bidders because they may have submitted a reserve price bid because the platform did not allow them to submit a lower bid. Thus, in the presence of reserve price bidders, the distribution of bids that we see is truncated at the reserve price. In such a situation, we can only infer the truncated distribution of valuations. In Appendix §A.9.1, we discuss how we can address this issue.

All the results presented in the main text are based on click valuations estimated based on Equation (2.15). However, in Appendix §A.9.1, we present six alternative methods to estimate click valuations that progressively relax the simplifications made to derive Equation (2.15). Table A.7 in Appendix §A.9.1 presents an overview of the simplifications relaxed in each of the alternative methods. In particular, the last alternative method employs [105]'s recently proposed estimator for quasi-proportional auctions. His method is fully non-parametric and does not make any of the simplifications listed earlier. We find that the main results remain the same (qualitatively) even when we use these more complex estimators. Therefore, we stick with the simpler estimator in the main text and refer interested readers to the Appendix for these robustness checks.[15]

Recovering Match Values

We now discuss how we can use our estimate of matrix M from our targeting framework to recover match values for any targeting regime. $m_{i,a}$ is ad a's match value for any impression i, if all impressions are distinguishable to her and she is competing for that impression. This follows naturally from our arguments on the accuracy of match value estimates in §2.4.2. However, if two impressions are not distinguishable, the advertiser needs to use the aggregate estimate for that bundle. That is, for any targeting regime $\mathcal{I} = \{I_1, I_2, \ldots, I_L\}$, we can write the match value of advertiser a for impression i in bundle $\mathcal{I}$, $m^{\mathcal{I}}_{i,a}$, as follows:

$$\hat{m}^{\mathcal{I}}_{i,a} = \sum_{j=1}^{L} \mathbb{1}(i \in I_j) \frac{\sum_{k \in I_j} \hat{m}_{k,a} e_{k,a}}{\sum_{k \in I_j} e_{k,a}} \;\; \forall \;\; i \in \mathcal{I}, \tag{2.16}$$

where $e_{k,a}$ are elements of the filtering matrix that allows us to disregard inaccurate estimates and take the average of the rest. Figure 2.10 illustrates how the bundling and aggregation

[15]The underlying theory behind our findings relates to match valuations, and not click valuations: with more granular targeting, advertisers have more accurate match valuations, which in turn, softens the competition and hurts platform's revenues. As such, the heterogeneity induced by allowing more granular targeting comes from the heterogeneity in match valuations, and click valuations are invariant to targeting scenarios. Therefore, using an approximate method to quantify the distribution of click valuations is sufficient for our purpose, and does not change the main findings.

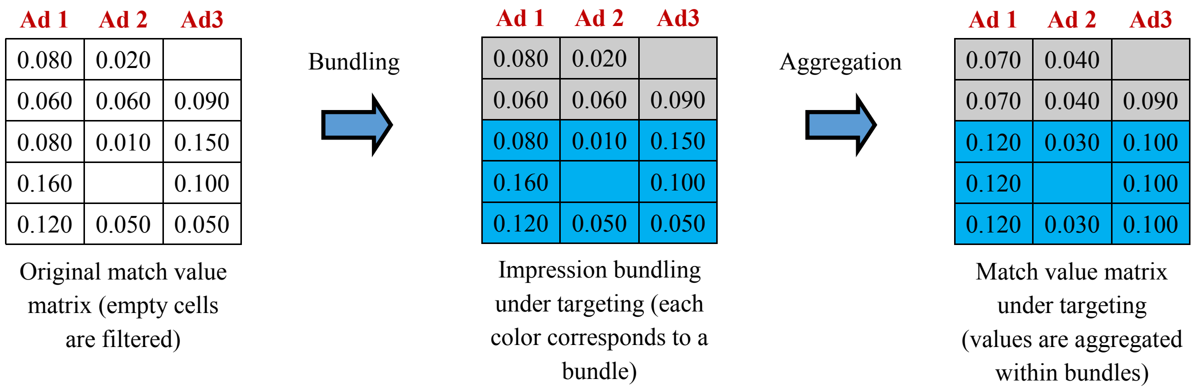

Figure 2.10: Construction of match value matrix under targeting in a simple example with five impressions and three ads.

are performed on the match value matrix in a simple example with five impressions and three ads.

Here, we assume that for any targeting regime $\mathcal{I}$, advertisers can infer their private match values for the bundles at that targeting regime $\hat{m}_a^{\mathcal{I}}$. This is reasonable because if the platform allows impression-level targeting, the platform would automatically share the impression-level data of each advertiser a with that advertiser (but not other advertisers). If a has sufficient data, then a can accurately estimate the match value vector $m_{i,a}$ for impression i from their own data. Similarly, if the platform only allows targeting at level $\mathcal{I}$, then advertisers would automatically have information on which bundle an impression belongs to as well as outcomes (whether impressions in a given bundle received clicks or not), and can therefore accurately infer their match values at the granularity of the bundle. While this assumption always holds from a theoretical standpoint, it may not hold in practice because advertisers need sufficient data to obtain accurate estimates of their match values. Thus, the match value estimates of smaller advertisers and/or new advertisers can be noisy (though they will be consistent). Later, in Appendix §A.9.2, we show that our findings are robust even in situations where advertisers' match value estimates are noisy/imperfect.

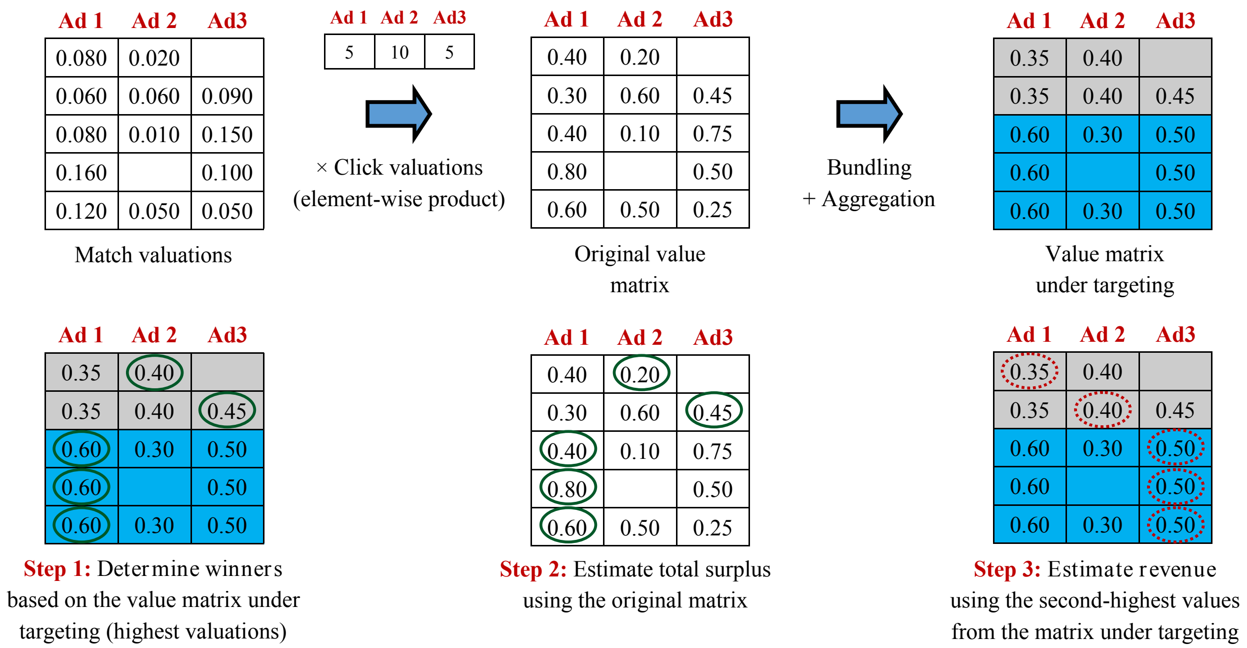

Figure 2.11: Step-by-step procedure to estimate revenue and surplus in a simple example.

Finally, the match value estimates derived from our quasi-proportional auction are assumed to remain the same under a second-price auction. This is reasonable because the match value simply indicates the click probability of a user in a given context for an ad. There is no economic rationale for users' click behavior to be a function of the auction, especially since users often do not know which auction is running at the back-end. Intuitively, click valuations and match valuations are treated as structural parameters.

Estimation of Revenue and Surplus

Given our estimates for click and match valuations, we can obtain estimates of the elements of the valuation matrix V as $\hat{v}_{i,a} = \hat{v}_a^{(c)} \hat{m}_{i,a}$. Further, we can estimate advertisers' expected value of impression i under targeting regime $\mathcal{I}$ as $\hat{v}_{i,a}^{\mathcal{I}} = \hat{v}_a^{(c)} \hat{m}_{i,a}^{\mathcal{I}}$. We now formally discuss our procedure to estimate revenue and surplus for any targeting regime $\mathcal{I}$.

First, we determine the winners of each impression as follows:

$$\hat{a}_i^*(\mathcal{I}) = \underset{a}{\operatorname{argmax}}\, \hat{v}_{i,a}^{\mathcal{I}} e_{i,a}, \tag{2.17}$$

where $\hat{a}_i^*(\mathcal{I})$ is the winner for impression i under targeting regime $\mathcal{I}$. Note that the multiplication by the element of the filtering matrix $e_{i,a}$ simply ensures that the ad is competing in the auction for impression i and that the counterfactual match value estimates are valid, as discussed in §2.4.2.

While the winner is determined using the advertisers' expected value of impression i under a specific targeting regime, surplus is calculated using the actual valuation matrix because it denotes the expected value that would be realized in the system if advertiser $\hat{a}_i^*(\mathcal{I})$ is allocated impression i. So we can write the surplus under targeting granularity $\mathcal{I}$ as:

$$\hat{S}^{\mathcal{I}} = \sum_{i=1}^{N^F} \hat{v}_{i,\hat{a}_i^*(\mathcal{I})} \tag{2.18}$$

To estimate the platform revenues, however, we need to use advertisers' expected values under targeting regime $\mathcal{I}$, as these values guide their bidding behavior. Further, we need to incorporate the fact that the revenue generated from impression i is the second highest bid (or valuation) for it. Thus, the revenue under $\mathcal{I}$ is:

$$\hat{R}^{\mathcal{I}} = \sum_{i=1}^{N^F} \max_{a \backslash \hat{a}_i^*(\mathcal{I})} \hat{v}_{i,a}^{\mathcal{I}} e_{i,a} \tag{2.19}$$

Finally, we can estimate advertiser a's surplus under targeting regime $\mathcal{I}$ as follows:

$$\hat{W}_a^{\mathcal{I}} = \sum_{i=1}^{N^F} \left(\hat{v}_{i,\hat{a}_i^*(\mathcal{I})} - \max_{a \backslash \hat{a}_i^*(\mathcal{I})} \hat{v}_{i,a}^{\mathcal{I}} e_{i,a}\right) \mathbb{1}\left(\hat{a}_i^*(\mathcal{I}) = a\right) \tag{2.20}$$

This estimation is carried out on the Filtered Sample to ensure that our match value estimates are accurate, and hence the averaging in the above equations is done over N_F. Figure 2.11 presents a step-by-step procedure to estimate revenue and surplus for the example case shown in Figure 2.10.

2.7 Counterfactual Results and Privacy Implications

While we can analyze market outcomes for any targeting regime, we focus on the following four targeting regimes that have a one-to-one correspondence with our analysis in §2.5.1:

- No targeting ($\mathcal{I}^{(\mathcal{N})}$) – The platform allows no targeting. As such, there is only one bundle (which constitutes all impressions) and advertisers cannot distinguish between any impressions.
- Contextual targeting ($\mathcal{I}^{(\mathcal{C})}$) – The platform only allows contextual targeting. Here, advertisers can distinguish between impressions in different contexts (app and time). However, impressions from different users in the same context is not distinguishable.
- Behavioral targeting ($\mathcal{I}^{(\mathcal{B})}$) – The platform allows behavioral targeting, thereby allowing advertisers to distinguish between users but not contexts. Here, advertisers can submit bids targeted at the user-level, but cannot distinguish two impressions by the same user in different contexts.
- Full targeting ($\mathcal{I}^{(\mathcal{F})}$) – Platform allows impression-level targeting, i.e., each impression is a bundle and therefore distinguishable. Advertisers can submit targeted bids for each impression.

Using Proposition 1, we can show that $S^{(\mathcal{F})} \geq S^{(\mathcal{N})}$, and that $S^{(\mathcal{C})}$ and $S^{(\mathcal{B})}$ lie in between since both contextual and behavioral targeting can be interpreted as imperfect targeting. However, we cannot theoretically pin down their relative magnitudes because these two types of information are orthogonal (one cannot be interpreted as more granular than the other). So we can only show that: $S^{(\mathcal{F})} \geq S^{(\mathcal{C})}, S^{(\mathcal{B})} \geq S^{(\mathcal{N})}$. Further, we have no theoretical guidance on which of these targeting regimes maximizes platform revenues. We therefore use the empirical framework described in §2.6.3 to derive estimates of platform revenue, advertisers' surplus, and total surplus under the four targeting regimes, and answer the question: "What is the optimal level of targeting that maximizes the platform's revenue?". The results from this exercise are shown in Table 2.5.

Targeting	Total Surplus	Platform Revenue	Advertisers' Surplus
Full	9.45	8.35	1.10
Behavioral	9.18	8.35	0.84
Contextual	8.99	8.44	0.55
No targeting	8.36	8.30	0.06

Table 2.5: Platform revenues, advertisers' surplus, and total surplus for different levels of targeting. The numbers are reported in terms of the average monetary unit per impression.

2.7.1 Platform's Revenue

Consistent with our theory model, our empirical results suggest that more granular targeting leads to higher efficiency in the market: the total surplus under full targeting is 13.02% higher than the no targeting case. Further, in line with our findings in §2.5.1, we find that the total surplus under behavioral targeting is 2.13% higher than the contextual targeting. However, platform revenues exhibit more of an inverted U-shaped curve. They are maximized when the platform restricts targeting to the contextual level. When the platform allows behavioral targeting, advertisers achieve a greater ability to target. While this increases the total surplus in the system, much of this surplus is appropriated by advertisers and the platform's revenue suffers.[16] Thus, the platform's incentives are not perfectly aligned with that of advertisers. Indeed, the platform's optimal targeting level is privacy preserving and aligned with consumers' preferences. We thus find support for the advertising industry's claim that the industry has natural economic incentives to limit user-tracking/targeting and that self-regulation is feasible.

Our findings give rise to many interesting suggestions/ideas on optimal mechanism design and information revelation from the platform's perspective. Limiting targeting to the

[16] Nevertheless, our findings are weaker than those predicted by theory models, i.e., while revenues reduce with more granular targeting, the drop is not very large. This suggests that the strong distributional assumptions on the match values in earlier theory papers (e.g., [65]) may not hold in real ad auctions.

contextual-level is an obvious strategy. However, this approach also reduces the total surplus and hence caps the platform's revenues. Thus, the optimal path for the platform may not be to restrict targeting, but instead to consider mechanisms that can do both – increase efficiency and extract the revenue from winning advertisers by shrinking the informational rent. For instance, the platform could allow behavioral targeting, and also adopt the standard theoretical solution proposed for revenue extraction – optimal reserve prices [96]. [100] validate these theoretical findings using field experiments for search ads. However, they only consider optimal reserve prices for broad sets of keywords and assume CTRs to be homogeneous across advertisers. In contrast, we have a setting where each impression is a unique product and advertisers' match values for an impression are heterogeneous. So, in our case, the platform has to develop a system that can set dynamic impression-specific optimal reserve prices.

2.7.2 Advertisers' Surplus

We begin by comparing total advertiser surplus across the four targeting regimes. As shown in Table 2.5, advertiser's surplus is increasing with more granular targeting. This validates our theoretical prediction that more granular targeting helps advertisers by allowing them to generate more accurate estimates of their match values and place targeted bids. Under full targeting, advertisers' surplus is 11.69% of the total surplus whereas this share drops to 0.69% when no targeting is allowed. Further, the share of advertisers' surplus under behavioral targeting is 8.99% which is considerably higher than 6.13%, their share under contextual targeting. Together these findings emphasize the value of behavioral information for advertisers.

Next, we explore whether all advertisers benefit when their ability to target is enhanced. In a competitive environment, greater ability to target does not necessarily translate into higher profits. Instead, it is the ability to target relative to competitors that matters. In Table 2.6, we show how many advertisers benefit as we move from one targeting regime (column) to another (row).

In general, more advertisers benefit when the platform allows more granular targeting,

To \ From	Full	Behavioral	Contextual	Baseline
Full	NA	23	33	35
Behavioral	14	NA	34	36
Contextual	4	3	NA	33
Baseline	2	1	4	NA

Table 2.6: Number of advertisers who benefit by moving from one targeting regime (column) to another (row) (out of 37 top advertisers).

especially when it allows behavioral targeting. Moving from behavioral, contextual, and no targeting to full targeting benefits 23, 33, and 35 advertisers respectively (first row of Table 2.6). However, more granular targeting is not uniformly better for all advertisers. The first column of Table 2.6 depicts situations where advertisers go from the most granular to less granular targeting regimes. Interestingly, it is populated with positive numbers, which suggests that some advertisers actually benefit from less granular targeting. For example, there are 14 advertisers who prefer behavioral targeting to full targeting. Similarly, while the majority of advertisers prefer behavioral targeting, there is a small portion of advertisers (3) who prefer contextual targeting. We present a simple example to highlight the intuition behind this – a nutrition supplement ad that advertises on a fitness app can get all the slots in that app at a low cost because other advertisers would place low bids when only app-level targeting is allowed. However, this ad would be worse off if only behavioral targeting is allowed, because the competition for users in this app becomes more intense and this ad will no longer be able to extract a large informational rent.

In sum, our findings offer some evidence that advertisers are likely to be differentially affected by privacy regulation on user-tracking and behavioral targeting. Further research on the sources of heterogeneity in advertisers' incentives can help regulators craft the appropriate privacy policies.

2.7.3 Robustness Checks and Limitations

We run a series of robustness checks on the two main components of our estimation – click valuations and match valuations. First, in Appendix §A.9.1, we consider alternative approaches to estimate click valuations from observed bids and show the robustness of our results. Second, in Appendix §A.9.2, we show that our results are robust to the addition of noise to all the match value estimates (to reflect the cases where advertisers realize a noisy version of match value estimates from our ML framework).

Finally, while we have tried to make our analysis as exhaustive and complete as possible, our results should nevertheless be interpreted as short-run counterfactuals with the necessary caveats. First, we assume that advertisers' enhanced ability to target is only reflected in their targeted bidding. In reality however, there might be value creation through other decision variables as well. Second, we assume that the set of ads competing for an impression will not change under different targeting regimes. This implies that there is no entry of new ads or exit of existing ads for an impression. While this assumption may not be realistic, it is unlikely to change the qualitative findings of this chapter. Third, we consider the case where the platform is a monopolist, which reflects our empirical setting. The question of how upstream competition affects privacy-preserving equilibrium outcomes is an important one, but outside the scope of our empirical setting. Finally, all our analysis is static. However, ad-networks can adopt a forward-looking approach in allocate and sell ads. We refer readers to the recent series of work on adaptive ad sequencing that provides frameworks to maximize user engagement [104] and platform revenues [105].

2.8 Conclusions

Mobile in-app advertising is now a dominant ad-format in the digital advertising eco-system. In-app ads have unique tracking properties: they allow advertisers and ad-networks to access the device ID of users' mobile devices, and thereby enable high quality behavioral targeting. While this has made them appealing to advertisers, consumers privacy advocates are concerned about their invasiveness. Therefore, marketers and policy-makers are interested in understanding the relative effectiveness of behavioral targeting compared to contextual targeting, the incentives of ad-networks to engage in behavioral targeting, and the role of regulation in preserving privacy.

We propose a unified framework that consists of two components – a machine learning framework for targeting and an analytical framework for targeting counterfactuals when considering the competition in the market. We apply our framework to data from the leading in-app ad-network of an Asian country. Our machine learning model achieves a RIG of 17.95% over the baseline when we evaluate it on test data. This translates to a 66.80% increase in the average CTR over the current system if we were to deploy an efficient targeting policy based on our machine learning framework. These gains mainly stem from behavioral information and the value of contextual information is relatively small. Next, we build an analytical model of targeting and theoretically prove that although total surplus grows with more granular targeting between the ad-network and advertisers, the ad-network's revenues are non-monotonic in the granularity of targeting. We then take our analytical model to data and conduct a series of targeting counterfactuals and show that the platform prefers to not allow behavioral targeting. There is also some heterogeneity among advertisers' on their preferred level of targeting. Our findings suggest ad-networks have economic incentives to preserve users' privacy in the mobile advertising domain.

This chapter makes several contributions to the literature. First, from a methodological standpoint, we propose a unified framework for targeting that provides counterfactual estimates of platform revenues and efficiency under various targeting regimes. Our framework is generalizable and can be applied to a wide variety of advertising platforms as long as we are able to recover both match valuations and click valuations. In our setting, this is facilitated by the quasi-proportional auction which induces randomness in the allocation of ads over impressions while preserving the linkage between observed bids and click valuations. However, other ad-networks that employ deterministic auctions can also use our framework as long as they randomize ad allocation for a small portion of their traffic.[17] In such cases – (1) the data from the auctions can be used to recover click valuations, and (2) the data from the randomized traffic would satisfy the unconfoundedness assumption because of the exogenous variation in the allocation of ads, and can be used to recover match valuations using our machine learning framework that combines ideas from causal inference and large-scale prediction tasks. Once these two primitives are available, our framework on revenue-efficiency analysis is directly applicable to evaluate market outcomes under different targeting scenarios.

Next, from a substantive perspective, this chapter provides new insights on contextual and behavioral targeting. To our knowledge, this is the first work to study both revenue and efficiency under these two types of targeting. Finally, from a policy point-of-view, we examine the incentives to target for the two major parties in the advertising eco-system: the platform and advertisers. We expect our model and findings to speak to the debate on privacy regulations in the advertising industry.

[17]Indeed, this is standard practice in large ad-networks; e.g., Bing always randomizes ads for a small portion of its traffic [88]. More broadly, all prominent ad-networks now run A/B tests on portions of their data, and this portion of the traffic can be used to infer match valuations.

Chapter 3

DYNAMIC EFFECTS OF ADVERTISING

3.1 Introduction

The smartphone industry has seen unprecedented growth over the past decade, with over three billion users reported having a smartphone at the beginning of 2020 [38]. In 2019, the average time spent on mobile by a US adult approached nearly four hours per day, surpassing the time spent on TV for the first time [37]. The steady growth in smartphone adoption and usage has made mobile an attractive medium for marketing interventions. These interventions are often short-lived in nature, which provides marketers with numerous opportunities to interact with users. As such, mobile users are exposed to a variety of marketing interventions, even in a short period of time.

As mobile users interact with a variety of marketing interventions every day, it is important to understand the consequences of this increased variety in mobile ecosystem. So far, the literature has viewed the increased variety of marketing interventions as a means to explore and learn consumers' tastes in a sequential setting [81]. However, it is not clear how the increased variety itself would affect consumer behavior – would consumers react differently to an intervention when the variety of past interventions is high vs. low? Knowing the answer to this question is of great value to marketers who want to target consumers, and has important implications for the intermediary platforms who design markets for consumers and marketers to interact.

In this chapter, we examine the effects of variety in the context of mobile in-app advertising. In-app advertising is the most popular type of marketing intervention on mobile, generating over $100 billion in ad spend, which makes it the dominant advertising channel in the US [39]. In-app advertising slots are refreshable: each ad exposure lasts a short amount of time (e.g., one minute), and the slot will then be refreshed with another ad. Hence, even in a 10-minute session, a user can be exposed to a variety of different ads. This feature makes in-app advertising particularly suitable for studying how the variety of the sequence of ads seen earlier affects a user's responsiveness to the current ad.

In the advertising context, the effects of variety on users can stem from potential differences in users' information processing when being exposed to high vs. low variety of ads. On the one hand, past behavioral research has shown that consumers' categorization ability increases when they are exposed to a higher variety of objects. This, in turn, can make them attend more to the differentiating aspects of the objects [110]. As such, when exposed to a higher variety of previous ads, users are more likely to actively engage with the next ad, since they can differentiate it more easily. Further, prior studies on eye-tracking have documented that ad repetition lowers consumers' attention during advertisements. Therefore, increasing variety can result in users paying more attention to the ad slot in general, thereby engaging more with the next ad [103]. On the other hand, the literature on information overload argues that excessive information supply makes users less receptive to new information, which in turn depresses their engagement and decision making ability [67, 68]. Thus, we do not have a clear prediction on the direction and source of variety effects in the context of advertising. In this chapter, we seek to address this gap by providing answers to the following key questions:

1. How does variety of previous ads influence users' response to the next ad?
2. Which aspects of variety do contribute more to the outcome? Is it solely the pure diversity in the set of prior ads, or the sequential organization of the set also matters?
3. What is the extent of heterogeneity in the effect of variety across user-level attributes?

To empirically answer these questions, we use impression-level data from the leading mobile advertising ad-network from a large Asian country. The ad-network uses a refreshable ad format where each ad lasts one minute and is followed by another ad exposure. Like other ad-networks, this ad-network runs an auction to allocate ads across impressions. However, the auction differs from the norm in this industry in that it uses a probabilistic allocation rule. More specifically, while the industry norm is to run a first- or second-price auction where the ad with the highest quality-adjusted bid always wins an impression, this ad-network runs a quasi-proportional auction whereby each ad has a non-zero probability of winning proportional to their quality-adjusted bid. Thus, this induces great variation in the allocation of ads, which facilitates studying the effects of variety in our setting.

We face a series of challenges in answering our research questions. First, to quantify the effects of variety on consumers' ad response, we need to measure variety. In particular, we need measures that capture different aspects of variety so we can identify which aspects contribute more to the outcome. Second, given the observational nature of our data, we have to control for the non-experimental variation in users' assignment to variety. There are two sources of this non-experimental variation – (1) *cross-sectional selection*, which is due to the lack of full randomization in users' assignment to different levels of variety, and (2) *dynamic selection*, which stems from sequential nature of variety in our study and users' compliance – the observed data only include the cases where the users have complied to receive variety treatment.

We propose a methodological framework that helps us overcome these challenges and causally measure the effects of variety. Our framework has two key components corresponding to the two challenges that we face. As such, the first component establishes our measures of variety, and the second component presents the identification strategy to address the selection issues. While we exploit the specifics of our setting in our identification strategy, the intuition is fairly general and can be used in settings other than ours, as long as the ad assignment is conditionally exogenous.

In the first component of our framework, we use three different measures of variety – (1) breadth of variety that captures the number of distinct ads shown in a sequence, (2) total number of consecutive changes, which reflects the sequential aspect of variety, and (3) Gini-Simpson index of diversity, which calculates the probability that two random prior exposures are different. Our innovation is in exploring the relationship between these measures, which allows us to decompose the effects of variety into different parts – the part that relates to the pure diversity of the set, and the part that stems from the sequential organization of the set. The ability to distinguish between these two sources is important in the mobile context, where the interventions are sequential and adaptive.

The main goal of the second component of our framework is to address both cross-sectional and dynamic selection problems. We focus on a key distinction between these two selection problems: while cross-sectional selection is caused by the ad-network's decision to show an ad through a known auction, dynamic selection is caused by the users through an unknown underlying mechanism. This implies that the *observed variety* in the data may be confounded by some user unobservables. Instead of imposing assumptions on the randomness of user compliance, we re-define our treatment variable as *intended variety*, which is not necessarily received by the user. Therefore, intended variety is only a function of the auction's ad assignment, but not users' compliance. In principle, this allows us to obtain consistent *Intent-to-Treat (ITT)* estimates of the effects of variety. However, there are still two challenges: (1) we need to control for the cross-sectional selection caused by the auction, and (2) we need to impute the intended variety, since we do not observe it unless users have actually received it.

Cross-sectional selection issue caused by the auction is similar to the main challenge in most observational studies measuring the effects of advertising, i.e., targeting of ads. In our setting, advertisers can target sessions based on their characteristics such as Province, Hour of the Day, Smartphone Brand, Mobile Service Provider (MSP), and Connectivity Type. As a result, sessions that differ in these characteristics may have different distribution to intended variety. We control for this selection problem by stratifying the session space such that each stratum has the same targeting characteristics. As a result, the ad assignment is identical for all the sessions within the strata, thereby making the assignment to the intended variety identical within the strata. Thus, we transform our problem into a stratified randomized experiment where users' assignment to the intended variety is fully random within the strata.

We impute the intended variety for users who have left the session by imputing each single ad assignment. Again, we use the logic behind our stratification: since ad assignment distributions are identical within the same stratum, we can impute the next ads for users who have left the session, using the actual ad assignment in other sessions of that stratum in the full data. This gives us the sequence of ads the user would have been assigned to had the user not left the session. Using the imputed sequence of intended ad assignments, we can impute the intended variety for all the sessions and construct the intended sample needed for obtaining ITT effects of variety.

We specify a regression model to estimate the ITT effects of variety of previous ads. We find that holding all else equal, an increase in the variety of previous ads will lead to higher probability of click on the next ad. The results are robust across different exposure numbers and measures of variety. We also find that holding all else equal, repeated exposures of an ad within the session increase the likelihood of clicking on that ad. We then only focus on the observed sample to evaluate the extent of dynamic selection in our model. We find some evidence for the existence of dynamic selection: users who are assigned to a lower variety are more likely to leave the session. Thus, if we assume random compliance, we will underestimate the ITT effects of variety. However, we find that the extent of bias due to dynamic selection is minimal.

Next, we explore the source for the effects of variety to find which aspects of variety contribute more to the main effect. In particular, we want to examine whether or not the effects of variety solely stem from the pure diversity of ads shown earlier irrespective of how they are sequentially ordered. If this statement is true, two prior sequences like $\langle A, B, A\rangle$ and $\langle A, A, B\rangle$ must generate the same impact on the user's CTR on the next ad. We use the first part of our framework to decompose our measures of variety into sequential and non-sequential elements and compare the contribution of each to the main effect. Through a series of comparisons, we find strong evidence that a larger fraction of the variety effects is attributed to the sequential organization of the set. Thus, our findings suggest that the sequential organization of exposures is a major contributing factor to the effects of variety.

Finally, we document the heterogeneity in the effects of variety across user-level observables. We specifically focus on variables based on users' past behavioral history. We first find that the effects of variety is stronger when users are new to the platform. Similarly, the effects of variety tend to be stronger on users who have seen the current ad more in prior sessions. Finally, we show heterogeneity in the effects of variety based on how users responded to prior ads: users who already expressed some interest in ads (by clicking at least once before) are more likely to be responsive to variety-based manipulations. Interestingly, for users who never clicked on an ad before, the effects of variety are sometimes negative.

In sum, this chapter contributes to the literature in several ways. Methodologically, we propose a framework for studying the causal effects of variety in sequential settings. A novel contribution of our framework is in recognizing two types of selection problems arising in observational data and offering a unified identification strategy based on stratification and imputation that only requires unconfoundedness assumption. Further, our framework presents different ways to decompose the effects of variety in order to understand what fraction of the main effects stems from the sequential organization of objects. Given the increase prevalence of sequential marketing interventions on mobile, we believe this feature of our framework is particularly interesting for practitioners. From a substantive perspective, we establish the effects of variety in the advertising context. In particular, we find that assignment to higher variety of ads in the past will increase users' engagement with ads in the future. To our knowledge, this is the first study on the effects of variety in the advertising context. We further explore the sources for variety effects and identify the sequential organization of exposure within the session as a major source. Overall, we believe that this chapter broadens our understanding of how variety shapes consumer behavior in the digital world.

3.2 Related Literature

This chapter relates and contributes to different streams of literature in marketing and economics.

Broadly, this chapter relates to the literature on measuring the causal effects of advertising. Selection problems induced by the ad allocation mechanism make it hard to obtain unbiased estimates for the effects of advertising [15, 52]. Besides, statistical power is usually very low requiring a data-rich environment for identification of the effects [83, 70]. A series of works have conducted large-scale randomized experiments to answer various research questions while mitigating the above challenges (see [83] for an excellent review of this stream of research). To address these challenges in observational data, however, researchers have employed novel identification strategies such as difference-in-difference strategy to examine the effects of privacy regulations on advertising effectiveness [48], border strategy to quantify the effects of TV advertising [125], regression discontinuity design (RDD) to determine the position effects in sponsored search advertising [99], and within-advertiser exogenous variation in positions due to randomization in Bing's search auctions to study the interplay between position and prominence [69]. This chapter contributes to this literature in two ways. First, we propose an identification strategy for settings where ads are placed through a probabilistic mechanism. Second, we bring behavioral literature on variety and establish the causal link between variety of previous ads shown to the user and her propensity to click on the next ad. To the best of our knowledge, this is the first work to study the effects of ad variety on consumers' ad response.

More specifically, this chapter relates to a somewhat narrower area of research on temporal effects of advertising. This stream of research includes the works on spillover and carryover effects in online advertising [117, 119, 70, ?] and attribution problem [85]. Our work closely relates to [118], as both focus on temporal features of ad exposures, the general idea of temporal sequencing of ads and its impacts on user's information processing. This chapter adds to this stream of literature in two ways. First, we focus on variety of ads as main construct and offer it as an important aspect of ad sequencing. Second, we develop a new approach for observational data to recover intent-to-treat estimates for temporal effects of advertising when compliance is influenced by the intervention. Conceptually, our approach is close to [71] whose predicted ghost ad approach is to transform the intended sample into that of compliers to estimate average treatment effects on the treated. However, we do the opposite by constructing the intended sample in order to recover intent-to-treat estimates.

Finally, this chapter relates to the marketing literature on the concept of variety. We can broadly categorize prior research on variety into two streams based on the role of variety in the analysis. In the first stream of work, variety serves as the main outcome of interest and consumers can actively choose different levels of variety. This includes studies on consumer's demand for variety [76, 25] and variety-seeking behavior [91, 109]. The second stream, however, views variety as a factor influencing the outcome, such as the link between variety of assortments and store choice [63], variety of episodes and consumer's engagement [110], and the dispersion of word-of-mouth and TV ratings [46]. Related to the second stream of research, we bring the notion of variety to the advertising context and examine whether variety of ads seen earlier affect user's ad response. To our knowledge, the only study related to variety in this context is [122] who show the variation of ad content over a repeated advertising schedule will increase user's responsiveness to that ad. While they only focus on variation in the content for one ad in a lab setting, our work extends it to the variety of potentially competing ads in a large-scale in-app advertising market. Further, this chapter adds to the marketing literature on variety by providing a framework for decomposing the effects of variety and developing a new measure of sequential variety that can be used for cases where objects are presented in a sequence.

3.3 Setting and Data

We now describe our setting and data. In §5.4.1, we discuss the setting, in §4.3.2 and §3.3.3 we describe the data, sampling, and the data generating process in detail. Finally, in §3.3.4, we present some summary statistics and descriptive analysis.

3.3.1 Setting

Our data come from the leading mobile in-app advertising network of a large Asian country, which had over 85% marketshare during the time of this study. The ad-network functions as a match-maker between advertisers and app developers and serves ads inside mobile apps.

The scope and scale of the ad-network are quite large; it serves more than 10,000 unique apps that generate a total of over 50 million ad impressions daily.

Main Players

We begin by describing the four main players in this marketplace:

- *Users:* Individuals who generate eyeballs or impressions by using their mobile apps. For each ad impression shown to her, the user can decide whether to click on the ad or not.
- *Publishers:* App developers who have joined ad-network. Their revenue comes from the clicks generated in the ads shown in their apps. This is the main monetization strategy for most of the apps in our data.
- *Advertisers:* Firms that buy ad slots in the mobile apps served by the ad-network. Advertisers create banner ads (in either JPEG or GIF formats), submit a bid that indicates their willingness to pay per click, and a maximum daily budget. Advertiser can also target their ads based on the following variables associated with the impression: province, app category, hour of the day, smartphone brand, mobile service provider, and connectivity type. We discuss targeting in greater detail in §3.3.3.
- *Ad-network:* The platform that facilitates the matching between ads and impressions generated by a user-app combination. The ad-network runs a real-time auction to select the ad to show in each impression. Like publishers, the ad-network's revenue comes from clicks: for each click generated in the system, the ad-network retains 30% of the price that the winning advertiser paid for it, and the gives the remaining 70% of the price to the publisher of the app where the impression came from.

Ad Refresh

Mobile in-app ad slots are distinct from standard desktop display ads in one important dimension – they are *refreshable*, i.e., the ad-network has the ability to refresh or change the ad during a user's app session. This is different from the standard display ads shown in websites, where the ad shown in an ad-slot is fixed for the duration of a user's visit. So the definition of an impression in mobile in-app advertising is different from that in desktop display advertising. In the latter case, one impression refers to one unique website visit. In the former, each impression refers to a fixed period of time during which the ad slot shows the same ad.[1]

Our ad-network employs refreshable ad-slots, where each impression lasts one minute. When a user starts using an app, the ad-network runs an auction to determine the winning ad and serves this ad for one minute. If the user continues using the app beyond one minute, the ad-network treats this as a new impression and runs another auction to determine the next ad to show the user.

The use of refreshable ad slots in mobile apps (unlike fixed ad slots in desktop settings) stems from a few contextual distinctions between mobile app usage and desktop-based browsing. First, users tend to spend long periods of time on mobile apps, which makes it possible to show them multiple ads within the same session. In contrast, users spend only a short time on a website before scrolling/clicking, which makes it unnecessary to change the ad on a given website during the browsing session. Second, mobile screens are generally small, and users tend to be closer to the screen. As a result, ads are harder to ignore in mobile apps. So it can be irritating for the user to see the same ad within the app that she is using for minutes at a stretch. In contrast, the user only stays on the website for a short time and is unlikely to be bothered by seeing the same ad for a short time on a large screen. Finally, publishers and ad-networks view eye-balls as a valuable resource; so replacing/refreshing ads within mobile app sessions allows them to monetize the app better.

Definition of a Session

We now define a session, which is the main unit of analysis in this chapter. A session is a set of consecutive impressions generated by a user within an app, such that the gap between

[1]Recently, some ad-networks like Google Display have started allowing publishers to refresh ads on their web page after a fixed period of time or after a specific action taken by the user [9].

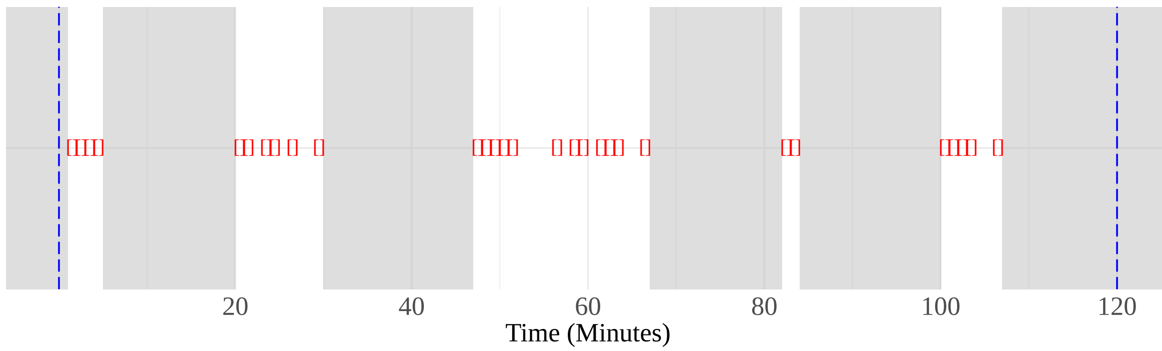

Figure 3.1: Definition of the Session

two consecutive impressions is less than 10 minutes. Figure 3.1 shows a two-hour snapshot of a user's exposure to ads to help illustrate the main idea behind the definition of sessions. Impressions are visualized by their starting and ending points, depicted using red [and] respectively. The grey areas show the time between the session, i.e., the length of each grey area exceeds 10 minutes. Overall, we see that this user has initiated five separate sessions of varying lengths. For example, the third session in this picture lasts for 20 minutes and contains 10 impressions, whereas the fourth one lasts only for two minutes (and contains two back-to-back impressions). Not all impressions within a session are back to back, i.e., there can be some gap in time ($<$ 10 minutes) between consecutive impressions within a session. For example, in the third session, there is a five minute gap between the fifth and sixth impressions.

3.3.2 Data

We have data on all the impressions and corresponding clicks (if any) in the platform for all the participating apps for a one month period from 30 September 2015 to 30 October 2015.

Raw Variables

For each impression in the data, we have access to the following information:

- Time and date: The time-stamp of the impression.
- AAID: Android Advertising ID is a user re-settable, unique, device ID that is provided by the Android operating system.[2] It is accessible to advertisers and ad networks for tracking and targeting purposes. We use it as the user-identifier in our main analysis.
- App ID: A unique identifier for apps that advertise through the platform.
- Ad ID: This is an identifier for ads that are shown to the users.
- Bid: The bid that the advertiser has submitted for her ad.
- Province: The province/state where the user is located when the impression occurs.
- Connectivity type: It refers to the user's type of connectivity (e.g., WiFi or cellular data)
- Smartphone brand: The brand of the user's smartphone (e.g., Samsung, Huawei, etc.)
- MSP: The user's mobile-phone service provider.
- Click indicator: This variable indicates whether or not the user has clicked on the ad.

Notice that all the variables that the advertiser can target on (province, app category, hour of the day, smartphone brand, MSP, and connectivity type) are also observable to us. Thus, we are able to avoid many of the common problems related to endogeneity in the measurement of ad effects due to targeting based on unobservables. We will discuss how we incorporate role of targeting in our identification strategy in greater detail in §3.5.3.

Sampling Procedure

The total data available is too large to work with, and all the data are not amenable to our analysis. So we draw a sample from our full data to conduct our study. We now discuss our sampling procedure and explain the rationale for it. Recall that our goal is to study the effects of the variety of ads shown earlier in the session on user's response to the current ad. So any sampling procedure we adopt should be at the session-level, i.e., include all the impressions within a session. With this basic criterion in mind, we sample a set of sessions

[2]Apple's app store is not available in the country where our data are sourced from. Hence, all smartphones use the Android operating system.

from our data as follows:

- *Users:* One of the goals of our analysis is to examine the heterogeneity in users' response (to the variety of previous ads) as a function of their observed behavioral history. To do so, we need the full history of users. The challenge is that, for some users, the first observed activity in our data goes back to 30 September 2015, which is the first day of our observation period. For these users, activity logs are truncated, and we cannot generate accurate measures of their past activity. To identify users with untruncated history, we split the data into two parts and make two sets of users: (1) $\mathcal{U}_1$ that consists of users generated at least one impression from 30 September 2015 to 21 October 2015, and (2) $\mathcal{U}_2$ that consists of users who generated at least one impression from 22 October 2015 to 30 October 2015. We then sample all users who are available in the second set but not in the first set, i.e., $\mathcal{U}_2 \setminus \mathcal{U}_1$. The fact that the user was not available in the first set suggests that we have the entire observed activity for that user.
- *App:* Within the set of users who satisfy the above condition, we then exclusively focus on their impressions in the most popular app. This is a messaging app that is widely used in the country, and generates over 30% of the total traffic observed in the ad-network. Just focusing on the top app allows us to hold the context of the app constant and cleanly derive the causal effect of the variety of previous ads.[3]

Overall, our working sample consists of 96,076 users who generate 1,291,477 sessions and 7,066,483 impressions in the messenger app. The average CTR in this data is 0.022. All our summary statistics and analysis are shown for this sample. However, we utilize the data from other users and apps to supplement our analysis. For the sample of users that we focus on, even though we use only data from the top app for our analysis, we keep track of the data generated by these users in other apps to segment them based on their behavioral history and explore the heterogeneity in their responsiveness to ad variety. Please see §3.6.4 for a detailed discussion of measures of user-level heterogeneity. Further, we also use the

[3] In principle, we can include more apps in the analysis as long as we have a sufficient number of impressions within each app so that we can stratify the data at the app level to avoid power issues [83].

impressions from other users not in our sample in the top app for imputation purposes in our estimation procedure. Please see §3.5.3 for details.

3.3.3 Data Generating Process

We now describe the data generating process in our setting. Let i denote a session and t denote an exposure number within a session. Each exposure t in session i in our data comes with three pieces of information: (1) impression-level characteristics ($X_{i,t}$), (2) ad shown in the exposure ($A_{i,t}$), and (3) the user's decision to click on the ad shown in the exposure ($Y_{i,t}$). We define our data as $\mathcal{D} = \{(X_{i,t}, A_{i,t}, Y_{i,t})\}_{i,t}$ where each element is an exposure in a session with the three sets of variables characterized above. We now describe the underlying data generating process behind each of these three components:

- *Impression-level characteristics ($X_{i,t}$):* They capture all the observable attributes associated with the user and context of the impression, e.g., the smartphone brand or province of the user, or the hour of the day when the impression occurs.
- *Ad allocation ($A_{i,t}$):* Whenever an impression is generated, a request is sent to the ad-network's server to fill that impression with an ad. The ad-network uses a *quasi-proportional* auction mechanism to allocate ads to impressions [94]. The main distinction between a quasi-proportional auction and other commonly used auction mechanisms (e.g., Second Price or Vickrey) is the use of a probabilistic winning rule, i.e., all ads participating in the auction for an impression have a probability of winning of the impression. In our setting, the probability or propensity that ad a wins the t^{th} impression in session i is denoted by $\pi_{i,t}(a)$ and defined as:

$$\pi_{i,t}(a) = \mathbb{1}(a \in \mathcal{C}_{i,t}) \frac{b_{i,a} m_{i,a}}{\sum_{k \in \mathcal{C}_i} b_{i,k} m_{i,k}}, \tag{3.1}$$

  where $\mathcal{C}_{i,t}$ is the set of ads participating in the auction for the t^{th} exposure in session i, and $b_{i,a}$ and $m_{i,a}$ are advertiser a's bid and quality score in session i respectively.[4] Thus,

[4]The bid and quality score are fixed within a session, and hence are not indexed by t.

the variation in $\pi_{i,t}(a)$ can stem from variation in $\mathcal{C}_{i,t}$, $b_{i,a}$, $m_{i,a}$, or some combination of these variables. The variation in $\mathcal{C}_{i,t}$ is mainly driven by advertisers' targeting decisions. For example, if ad a chose not to target the province where the user generating session i is located, then ad a will not belong to $\mathcal{C}_{i,t} \Rightarrow \pi_{i,a} = 0$. The quality score $m_{i,a}$ is a measure of profitability that the platform assigns to each ad a in session i. The extent of customization in the quality scores is quite low: the ad-network simply assigns one aggregate quality score to each ad and only updates it once in a day. So while $m_{i,a}$ can vary across sessions, the extent of variation is quite low. Finally, $b_{i,a}$ is ad a's bid for session i. It is important to notice that each ad could submit only one bid at a given time of the day.[5] If the advertiser changed the bid at some point in time, it was updated for all the sessions that started in the next hour of the day.

Together, these three components determine the distribution of propensity scores ($\pi_{i,t}(a)$s) for ads at the t^{th} exposure in session i in our problem setting. As such, $A_{i,t}$ is a draw from this probability distribution.

- *Click outcome ($Y_{i,t}$):* When an impression with $X_{i,t}$ is served with ad $A_{i,t}$, the user generating the impression may choose to click on the ad. We denote the click outcome using $Y_{i,t}$, which is a binary variable that takes value one when a click happens and zero otherwise. Our goal is to learn how interventions in the prior variety of ad shift $Y_{i,t}$.

A key advantage of our quasi-proportional auction mechanism is the use of a probabilistic allocation rule, which is different from the common practice of using deterministic auctions (e.g., second- or first-price auctions) in this industry. This gives us in considerable variation in the ad assignment across and within sessions. More importantly, this variation is a function of simply observables. Thus, in light of Equation (3.1), we can show the following proposition.

Proposition 3 *The distribution of propensity scores for ads, $\pi_{i,t}(a)s$, for any impression is only a function of impression-specific observables, $X_{i,t}$, in the data.*

[5]Advertisers could not customize bids by targeting variable. For example, they could not submit different bids for two different provinces at the same time, even if their willingness to pay for the clicks in the two provinces was different.

Proof. *See Appendix B.1.* ■

This is a crucial result since it ensures that there are no user or impression-specific unobservables that affect ad assignment that are observable to advertisers (and the ad-network), but not to the researcher. An immediate corollary of Proposition 3 is that we can fully control for any non-random selection in the assignment of ads. Our setting is thus similar to a randomized experiment where the treatments (or ads in this case) are assigned using a known propensity score distribution. Later in §3.5.3, we use this result to transform our problem into a stratified randomized experiment, where the probability of being allocated a specific ad is fixed within a stratum and is only a function of observables.

3.3.4 Summary Statistics

We now describe present some summary statistics for the key variables of interest and some simple descriptive analysis.

Targeting Variables

As discussed in §3.3.1, targeting variables are the dimensions that advertisers can target their ads based upon. In our setting, these consist of: province, app category, hour of the day, smartphone Brand, MSP, and connectivity type. All these variables are categorical and advertisers can specify which sub-categories they want to show their ads in, e.g., an advertiser can indicate she want her ads to be shown only from 6 pm to 9 pm every day on Samsung phones in one specific province.

We first report the impression share of the top three sub-categories within each targeting variable in Table 5.1. All these numbers are shown for the sample that we use for analysis.[6]

Next, we illustrate whether the impression share of a targeting subcategory (e.g., Samsung in the category of Smartphone Brand) is correlated with the extent to which that subcategory is targeted. We know that if a subcategory is excluded in an ad's targeting, no impression

[6] Since our sample only consists of one app, we exclude it from Table 5.1.

Variable	Number of categories	Share of top categories 1^{st}	2^{nd}	3^{rd}	Number of impressions
Province	31	24.55%	9.52%	7.42%	7,066,483
Hour of the day	24	8.47%	8.02%	7.26%	7,066,483
Smartphone brand	7	44.72%	38.12%	6.65%	6,416,152
Connectivity type	2	50.23%	49.47%		7,066,483
MSP	3	50.14%	44.04%	5.81%	6,890,873

Table 3.1: Summary statistics of the targeting variables. The number of non-missing observations for each variable are shown. While the information about province and hour of the day is always available, other variables are missing for some impressions. The shares shown are computed after excluding the missing observations for each variable.

of that ad will be shown in that subcategory. As such, for each targeting subcategory, we calculate the number of distinct ads that show at least one impression in that subcategory. For all the subcategories, we plot the number of distinct ads targeting that subcategory against the impression share of that category in Figure 3.2. Subcategories within each targeting category are shown as points with the same color and shape. The grey dashed line in the top shows the total number of distinct ads in our study. As such, the points that are very close to the dashed line are almost included in all ads' targeting.

Three important patterns emerge from Figure 3.2. First, we find that some variables are less used for targeting. In particular, all the subcategories within the Connectivity Type, MSP, and Smartphone Brand are close to the dashed line, which means that advertisers simply show their ads irrespective of these variables. While smartphone brands slightly differ in the number of ads targeting them, the extent of targeting is limited. On the other hand, Province and Hour of the Day are the main variables used for targeting: all subcategories within these variables are considerably different in terms of the number of distinct ads targeting them. The second important insight form Figure 3.2 relates to this

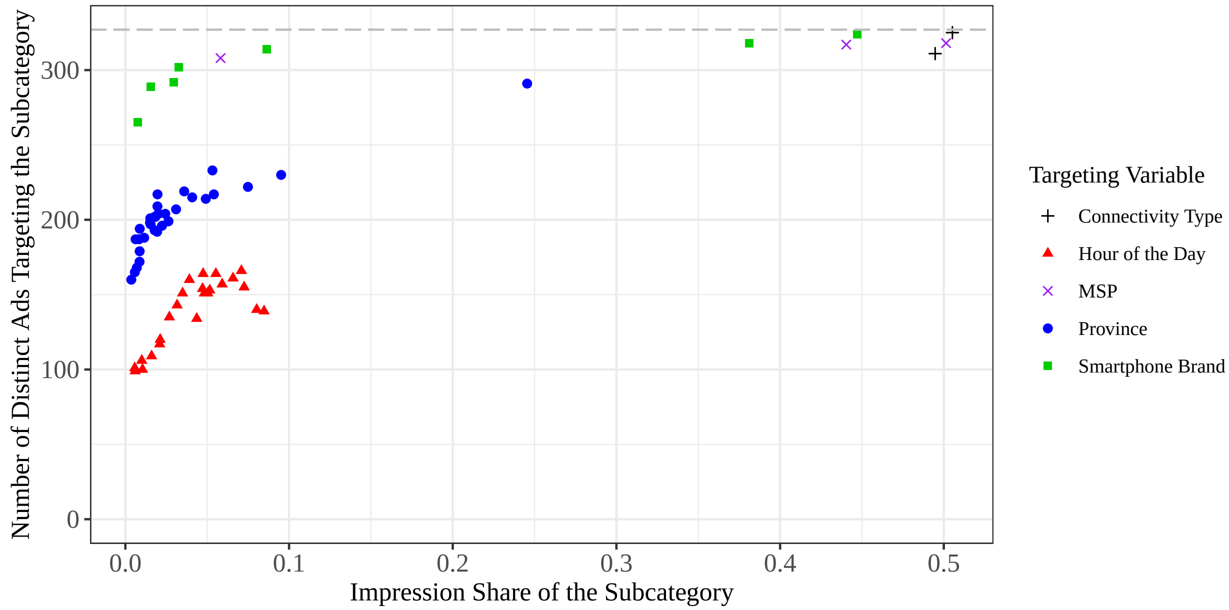

Figure 3.2: Relationship between the number of distinct ads targeting a subcategory and the impression share of that subcategory within the targeting variable. All subcategories within each targeting variable are in the same color and shape. The dashed grey line on the top is the total number of distinct ads available in our study.

difference: subcategories with greater impression shares seem to be more popular among advertisers. For example, the blue circle in the middle of Figure 3.2 refers to the largest province in the country that has the highest impression share among provinces as well as the number of distinct ads. On the other hand, the red triangles in the bottom left of the figure are midnight hours that have lower impression shares and fewer advertisers. However, it is worth noting that even these unpopular hours attract a lot of ads. This bring us to the third important insight from Figure 3.2 that shows there is no niche targeting in this market. This feature results in great within-session variation in ads in almost all sessions, which allows us to study the effects of variety of ads.

Variation in Click-through Rate (CTR)

Click is the main outcome of interest in our study. Here we want to show how CTR varies within and across sessions. We start with the most basic graph that shows within-session variation in CTR. In Figure 3.3a, we show average CTR in each exposure number in a session. It reveals a downward trend in CTRs over exposure number, suggesting that users who have continued for long periods of time are less likely to click. Hence, it is essential to control for the exposure number when studying the effects of within-session interventions such as variety.

Next, we show how CTR varies across sessions. In Figure 3.3b, we show the plot of average CTR in each targeting subcategory against the number of distinct ads targeting that subcategory. First, by looking only at the y-axis, we find quite a bit of variation in CTR across provinces and hours of the day. However, the variation in CTR is minimal in subcategories of other targeting variables such as Smartphone Brand, MSP, and Connectivity Type. In general, the variation in CTR shrinks as we move towards the right of Figure 3.3b, where subcategories are targeted by default by all ads. This intuitively makes sense because advertisers want to use variables for targeting that have considerable within-variation in CTR. Finally, we observe a slightly positive correlation between the number of distinct ads targeting a subcategory and its CTR, which suggests that subcategories with higher CTR are more popular among advertisers. This can potentially create cross-sectional selection problems in studying the effects of variety on users' CTR. Later in §3.5, we discuss this selection issue in greater detail and provide our solution to address it.

Session-level Metrics

Sessions are of great importance in our study since we want to measure the causal effects of variety of ads shown within a session on users' ad response. Here we show the distribution of two key metrics across sessions: (1) session length, which is defined as the number of impressions (or exposures) that were served during a session, and (2) number of distinct ads

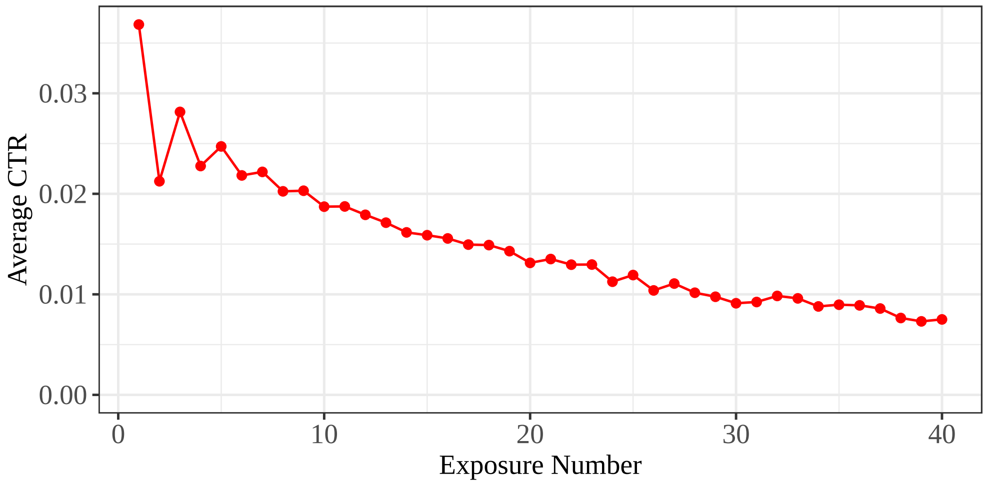

(a) Within-session variation in CTR: average click-through rate (CTR) across different exposure numbers.

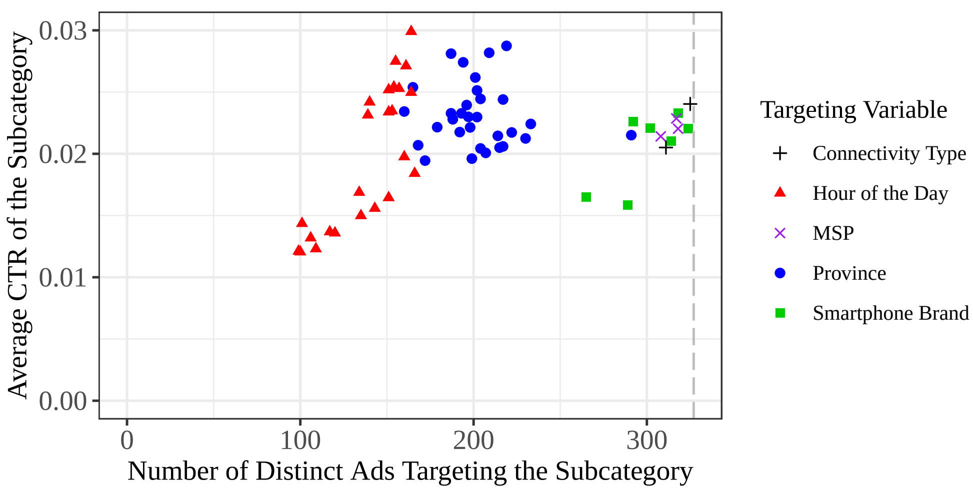

(b) Across-session variation in CTR: relationship between the subcategory's CTR and the number of distinct ads targeting that subcategory.

Figure 3.3: Variation in CTR.

shown in a session. We show the empirical CDF of both in Figure 3.4. We see significant variation in the length of the sessions in our data, as shown in Figure 3.4a. Nearly half of the sessions end after the first two impressions/exposures. Further, a vast majority of sessions last for less than 10 impressions. However, we see some very long sessions during which the user sees over 60 ad exposures. This implies that the distribution of session length has a long tail.

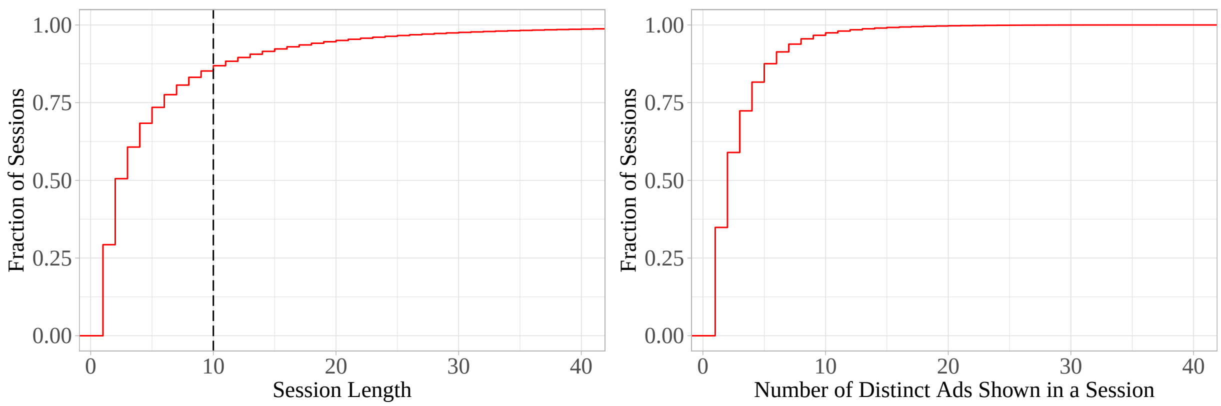

(a) Session length (truncated at 40). (b) Number of distinct ads shown in a session.

Figure 3.4: Empirical CDF of two session metrics.

Next, we focus on Figure 3.4b that shows the empirical CDF of the number of distinct ads shown within a session. This distribution also exhibits great variation, which highlights the unique feature of the auction mechanism discussed in §3.3.3. This variation in the number of distinct ads is necessary for our goal of estimating the effects of variety of ads.

Share of Ads

As mentioned earlier, the platform runs a probabilistic auction that determines the winner for each impression using a probabilistic rule. As such, all ads participating in the auction for an impression have a chance to win. However, the propensity to win the auction for an impression varies significantly across ads due to three main factors – bid, budget, and targeting. First, ads that have a higher bid have a higher probability of winning the auction as shown in Equation (3.1). However, if an ad has a high bid but a fixed budget, then it would win impressions earlier in the day, but may run out of the budget soon as a result of getting clicks. Thus, they will be out of competition for later impressions and may only get a few impressions overall. In some cases, the ad may have a high bid and also set a high budget, but target more narrowly. Then, in spite of the high bid and budget, it will still get a smaller share of the impressions in the ad-network because it only competes for (or targets) a small set of impressions. Thus, the bid, budget, and targeting decisions of the ad together determine the share of total impressions awarded to it.

We see a total of 327 distinct ads in our data. We first calculate the shares of a given ad as the fraction of impressions showing that specific ad in the total set of impressions in our data. We then sort the ads by their share and present the CDF of the shares in Figure 3.5. Note that most ads only get a tiny fraction of impressions due to the reasons outlined above. However, a few ads account for the vast majority of impressions. Specifically, the top 50 ads account for over 90% of the total impressions.

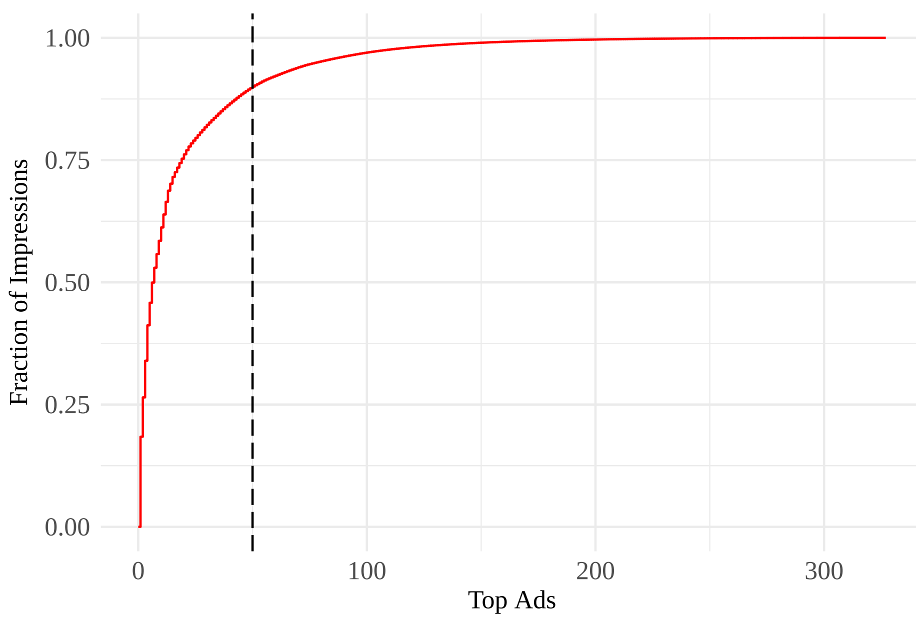

Figure 3.5: Cumulative fraction of impressions associated with ads.

3.4 Variety of Previous Ads

Our main goal is to examine how the variety of previous ads affects user's clicking behavior on the next ad. As such, variety of previous ads is the key independent variable of interest in this study, and the focus of this section. We first discuss the conceptualization of variety over a sequence of ads and introduce a few measures of variety in 3.4.1. Next, in §3.4.3, we present some summary statistics on the different variety measures in our data. Finally, in §3.4.4, we offer some model-free evidence from our data linking variety of previous ads to user's clicking behavior.

3.4.1 Measures of Variety

Quantifying variety is a challenging task because consumers' perception of variety can vary depending on the context and the information structure of the assortment [63]. Therefore, we do not have a universal measure of variety that is applicable to all settings. Instead, researchers have developed a number of metrics that capture different notions of variety such as the breadth of variety, diversity, and concentration [56]. In our setting, variety is defined over the prior sequence of ads shown in the session. Figure 3.6 presents a example of such a sequence, where a user at the ninth exposure in the session, and has seen a sequence of eight ads prior to the current exposure.

We now introduce some notation. Let i denote sessions and t denote the exposure number within a session. We define the sequence of ads shown in session i as $\{A_{i,t}\}_{t=1}^{T_i}$, where $A_{i,t}$ is the ad shown in t^{th} exposure in session i, and T_i is the total number of exposures shown in session i. We now present measures that capture the variety of a sequence of $t-1$ ads.

- *Breadth of variety*: This metric counts the number of distinct or unique ads shown within the session so far. As such, it captures the breadth of variety. For exposure number t in session i, we define the breadth of variety as:

$$Breadth_{i,t} = |\{A_{i,1}, \ldots, A_{i,t-1}\}|, \tag{3.2}$$

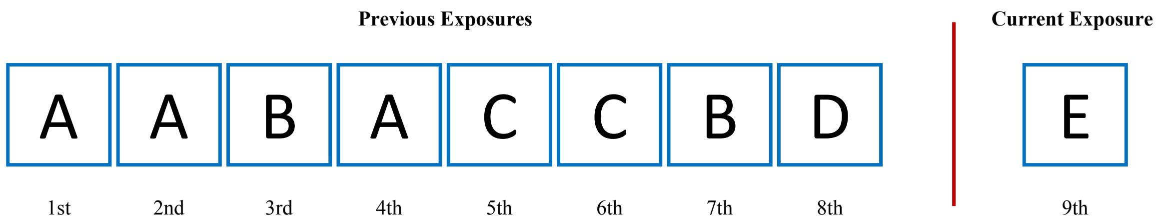

Figure 3.6: An example of a session with where the user is at the ninth exposure. The numbers represent exposure number t and each rectangle represents the ad shown in that exposure. The letter coding refers to ad IDs, i.e., each letter represents one unique ad. For instance, the user is shown the same ad (coded in letter A) during the first, second, and fourth exposures.

where $A_{i,t}$ denotes the ad shown in the t^{th} of session i. For example, the user has seen four distinct ads $\{A, B, C, D\}$ prior to the current exposure in Figure 3.6; so $Breadth_{i,9} = 4$.

- *Consecutive changes*: Since ads are replaced sequentially in our setting, users never see the full sequence at the same time. However, the sequential organization of exposures may still influence users' perception of variety. This is in line with prior research, which shows that show the organization and structure of the assortment play an important role in the user's perceptions of variety [63, 72]. One measure of variety that captures the sequential organization of the session is the total number of consecutive exposures that show different ads. Formally, consecutive changes is defined as:

$$Change_{i,t} = \sum_{j=2}^{t-1} \mathbb{1}(A_{i,j} \neq A_{i,j-1}), \tag{3.3}$$

  In Figure 3.6, we see five consecutive changes (the identity of the ad shown changes at each exposure, except the second and sixth exposures); so $Change_{i,9} = 5$.

- *Gini–Simpson index*: Next, to capture the diversity of the set of ads shown earlier, we use the Gini-Simpson index, which simply calculates the probability that two random prior exposures showed different ads. We use this measure because it has a nice mathematical interpretation. For exposure number t in session i, we can define the Gini-Simpson index as follows:

$$GiniSimpson_{i,t} = \frac{\sum_{s<s'<t} \mathbb{1}(A_{i,s'} \neq A_{i,s})}{\frac{(t-1)(t-2)}{2}}, \tag{3.4}$$

where the numerator calculates all the pairs with two different ads, and the denominator is the total number of pairs of prior exposures. As such, $GiniSimpson_{i,t}$ measures the probability that a random pair of prior exposures showed different ads. This measure is closely linked to the Herfindahl-Hirschman Index (HHI), which is the most commonly used concentration metric in marketing and economics literature.[7] Moreover, the Gini-Simpson index is only a function of the set of the ads; the actual sequence does not matter (unlike *Change*).[8]

In Figure 3.6, $GiniSimpson_{i,9} = 0.82$, and it denotes the probability that two random exposures show different ads. Gini-Simpson is higher if the shares of ads are more evenly distributed. For example, if we change the fourth exposure in Figure 3.6 from A to D, all ads will have equal share (0.25) and Gini-Simpson index would increase to ≈ 0.86. If we show a distinct ad in each exposure, the Gini-Simpson measure would take the highest possible value of 1.

Together, these three measures capture the notions of ad uniqueness and the spatial spread of ads over the sequence; both of which have been shown to have a significant effect of consumers' perception of variety [63]. Finally, note that our variety measures are based on the identity of ads and not ad attributes. There are two reasons for this. First, we do not have access to ad creatives. Second, unlike a standard super-market product setting, where we can define similarity measures on standard product attributes (e.g., price, on sale or not, flavor, local/national brand, product size), there are no standard features to compare banner ads on. Moreover, the banner ads in our setting quite small with simple ad creatives (a short

---

[7] Gini-Simpson index is based on the original Simpson's entropy measure that calculates the probability of the complementary event, i.e., the probability that two random exposures show the same ad [**?**]. HHI also calculates this probability when we have infinite samples.

[8] For instance, the Gini-Simpson index is the same even if we change the sequence shown in Figure 3.6 to $\{A, A, A, B, B, C, C, D, D\}$. However, in this case $Change_{i,9} = 3$.

text with some JPEG or GIF). Therefore, variety measures based on ad features are unlikely to provide additional information over and beyond those based on ad identities. Nevertheless, all our variety measures can be easily extended to ad features if the researcher/manager has access to them.

3.4.2 Relationship Between Different Measures of Variety

Each one of the measures in §3.4.1 captures certain aspects of variety. For example, *Breadth* focuses on the distinctiveness of the sequence of prior ads, whereas *GiniSimpson* also takes into account how evenly distributed are these distinct ads. In order to understand how different components of variety contributes to its overall impact, we need to explore the relationship between these measures. Here we focus on our three measures of variety and show how each is related to the other ones:

- *Relationship between Breadth of Variety and Consecutive Changes:* Breadth of variety is the number of distinct ads shown in a sequence. Obviously, the first time we show each one of these distinct ads, we make a consecutive change. Hence, the total number of such consecutive changes equals the breadth of variety minus one, because we do not count the first exposure number in the total number of changes. As such, we can decompose the total number of consecutive changes into two parts – *breadth-increasing changes* and *breadth-constant changes*. The former is the total number of consecutive changes that increase the variety of the sequence (i.e., the second ad in the pair has not been shown before within the sequence), whereas the latter is the total number of consecutive changes that do not increase the variety (i.e., the second in the pair has been shown before within the sequence). Let $BIC_{i,t}$ and $BCC_{i,t}$ denote the total number of *breadth-increasing changes* and *breadth-constant changes* respectively. We can write:

$$Change_{i,t} = BIC_{i,t} + BCC_{i,t} = (Breadth_{i,t} - 1) + BCC_{i,t} \tag{3.5}$$

  The relationship in Equation (3.5) allows us to disentangle the effects of consecutive change from the breadth of variety.

- *Relationship between Consecutive Changes and Gini-Simpson Index:* The Gini-Simpson index measures the probability that two random exposures from a sequence show different ads. As shown in Equation 4.12, it takes all the pairs of prior ads and calculates what fraction of them that show different ads. We now decompose these pairs into consecutive and non-consecutive pairs, and re-write the Gini-Simpson index as follows:

$$GiniSimpson_{i,t} = \underbrace{\frac{\sum_{s+1=s'<t} \mathbb{1}(A_{i,s'} \neq A_{i,s})}{\frac{(t-1)(t-2)}{2}}}_{\text{consecutive diversity } (CD_{i,t})} + \underbrace{\frac{\sum_{s+1<s'<t} \mathbb{1}(A_{i,s'} \neq A_{i,s})}{\frac{(t-1)(t-2)}{2}}}_{\text{non-consecutive diversity } (NCD_{i,t})} \tag{3.6}$$

This decomposition allows us to understand which aspect of diversity is more important. It further reveals the relationship between $GiniSimpson_{i,t}$ and $Change_{i,t}$, as the numerator in consecutive diversity is $Change_{i,t}$.

3.4.3 Variation in Variety

A necessary condition to identify the effects of variety is to have sufficient variation in the variety measures. To visualize the extent of this variation, we first plot the empirical CDF of $Breadth_{i,t}$ at different exposure numbers in Figure 3.7. Notice that there is substantial variation in $Breadth_{i,t}$ at all exposure numbers.

This variation can stem from two possible sources. First, the proportional auction mechanism naturally induces variation in variety (if there are enough ads competing in a given session) and this variation is random conditional on the set of competing ads and helps with identification. Second, this variation could stem from variation in the set of ads competing for any given session. That is, some sessions have more ads competing for them compared to others, which naturally results in a higher variety of ads in those session. This type of variation is non-random since the ads competing in a given session are self-selected. Thus, we have to control for this selection in our modeling exercise. We discuss the latter issue in greater detail §3.5.2.

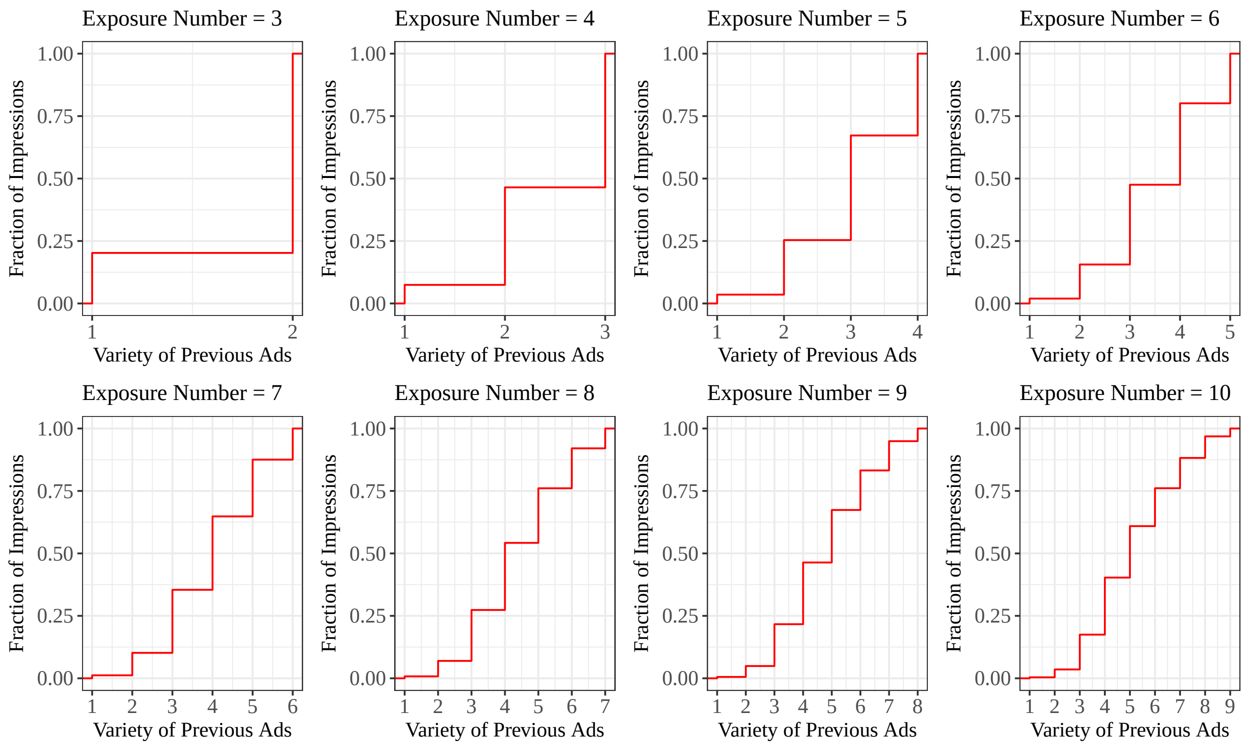

Figure 3.7: Empirical CDF of the breadth of variety for different exposure numbers. Note that by definition *Breadth* is less than exposure number.

3.4.4 Model-free Evidence

We now examine whether our data show any patterns that indicate a link between variety of previous ads and the click outcome on the current ad. For illustrative purposes, we focus on impressions of the top two ads in our data and consider their average CTR when they are shown at the eighth exposure. We plot each ad's average CTR against the variety of previous ads (as quantified by $Breadth_{i,t}$) and present the results in Figure 3.8. The plots suggests a positive relationship between the variety of previous ads and the average CTR of the current ad.[9]

[9]For the model-free analysis to be interpretable, we need to ensure that the plot is not muddled by strong ad fixed-effects or exposure number fixed effects. That is why show this analysis at the ad level and at a specific exposure number.

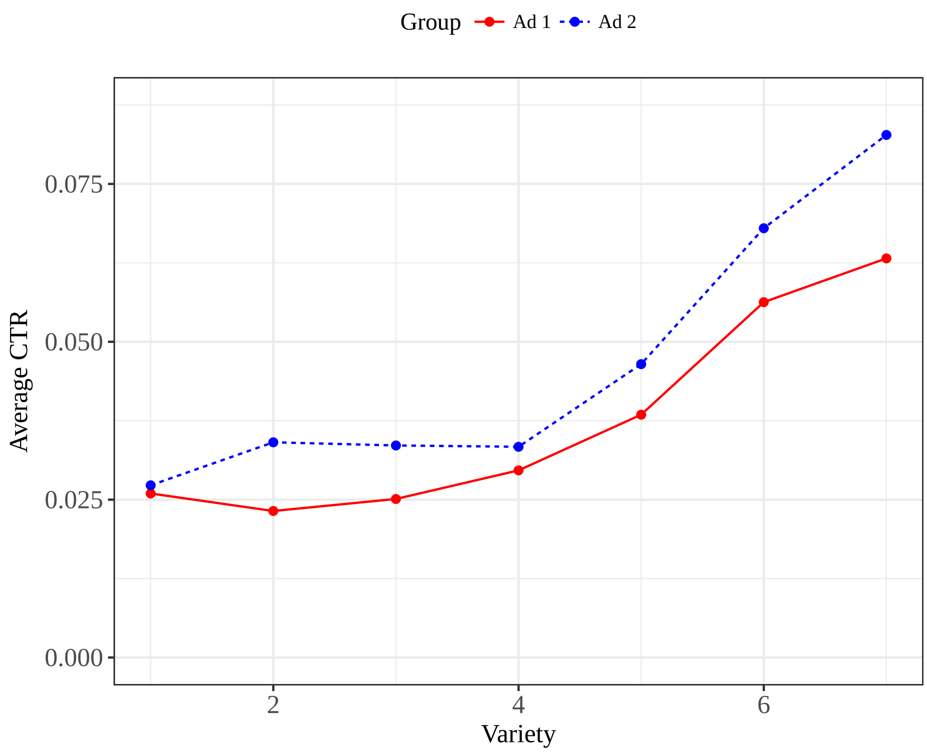

Figure 3.8: Average CTR of top two ads at the eigth exposure number 8 when exposed to different levels of variety

Nevertheless, we should be cautious in interpreting these patterns as causal. They can stem from potential selection issues rather than the effect of variety. For example, users with a higher propensity to click may be targeted by more advertisers, and therefore shown a higher variety of ads. Further, users shown a low variety of ads may get annoyed and leave the session before getting treated. In §3.5, we formalize these selection issues and outline our identification strategy in greater detail.

3.5 Empirical Framework

In this section, present our empirical framework.

3.5.1 Problem Definition

At each exposure t in session i, our main outcome of interest is the click indicator, $Y_{i,t}$, which is a binary variable that is equal to one if the user clicks on the ad shown during this exposure, and zero otherwise. Let $V_{i,t}$ be a variety measure that captures the variety of ads shown in the last $t-1$ exposures, with support distribution $\mathcal{V}_{i,t}$ that determines the assignment to different levels of variety. Note that $V_{i,t}$ can refer to any measure of variety (e.g., $Breadth_{i,t}$, $Change_{i,t}$, or $GiniSimpson_{i,t}$) or indeed any temporal measure that reflects some information on the prior sequence of ads ($\{A_{i,1}, \ldots, A_{i,t-1}\}$). For simplicity we refer to $V_{i,t}$ as variety of previous ads. We can now characterize our observed data as $\mathcal{D} = \{X_{i,t}, A_{i,t}, V_{i,t}, Y_{i,t}\}$, where $V_{i,t}$ is our treatment variable and $Y_{i,t}$ is the outcome.

In this chapter, we want to measure the causal effect of the variety of previous ads ($V_{i,t}$) on the user's decision as to click ($Y_{i,t}$) on the current ad ($A_{i,t}$). To further characterize the problem, let $X_{i,t}$ denote other covariates that we observe at the t^{th} exposure in session i that can affect the user's click decision (e.g., smartphone brand). For any exposure t, we consider a generic partially linear model to capture the relationship between $Y_{i,t}$ and $V_{i,t}$ as follows:

$$Y_{i,t} = \beta_t V_{i,t} + f(X_{i,t}, A_{i,t}) + \epsilon_{i,t}, \tag{3.7}$$

where β_t captures the marginal effect of variety of previous ads on the user's probability of click on the ad shown in the t^{th} exposure, and $f(X_{i,t}, A_{i,t})$ can be any non-parametric function that separates out the effects of the current ad and other covariates on the outcome from the effects of variety. Our main goal is to consistently estimate β_t for each exposure number t.

3.5.2 Selection Issues

In an ideal case where the distribution of variety is probabilistic and identical across sessions (i.e., $\mathcal{V}_{i,t} = \mathcal{V}_{j,t}$ for any two sessions i and j), we can consistently estimate β_t in Equation 3.7 using conventional regression approaches. However, the distribution of variety is not

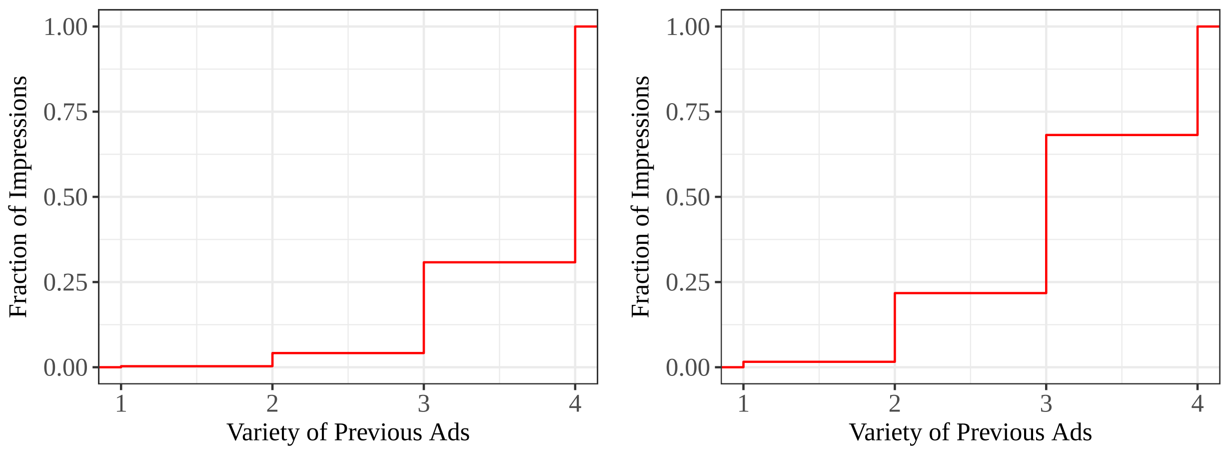

(a) Large Province, Evening

(b) Small Province, Morning

Figure 3.9: Empirical CDF of the breadth of variety for two slices of the data at the fifth exposure.

identical across sessions. We first empirically highlight this point from our data. We then present two types of selection problems that might cause this discrepancy in the distribution.

To illustrate the discrepancy in the distribution of variety across sessions, we focus on two slices of our data: (1) sessions in a large province in the evening, and (2) sessions in small province in the morning of the same day. We focus on the fifth exposure in these sessions and calculate the breadth of variety of the first four exposures in each session. Figure 3.9 plots the empirical CDF of the breadth of variety in these two slices of the data. As shown in this figure, users in the larger province in the evening are more likely to be exposed to a higher variety than users smaller province in the evening. A Kolmogorov-Smirnov test rejects the null hypothesis of identical distributions, thereby providing empirical evidence that the distribution of variety is not identical across all sessions, i.e., $\mathcal{V}_{i,t} \neq \mathcal{V}_{j,t}$ for some i and j.

The difference in the distribution of variety across sessions may originate from two potential selection problems:

1. *Cross-sectional selection:* As highlighted in §3.3.4, advertisers use some session-level characteristics to target their ads. As such, the number of distinct ads available for two different sessions can be different. For example, there are 73 distinct ads shown in impressions shown in the first slice in Figure 3.9, whereas only 25 distinct ads are shown in the second slice. This is because sessions in the first slice are more attractive to advertisers because they come from users in a large affluent province at a time of the day with enough leisure. With more ads competing for each exposure in the session, sessions in more popular areas will naturally be assigned to higher variety levels. Therefore, two sessions with different targeting variables can have different distribution of variety, i.e., $\mathcal{V}_{i,t} \neq \mathcal{V}_{j,t}$ for some i and j. We call this cross-sectional selection.

   In particular, cross-sectional selection creates estimation challenges when the popularity of targeting characteristics is also associated with higher CTR. In Figure 3.3b, we show a positive correlation between the number of distinct ads targeting each subcategory and the average CTR of that category. Hence, failure to control for this cross-sectional selection will result in selection bias in the estimates of the effects of variety. This is similar in spirit to the endogenous targeting problem, documented in the literature on measurement of ad-effectiveness – users who are more likely to respond to an ad are also more likely to be targeted by that ad [**?**]. However, it is worth emphasizing that in light of Proposition 3, this type of selection in assignment to variety comes form an unconfounded ad assignment that can be fully learned from the data. So we can technically view it as a *known* part of the distribution of variety.

2. *Dynamic selection:* A separate confounding factor that can lead to the discrepancy in the distribution of variety is the lack of user-compliance. This issue is present even in a fully randomized assignment: regardless of how we want to assign users to different levels of variety, it is the user's decision to comply with the assignment and stay in the session. Figure 3.10 depicts this dynamic selection issue in our problem in a case

where there are six sessions that are randomly assigned to two different treatment arms – high ($Breadth_{i,t} = 4$) vs. low variety ($Breadth_{i,t} = 1$). In all these cases, the user is supposed to see ad X in the fifth exposure. The dynamic selection in our case stems from users' compliance: as shown in Figure 3.10, only three out of six sessions fully received the intended variety treatment in the fifth exposure. As such, we only observe three sessions that actually received the treatment in our data: one in the high variety condition and two in the low variety condition. Therefore, $\mathcal{V}_{i,t} \neq \mathcal{V}_{j,t}$ for some i and j. We refer to it as dynamic selection since this stems from the sequential nature of the treatment assignment. In particular, this can interfere with inference if users' decision

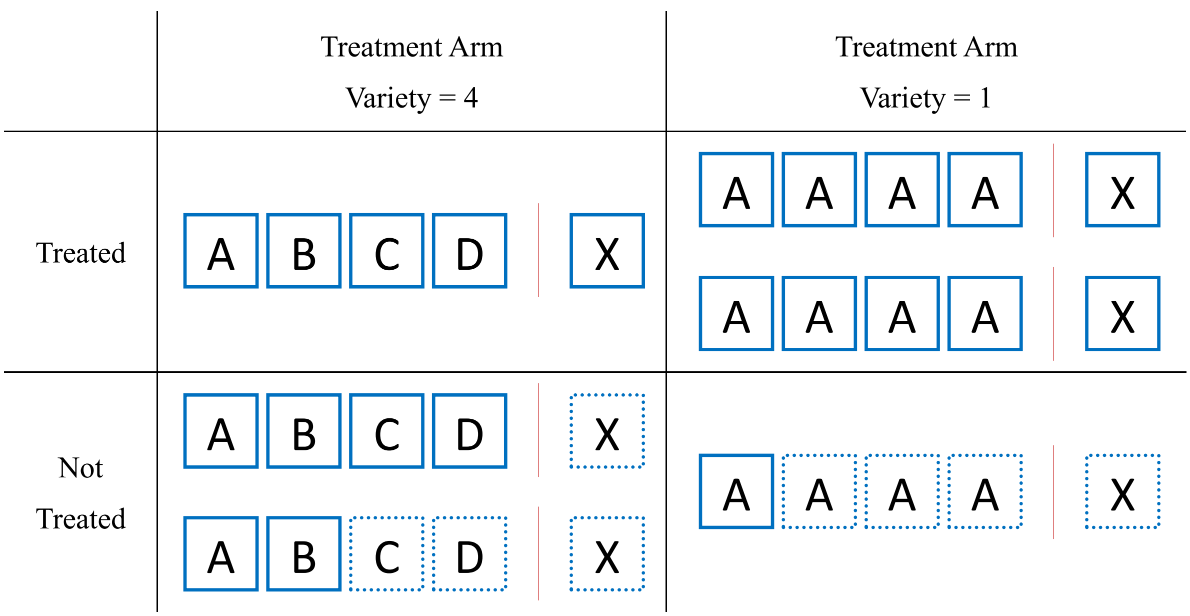

Figure 3.10: Compliance and non-compliance in receiving the treatment

to stay in the session is a function of their variety assignment. To see how it can result in incorrect inference, suppose that only one impression in each treatment arm has been clicked. As such, there is no difference in the CTR between the two treatment arms as both have a CTR of $\frac{1}{3}$. However, if we solely rely on the observed data, we find a CTR of 1 for the high variety assignment compared to a CTR of $\frac{1}{2}$ for the low variety assignment. This results in an incorrect inference drawn from the observed data that higher variety of previous ads results in a higher CTR, while this is only because higher variety of previous ads resulted in fewer users who stayed in the session to fully receive this variety treatment. Since we do not fully observe all user-level characteristics that drive users' compliance rate, this is an *unknown* part of the distribution of variety.

Together, both cross-sectional and dynamic selection problems make the main estimation task more challenging. However, it is essential to distinguish between these two types of selection problems: while cross-sectional selection is caused by the ad allocation mechanism, dynamic selection is caused by the user. Later in §3.5.3, we explain how we use this distinction to address each selection problem.

3.5.3 Identification Strategy and Solution to Selection Issues

In this section, we present our identification strategy to address both cross-sectional and dynamic selection issues outlined in §3.5.2. We start by reformulating the problem in, which allows us to define our main goal and establish the feasibility of our identification. Next, we present our solution to the cross-sectional selection problem. Finally, we describe our strategy to address the dynamic selection problem.

Reformulated Problem Definition

As discussed in §3.5.2, there are two types of selection in the distribution of variety – cross-sectional selection and dynamic selection. Since cross-sectional selection originates only from the ad assignment mechanism, we can use Proposition 3 and adjust for the non-random discrepancies in the distribution of variety across the sessions. However, even if we can fully address the cross-sectional selection problem, the unconfoundedness assumption is still likely to fail for $\mathcal{V}_{i,t}$ because of dynamic selection: while we observe all the covariates that determine ad assignment, it is quite possible that we miss variables that drive users' decision to stay in a session. As such, unless we impose assumption on the unconfoundedness of users' compliance, we cannot use the variation in $\mathcal{V}_{i,t}$ to identify β.[10]

Instead of imposing restrictive assumptions on users' compliance, we take an alternative approach based on the following logic: while the *observed variety* is endogenous due to users' compliance issue, the distribution of *intended variety* is only a function of the auction mechanism and not users' compliance. We formalize this intuition in Remark 3. However, we first define some notation. Let $A^*_{i,t}$ is the intended ad assignment for exposure t in session i; i.e., the ad that would have been shown had the user stayed in the session to see exposure t. Of course, for observed exposures in the data, we have $A^*_{i,t} = A_{i,t}$. Using the intended assignment to ads, we can define the intended assignment to variety as $V^*_{i,t}$, and its distribution as $\mathcal{V}^*_{i,t}$. For example, in Figure 3.10, the intended breadth of variety in the fifth exposure is four for all the sessions in the left and one for all the sessions in the right, irrespective of whether the user has stayed in the session. It is easy to show that the distribution of the intended variety is fully determined by the distribution of ad assignments. In light of this and Proposition 3, we can write the following remark:

Remark 3 *Conditional on the set of observed covariates, the intended assignment to variety is random.*

Remark 3 essentially gives us unconfoundedness or conditional exogeneity in intended variety. This allows us to re-formulate the problem from Equation 3.7 and focus on intended variety as our main treatment variable. Let $Y^*_{i,t}$ be the corresponding click outcome for the intended exposure. It is is the observed outcome for exposures shown in the data ($Y^*_{i,t} = Y_{i,t}$), and zero for all the intended exposures not shown in the data, because the user does not click on these unrealized exposures. Now we can re-write the model in Equation (3.7) as follows:

$$Y^*_{i,t} = \beta^*_t V^*_{i,t} + f^*(X^*_{i,t}, A^*_{i,t}) + \epsilon^*_{i,t}, \tag{3.8}$$

[10]It is worth emphasizing that this dynamic selection issue is present even in the case of fully randomized experiment where we users are randomly assigned to assign to different levels of variety. Indeed, it is present in any experiment where the treatment assignment spread over multiple time periods.

where β_t^* is the main estimand of interest and $X_{i,t}^*$ is the set of covariates for both observed and intended exposures. The β_t^* is essentially equivalent to the Intent-To-Treat (ITT) estimate since the treatment variable is the intended variety ($V_{i,t}^*$), and not the observed variety of ads users received in the data ($V_{i,t}$).

The main benefit of the formulation in Equation (3.8) over that in Equation (3.7) is the fact that assignment to the treatment variable is not confounded by unobserved covariates that drive users' compliance in Equation (3.8). In light of Remark 3, we can therefore consistently estimate β_t^* if: (1) we control for all the confounding factors in the intended assignment to variety through $f^*(X_{i,t}^*, A_{i,t}^*)$ in Equation (3.8), and (2) we construct the intended sample, i.e., $\{X_{i,t}^*, A_{i,t}^*, V_{i,t}^*, Y_{i,t}^*\}$, for any intended exposure t in session i. The former addresses the cross-sectional selection problem, whereas the latter addresses the dynamic selection by obtaining ITT estimates.

Solution to Cross-Sectional Selection: Stratification

In this section, we focus on the cross-sectional selection problem highlighted in §3.5.1. In a fully randomized experiment, we would have $\mathcal{V}_{i,t}^* = \mathcal{V}_{j,t}^*$ for all pairs of i and j by construction, i.e., the distribution of intended assignment to variety would be identical across sessions. In that case, we can consistently estimate β_t^* in Equation (3.8) without controlling for $f^*(X_{i,t}^*, A_{i,t}^*)$. However, as discussed earlier, in our setting there may be discrepancies in the distributions of intended assignment to variety across sessions, i.e., there exist i and j such that $\mathcal{V}_{i,t}^* \neq \mathcal{V}_{j,t}^*$. However, from Remark 3, we know that the intended assignment to variety is exogenous conditional on observed covariates. This gives rise to the following challenge:

Challenge 3 *We need to specify $f^*(X_{i,t}^*, A_{i,t}^*)$ in Equation (3.8) such that it controls for all the confoundedness in the intended assignment to variety, i.e., conditional on $f^*(X_{i,t}^*, A_{i,t}^*)$, the intended assignment to variety is fully random.*

To find the solution to this challenge, we start with a rather trivial remark as follows:

Remark 4 *If the ad allocation to each single exposure is identical across sessions i and j, the intended assignment to variety is identical, i.e.,* $\mathcal{V}^*_{i,t} = \mathcal{V}^*_{j,t}$.

Now, our goal is to find a stratification of sessions where the sessions within the same stratum have identical distribution of intended assignment to variety. To find such a stratification, we turn to Remark 4 and explore under what conditions the allocation to each single exposure is identical across two sessions. Using the definition of allocation function in Equation (3.1) in §3.3.3, for two sessions i and j to have the same distribution of ad allocation, the following condition must hold:

$$\forall a \in \mathcal{C}, \quad \mathbb{1}(a \in \mathcal{C}_i)\frac{b_{i,a}m_{i,a}}{\sum_{k\in\mathcal{C}_i} b_{i,k}m_{i,k}} = \mathbb{1}(a \in \mathcal{C}_j)\frac{b_{j,a}m_{j,a}}{\sum_{k\in\mathcal{C}_j} b_{j,k}m_{j,k}} \tag{3.9}$$

This condition guarantees that the allocation probability for each ad is the same in both session i and j. Now, we can identify the potential sources that can violate the condition in Equation (3.9):

- *Difference in bids:* If $b_{i,a} \neq b_{j,a}$ for some a, then $\mathcal{V}^*_{i,t} \neq \mathcal{V}^*_{j,t}$. In our setting, ads are only allowed to submit a single bid at any time. Therefore, $b_{i,a} \neq b_{j,a}$ only happens if the ad has changed his bid between these two sessions. If the ad decides to change his bid, it will only be effective in the next hour. Thus, for sessions i and j around the same time (e.g., within the same hour), we have $b_{i,a} = b_{j,a}$.
- *Difference in quality scores:* If $b_{i,a} \neq b_{j,a}$ for some a, then $\mathcal{V}^*_{i,t} \neq \mathcal{V}^*_{j,t}$. A unique feature of our setting is that unlike most platforms, this platform does not customize quality scores and only uses an aggregate measure for every ad. Every once in a while (often daily), the platform updates these quality scores. Thus, for sessions i and j around the same time, we have $q_{i,a} = q_{j,a}$.
- *Difference in the set of participants:* If $a \notin \mathcal{C}_i$ but $a \in \mathcal{C}_j$ (or vice versa) for some a, then $\mathcal{V}^*_{i,t} \neq \mathcal{V}^*_{j,t}$. In words, if there exists an ad that participates in auction for session i but not session j (or vice versa), the intended assignment to variety may differ across these two sessions. There are two sources that can lead to the difference in the set of participants in two sessions:

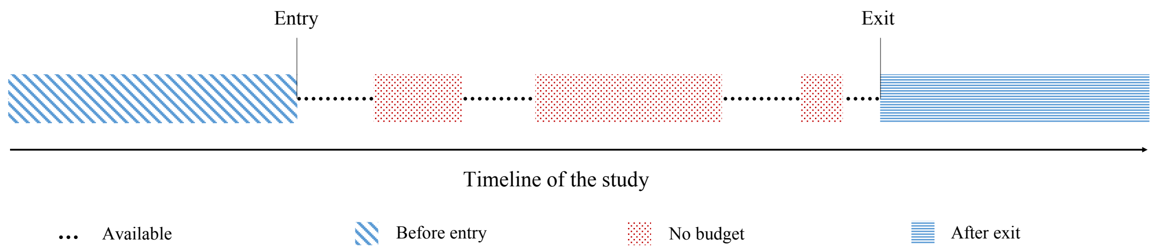

Figure 3.11: Availability of an ad in the timeline of the study due to entry, exit, and budget exhaustion

1. *Difference in targeting:* Advertisers can target their ads based on app category, province, connectivity type, time of the day, MSP, ISP, and smartphone brand. As such, each session has a set of targeting characteristics. Hence, we may have $a \notin \mathcal{C}_i$ and $a \in \mathcal{C}_j$ because ad a decided to target session i but not session j (or vice versa). For example, if an advertiser selects Huawei and LG as the set of smartphone brand categories he wants to target, his ad will not be shown to any Samsung users, as Samsung is excluded from her targeting set. Now, if the smartphone brand is Huawei in session i and Samsung in session j, this ad will be in $\mathcal{C}_i$ but not $\mathcal{C}_j$. Thus, for sessions i and j with the identical targeting characteristics, we have no difference in targeting.

2. *Difference in availability over time:* Ads' campaign availability over time may change due to three reasons: entry, exit, and budget. Figure 3.11 illustrates this point by showing an ad's availability (vs. unavailability) during the time of study. As such, if session i is in areas where ad a is unavailable (due to either entry, exit, and budget) while session j is in areas where this ad is available, we have $a \notin \mathcal{C}_i$ but $a \in \mathcal{C}_j$ (or vice versa). Thus, for sessions i and j around the same time, there is no difference in ads' availability.

Together, for sessions i and j around the same time with the same targeting characteristics, we have $\mathcal{C}_i = \mathcal{C}_j$.

Now, we can show that the condition in Equation (3.9) holds if we have the following requirements for two sessions i and j: (1) $b_{i,a} = b_{j,a}$ for all a, (2) $q_{i,a} = q_{j,a}$ for all a, and (3) $\mathcal{C}_i = \mathcal{C}_j$. We derived the sufficient condition for each of these requirements separately above. Now to have a sufficient condition that satisfies all these requirements together is the intersection of all separate sufficient conditions. This brings us to our main remark in this section:

Remark 5 *For sessions i and j around the same time with the same targeting characteristics, we have $\mathcal{V}^*_{i,t} = \mathcal{V}^*_{j,t}$.*

In light of Remark 5, we define each stratum $G_k \in \mathcal{G}$ as a subset of sessions in our data where all the targeting variables are the same. Since hour of the day is one of the targeting variables, sessions i and j in stratum G_k are within the same hour by definition. Later, as a robustness check, we show that the results will not change even for shorter time intervals.

In sum, our solution to address cross-sectional selection in our problem is to transform the problem into a stratified randomized experiment. As such, within the stratum $G_k \in \mathcal{G}$, the intended assignment to variety is random. Thus, if $f^*(X^*_{i,t}, A^*_{i,t})$ in captures all $G_k \in \mathcal{G}$, the intended assignment to variety will be conditionally exogenous in Equation (3.8).

Solution to Dynamic Selection: Imputation

In this section, we present our solution to address dynamic selection problem in great detail. In our solution to cross-sectional selection, we stratify the covariate space such that the intended assignment to variety is fully random within each stratum G_k. As such, we can estimate the effects of the intended assignment to variety also known as intent-to-treat estimates of the effects of variety. However, as shown in Equation (3.8), we need to have the values of all variables for the intended exposures. For example, if session i ended in 5 exposures, we do not directly observe what level of variety this session would have been assigned in 8^{th} exposure had it been continued to fully receive the intended variety treatment. This brings us to our main empirical challenge in this section:

Challenge 4 *We need to construct the values of the main variables for all the intended exposures. That is, for any session i that ended in T_i exposures, we need to impute the set of $\{Y^*_{i,t}, A^*_{i,t}, V^*_{i,t}\}$ for all $T_i < t \leq T$.*

As discussed earlier, we impute the outcomes for missing exposures by zero: for any session i and exposure t such that $T_i < t \leq T$, we have $Y^*_{i,t} = 0$. The key challenge is to impute intended ad assignments. For each session i, we observe the sequence $\{A_{i,1}, A_{i,2}, \ldots, A_{i,T_i}\}$. Our goal is to obtain the full set of all intended ad assignments after the session ended, i.e., $\{A^*_{i,T_i+1}, \ldots, A^*_{i,T}\}$.

We first characterize the distribution each ad is drawn from. Let τ denote the exact timestamp of an exposure in session i. We define $\mathcal{A}_i(\tau)$ as the distribution of ad allocation at timestamp τ in session i. The following remark about the distribution of ad allocation helps characterize our imputation strategy:

Remark 6 *For any two sessions i and j with the same targeting characteristics ($i, j \in G_k$), the distribution of ad allocation is the same at any arbitrary timestamp, i.e., $\mathcal{A}_i(\tau) = \mathcal{A}_j(\tau)$.*

In light of Remark 6, we can use the actual ad assignment in other sessions to impute the intended ad assignment for a specific session. For example, in order to impute A^*_{i,T_i+1}, we need to find the timestamp τ that indicates when this exposure would have happened had the session not ended, as well as another session j with the same targeting characteristics that served an ad at timestamp τ. We then impute A^*_{i,T_i+1} by the ad shown in session j at timestamp τ, since the ad allocation mechanism is identical across session i and j. It is worth noting that unlike most imputation approaches that use models to approximate the original distribution and simulate missing data from this approximate distribution, our approach is model-free and guarantees that the imputed exposures are drawn from the original distribution, e.g., $A^*_{i,T_i+1} \sim \mathcal{A}_i(\tau)$ in the example above.

We now present the step-by-step algorithm that allows us to impute all the sessions in our data. We first set $T = 10$ because the vast majority of sessions ends in the first 10 sessions.[11] As such, for any session i, we want to obtain $\{A^*_{i,T_i+1}, \ldots, A^*_{i,10}\}$. We do not use sessions in our sample for imputation purpose. Rather, we use a complementary data set of all the impressions that are not in our sample. We call this data set $\tilde{D}$ and it contains the same characteristics as our original sample. Below we present the step-by-step imputation algorithm that we use:

- *Step 1:* For any session i that ended in T_i exposures, we first impute the timestamps for the intended exposures $\{A^*_{i,T_i+1}, \ldots, A^*_{i,10}\}$. Each exposure lasts one minute, so we just need to add one minute to the timestamp for the last exposure. We denote these timestamps by $\{\tau^*_{i,T_i+1}, \ldots, \tau^*_{i,10}\}$.
- *Step 2:* For any $s \geq 1$ that $T_i + s \leq 10$, we take τ^*_{i,T_i+s} and session i, and find a session j in the complementary data ($j \in \tilde{D}$) with the same targeting characteristics as i that served ads at timestamp τ^*_{i,T_i+s}. Let $\tilde{A}_j(\tau^*_{i,T_i+s})$ denote the ad shown in session j at timestamp τ^*_{i,T_i+s}.
- *Step 3:* We impute the intended exposure in session i at exposure $T_i + s$ with the ad found in the complementary data as it would have been the ad shown had the session not ended. Hence, $A^*_{i,T_i+s} = \tilde{A}_j(\tau^*_{i,T_i+s})$.

Once we imputed the ad assignment for all the intended exposures, we can impute the intended assignment to variety for all these exposures since the variety definition is solely based on the sequence of prior ads.

3.5.4 Model Specification

In this section, we present our model specification. This model is based on the generic form presented in Equation (3.8). The only part from this equation that we need to specify is $f^*(X^*_{i,t}, A^*_{i,t})$. While we know that $f^*(X^*_{i,t}, A^*_{i,t})$ must incorporate the stratification defined in §3.5.3, we may need other controls depending on our experimental goal. In particular, we need to specify how we want to control for the information in the current ad $A^*_{i,t}$. This

[11] We do not go beyond exposure number 10 as the vast majority of sessions end in less than 10 exposures (see Figure 3.4a).

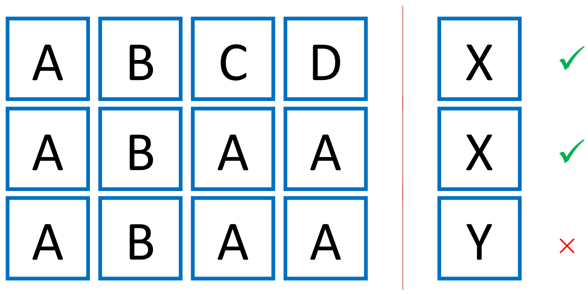

(a) Controlling for ad fixed effects

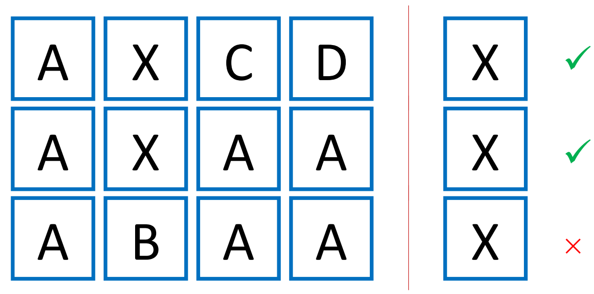

(b) Controlling for previous exposures of the focal ad

Figure 3.12: Controls for ad-specific features and their experimental equivalents

specification also defines the causal interpretation that our main parameter of interest has.

Ideally, we want to compare cases where the only difference comes from the assignment to variety. As such, failure to control for the differences in the effects of the current ad may result in some confounding issues. In Figure 3.12, we visualize the experimental comparisons that we wish to make in terms of the effects of variety. In Figure 3.12a, we want to compare the first two cases where the current ad is the same and sessions are assigned to different levels of variety. Comparison of the first and third case, however, is not of our interest because not only is assignment to variety different between these two cases, but also the assignment to the current ad is different. Hence, it is not clear if the difference in the outcome comes from the assignment to variety or the current ad. To address this issue, we include ad fixed effects allows us to control for the ad that is currently being shown.

While controlling for ad fixed effects helps compare the effects of variety within a specific ad assignment, but there might still be more information about the current ad that is missing. Figure 3.12b shows a case where the current ad is the same in all three sessions, but the number of times this ad has been shown within the session is different: the first two sessions have shown the current ad once before within the session, whereas the third session has not shown this ad at all. Since the first two sessions only differ in variety of previous ads, we want to compare these two with each other, not with the third session. We control for this factor by controlling for the number of times the current has been shown. We define this variable as $I^*_{i,t}$ as follows:

$$I^*_{i,t} = \sum_{j=1}^{t-1} \mathbb{1}(A^*_{i,j} = A^*_{i,t}) \tag{3.10}$$

Now, we can specify our regression model for each exposure number as follows:

$$Y^*_{i,t} = \beta_t^{ITT} V^*_{i,t} + \gamma_t I^*_{i,t} + \underbrace{\eta^*_{a,t}}_{\text{ad fixed effects}} + \underbrace{\zeta^*_{g,t}}_{\text{strata fixed effects}} + \epsilon_{i,t}, \tag{3.11}$$

where β_t^{ITT} is the intent-to-treat (ITT) estimate of the effects of variety and the combination of all controls forms our specification of $f^*(X^*_{i,t}, A^*_{i,t})$ in Equation (3.8). For our treatment variable $V^*_{i,t}$, we use the measures defined in §3.4.1. We present our estimates for each exposure number in §3.6.

3.6 Results

In this section, we present our results. We first present the ITT estimates of the effects of variety in §3.6.1. In §3.6.2, we establish the extent of dynamic selection bias if we mainly rely on the observed sample and ignore imputation. Next, in §3.6.3, we show our findings on the source for the effects of variety. Finally, in §3.6.4, we document the heterogeneity in the effects of variety across user-level observables.

3.6.1 Intent-to-Treat (ITT) Estimates of the Effects of Variety

We start by estimating the parameters in Equation (3.11) using the intended sample. For each exposure number t, we use each one of the variety measures in §3.4.1 to estimate the model. Overall, we estimate $7 \times 3 = 21$ models and present all the results in Table 3.2. With no exception, all estimates for the coefficient of variety indicate a positive association between the variety of previous ads and clicking on the next ad. That is, when assigned to a higher variety of previous ads in a sessions, users are more likely to click on the next ad.

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
				A. Breadth of Variety			
$Breadth_{i,t}$	0.00066***	0.00078***	0.00045***	0.00024**	0.00028***	0.00028***	0.00024***
	(3.67)	(5.59)	(4.28)	(2.71)	(3.74)	(4.20)	(4.21)
Previous	0.00042**	0.00073***	0.00038***	0.00032***	0.00015*	0.00021***	0.00021***
Exposure ($I_{i,t}$)	(2.74)	(5.77)	(3.93)	(3.90)	(2.23)	(3.47)	(4.00)
R^2	0.099	0.088	0.087	0.077	0.071	0.063	0.061
Adjusted R^2	0.028	0.024	0.029	0.024	0.022	0.017	0.019
				B. Consecutive Changes			
$Change_{i,t}$	0.00056**	0.00079***	0.00043***	0.00027**	0.00032***	0.00037***	0.00024***
	(3.18)	(5.83)	(4.29)	(3.19)	(4.58)	(5.94)	(4.57)
Previous	0.00040**	0.00072***	0.00037***	0.00033***	0.00016*	0.00023***	0.00021***
Exposure ($I_{i,t}$)	(2.58)	(5.69)	(3.84)	(3.96)	(2.31)	(3.80)	(4.00)
R^2	0.099	0.088	0.087	0.077	0.071	0.063	0.061
Adjusted R^2	0.028	0.024	0.029	0.024	0.022	0.017	0.019
				C. Gini-Simpson Index for Diversity			
$GiniSimpson_{i,t}$	0.00133**	0.00261***	0.00245***	0.00125*	0.00226***	0.00285***	0.00283***
	(3.23)	(5.09)	(4.64)	(2.18)	(3.87)	(4.66)	(4.68)
Previous	0.00041**	0.00073***	0.00041***	0.00032***	0.00017*	0.00024***	0.00024***
Exposure ($I_{i,t}$)	(2.67)	(5.74)	(4.17)	(3.81)	(2.46)	(3.83)	(4.39)
R^2	0.099	0.088	0.087	0.077	0.071	0.063	0.061
Adjusted R^2	0.028	0.024	0.029	0.024	0.022	0.017	0.019
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
Note:							*p<0.05; **p<0.01; ***p<0.001

Table 3.2: Intent-to-Treat (ITT) estimates of the effects of variety previous ads and number of prior ads shown on CTR

Further, we find that the number of previous exposures of the current ad also leads to higher likelihood of click. That is, the more an ad has been shown within the session, the more likely the user is to click on that ad. It reinforces the prior findings of the literature regarding the positive effects of repeated exposure [118]. It also gives rise to an inter-temporal trade-off in ad placement. On the one hand, repeated exposures of the same ad increases the user's likelihood of clicking on it. On the other hand, showing the same ad repeatedly reduces variety. Thus, it suggests that employing a mix of both would be optimal in placement of ads.

3.6.2 Dynamic Selection Bias

Our results in Table 3.2 are ITT estimates of the effects of variety accounting for possible dynamic selection in users' assignment to variety. In this section, we quantify the extent to which the main estimates would be biased if we only use the observed sample and do not adjust for dynamic selection. In principle, if users' decision to leave the session is fully exogenous, we can easily link our estimates from the observed sample with those from the intended sample. In this section, we first present a simple regression model to test how exogenous user compliance is, and then quantify the extent of dynamic selection bias if we only rely on the observed sample.

How Reasonable is the Assumption of Exogenous Compliance?

As discussed above, if users decide to leave a session irrespective of their variety assignment, the observed sample reflects the effects of variety with no selection bias. On the other hand, if users' decision to leave is a function their variety assignment, we disproportionately observe exposures of each variety treatment, which would bias our main estimates. In this section, we test the assumption that compliance is not a function the variety assignment. To do so, we define *leave* variable indicating whether the user has left after that exposure. Let $L_{i,t}$ denote the user's decision to leave after seeing exposure t in session i. We specify the following regression model:

$$L_{i,t} = \alpha V_{i,t} + \theta_t I_{i,t} + \eta'_{a,t} + \zeta'_{g,t} + \epsilon'_{i,t} \tag{3.12}$$

It is worth noting that the right-hand side of Equation (3.12) is the equivalent of that in Equation (3.11) on the observed sample. If the assumption that users' decision to leave is independent of variety assignment is true, α in Equation (3.12) must be insignificant. We present the estimates of this model in Table 3.3. Our results indicate that at some exposure numbers, there is a statistically significant link between variety assignment and the user's decision to leave the session. This finding provides evidence for the existence of dynamic selection in our problem. In particular, given the negative association between variety and users' decision to leave, users are more likely to leave the session when the assignment to variety is low. As such, the number of observed cases in low variety assignment is lower. Thus, we expect the dynamic selection to underestimate the positive effects of variety.

Magnitude of Dynamic Selection Bias

As shown in §3.6.2, user compliance is a function of the variety assignment. Hence, if we only rely on the observed sample to estimate the effects of variety, our estimates will be biased. In this section, we want to quantify the magnitude of this dynamic selection bias. To do so, we need to first estimate the coefficient of variety using the observed sample. We use the exact specification in Equation 3.11 and estimate the parameters on the observed sample. The results are shown in Table 3.4. The estimates qualitatively reveal the same pattern as our ITT estimates in Table 3.2: there is a positive association between CTR on the current ad and both the variety of previous ads and the number of previous exposures of current ad. Thus, the existence of dynamic selection bias does not directionally change our main findings.

However, our goal is to quantitatively compare the variety estimates in Tables 3.2 and 3.4. Of course, a plain comparison between estimates is not correct, since the estimands are fundamentally different: the estimand in Table 3.2 is ITT estimates obtained using the entire

	Dependent variable: Leave ($L_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
				A. Breadth of Variety			
$Breadth_{i,t}$	-0.00715***	-0.00254*	-0.00268*	-0.00171	-0.00124	-0.00039	-0.00005
	(-5.20)	(-2.11)	(-2.33)	(-1.55)	(-1.15)	(-0.36)	(-0.05)
Previous	-0.00284*	-0.00323**	-0.00084	-0.00223*	-0.00115	-0.00383***	0.00030
Exposure ($I_{i,t}$)	(-2.39)	(-2.96)	(-0.79)	(-2.21)	(-1.17)	(-3.98)	(0.32)
R^2	0.222	0.250	0.279	0.305	0.329	0.348	0.382
Adjusted R^2	0.014	0.010	0.013	0.012	0.012	0.006	0.023
				B. Consecutive Changes			
$Change_{i,t}$	-0.00402**	-0.00126	-0.00144	-0.00062	0.00136	0.00027	-0.00044
	(-2.97)	(-1.08)	(-1.32)	(-0.60)	(1.36)	(0.28)	(-0.46)
Previous	-0.00217	-0.00292**	-0.00052	-0.00196	-0.00053	-0.00367***	0.00020
Exposure ($I_{i,t}$)	(-1.83)	(-2.69)	(-0.50)	(-1.95)	(-0.54)	(-3.84)	(0.22)
R^2	0.222	0.250	0.279	0.305	0.329	0.348	0.382
Adjusted R^2	0.014	0.010	0.013	0.012	0.012	0.006	0.023
				C. Gini-Simpson Index for Diversity			
$GiniSimpson_{i,t}$	-0.01529***	-0.00809	-0.01489**	-0.00457	0.00322	-0.01270	-0.01093
	(-4.86)	(-1.84)	(-2.59)	(-0.66)	(0.39)	(-1.34)	(-1.03)
Previous	-0.00283*	-0.00322**	-0.00102	-0.00203*	-0.00071	-0.00418***	-0.00003
Exposure ($I_{i,t}$)	(-2.37)	(-2.92)	(-0.96)	(-1.98)	(-0.71)	(-4.25)	(-0.03)
R^2	0.222	0.250	0.279	0.305	0.329	0.348	0.382
Adjusted R^2	0.014	0.010	0.013	0.012	0.012	0.006	0.023
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	365,166	285,373	234,109	194,485	165,558	142,878	124,864
Note:							*p<0.05; **p<0.01; ***p<0.001

Table 3.3: OLS estimates of the effects of variety on users' decision to leave a session using the observed sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
				A. Breadth of Variety			
$Breadth_{i,t}$	0.00151**	0.00242***	0.00152***	0.00118*	0.00157***	0.00140**	0.00156***
	(2.99)	(4.82)	(3.34)	(2.51)	(3.43)	(2.95)	(3.34)
Previous	0.00138**	0.00264***	0.00154***	0.00158***	0.00095*	0.00128**	0.00150***
Exposure ($I_{i,t}$)	(3.17)	(5.82)	(3.68)	(3.68)	(2.28)	(2.99)	(3.61)
R^2	0.229	0.258	0.296	0.316	0.337	0.354	0.375
Adjusted R^2	0.023	0.020	0.035	0.027	0.024	0.015	0.013
				B. Consecutive Changes			
$Change_{i,t}$	0.00133**	0.00239***	0.00163***	0.00150***	0.00170***	0.00209***	0.00166***
	(2.68)	(4.93)	(3.74)	(3.41)	(4.01)	(4.82)	(3.94)
Previous	0.00133**	0.00259***	0.00154***	0.00164***	0.00096*	0.00144***	0.00153***
Exposure ($I_{i,t}$)	(3.06)	(5.74)	(3.71)	(3.85)	(2.32)	(3.39)	(3.70)
R^2	0.229	0.258	0.296	0.316	0.337	0.354	0.375
Adjusted R^2	0.023	0.020	0.035	0.027	0.024	0.015	0.013
				C. Gini-Simpson Index for Diversity			
$GiniSimpson_{i,t}$	0.00302**	0.00794***	0.00874***	0.00645*	0.01158***	0.01421***	0.01724***
	(2.62)	(4.34)	(3.84)	(2.18)	(3.31)	(3.37)	(3.63)
Previous	0.00136**	0.00264***	0.00166***	0.00159***	0.00102*	0.00143**	0.00166***
Exposure ($I_{i,t}$)	(3.11)	(5.78)	(3.93)	(3.65)	(2.42)	(3.27)	(3.88)
R^2	0.229	0.258	0.296	0.316	0.337	0.354	0.375
Adjusted R^2	0.023	0.020	0.035	0.027	0.024	0.015	0.013
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
Avg. CTR	0.0234	0.0260	0.0227	0.0233	0.0215	0.0221	0.0203
No. of Obs.	365,166	285,373	234,109	194,485	165,558	142,878	124,864
Note:							*p<0.05; **p<0.01; ***p<0.001

Table 3.4: OLS estimates of the coefficient of variety of previous ads and number of previous exposures of the current ad using the observed sample

sample, whereas the estimand in Table 3.4 only focuses on the observed sample of compliers. To link these two estimands, we turn to a well-known decomposition of ITT estimates:

$$\beta_t^{ITT} = \pi_t^{co}\beta_t^{ITT,co} + (1-\pi_t^{co})\beta_t^{ITT,nc}, \tag{3.13}$$

where $\beta_t^{ITT,co}$ and $\beta_t^{ITT,nc}$ are ITT estimates on the sample of compliers vs. non-compliers respectively, and π_t^{co} denotes the share of compliers for exposure number t. Since the outcome is zero for any non-treated user regardless of their intended treatment assignment, we have $\beta_t^{ITT,nc} = 0$. If there is no dynamic selection, the estimates in Table 3.4 are equivalent to $\beta_t^{ITT,co}$. Therefore, the ITT estimates we would get from our estimates in Table 3.4 under no dynamic selection assumption is as follows:

$$\tilde{\beta}_t^{ITT} = \frac{N_t}{N_1}\hat{\beta}_t^{observed}, \tag{3.14}$$

where N_t and N_1 are the number of sessions available at exposure number t and 1 respectively so the fraction is an empirical estimate of π_t^{co}, and $\hat{\beta}_t^{observed}$ is simply the estimates on the observed sample presented in Table 3.4. As such, $\tilde{\beta}_t^{ITT}$ gives us the ITT estimates if there is no dynamic selection, and the difference between $\tilde{\beta}_t^{ITT}$ and our original ITT estimates in Table 3.2 quantifies the magnitude of dynamic selection bias. We plot both estimates for the breadth of variety in Figure 3.13. As shown in this figure, the actual ITT estimates are higher than those under the random compliance assumption. This is in line with the results presented in Table 3.3 that show both variety and previous exposures of the focal ad have negative effects on user's decision to leave, at some exposure numbers. This lower compliance rate for higher variety conditions explains why the ITT estimates under random compliance is biased downward.

3.6.3 What is the Source for the Effects of Variety?

Our findings in §3.6.1 establish the main effects of variety of previous ads on the likelihood of clicking on the next ad. We now want to explore the sources for the effects of variety. In particular, we examine the role that the sequential organization of the previous ads play in

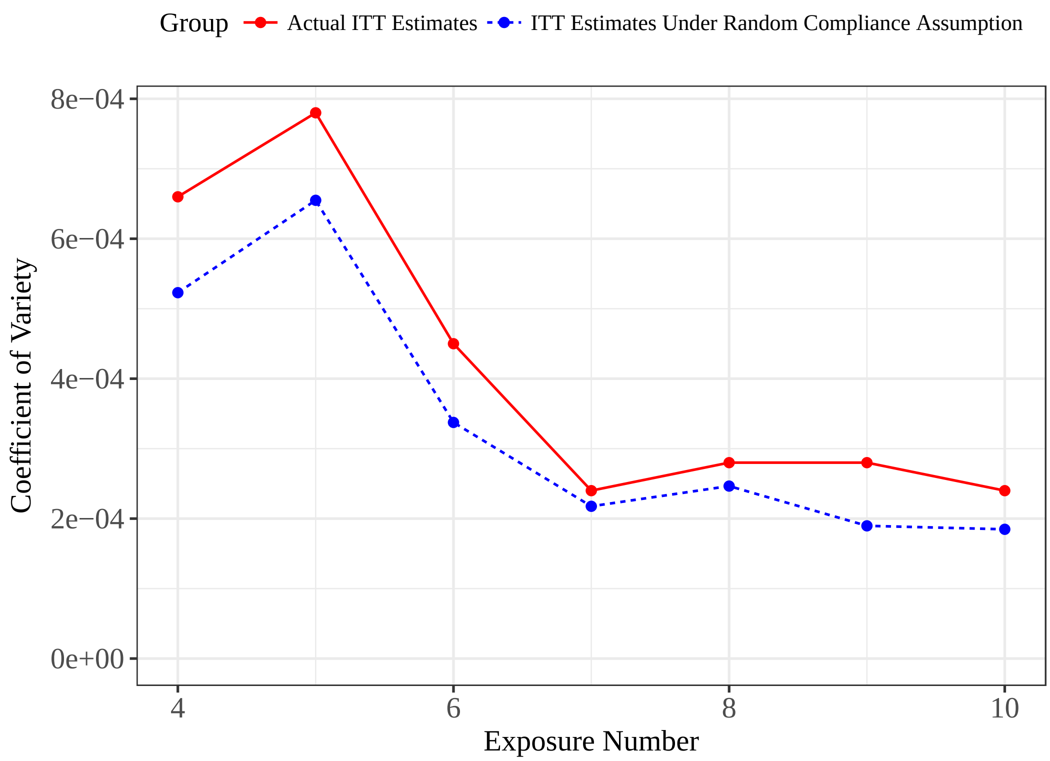

Figure 3.13: Comparison of ITT estimates using the intended sample vs. ITT estimates under random compliance assumption

driving the main effects of variety. As such, we mainly use the relationships between our three measures of variety and decompose the main effects into sequential and non-sequential elements, as presented in §3.4.2.

Breadth-Increasing vs. Breadth-Constant Changes

We start with the relationship between two of our variety measures – breadth of variety ($Breadth_{i,t}$) and consecutive changes ($Change_{i,t}$). While the former views the sequence of prior ads as a set and ignores the order in the sequence, the latter relates to the sequential ordering of previous ads. Recall our decomposition of consecutive changes in Equation 3.5

in §3.4.2:

$$Change_{i,t} = BIC_{i,t} + BCC_{i,t} = (Breadth_{i,t} - 1) + BCC_{i,t},$$

where we decompose the total number of consecutive changes into two parts – *breadth-increasing changes* ($BIC_{i,t}$) and *breadth-constant changes* ($BCC_{i,t}$). The former is the total number of consecutive changes that increase the variety of the sequence (i.e., the second ad in the pair has not been shown before within the sequence), whereas the latter is the total number of consecutive changes that do not increase the variety (i.e., the second in the pair has been shown before within the sequence). We can now include both $BIC_{i,t}$ and $BCC_{i,t}$ in Equation (3.11) as $V_{i,t}$ to see which one contributes more to the outcome. We can re-write our main model as follows:

$$Y^*_{i,t} = \beta_t^{BIC} BIC^*_{i,t} + \beta_t^{BCC} BCC^*_{i,t} + \gamma_t I^*_{i,t} + \eta^*_{a,t} + \zeta^*_{g,t} + \epsilon_{i,t}, \quad (3.15)$$

where β_t^{BIC} and β_t^{BCC} demonstrate the impact of breadth-increasing changes and breadth-constant changes respectively. If $\beta_t^{BIC} = \beta_t^{BCC}$, we can then argue that the source is mainly consecutive changing of ads. In general, the ration $\frac{\beta_t^{BCC}}{\beta_t^{BIC}}$ shows what fraction in the effects of breadth of variety stems from consecutive changing of ads. We present our results in Table 3.5. As shown in this table, we find support for both sources. The point estimates for breadth-increasing changes are higher than those of breadth-constant changes, indicating that the distinctiveness of previous ads is a source for the effects. However, the estimates of breadth-constant changes are also significant for some exposure numbers, suggesting that a large fraction of the effects of variety stems just from the consecutive changing of ads. Further, the estimates of breadth-constant changes become significant from exposure number eight onward. The fact that the gap between the estimates of breadth-increasing and breadth-constant changes shrink in later exposure numbers might be due to diminishing returns to breadth of variety. A potential explanation for the diminishing returns to breadth of variety is that users may not have the full memory of the session. To summarize, while we find the positive effects of breadth of variety in all exposures, a large fraction of the effects of variety is attributed to consecutive changing of ads, especially in later exposures.

	Dependent variable: Click ($Y_{i,t}$)						
	($t=4$)	($t=5$)	($t=6$)	($t=7$)	($t=8$)	($t=9$)	($t=10$)
Breadth-Increasing	0.00067***	0.00091***	0.00053***	0.00031**	0.00038***	0.00041***	0.00032***
Changes	(3.61)	(6.12)	(4.68)	(3.23)	(4.60)	(5.61)	(4.97)
Breadth-Constant	0.00009	0.00051*	0.00026	0.00021	0.00026**	0.00033***	0.00017**
Changes	(0.29)	(2.54)	(1.89)	(1.92)	(2.95)	(4.33)	(2.74)
Previous	0.00042**	0.00075***	0.00040***	0.00034***	0.00017*	0.00024***	0.00023***
Exposure ($I_{i,t}$)	(2.75)	(5.89)	(4.07)	(4.05)	(2.50)	(3.92)	(4.29)
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
R^2	0.099	0.088	0.087	0.077	0.071	0.063	0.061
Adjusted R^2	0.028	0.024	0.029	0.024	0.022	0.017	0.019
Note:							*p<0.05; **p<0.01; ***p<0.001

Table 3.5: Intent-to-Treat (ITT) estimates of the effects of Breadth-Increasing Changes vs. Breadth-Constant Changes

Consecutive vs. Non-Consecutive Diversity

We now turn to the second decomposition that we present in §3.4.2. Recall that the Gini-Simpson index for diversity can be decomposed into consecutive vs. non-consecutive elements:

$$GiniSimpson_{i,t} = \frac{\sum_{s+1=s'<t} \mathbb{1}(A_{i,s'} \neq A_{i,s})}{\frac{(t-1)(t-2)}{2}} + \frac{\sum_{s+1<s'<t} \mathbb{1}(A_{i,s'} \neq A_{i,s})}{\frac{(t-1)(t-2)}{2}} = CD_{i,t} + NCD_{i,t},$$

where the first element in the right-hand side only considers the consecutive pairs that show different ads ($CD_{i,t}$), and the second element takes the non-consecutive pairs of prior ads that show different ads ($NCD_{i,t}$). As such, the former quantifies the contribution of consecutive pairs to Gini-Simpson index for diversity, whereas the latter quantifies the contribution of non-consecutive pairs. We include both consecutive and non-consecutive diversity in the main specification and estimate the following model:

$$Y_{i,t}^{*} = \beta_t^{CD} CD_{i,t}^{*} + \beta_t^{NCD} NCD_{i,t}^{*} + \gamma_t I_{i,t}^{*} + \eta_{a,t}^{*} + \zeta_{g,t}^{*} + \epsilon_{i,t}, \tag{3.16}$$

where β_t^{CD} and β_t^{NCD} demonstrate the impact of breadth-increasing changes and breadth-constant changes respectively. We present our estimates in Table 3.6. We find that consecutive diversity plays a more important role in driving users' clicks. Interestingly, for exposure number 7 and 10, we find that the coefficients for non-consecutive diversity are also significant. This finding provides more support for the hypothesis that users' memory of the session is limited. Overall, our results suggest that the sequential organization of exposures is critical since there is a difference between consecutive and non-consecutive diversity. This is similar to the findings of [63] and [72] who find that the organization and structure of the assortment affect users' perceptions of variety.

3.6.4 Heterogeneity Across User History

As mentioned in §3.3.2, we sample our users such that their activity logs are not truncated. The fact that we observe the entire history for our sample allows us to distinguish between sessions based on their user-level observables and examine the effects of variety conditional on these observables. For instance, users with longer history may react differently to manipulation strategies compared to those who are new to the platform. This motivates us to explore and document heterogeneity in the main effects across sessions with different user-level observables.

We take a simple descriptive approach to show heterogeneity in the effects of variety. This involves splitting our data into two parts based on users' historical features (e.g., users' past activity). We then obtain estimates for the breadth of variety ($V_{i,t}$) on both splits of the data for all exposure numbers. Below we present three dimensions of users' past history

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
Consecutive	0.00160**	0.00459***	0.00354***	0.00422**	0.00621***	0.00984***	0.00662**
Diversity	(3.00)	(5.43)	(3.30)	(3.04)	(3.79)	(4.99)	(3.02)
Non-Consecutive	0.00079	0.00063	0.00173*	-0.00023	0.00068	0.00054	0.00175*
Diversity	(1.00)	(0.75)	(2.12)	(-0.27)	(0.80)	(0.62)	(2.06)
Previous	0.00042**	0.00073***	0.00041***	0.00032***	0.00017*	0.00024***	0.00024***
Exposure ($I_{i,t}$)	(2.68)	(5.74)	(4.18)	(3.82)	(2.44)	(3.84)	(4.39)
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
R^2	0.099	0.088	0.087	0.077	0.071	0.063	0.061
Adjusted R^2	0.028	0.024	0.029	0.024	0.022	0.017	0.019
Note:							*p<0.05; **p<0.01; ***p<0.001

Table 3.6: Intent-to-Treat (ITT) estimates of the effects of Consecutive Diversity vs. Non-Consecutive Diversity

that we are interested in and determine the splits of the data accordingly:

- *Length of history*: Users in our sample vary by the number of impressions they have seen in the past. In general, we expect this variable to be a major source of heterogeneity, as it shapes user's behavior towards ads. We call a user *experienced* if he/she has seen over 50 impressions. We split the data into two parts each containing impressions shown to experienced vs. inexperienced users. The effects on the sample of tenured and non-tenured users are shown in Figure 3.14a and 3.14b respectively. While the effects are significant both economically and statistically for inexperienced users, experienced users seem to be unaffected by variety interventions. This may be because experienced users are less susceptible to session-level manipulations such as variety.

- *Clicking history*: Users' clicking record is usually a good indicator of their interest in ads. We call a user *responsive* if he/she has at least clicked on one ad in his/her record. We estimate our model on the sample of responsive vs. non-responsive users and present the results in Figure 3.14c and 3.14d respectively. The patterns reveal a stark contrast: while variety of previous ads leads to a higher CTR on the sample of responsive users, the effects flip on the non-responsive sample. This can stem from the fundamental difference between these two groups: the responsive users are more interested in ads, thereby variety increases their engagement with ads. However, non-responsive users who have not shown any interest in ads seem not to enjoy higher variety of ads.
- *Past interaction with the focal ad*: Another historical feature of an impression that can be an important source of heterogeneity is the number of previous exposures of the focal ad that the user has seen prior to the session. Clearly, the number of previous exposures of the current ad affects user's information processing of the ad. Thus, the effects of variety are likely heterogeneous conditional on the previous exposures of the current ad. We split the data into two parts based on the previous exposures of the current ad. We call the condition *high repetition* if the user has seen over 5 exposures prior to the session and split the data based on this condition. As shown in Figure 3.14e and 3.14f, we find that the effects of variety are stronger if the user has seen fewer exposures of the current ad prior to the session. However, variety of prior ads is less effective if the current ad has been shown over 5 times before.

3.6.5 Robustness Checks

In this section, we present some robustness check to examine whether our findings hold under different specifications. We outline different checks below:

- *Temporal Spacing:* One confounding factor for temporal features can be temporal spacing.

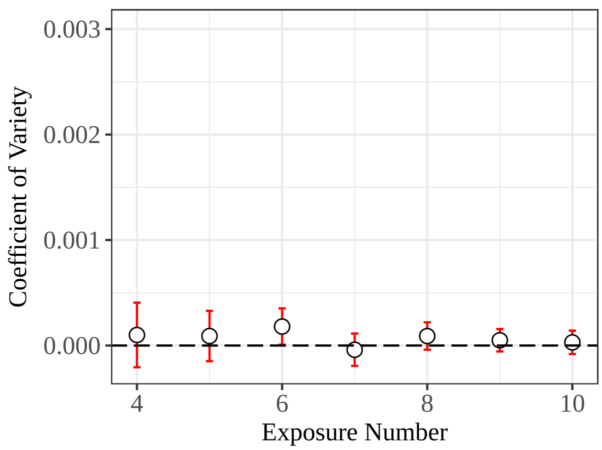

(a) Experienced

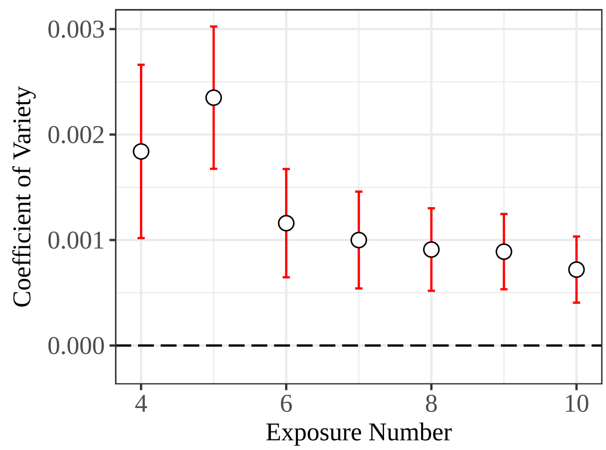

(b) Inexperienced

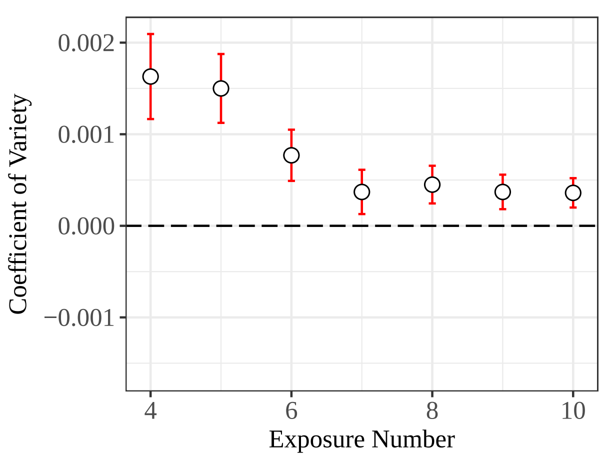

(c) Responsive

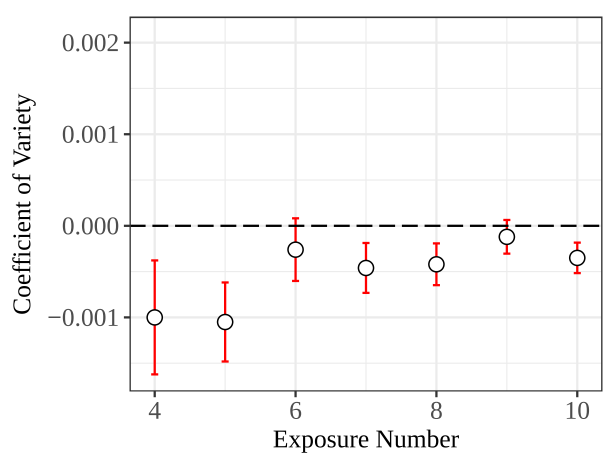

(d) Non-responsive

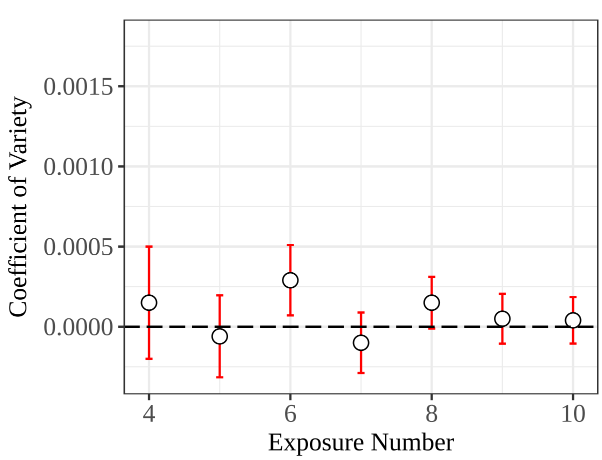

(e) High Repetition

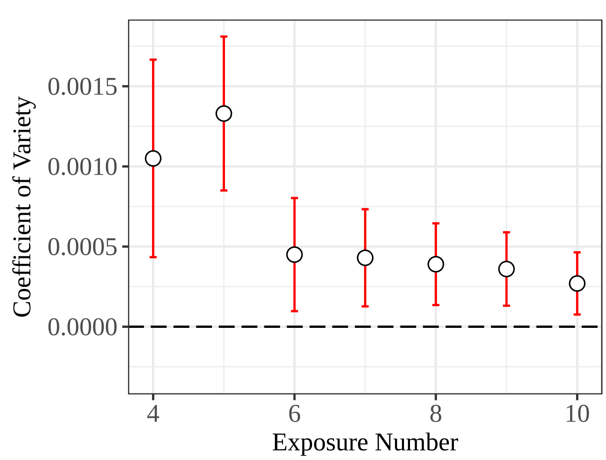

(f) Low Repetition

Figure 3.14: Heterogeneity in the effects of variety of previous ads

[118] finds that temporal spacing between the current ad and the last time the ad was shown has a positive effect on the clicking outcome. One would argue that omitting this variable would result in biased estimates for variety, as higher spacing might be correlated with higher variety. To control for this potential confounding factor, we create the variable $space_{i,t}$, which is the number of impressions the user has seen in session i, since she last saw ad $A_{i,t}$. Obviously, $space_{i,t}$ is not defined if this is the first time showing $A_{i,t}$ within the session. We run two different models to address this issue. First, we add the variable $space_{i,t}$ to the specification presented in Equation (3.11) and estimate the model. However, this might reduce our power as we drop many data points. This brings us to our second strategy where we define the dummy variable for first exposure of an ad. The result of these models are presented in Table B.1 and B.2 in Appendix. By and large, the coefficient for variety is positive and significant, implying positive effects of variety on click. Further, we find positive effects for temporal spacing in most exposures, which confirms the findings of [118].

- *Other Measures of Variety:* Besides measures that we used for variety, Shannon entropy is another metric for dispersion and diversity which is widely used in information theory literature [124].[12] Below is a formal definition of this metric:

$$Shannon_{i,t} = -\sum_{a} p_{i,a,t} \log_2 p_{i,a,t}, \tag{3.17}$$

  where $p_{i,a,t}$ is the share of ad a in the session i prior to exposure number t, i.e., $p_{i,a,t} = \frac{I_{i,a,t}}{t-1}$. This measure captures the *amount of information* in the past sequence.[13] We use this measure of variety as $V_{i,t}$ in (3.11) and estimate the model on the intended sample. As presented in Table B.3, our estimates for the effects of variety are robust to Shannon measure of variety.

- *Different Stratification:* In our main model, we used strata fixed effects to control for cross-sectional selection. Each stratum is a combination of a targeting area g and hour h.

[12] For a marketing use and discussion of this metric, please see [46].

[13] More practically, it translates into number of bits required to store the sequence.

Clearly, such strata can be very narrow containing only a few corresponding observations. While it helps us avoid the cross-sectional selection bias, it may reduce our power of identifying the effects. On the other hand, one would argue that the current stratification may not fully eliminate the selection bias, as the time intervals can be wide and we pool all the ads and control for ad effects separately. We conduct a series of different specifications with more or fewer fixed effects. The results are shown in the appendix. Table B.4, B.5, and B.6 present the estimates when using less fixed effects. By and large, the substantive results remain the same, though the effects are underestimated compared to the main results. Such discrepancy suggests that advertisers take into account the interactions between targeting variables when targeting their ads, thereby inducing bias in the models that do not capture such interactions.

Further, we estimate models with even narrower strata. As discussed in §3.5.3, to control for time-dependent factors that change the auction, we used hourly intervals. Here we use half-an-hour intervals to see whether the results are robust to tighter time intervals. As shown in Table B.7, the results are almost the same as our main model. Finally, we incorporate ad fixed effects in our stratification and estimate the model. That is, for any auction-invariant stratum and ad, we define a separate group and control for such group fixed effects. The results, as presented in Table B.8, show the same pattern as our main model. However, for later exposures the coefficients lose significance which is likely due to lack of power.

- *Top Ad:* In our main model, we pool all ads together and control for their fixed and time-varying effects. An alternative approach is to run models for each ad separately to show the effects on each individual ad and explore the heterogeneity in the effects across ads. However, we do not have enough power for each ad to identify the effects. Here we estimate the model on the top ad for which we have enough observations. As show in Table B.9, we see the same pattern as our main model.
- *Controlling for User's Panel Length:* In our main analysis, we take session as the unit of analysis. However, we may observe multiple sessions for only one users. As a result, our

main results may be driven by some users for whom we see many sessions. Further, one would argue that user's behavior may evolve over time with respect to session-level temporal features, which in turn, violates *Stable Unit Treatment Value Assumption* (SUTVA) [24, 116]. To address this issue, we only focus on user's first assignment to an experimental condition. That is, for any exposure number t, we focus on the first time users get to see that specific exposure number and estimate our model. This changes our unit of analysis from sessions to users. The only problem with that approach is the significant loss in the power of our analysis, as we have to drop many sessions. We estimate our model using this approach and present the results in Table B.10. The estimates of the effects of variety show the same patterns as our main results.

3.7 Conclusion and Future Directions

Mobile in-app advertising is now a major source of revenue for many app developers. Given some specific features of this medium of advertising, sequential ad placement is a common practice among publishers who serve ads. That is, users are exposed to a sequence of potentially different ads within a session. This motivates a series of questions related to the sequencing of ads. We particularly focus on variety as an important feature of ad sequences and study how variety of previous ads is linked to user's clicking behavior on the next ad. To answer this question, we use data from the leading in-app ad-network of an Asian country to examine this question. A unique feature of our data is the use of probabilistic auction for ad placement that creates a great degree of randomness in the sequence of ads users are exposed to within the session. We develop an identification strategy that allows us to exploit the exogenous variation in users' assignment to variety and obtain intent-to-treat estimates. We find that variety of previous ads leads to higher likelihood of click on the next ad. We then examine the source for such effects using different measures of variety and identify sequential organization of exposures as a major source. Finally, we use a descriptive approach to document heterogeneity in variety effects across user's past history.

In sum, this chapter contributes to two broad streams of literature on advertising and variety. First, this chapter adds to the methodological literature on quantifying causal effects of advertising. We propose an identification strategy that can be applied to cases where ads are allocated through a probabilistic mechanism. Specifically, our approach is useful for identification of temporal effects of advertising using observational data, where users may drop out before fully assigned to the treatment. Second, this chapter adds to the marketing literature on variety by establishing the role of variety in advertising context. Our findings provide some managerial insights for both advertisers and publisher in mobile in-app advertising industry. For advertisers, our findings would

Our findings have managerial implications for both advertisers and publishers. One implication for advertisers is to participate in auctions with many competitors to be assigned to higher variety conditions. Further, our results on heterogeneity in variety effects would guide advertisers' targeting decisions. For publishers, our results highlight the importance of questions related to mechanism design – whether they must use deterministic vs. probabilistic auctions and how they should design the sequence of ads.

Nevertheless, there are certain aspects of the problem that this chapter overlooks but can serve as avenues of future research. First, we only focus on click as the main outcome of interest. However, click is not the ultimate outcome that advertisers care about. Examining whether this increment in clicks will lead to higher conversion would be an interesting area for future research. Second, while our results establish the causal effects of variety, it is not clear what the optimal policy would be. Increasing variety may come with the cost of showing irrelevant ads giving rise to an inter-temporal trade-off in sequencing of ads. Future research can explore how the publisher can optimally design the sequence of ads. Third, we show that assignment to variety affects user's decision to stay in the app. Questions related to the effects of advertising interventions on app usage are particularly important yet understudied in the literature. Thus, such question can open fruitful avenues for future research.

Chapter 4

ADAPTIVE AD SEQUENCING: VALUE CREATION

4.1 Introduction

4.1.1 Adaptive Interventions in Mobile Advertising

Consumers now spend a significant portion of their time on mobile devices. The average time spent on mobile devices by US adults has grown steadily over the last few years [37]. This demand expansion, in turn, has amplified marketing activities towards mobile users. In 2018, mobile advertising generated over $71 billion in the US, accounting for roughly double the share of its digital counterpart, desktop advertising [36]. Most of this growth in mobile advertising is attributed to in-app ads, i.e., ads shown inside mobile apps. Indeed, in-app advertising is now the dominant channel for mobile advertising, generating over 80% of ad spend in the mobile advertising category [33].

Two key features of mobile in-app ads have contributed to their growth. First, the mobile app ecosystem has excellent user tracking ability, thereby allowing "personalization" of ad interventions, i.e., targeting of users based on their prior behavioral history [55]. Second, in-app ads are usually short-lived and dynamic in nature: each ad intervention is shown for a fixed amount of time (e.g., 30 seconds or one minute) inside the app, and is then followed by another ad intervention. As such, a user can see multiple ad exposures within a session.[1] This is in contrast with the common practice in desktop advertising, where ads remain fixed throughout a session. Short-lived ads together with personalization potential make in-app advertising amenable to "adaptive interventions", i.e., targeting of ads based on time-varying behavioral information about users to maximize user engagement.

---

[1]A session is an uninterrupted time that a user spends inside an app.

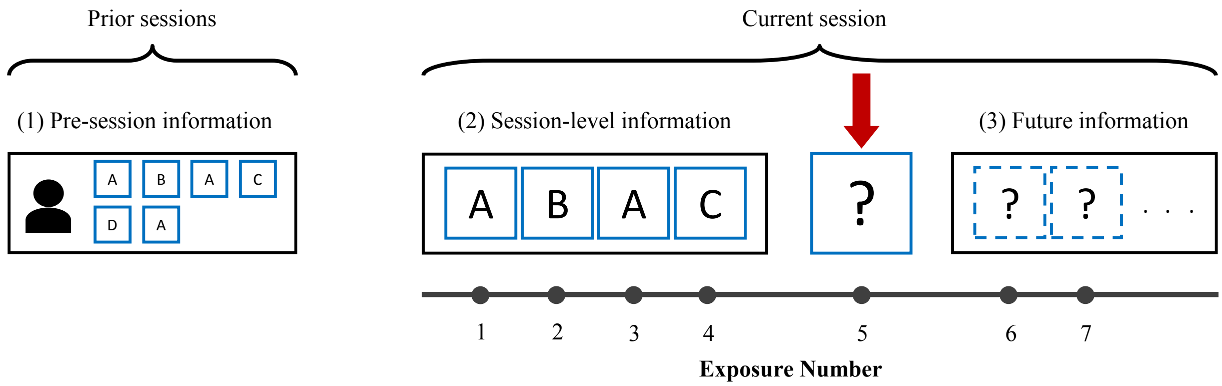

Figure 4.1: A visual schema the publisher's ad sequencing decision and different types of information available.

We illustrate the general schema for a mobile publisher's ad sequencing problem in Figure 4.1.[2] In this problem, the publisher has to decide which ad to serve in each period that the user stays in the app. We characterize three separate pieces of information that the publisher can incorporate when deciding which ad to show in any impression: (1) the pre-session information, which consists of user characteristics as well as the behavioral history of the user up until the current session, (2) the session-level information, which is the sequence of ads the user has seen so far within the session, as well as his response to each, and (3) the future information, which captures the publisher's expectations on how the sequence is likely to evolve in the next periods. The current research and practice mostly focus on the first two pieces and overlook the future information [86, 89, 107]. One reason is that capturing future information often adds significantly to the complexity of the problem. In addition, the returns from adopting a forward-looking model are not clear. Thus, the publisher's decision on whether to use a dynamic framework boils down to whether incorporating future information helps her achieve a better outcome.

In principle, if there is no interaction between sequential ad interventions, incorporating

[2] In this chapter, we use the publisher, ad-network, and platform interchangeably, when we refer to the agent who makes the ad placement decision.

future information will not improve the ad sequencing policy. However, at the advertiser level, the extant literature has documented various empirical effects that highlight such inter-temporal trade-offs. Prior findings on spillover and carryover effects of advertising rule out the independence of ads shown in a sequence [**?**, 118]. In particular, in the context of mobile in-app advertising, [106] find that a higher variety of previous ads shown within a sequence leads to a higher likelihood of click on the next ad. At the same time, they show that an ad will have a higher chance of being clicked if it has been shown more within the sequence of prior ads. These findings suggest a natural inter-temporal trade-off: higher variety by definition implies fewer repeat exposures of each ad. While these effects are well-established at the advertiser level, neither research nor practice has looked into how to collectively incorporate these findings to dynamically sequence ads to optimize publishers' outcomes. Thus, to the extent that these effects have significant impact on user's engagement with ads, it is important to develop a dynamic framework for ad sequencing and examine the gains from adopting such a framework.

4.1.2 Research Agenda and Challenges

In this chapter, we propose a dynamic framework that incorporates the inter-temporal trade-offs in ad sequencing and develops an optimal policy that optimizes user engagement with ads in a session. User engagement with ads or the match between users and ads is particularly an important outcome for publishers, as it is the main channel through which the publisher can create value by ad sequencing. We can use different metrics to measure user engagement with ads depending on the context of advertising. In the context of mobile in-app advertising, click is a reasonably good metric for the user engagement with ads because of two reasons. First, all ads are mobile apps whose objective is to get more clicks and installs (e.g., performance ads). Second, clicks are directly linked to revenues as most in-app ad publishers employ cost-per-click or cost-per-install as their monetization strategy.

Therefore, we use a click-maximizing objective and seek to answer the following three questions in this chapter:

1. How can we develop a dynamic theoretical framework that incorporates the inter-temporal trade-offs in ad sequencing and designs a policy that maximizes the expected number of clicks per session?

2. How can we empirically evaluate the performance of this dynamic ad sequencing policy relative to other benchmark policies?

3. What are the gains from using a dynamic framework to allocate ads? What explains the differences in outcomes and interventions under a fully dynamic sequencing policy and other benchmark policies?

We need to overcome three major challenges to satisfactorily answer this set of questions. First, to capture the dynamics of this problem, we need a theoretical framework that incorporates inter-temporal trade-offs in the publisher's ad allocation decision. Second, to develop a click-maximizing dynamic sequencing policy, we need to obtain accurate personalized counterfactual estimates of user behavior that allows us to evaluate all feasible policies (and not only those implemented in the data). As such, we need a setting with enough randomization in the ad allocation process that enables us to exploit this randomization and obtain accurate counterfactual estimates of user behavior. Finally, to measure the gains from the dynamic framework and explore the differences in sequencing strategies across sessions, we need to have a solution concept that determines the optimal dynamic policy and an evaluation approach that accurately estimates the outcomes under the new optimal policy for each session.

4.1.3 Our Approach

In this chapter, we present a unified three-pronged framework that addresses these challenges and develops a forward-looking adaptive sequencing policy to maximize user engagement with ads. We present an overview of our approach in Figure 4.2. As shown in the top row of this figure, we start with a theoretical framework that models the domain structure

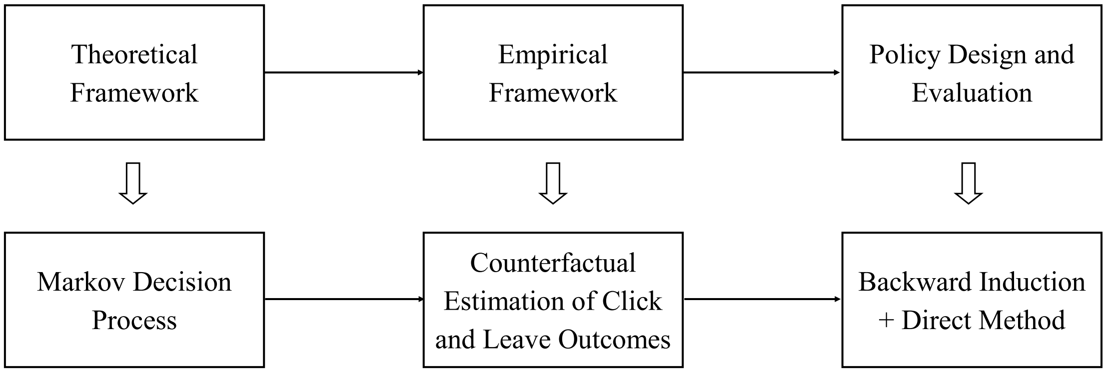

Figure 4.2: An overview of our approach. The top row presents our general framework and the bottom row shows the specific approach we take in this chapter.

of our problem and allows us to identify key empirical tasks required for the policy design and evaluation. The bottom row follows the same flow and illustrates the specifics of our approach to the ad sequencing problem in mobile in-app advertising:

Theoretical Framework: We start with a Markov Decision Process (MDP henceforth) that characterizes the structure of adaptive ad interventions. Starting from a theoretical framework allows us to address all the challenges discussed above. First, to capture the dynamics of the problem, we specify a domain-specific MDP with a rich set of state variables that incorporates the inter-temporal trade-offs identified in the literature. Our MDP characterizes the reward at any time period as well as how the state evolves in future periods, given any action taken by the publisher. Since our goal is to optimize the number of clicks per session, we define the reward as the expected probability of click on any ad selected. This probability is also informative of the available history in the next period, and thereby captures a probabilistic component of state transitions in our MDP. Another probabilistic factor that affects the future state is the expected probability of the user leaving the session after an intervention. It determines with what probability the user will be available to receive the next intervention. Together, these two probabilistic outcomes help us define the empirical tasks required in our problem: personalized counterfactual estimation of click and leave outcomes.

Empirical Framework: To obtain accurate counterfactual estimates for the click and leave outcomes, we use the filtering strategy proposed by [107] that filters out the ads that *could have never been shown* in a session. This is because we cannot estimate the outcome reliably for these ads as they are not within the joint distribution of the training set that is used for model fitting. However, if an ad *could have been shown* within a session (i.e., has a non-zero propensity score), our outcome estimates will be accurate for this ad in this session as long as we control for all other covariates affecting the propensity of seeing this particular ad. Further, to personalize these counterfactual estimates, we employ machine learning methods that can capture more complex relationships between the covariates and the outcome. In particular, we use the Extreme Gradient Boosting (XGBoost henceforth) method developed by [22], which is a fast and scalable version of Boosted Regression Trees [41].

Policy Design and Evaluation: In order to develop the optimal dynamic policy, we can use our counterfactual estimates for click and leave probabilities and numerically solve for the value functions, given the state variables. In our case, we employ a non-stationary finite horizon MDP and determine the optimal dynamic policy using backward induction. Finally, we use a direct method evaluation method to evaluate the performance of any policy at the session level. This method uses the outcome estimates for both click and leave to simulate the stochasticity within the session and evaluate the outcome/outcomes of interest.

4.1.4 Findings and Contribution

We apply our framework on the data from a leading mobile in-app ad-network of a large Asian country. Our setting has some notable features that make it amenable to our research goals. First, the ad-network uses a short-lived ad format where ad interventions last for a short period of time and change within the session. Second, the extent of randomization in ad allocation is quite high because of two reasons: (1) the ad-network runs a quasi-proportional auction that employs a probabilistic allocation rule, and (2) the ad-network only allows limited targeting on broad categories. These two features help us estimate click and leave outcomes for a wide range of counterfactual ads. This task would not have been possible if the ad-network had run a deterministic mechanism such as the second-price auction or allowed micro-level targeting.

We first discuss the results from our first-stage machine learning models for both click and leave outcomes. We evaluate these predictive models on a hold-out test set, using various goodness-of-fit measures. We show that both the click and leave models achieve high out-of-sample predictive accuracy. We then show that dropping variables that are defined either at the ad- or session-level from the predictive model leads to a significant drop in the performance of both click and leave models. These results provide preliminary evidence for the gains from the dynamic framework since the publisher can actively influence these variables by adopting an adaptive forward-looking policy.

Next, we focus on the main goal in this chapter – evaluating the gains from *adaptive forward-looking sequencing policy*. This policy incorporates all three pieces of information in Figure 4.1. To establish the performance of our adaptive forward-looking sequencing policy and identify where these gains come from, we define a set of benchmark policies that use a combination of different pieces of information:

- *Random sequencing policy*: This policy allocates exposures to ads randomly. It serves as a baseline for capturing how well we can do without any model.
- *Non-adaptive single-ad policy*: This policy uses the pre-session information to allocate ads within the session. Since this information does not change within the session, the optimal allocation based on only this information is to select a single ad that maximizes publisher's rewards. As such, this policy simulates the case where the ad slot is fixed throughout the session and helps us identify the opportunity costs of using a fixed ad slot.
- *Adaptive myopic sequencing policy*: This policy uses both pre-session and session-level information and allocates each impression to the ad with the highest probability of click in that impression, regardless of how it affects the expected future rewards.

We find that all model-based policies lead to substantial gains compared to the random sequencing policy, in terms of the expected number of clicks per session. In particular, the adaptive forward-looking sequencing policy increases the expected number of clicks by 80.36% relative to the random sequencing policy. We then show that the expected number of clicks is 7.87% higher under the adaptive forward-looking sequencing policy as compared to the non-adaptive single-ad policy. This finding demonstrates the opportunity cost of using a fixed ad slot throughout the session, which supports the current industry trend of using short-lived ad slots. Finally, we show that the adaptive forward-looking sequencing policy results in 1.50% increase in the expected number of clicks per session, compared to the adaptive myopic sequencing policy. This suggests that choosing the best match at any point will not necessarily create the best match outcome at the end of the session. Rather, the right action sometimes is to show the ad that is not necessarily the best match at the moment but transitions the session to a better state in the future. Together, these findings establish the benefits of adopting an adaptive and forward-looking approach to allocate ads. This has important implications for publishers and ad-networks, especially since the current practice in the industry overlooks the dynamics of ad sequencing.

Next, we explore the heterogeneity in the gains from the adaptive forward-looking sequencing policy compared to both non-adaptive single-ad and adaptive myopic sequencing policies, across the sessions. We examine how the relative gains from the adaptive forward-looking sequencing changes, as the number of prior sessions a user has participated in increases. On the one hand, we expect more data on the user to benefit the adaptive forward-looking sequencing policy more than other policies, as the publisher can take prior usage patterns of the user into account to improve the adaptive forward-looking sequencing policy. On the other hand, sequencing effects are shown to become smaller as the user becomes more experienced [106]. We find evidence for the latter: the relative gains from adopting adaptive forward-looking sequencing policy diminishes as the user participates in more sessions. We further provide some descriptive evidence that the source for these diminishing returns seems to be the number of distinct ads the user has seen: if a user has seen many distinct ads over time, it is harder to affect their decision through adaptive forward-looking sequencing of ads.

Finally, we present some descriptive analysis to better understand how and why the adaptive forward-looking and adaptive myopic sequencing policies differ. We find that the adaptive forward-looking sequencing policy tends to repeat the same ad in consecutive exposures more than the adaptive myopic sequencing policy. This is likely because repeating an ad consecutively is not the optimal decision if the publisher only takes the reward at that moment into account. However, it is the right decision if the publisher is forward-looking, as it increases the number of prior exposures for an ad and strengthens the session-level carryover effects.

In sum, this chapter makes three contributions to the literature. First, from a methodological standpoint, we propose a unified dynamic framework that theoretically characterizes the domain structure of the mobile in-app advertising environment, and an empirical approach that allows us to break the problem into composite machine learning tasks. To our knowledge, this is the first work to collectively incorporate temporal effects of advertising documented in the literature at the advertiser level, and propose a dynamic framework for the sequential allocation that characterizes optimal policy design for publishers. The generality of our framework makes it applicable to the contexts where advertisers have to make dynamic decisions such as the attribution problem. Second, from a substantive point-of-view, we establish the gains from an adaptive forward-looking sequencing policy as compared to other benchmarks that are often used in research and practice. This finding is of importance, as the current practice in this industry ignores the dynamics of ad allocation problem. Third, from a managerial perspective, we quantify the opportunity costs of using fixed ad slots. Our framework can help marketing practitioners and ad-networks to design the ad slot that is optimal in their context.

4.2 Related Literature

This chapter relates and contributes to several streams of literature.

First, this chapter relates to the marketing literature on personalization and targeting in digital platforms. Early papers in this stream build Bayesian frameworks that exploit behavioral data and personalize marketing mix variables [?, ?, 4]. Other recent papers in this area use a combination of feature generation and supervised machine learning frameworks to provide more scalable solutions for personalization and targeting policies for large-scale data [137, 107]. While all these papers focus on prescriptive or substantive frameworks to study personalization, they all study this phenomenon from a static point-of-view. This chapter extends this literature by offering a scalable and dynamic framework to develop personalized targeting policies.

Second, this chapter relates to the literature on the temporal effects of advertising, such as spillover effects in search advertising [?, 119], carryover effects in display advertising [70], temporal interactions between multiple advertising channels [85], effects of temporal spacing in search advertising [?], and the effects of variety of previous ads in mobile in-app advertising context [106]. While these papers establish the presence of these temporal effects from advertisers' perspective, they do not address how a publisher can use this information to optimally show ads in sequences. In this chapter, we view this problem from a publisher's perspective who wants to maximize users' engagement with ads and develop a dynamic framework that incorporates all possible temporal effects that have been documented in the literature.

Third, this chapter relates to the literature on dynamic policy design in digital advertising. Given the complexity of solving a dynamic policy, prior works often simplify the problem to avoid the curse of dimensionality. [130] focus on 16 cognitive-style segments and combine dynamic programming with a Bayesian framework to infer segment membership. In the context of linear video ads, [74] incorporate ad-specific leave probability as the only source of inter-temporal trade-off and use a cascade model to obtain a dynamic policy. Using the context of mobile in-app advertising, [126] theoretically examine the problem of short-lived ads in mobile in-app advertising and theoretically derive the dynamic policy by only focusing on certain aspects of the dynamics in this problem. Closely related to our problem, [127] present a framework to develop and evaluate dynamic policies and study the gains at the advertiser level. This work differs from these papers since we are the first to present a dynamic framework that incorporates all established temporal effects of advertising and examines the gains from adopting an adaptive forward-looking sequencing policy, from the publisher's perspective.

Finally, this chapter relates to the growing literature on machine learning applications in marketing. The vast majority of papers in this stream focus on prediction tasks and use various supervised or unsupervised learning algorithms to achieve a better predictive accuracy [128, 59, 29]. A narrower body of works in this area brings machine learning methods to policy design questions. Using a static approach, some recent papers develop machine learning methods to design optimal policies in various contexts such as pricing [27], ad placement [107], and CRM campaigns [62]. Incorporating the dynamics of exploration-exploitation trade-off, [123] offer a multi-armed bandit approach in a display advertising context. This chapter adds to this literature by fully incorporating the dynamics of the ad sequencing problem through an MDP and linking it to smaller machine learning tasks.

4.3 Setting and Data

4.3.1 Setting

Our data come from a leading mobile in-app advertising network of a large Asian country that had over 85% of the market share around the time of this study. Figure 4.3 summarizes most key aspects of the setting. We number the arrows in Figure 4.3 and explain what each step of the ad allocation process in details below:

1. The ad-network designs an auction to sell ad slots. In our setting, the ad-network runs a quasi-proportional auction with a cost-per-click payment scheme. As such, for a given ad slot and a set of participating ads $\mathcal{A}$ with a bidding profile $(b_1, b_2, \ldots, b_{|\mathcal{A}|})$, the ad slot is allocated to ad a with the following probability:

$$q_a^p(b; z) = \frac{b_a z_a}{\sum_{j \in \mathcal{A}} b_j z_j}, \tag{4.1}$$

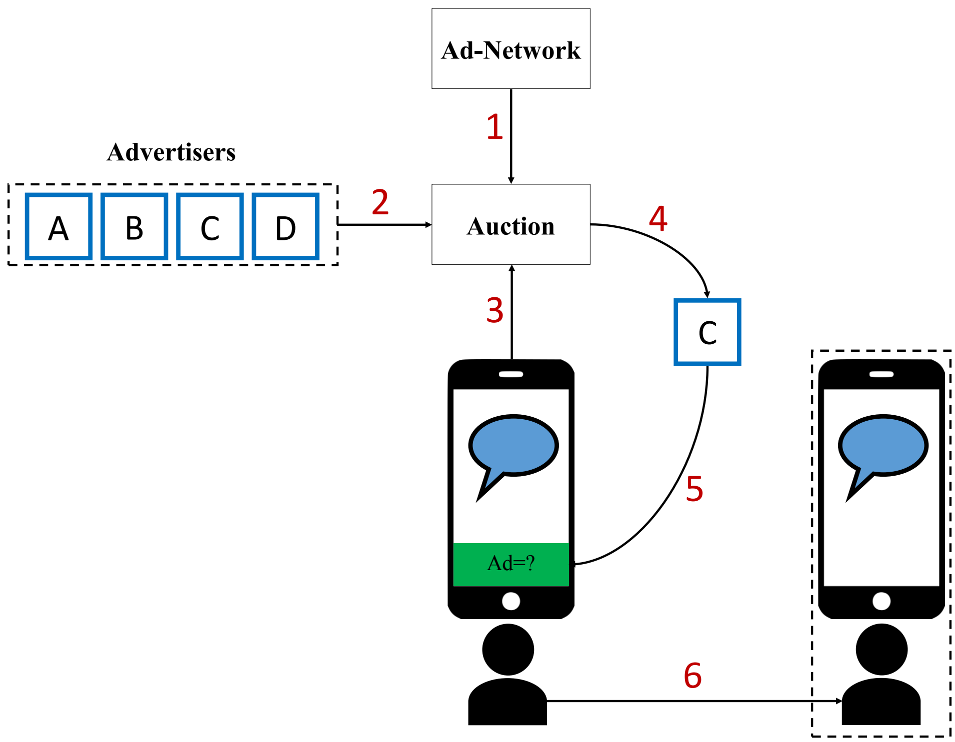

Figure 4.3: A visual schema of our setting

where z_a is ad a's quality score, which is a measure reflecting the profitability of ad a. The ad-network does not customize quality scores across auctions.[3] The payment scheme is cost-per-click and is similar to Google's sponsored search auctions. That is, ads are first ranked based on their product of bid and quality score, and the winning ad pays the minimum amount that guarantees their rank if a click happens on their ad.

2. Advertisers participating in the auction choose: (a) design their banner, (b) specify the areas in which they want to show their ad, and (c) submit their bid. Figure 4.3 shows an example of auction with four different ads.

[3]In our data collection period, each ad just had one quality score.

3. Whenever a user starts a new session in an app (in Figure 4.3, we use a messaging app as an example), a new impression is being recognized, and a request is sent to the publisher to run an auction.

4. The auction takes all the participating ads into account and selects the ad probabilistically based on the weights shown in Equation (5.23). It is worth noting that all the participating ads have the chance to win the ad slot. This is in contrast with more widely used deterministic mechanisms like second-price auctions, where the ad with the highest product of bid and quality score always wins the ad slot.

5. The selected ad is placed at the bottom of the app, as shown in Figure 4.3.

6. Each ad exposure lasts one minute. During this time, the user makes two key decisions: (a) whether to click on the ad, and (b) whether to stay in the app or leave the app and end the session. If the user clicks on the ad, the corresponding advertiser has to pay the amount determined by auction. After one minute, if the user continues using the app, the ad-network treats the continued exposure as a new impression and repeat steps 3 to 6 are repeated until the user leaves the app. We assume that a user has left the app when the time gap until the next exposure exceeds 5 minutes. Consistent with this definition, we define a session as the time interval between the time a user comes to an app and the time she leaves the app.[4]

4.3.2 Data

We have data on all impressions and clicks for the one month period from 30 September 2015, to 30 October 2015. Overall, we observe 1,594,831,699 impressions along with 14,373,293 clicks in the data, implying a 0.90% CTR. We now describe our raw variables and sampling procedure.

[4]There are obviously various ways to define a session based on the time gap between two consecutive exposures. We show that our results are robust to different definitions.

Raw Variables

Each impression contains the following raw information:

- *Time and date:* The exact time-stamp of the impression.
- *App information:* The identifier for a mobile app that shows ads through the ad-network.
- *User information:* An identifier that is unique to each mobile device and serves as our user ID.
- *GPS information:* The exact latitude and longitude of the user at the time of the impression.
- *Targeting variables:* The set of variables that advertisers can target on. There are five main categories that advertisers can target: province, hour of the day, smartphone brand, connectivity type, and Mobile Service Provider (MSP). If an advertiser decides to exclude a certain sub-category within these variables (e.g., Samsung smartphones), his ad will not be shown in impressions in that sub-category.
- *Ad information:* The set of variables related to the ad shown in the impression. It consists of an ad identifier and the potential cost-per-click.[5][6]
- *Click outcome:* A binary variable indicating whether the user clicked on the ad. This is our primary outcome of interest to measure the match between users and ads. While click is generally an imperfect outcome for the effects of ads, it is reasonable to use it as our primary outcome for two reasons. First, all ads are mobile apps whose objective is to get more clicks or installs (e.g., performance ads). Second, the ad-network uses a cost-per-click auction, which directly links the click outcome to the ad-network's revenues.

---

[5] The potential cost-per-clik is the amount that the ad would have paid if the user had clicked on their ad.

[6] We do not have the data on the banner creatives and its format, i.e., whether it is a jpeg file or an animated gif.

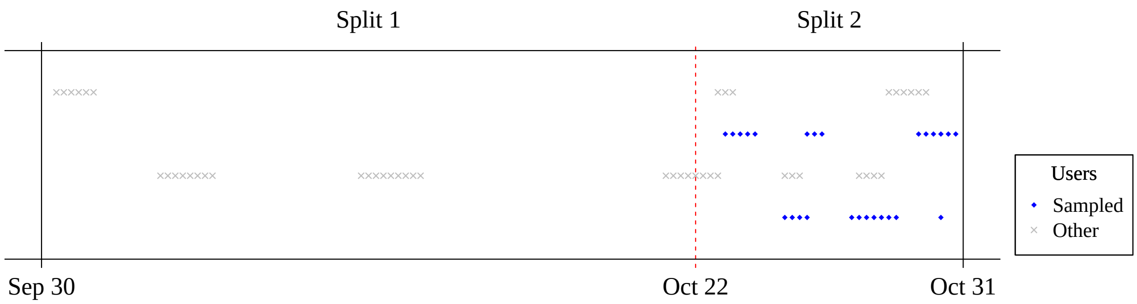

Figure 4.4: Schema for identification of new users.

Sampling Procedure

Our sampling procedure consists of two essential steps – (1) user sampling, and (2) app sampling. We describe each step in greater details below:

- *User sampling:* Since we want to optimally sequence ads within the session, our optimal intervention depends on users' past history. As such, we only focus on users for whom we can exploit their entire history. The challenge is that there is no variable in our data identifying new users. As illustrated in Figure 4.4, our approach is to split our data into two parts based on a date (October 22), and keep users who are active in the second part of the data (October 22 to October 30), but not in the first part (September 30 to October 22). This sampling scheme guarantees that the users who are identified as new users have not had any activity in the platform at least for the last three weeks. We drop all the other users from our data.
- *App sampling:* We only focus on the most popular mobile app in the platform, which is a messaging app that has over 30% share of total impressions. As such, we drop new users who do not use this app. There are a few reasons why we focus on this app. First, this is the only app whose identity is known to us. Second, we expect the sequencing effects to be context-dependent, and focusing on one app helps us perform a cleaner analysis.

Finally, it takes users a relatively long time to learn how to use certain apps (e.g., games), and learning effects can interact with sequencing effects. However, this messaging app is widely popular in the country and easy to use, so we expect users to pay more attention to ads from the beginning.

Overall, our sampling procedure gives us a total of 8,323,778 impressions shown to a set of 94,884 unique new users. Over 40% of these users use other apps in addition to the messaging app. In our data, there are 6,955,995 impressions shown inside the focal messaging app that corresponds to 1,271,068 unique sessions. For our analysis, we only focus on the impressions shown in the messaging app. However, we use impressions shown in other apps for feature generation.[7]

4.3.3 Summary Statistics

We now present some summary statistics on the data. As mentioned earlier, these statistics correspond to the impressions shown in the messaging app.

Shares of Categorical Variables

Since most raw variables presented in §4.3.2 are categorical, we cannot show the mean and standard deviation for them. Instead, we present the number of categories as well as the shares of the top three sub-categories within each variable in Table 4.1.

As shown in Table 4.1, one province accounts for a quarter of all impressions. The second row shows that there are certain hours of the day with more user activity. We find that these hours are late at the night when users are not at work. Table 4.1 shows that two major smartphone brands constitute over 80% of all impressions. Finally, we find that ads have different shares: the top three ads account for roughly 35% of all impressions. On the other hand, we observe that most ads have a very small share. This is mostly because these

[7] Our sampling procedure is almost identical to that of [106]. However, the number of impressions and sessions is slightly different, because we need to drop users with missing information on the latitude and longitude. [106] use those impressions because latitude and longitude do not play a role in their analysis.

Variable	Number of categories	Share of top categories		
		1^{st}	2^{nd}	3^{rd}
Province	31	24.58%	9.54%	7.50%
Hour of the Days	24	8.48%	8.03%	7.25%
Smartphone Brand	7	44.72%	38.07%	10.11%
Connectivity Type	2	50.54%	49.46%	
MSP	3	50.15%	44.07%	5.77%
Ad	327	18.43%	8.03%	7.50%

Table 4.1: Summary statistics of the categorical variables. This includes the number of categories and the percentage shares for the top sub-categories within each variable.

ads ran short campaigns. Later in §4.4.2, we show the full distribution of ad shares and provide a more detailed discussion.

Distribution of Session-Level Outcomes

Our goal in this chapter is to examine how much we can improve session-level user engagement with ads through optimal sequencing of ads. As such, the key outcomes are defined at the session level. Figure 4.5 shows the empirical CDF of two main outcomes of interest in this study – session length, and the total number of clicks made in a session, which is our primary outcome of interest. We measure session length by the number of exposures shown within any session. Figure 4.5a shows how this outcome varies across sessions. As shown in this figure, around 50% of all sessions end in only two exposures. Further, the empirical CDF in Figure 4.5a shows that the vast majority of sessions do not last for more than 10 exposures and only a small fraction of them last for 30 or more exposures.

In Figure 4.5b, we show the empirical CDF for our primary outcome of interest – the total number of clicks per session. As expected, most sessions end with no clicks being made on ads shown within the session, and the percentage of sessions with at least one click amounts to 7.77%. This is a reasonably high percentage in this industry. Interestingly, there

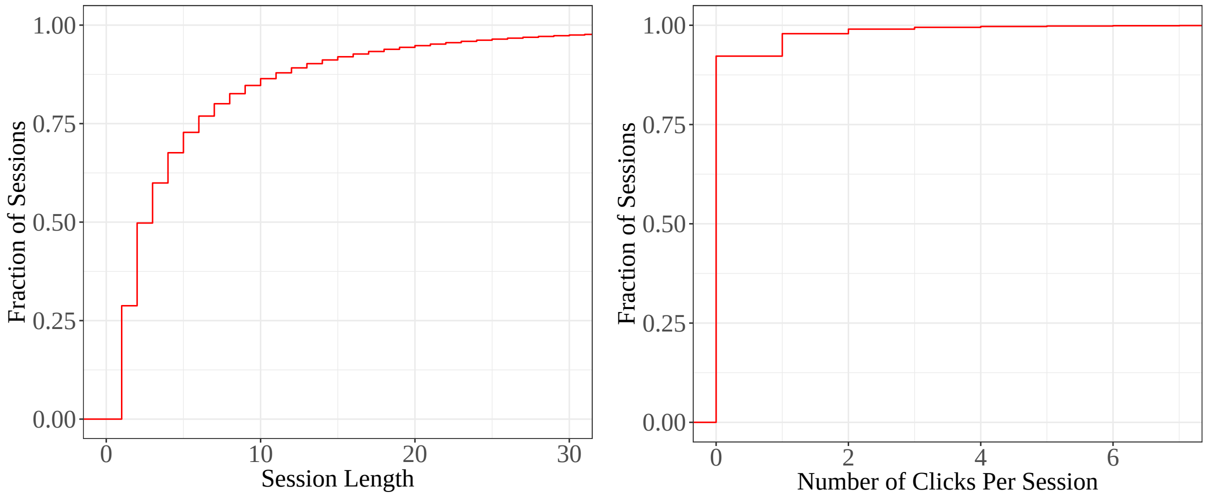

(a) Session Length

(b) Number of Clicks Per Session

Figure 4.5: Empirical CDF of the session length and total number of clicks per session.

are sessions with more than one click. Further exploration suggests that these sessions are typically much longer than other sessions, with an average length of over 15 exposures.

Distribution of Micro-Interventions

A central piece of our study is the sequence of ads shown within the session. These sequences are determined by the publisher's ad placement decision at a given exposure. As such, each ad exposure can be treated as a micro-intervention that forms the whole sequence. We focus on three binary micro-interventions at any time period given the history of prior ads shown within the session – (1) *repeat*, (2) *breadth-increasing change*, and (3) *breadth-constant change*. Figure 4.6 illustrates these micro-interventions. This figure presents a case where the publisher wants to select the fourth ad. We define a micro-intervention as *repeat* if the publisher selects the last ad shown within the session. On the other hand, if the publisher shows any other ad, it is called a *change*, meaning that the current exposure shows a different ad from the last ad. In line with [106], we then decompose *change* into two parts – (1)

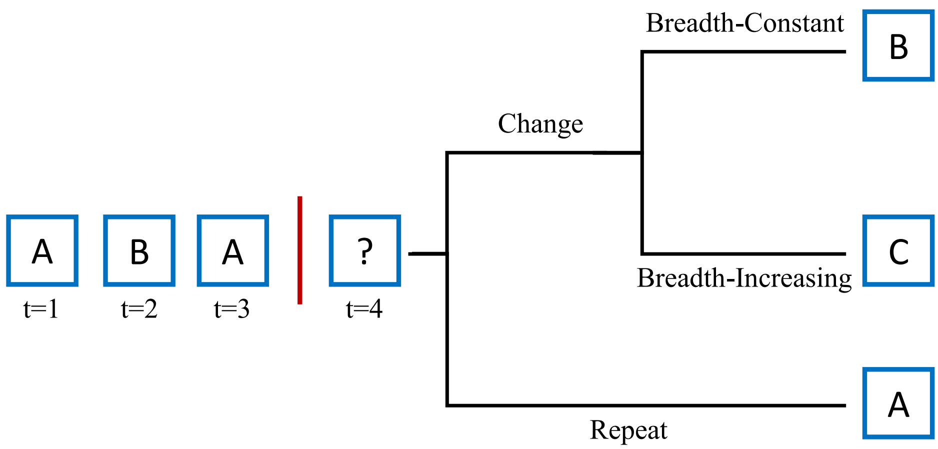

Figure 4.6: An example of three micro-interventions at exposure number 4

breadth-constant change, where the publisher changes the last ad but shows an ad that has been shown before in the session, and (2) *breadth-increasing change*, where the publisher changes the last ad and shows an ad that has not been shown before, thereby increasing the breadth of variety in the session. These three micro-interventions are mutually exclusive, i.e., only one of them takes the value one in any time period.

An important point to notice is that these binary micro-interventions together determine many characteristics of a sequence. For example, if prior interventions in a sequence are mostly breadth-increasing changes, we will have a high breadth of variety in that sequence. We present the distribution of these micro-interventions in our data in Figure 4.7. Each line in this figure represents the percentage of a specific micro-intervention at different points within the session. A few patterns emerge from Figure 4.7. First, the lines for both *change* and *repeat* are flat over time, illustrating the independence and stability of the auction across different exposures. Further, repeat accounts for only 20% of micro-interventions, which is mostly due to the probabilistic nature of the auction: each ad has a chance of being shown proportional to the product of its bid and quality score. This creates a great degree of randomization in the ad allocation process. Finally, we observe a decreasing pattern in

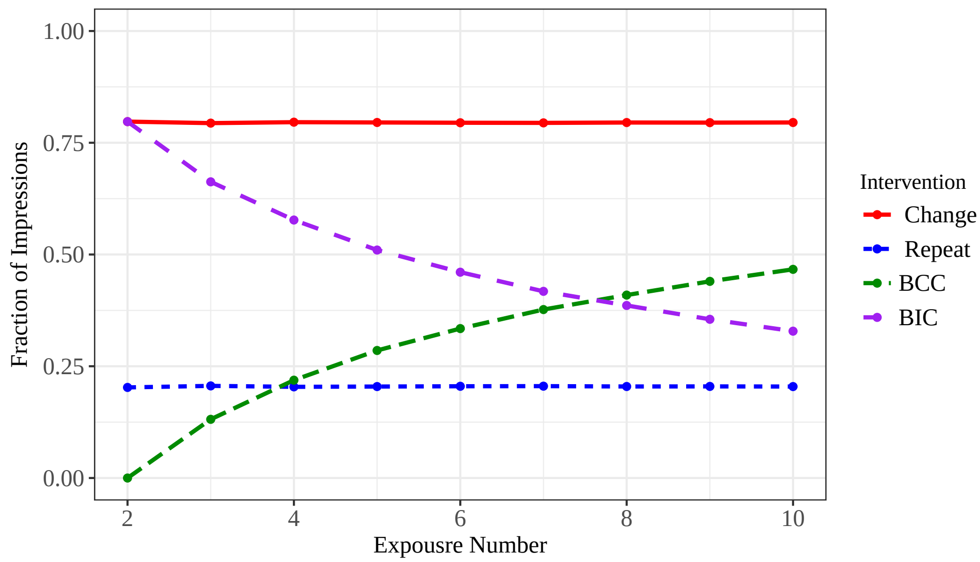

Figure 4.7: Distribution of micro-interventions at different exposure numbers within the session

breath-increasing changes. By definition: a random change is less likely to show a completely new ad because we have a finite set of ads competing in a given auction. The fraction of breadth-increasing changes is over 30% across all exposures, indicating the extent of variation in ad allocation within the session.

4.4 Dynamic Framework for Sequencing of Ads

We now present our dynamic framework for sequencing of ads. This section proceeds as follows. We start with the motivation for the use of a dynamic framework in our context in §4.4.1. Next, in §4.4.2 we define the primitives of our framework and then specify an MDP that incorporates publishers' current and expected future rewards.

4.4.1 Motivation

The micro-interventions presented in §4.3.3, by and large, depend on the history of the sequence. Further, they determine how a sequence evolves, thereby shaping the history of the sequence for future exposures. Thus, in principle, if the within-sequence history affects which micro-intervention is most effective at any point, it is not clear whether that intervention is the optimal action from a dynamic point-of-view. In fact, the optimal action will be the one with the right balance between how effective it is now and where it transitions the entire sequence. For example, suppose that the publisher wants to fill two impressions with two ads A and B. Now, consider a case where ad A is generally a better ad, so the publisher's optimal decision in both impressions is to allocate them to ad A. However, showing ad A after ad B generates the best overall outcome, as it differentiates ad A and generates a significantly higher click probability. In that sense, the best possible myopic action at a point may not be the optimal decision from a forward-looking perspective.

The main question is whether the within-sequence history affects the outcomes that the publisher cares about. Prior literature has offered some evidence on such effects, including the effects of multiple exposures and temporal spacing of ads [118]. In mobile in-app advertising setting, [106] show that users' clicking behavior on the next ad depends on the variety of the prior sequence, as measure by the total number of changes and the breadth of variety (sum of all breadth-increasing changes). While they document the positive effects of variety, they also find that the number of prior exposures of an ad within the session positively affects users' clicking behavior. Given these two findings, it is not clear how the publisher should manage these micro-interventions at any given point. On the one hand, more change creates a higher variety in the sequence, thereby increasing the probability of click on the next ad. On the other hand, repeating an ad will increase the number of prior exposures of that ad which is shown to positively affect the probability of click. Thus, the publisher faces various inter-temporal trade-offs when allocating ads within the session.

We address this challenge by developing a dynamic framework that: (1) captures the inter-temporal trade-offs in publishers' ad placement decision in the session, and (2) uses both pre-session and adaptive session-level information to personalize the sequence of ads for the user in any given session. Our framework incorporates both how an action affects the outcome in the current period, as well as the externalities that influence future exposures. In the next sections, we present the details of our framework.

4.4.2 Model Setup

We specify an MDP that captures the inter-temporal trade-offs in publishers' decision problem by taking into account both current and expected future rewards. An MDP is characterized by a set of primitives that serve as inputs into the objective function that the decision-maker seeks to maximize. We present a generic definition of these primitives below and discuss the specifics of each in the next sections.

1. Time Period (t): The first component we need to specify is the time unit. Since exposures are shown sequentially in our case, we treat each ad exposure as a time period wherein the publisher needs to decide on her actions. $t = 1$ indicates the first exposure in a session.

2. State Space ($\mathcal{S}$): The state space consists of all the information the publisher has about an exposure, which affects her decision-making process.

3. Action Space ($\mathcal{A}$): The action space contains the set of actions the publisher can take. In our case, this action is to show one ad from the ad inventory every time an impression is recognized. As such, $\mathcal{A}$ is the full ad inventory in our problem.

4. Transition Function (P): This function determines how the current state transitions to the future state given the action made at that point. As such, we can define $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ as a stochastic function that calculates the probability $P(s' \mid s, a)$ where $s, s' \in \mathcal{S}$ and $a \in \mathcal{A}$. Note that this is a crucial component of an MDP since

publishers cannot control the dynamics of the problem if the next state is not affected by the current decision.

5. Reward Function (R): This function determines the reward for any action a at any state s. As such, we can define this function as $R : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$. This function can take different forms depending on the publisher's objective. In our case, since the publisher is interested in optimizing user engagement, she can use different metrics that reflect user engagement such as the probability that the user clicks on the ad.

6. Discount Factor (β): The rate at which the publisher discounts the expected future rewards. In other words, it is the weight that the publisher assigns to the future relative to the current period.

With all these primitives defined, we can now write the publisher's maximization problem as follows:

$$\underset{a}{\operatorname{argmax}} \left[R(s,a) + \beta \mathbb{E}_{s'|s,a} V(s') \right], \tag{4.2}$$

where $V(s')$ is the value function incorporating expected future rewards at state s' if the publisher selects ads optimally. Following [12], we can write this value function for any state $s \in \mathcal{S}$ as follows:

$$V(s) = \max_{a} R(s,a) + \beta \mathbb{E}_{s'|s,a} V(s') \tag{4.3}$$

As shown in Equation (4.2), the optimization problem consists of two key elements – the current period reward and the expected future rewards. The publisher chooses the ad that maximizes the sum of these two elements.

State Variables

The state variables contain all the information that the publisher can use for any given exposure in a session. As discussed earlier, the publisher can take two pieces of information into account: (1) pre-session information, and (2) session-level information. Pre-session information contains any data on the user up until the current session, including his demographic variables and behavioral history. For any session i, we denote the pre-session state variables by X_i. It is important to notice that the pre-session variables are not adaptive, i.e., it does not change within the session and hence not have t subscript.

On the other hand, session-level variables are adaptive and change within the session. At any given point in a session, this information captures the information about the prior sequence of ads shown to the user as well as user's actions after seeing each ad. Let $G_{i,t}$ denote the session-level state variables. We can write:

$$G_{i,t} = \langle A_{i,1}, Y_{i,1}, A_{i,2}, Y_{i,2}, \ldots, A_{i,t-1}, Y_{i,t-1} \rangle, \tag{4.4}$$

where $A_{i,s}$ denotes the ad shown in exposure number s and $Y_{i,s}$ denotes whether the user clicked on this ad. As a result, $G_{i,t}$ is the sequence of all ads and actions within the session up to the current time period. Overall, we define the state variables as $S_{i,t} = \langle X_i, G_{i,t} \rangle$, i.e., a combination of both pre-session and session-level variables.

Action Space: Ad Inventory

As mentioned in §4.4.2, the publisher's action space in our case is the ad inventory. At each point of time, the publisher chooses one ad from the inventory to show to the user. We only focus on the top 15 ads in the ad inventory due to four key reasons. First, as shown in Figure 4.8, the top 15 ads generate over 70% of all impressions in the focal messenger app. So they collectively account for a significant portion of our observed data. Second, these top ads are more stable in terms of their budget. Since we want to run counterfactual policies, it is important to make sure that advertisers' budgets do not cause any problem for the reliability of our results. Third, we have more data for these ads as these are shown in a wide variety of state variables. This makes our estimates for their outcomes more accurate and less noisy. Finally, limiting the action space makes the problem computationally more tractable when we want to numerically solve for the optimal policy.

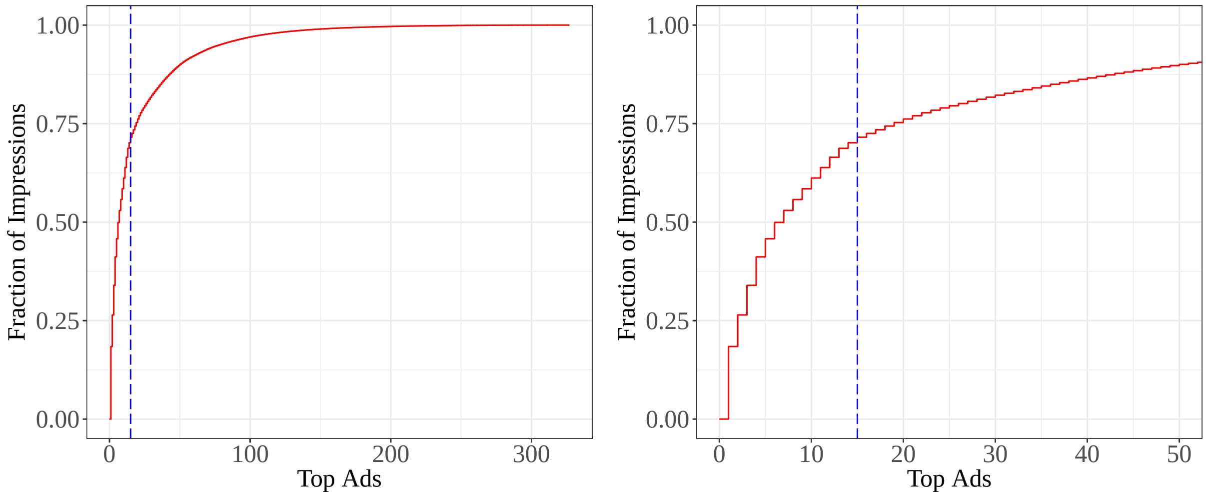

(a) All Ads (b) Top 50 Ads

Figure 4.8: Cumulative fraction of impressions associated with the top ads. The figure in the left shows the distribution for all 327 ads. The figure in the right zooms in the top 50 ads.

Transition Function

We now characterize the law-of-motion, i.e., how state variables transition given the publisher's action at any point. As mentioned earlier, we are interested in the probability of the next state being s', given that action a is taken in state s in the current period, i.e., $P(s' \mid a, s)$. Suppose that the user is in state $S_{i,t} = \langle X_i, G_{i,t} \rangle$ at exposure t in session i. The only time-varying factor in $S_{i,t}$ that can transition is $G_{i,t}$, which is the history of the sequence. Given the definition of $G_{i,t}$ in Equation (4.4), we can determine the next state if we know user's decision to click on the current ad and/or continue staying in the session. There are three mutually exclusive possibilities for state transitions:

1. *Click and stay:* If the user clicks on ad $A_{i,t}$ and stays in the session, we can define the

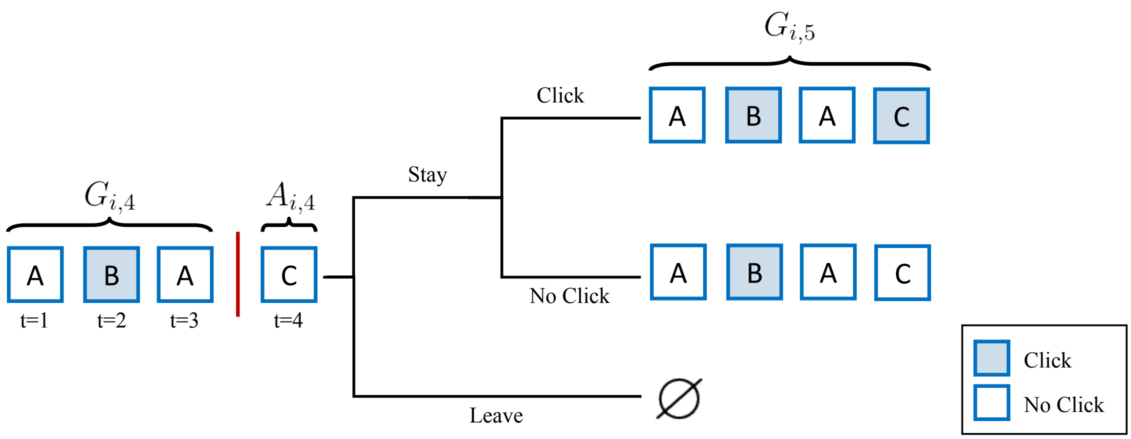

Figure 4.9: An example illustrating the state transitions.

next state as follows:

$$S_{i,t+1} = \langle X_i, G_{i,t}, A_{i,t}, Y_{i,t} = 1 \rangle, \tag{4.5}$$

where $Y_{i,t} = 1$ indicates that the user has clicked on the ad shown in exposure number t.

2. *No click and stay:* If the user does not click on ad $A_{i,t}$ and stays in the session, we can similarly define the next state as follows:

$$S_{i,t+1} = \langle X_i, G_{i,t}, A_{i,t}, Y_{i,t} = 0 \rangle, \tag{4.6}$$

where $Y_{i,t} = 0$ indicates that the user has not clicked on the ad shown in exposure number t.

3. *Leave:* Regardless of user's clicking outcome, if the user decides to leave, the entire session is terminated and there is no more decision to be made. Thus, we can write:

$$S_{i,t+1} = \varnothing \tag{4.7}$$

Figure 4.9 visually presents the three possibilities presented above. This figure illustrates an example where the publisher shows an ad in the fourth exposure in a session. It shows three possibilities and how each forms the next state. Based on this characterization, we can now define the transition function for any pair of action and state as follows:

$$P(S_{i,t+1} \mid a, S_{i,t}) = \begin{cases} \big(1 - P(L_{i,t} \mid a, S_{i,t})\big) P(Y_{i,t} \mid a, S_{i,t}) & \text{Equation (4.5)} \\ \big(1 - P(L_{i,t} \mid a, S_{i,t})\big)\big(1 - P(Y_{i,t} \mid a, S_{i,t})\big) & \text{Equation (4.6)} \\ P(L_{i,t} \mid a, S_{i,t}) & \text{Equation (4.7)} \\ 0 & \text{otherwise} \end{cases} \tag{4.8}$$

Equation (4.8) illustrates that the publisher needs to accurately estimate two user-level outcomes given any ad shown – click and leave probabilities. In §4.5, we discuss our approach to obtain these estimates.

Reward Function

Another piece of an MDP that needs to be defined is the reward function. The reward function can take different forms that vary with the publisher's objective and main outcome of interest. We primarily focus on the number of clicks per session as our main objective. In our case, clicks are particularly good measures of the user engagement with ads because of two reasons. First, all ads in our study are mobile apps that want more clicks and installs. In the literature, this type of ads is called performance ads and their match value is generally assumed to be the probability of click [3]. Second, click is the main source of revenue for the publisher, since the advertiser only pays when a click happens.

There are other advantages in using the click as the main outcome of interest. First, it is a well-recorded outcome that is realized immediately in the data. Second, click is a function of users' behavior, whereas other outcomes usually involve other players such as advertisers (e.g., publisher's revenues). Thus, the publisher's optimization problem only depends on inferring users' behavior, which is a feasible task in a data-rich environment.

Given that publishers want to maximize the number of clicks made per session, we can define the reward function as the probability of click for a pair of state and action. For exposure number t in session i, we can write:

$$R_t(a; S_{i,t}) = P(Y_{i,t} \mid a, S_{i,t}) \tag{4.9}$$

This is the probability of click on ad a if shown in the current state.

Discount Factor

The discount factor in Equation (4.2) reflects the relative importance of the expected future rewards compared to the current period reward from the publisher's perspective. As such, it generally takes a positive value close to one, if the publisher is forward-looking, i.e., they incorporate the expected future rewards in their decision problem. If the discount factor is zero, it means that the publisher is myopic and only cares about the reward in the current period.

In our dynamic framework, the publisher's decision is to select an ad for each impression in a session. The entire session happens in just a few minutes, depending on the number of exposures the user chooses to stay for. Given the short time horizon of the optimization problem, a risk-neutral publisher must value the current and expected future rewards equally, indicating that β is very close to 1.

Policy Definition

We characterize a dynamic policy as a mapping $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, that assigns a probability to any action $a \in \mathcal{A}$ taken in any given state $s \in \mathcal{S}$. For a deterministic policy, $\pi(a \mid s)$ will take value one only for one ad for any given state. However, we define the policy as a probability function to allow for non-deterministic policies as well.

In sum, the main goal of our framework is to develop a policy π^* that maximizes the expected reward given the initial set of state variables and the transition function. We later discuss how we develop and evaluate this policy in §4.6.2 and §4.6.3, respectively.

4.5 Empirical Strategy

In this section, we present our empirical strategy to estimate the primitives of the dynamic framework defined in §4.4. This involves the estimation of probabilistic components of the transition function as characterized in §4.4.2, as well as estimation of unknown components in the reward function as presented in §4.4.2. Together, this gives us two estimands – (1) leave outcome, and (2) click outcome. The task in both cases is to accurately predict the outcome for a pair of state and action. However, given the broader task of taking the best action at any state that we are interested in, we need to obtain these outcome estimates not only for the ad that is shown in the data but also for the counterfactual ads that are not shown in the data. Below is a formal definition of these two tasks:

Task 1: For any set of state variables observed in the data, we want to accurately estimate the click probability for all ads if shown in that impression. That is:

$$\hat{y}_{i,t}(a; S_{i,t}) = P(Y_{i,t} \mid a, S_{i,t}), \forall a \in \mathcal{A}_i \tag{4.10}$$

Task 2: For any set of state variables observed in the data, we want to accurately estimate the leave probability for all ads if shown in that impression. That is:

$$\hat{l}_{i,t}(a; S_{i,t}) = P(L_{i,t} \mid a, S_{i,t}), \forall a \in \mathcal{A}_i \tag{4.11}$$

We discuss our empirical strategy for Tasks 1 and 2 and helps us set the scope of our predictive model in §4.5.1. We then describe the learning algorithm that we use and its advantages in §4.5.3. Finally, in §4.5.4, we present the estimation results on our click and leave probability model.

4.5.1 Filtering Strategy for Counterfactual Estimation

For the task of outcome prediction like ours, we can use machine learning methods that capture more complex relationships between the covariates and outcomes [95]. However, these methods are only able to predict the outcome for observations within the joint distribution of the observed data. This implies that if an ad could have never been shown in an observation in the actual data, our outcome estimate for that ad may not be reliable. Thus, if there is no randomization in the ad allocation process, we cannot reliably estimate the outcome for counterfactual ads.

This informs our empirical strategy for counterfactual estimation of click and leave outcomes. While ads are selected through a deterministic allocation rule in most commonly used auctions such as second-price, the quasi-proportional auction in our setting has the advantage of using a probabilistic rule that induces randomization in the ad allocation process. The key issue with a deterministic allocation rule is that the ad that is actually shown in the data is the only ad that *could have been shown*, and other counterfactual ads *could have never been shown*. We argue that if an ad could have been shown in a session (i.e., has non-zero propensity score), this observation could have been generated within the joint distribution of the training data. Thus, we are able to accurately estimate the outcome for these ads in both factual or counterfactual situations.

To reflect this idea, we employ a filtering strategy similar to that in [107]. Here, for each session i, we identify the set of ads that could have been shown in that session and call it $\mathcal{A}_i$. We filter out the set of ads that are not in $\mathcal{A}_i$. This step sets the scope of our ability to generate counterfactual estimates. Figure 4.10 shows the empirical CDF of the number of ads participating in the auction for a session. As shown in this figure, the number of ads competing for each impression is quite variable across sessions.

Further, the extent of randomization in our problem allows us to make the unconfoundedness assumption: for any exposure number t in session i, the set of potential outcomes for all ads in $\mathcal{A}_i$ is independent of the actual ad that is shown in that exposure, conditional on the state variables. We can write:

$$\{Y_{i,t}(a)\}_{a \in \mathcal{A}_i} \perp\!\!\!\perp A_{i,t} \mid S_{i,t},$$

where $Y_{i,t}(a)$ is the potential outcome at the exposure t in session i when ad a is being shown. We can even weaken this assumption as only conditioning on demographic variables D_i yields

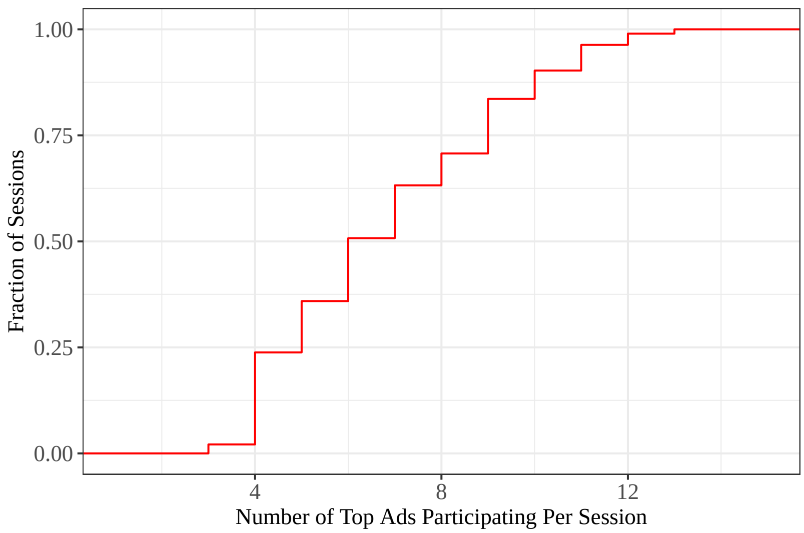

Figure 4.10: Empirical CDF of the number of ads competing per session.

the conditional independence presented above. This is because the ad allocation is random, controlling for the propensities determined by auction. Since only targeting variables can affect propensities in an auction, we only need to control for them. Thus, if we properly control for D_i in our predictive model, our outcome estimates preserve consistent treatment effects.

4.5.2 Feature Generation

As discussed earlier, our goal is to estimate click and leave outcomes for any combination of ad and state variables, as shown in Equations (5.21) and (5.22). A major challenge in estimating these equations is that the set of inputs is quite large, containing the entire sequence of prior ads shown to the user. In this section, we present a feature generation framework that maps a combination of state variables and ads ($\langle S_{i,t}, a \rangle$) to a set of meaningful features that we can give as inputs to our learning algorithm. Ideally, we need our final set of features to fully represent $\langle S_{i,t}, a \rangle$ in a lower dimension without any information loss. Thus, we generate a set of features that help us predict users' clicking behavior and app usage based on the prior

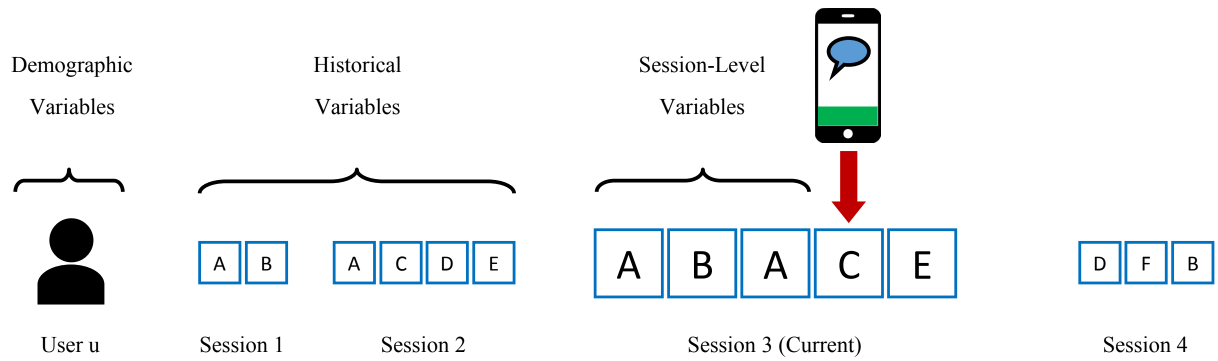

Figure 4.11: A visual schema for our feature generation and categorization.

literature on advertising.

We categorize these features into three groups: (1) demographic features, (2) historical features, and (3) session-level features. Demographic and historical features relate to the pre-session state variables (X_i), whereas session-level features relate to the session-level variables ($G_{i,t}$). Figure 4.11 provides an overview of our feature generation and categorization. In this example, the user is at her fourth exposure in her third session. The features for this particular exposure include the observable demographic features, historical features generated from the prior sessions, and session-level features that are generated from the first three exposures shown in the current session. Clearly, we do not use any information from the future to generate a feature: at any point, we only use the prior history up to that point. In the following sections, we describe all these features in detail.

Demographic Features

This includes the variables that we already observe in our data (see §4.3.2), such as the province, latitude, longitude, smartphone brand, mobile service provider (MSP), and connectivity type. For any session i, we use D_i to denote the set of demographic features. These features do not transition based on the ad that the publisher shows at any time period. As such, we do not use subscript t for them.[8] We include these features because of two reasons. First, these features help predict both users' clicking behavior and app usage. Second, the targeting variables are the main confounding source, and controlling them guarantees that we control for propensity score of ads when estimating the outcomes.

Historical Features

Historical features reflect the user's past activity prior to the current session. While demographic features are available in the data, we need to generate historical features based on the pre-session information. These features are not adaptive, i.e., they remain constant within the session. We follow the approach in [107] to generate these features. We refer the reader to that paper for details on why these features help predict user-level outcomes.

Let i, u, t, p, and a denote the session, user, exposure number, app, and ad respectively. Since we only focus on the top app, using the subscript p indicates that we only use impressions in that app to generate a feature. Otherwise, we calculate the feature using all impression. Below, we present the detailed set of our historical features along with their definition:

- $Imp_{i,u}$: The total number of impressions user u has seen prior to session i.
- $Click_{i,u}$: The total number of clicks user u has made prior to session i.
- $Imp_{i,u,p}$: The total number of impressions user u has seen in the top app prior to session i.
- $Click_{i,u,p}$: The total number of clicks user u has made in the top app prior to session i.
- $Imp_{i,u,t}$: The total number of impressions user u has seen at exposure number t prior to session i.
- $Click_{i,u,t}$: The total number of clicks user u has made at exposure number t prior to session i.

[8] One could argue that features such as latitude and longitude may change within the session. While this is possible, it is unlikely to happen as a result of the publisher's ad interventions. Further, the sessions are usually short, and we rarely observe such a change in our data.

- $Imp_{i,u,a}$: The total number of impressions of ad a that user u has seen prior to session i.
- $Click_{i,u,a}$: The total number of click that user u has made on ad a prior to session i.
- $Space_{i,u,a}$: The time space (in minutes) between session i and the last time ad a is shown to user u prior to session i. It takes value zero if there was no prior exposure of ad a.
- $LastSessionLength_{i,u}$: The length of last session (in number of exposures) that user u was exposed to prior to session i.
- $AvgSessionLength_{i,u}$: The average length of the sessions (in number of exposures) that user u was exposed to prior to session i.
- $LastFreeTime_{i,u}$: The free time (in minutes) user u has had between her last session and session i.
- $AvgFreeTime_{i,u}$: The average free time (in minutes) user u has had between her sessions prior to session i.
- $Breadth_{i,u}$: The total number of distinct ads that user u has seen prior to session i.
- $GiniSimpson_{i,u}$: The Gini-Simpson index for ads that user u has seen prior to session i [?]. This metric captures the diversity of prior ad exposures by calculating the probability that two random exposures from the past were of different ads. A higher Gini-Simpson index means that the user has seen a more diverse set of ads. We can write the Gini-Simpson index as follows:

$$GiniSimpson_{i,u} = 1 - \sum_{a \in \mathcal{A}} \frac{Imp_{i,u,a}(Imp_{i,u,a} - 1)}{Imp_{i,u}(Imp_{i,u} - 1)} \tag{4.12}$$

For a session i, we denote all these features as H_i. Like demographic features, publisher's actions will not change historical features within the session. However, for the next session that the user participates in, the historical features are updated. Further, it is worth noting that three historical features, $Imp_{i,u,a}$, $Click_{i,u,a}$, and $Space_{i,u,a}$, are ad-specific. It means that these features will change when different ads are selected.

Session-Level Features

The session-level features are key to our analysis as we are interested in optimal sequencing of ads within the session. Further, we allow these features to change by the exposure number. That is, depending on the prior exposures within the session, these features will evolve. Below is a list of session-level temporal features:

- $Imp_{i,t}$: The total number of impressions the user has seen in session i prior to exposure number t. For any exposure number t, this feature is $t-1$.
- $Click_{i,t}$: The total number of clicks the user has made in session i prior to exposure number t.
- $Imp_{i,a,t}$: The total number of impressions of ad a that user has seen in session i prior to exposure number t.
- $Click_{i,a,t}$: The total number of clicks that the user has made on ad a in session i prior to exposure number t.
- $Space_{i,a,t}$: The number of exposures between exposure number t and the last time ad a was shown in session i. It takes value 0 when there is no prior exposure of ad a in the session.
- $Breadth_{i,t}$: The total number of distinct ads that the user has seen within session i prior to exposure number t. We can define this feature as follows:

$$Breadth_{i,t} = \sum_{a \in \mathcal{A}} \mathbb{1}(Imp_{i,a,t} > 0) \tag{4.13}$$

- $Changes_{it}$: The total number of consecutive changes of ads prior to the exposure number t within the session i. We can write:

$$Change_{i,t} = \sum_{j=2}^{t-1} \mathbb{1}(A_{i,j} \neq A_{i,j-1}), \tag{4.14}$$

  where $A_{i,j}$ is the ad shown at exposure number j in session i.
- $GiniSimpson_{i,t}$: The Gini-Simpson index for the ads shown within session i prior to

exposure number t. Following the same logic in Equation (4.12), we can write:

$$GiniSimpson_{i,t} = 1 - \sum_{a} \frac{Imp_{i,a,t}(Imp_{i,a,t} - 1)}{(t-1)(t-2)} \tag{4.15}$$

For any session i and exposure number t, we denote all session-level features by $O_{i,t}$. As such, this is the only set of features that has subscript t, indicating that it changes within the session. Therefore, the publisher's actions affect the transition of these features in the session. One could argue that historical features also change within the session as user's history accumulates after each exposure. It is worth noting that we do not update the history within the session because session-level temporal features capture that information. As a result, not updating historical features will not result in any information loss.

4.5.3 Learning Algorithm

Here we describe the learning algorithm that we use to estimate the click and leave models. Since our goal is to estimate these outcomes as accurately as possible, we need to build a model that is able to capture complex relationships between covariates and outcomes. For both outcomes, we use Extreme Gradient Boosting (XGBoost henceforth) method developed by [22], which is a fast and scalable version of Boosted Regression Trees [41]. There are some key reasons why we use XGBoost as our main optimization method. First, it has been shown to outperform most existing methods in most prediction contests, especially those related to human decision-making like ours [22]. Second, [107] show that in the same context, XGBoost achieves the highest predictive accuracy compared to other methods.

There are important implementation details in making an XGBoost model that we summarize as follows. First, following the arguments in [107], we use logarithmic loss as our loss function. Second, we split our data into two parts – training and test sets. To tune parameters of XGBoost, we split the training set into two parts and use one as a hold-out validation set to prevent the model from over-fitting. Finally, to select the hyper-parameters accurately, we conduct a grid search over a large set of hyper-parameters and select those that give us the best performance on a hold-out validation set. Finally, we evaluate the performance of our model on both training and test set. Note that the test set is not used at any stage to select the model. Thus, the evaluation on the test set demonstrates an out-of-sample predictive accuracy.

4.5.4 Results

Results from the Click Estimation Model

We use three different evaluation metrics to evaluate the predictive performance of our model:

- *Relative Information Gain (RIG):* This is defined based on the log loss, which is our loss function in the XGBoost model. This metric reflects the percentage improvement in log loss compared to a baseline model that simply predicts the average CTR for all impressions.
- *R-Squared (R^2):* This is the most commonly used metric in marketing and economics, and intuitively calculates the percentage of variance in the outcome that our model can explain.
- *Area Under the Curve (AUC):* It determines how well we can identify *true positives* without identifying *false positives*. This score ranges from 0 to 1 and a higher score indicating better performance.

The results of these three metrics are shown in Table 4.2. We find that our model achieves over 28.47% and 27.68% RIG on training and test sets, respectively. This predictive accuracy is quite substantial compared to the literature [**?**, 107]. Further, the results on R^2 also indicate that our model reaches an excellent predictive performance. Given the inherent noise and variability in clicks, explaining over 18% of the variance in this outcome requires a very powerful model. Finally, the results on the last metric document a very good classification power.

Evaluation Metric	Training	Test
RIG	0.2847	0.2768
R^2	0.1899	0.1872
AUC	0.8560	0.8531

Table 4.2: Performance of the click estimation model on training and test sets

Evaluation Metric	Training	Test
RIG	0.0975	0.0949
R^2	0.0942	0.0918
AUC	0.7164	0.7135

Table 4.3: Performance of the leave estimation model on training and test sets

Results from the Leave Estimation Model

We now discuss the results from the leave model. The leave outcome is an important piece of the transition function, since the leave outcome by the user terminates the entire session. We use the same evaluation metrics to evaluate the predictive performance of our leave estimation model. Table 4.3 summarizes the results in terms of these metrics on both the training and test sets. We find that our leave estimation model explains over 9% of the variance in the leave outcome on both training and test sets.

Overall, it is worth noting that we should not expect the leave model to to have high predictive ability for a variety of reasons. First, clicking decision is mainly related to the ad shown, whereas leave decision is somewhat independent of ad exposures, especially in the messenger app where people's usage stems from users' messaging behavior that is unobserved to the researcher. Second, we focus on new users for whom we do not have a long panel. Hence, our features reflecting their past usage may not be very informative.

	Click Model			Leave Model		
Model	RIG	R^2	AUC	RIG	R^2	AUC
Full model	0.2768	0.1872	0.8531	0.0949	0.0918	0.7135
Without session-level	0.2245	0.1280	0.8334	0.0881	0.0889	0.7055

Table 4.4: Performance of models with/without session-level features

Preliminary Evidence on the Gains from the Adaptive Framework

An adaptive framework allows the publisher to use real-time session-level information to make decisions. Intuitively, the publisher can benefit from an adaptive framework for decision-making if session-level features play an important role in driving the main outcome of interest. In other words, if the user's decision to click on ads is only a function of his demographic or historical features and does not change based on session-level features, using an adaptive framework becomes ineffective. Thus, an intuitive test that can provide some preliminary evidence on the effectiveness of an adaptive framework is to drop session-level features from the main click and leave models and see whether it results in a drop in the predictive performance of these models.

The model without session-level features exclude the following features: $Imp_{i,t}$, $Click_{i,t}$, $Breadth_{i,t}$, $Changes_{i,t}$, $GiniSimpson_{i,t}$, $Imp_{i,a,t}$, $Click_{i,a,t}$, $Space_{i,a,t}$. We compare the performance of this model with the full model for both click and leave estimation models. The results on the performance of these models on the same test set are presented in Table 4.4. As shown in this table, the performance of the model drops when we exclude session-level features. This finding suggests and publishers can benefit from using adaptive frameworks for ad sequencing.

4.6 Ad Sequencing Policies

As discussed earlier, our main goal is to use our dynamic framework and develop a sequencing policy that maximizes user engagement in the session. We primarily measure user engagement by the number of clicks per session. To develop sequencing policies, we can use our estimates for the transition and reward functions to solve for the optimal policy in the MDP specified in §4.4.

In this section, we first present a series of other baseline policies that are commonly used in practice in §4.6.1. It allows us to compare the performance of the adaptive forward-looking policy developed by our framework with other benchmarks and establish the gains from adopting a dynamic framework. We then discuss the empirical evaluation of all these policies in §4.6.2. Finally, in §4.6.4, we present our results on different sequencing policies, explore heterogeneity in gains from the dynamic sequencing policy across sessions, and offer some explanation for the performance of different methods.

4.6.1 Definition of Sequencing Policies

We present a series of benchmark policies to examine the publisher's gains from adopting a dynamic framework for ad sequencing. We consider benchmark policies that drop modeling components in our dynamic framework and/or reflect the current norm in research and practice. As such, our comparison allows us to pin down how valuable each modeling component is and how much we can improve over the current practice. Starting with the adaptive forward-looking policy, which is based on our dynamic framework, we present the list of policies that we consider as follows:

- **Adaptive forward-looking sequencing policy:** This sequencing policy uses the dynamic framework in §4.4 with the expected probability of click being the reward function. Using our generic formulation of an MDP in Equation (4.2), we can write the publisher's optimization problem specific for a click-maximizing objective as follows:

$$\underset{a \in \mathcal{A}_i}{\operatorname{argmax}} \left[R_t(a; S_{i,t}) + \beta \mathbb{E}_{S_{i,t+1}|S_{i,t},a} V_{t+1}(S_{i,t+1}) \right], \tag{4.16}$$

where the value function for any state variable at exposure number t can be written as follows:

$$V_t(S_{i,t}) = \max_{a \in \mathcal{A}_i} \left[R_t(a; S_{i,t}) + \beta \mathbb{E}_{S_{i,t+1}|S_{i,t},a} V_{t+1}(S_{i,t+1}) \right] \tag{4.17}$$

This policy is both adaptive and forward-looking, i.e., it takes the session-level information (adaptive) and future information (forward-looking) into account. We also call this policy fully dynamic sequencing policy.

- **Adaptive myopic sequencing policy:** This sequencing policy does not take into account the expected future rewards when making the decision at any point. This is equivalent to the adaptive forward-looking sequencing with $\beta = 0$ that turns off the weight on the future rewards. Thus, we can write the objective function for adaptive myopic sequencing as follows:

$$\operatorname*{argmax}_{a \in \mathcal{A}_i} R_t(a; S_{i,t}) \tag{4.18}$$

  In this policy, the publisher selects the ad that maximizes CTR in the current period. It is worth noting that this policy is adaptive, as it uses the session-level information that is time-varying. However, it is myopic in the sense that it ignores future information. This case reflects the common practice of using contextual bandits in the industry.

- **Non-adaptive single-ad policy:** This policy only uses the pre-session information. Since it does not use adaptive information, this policy allocates all the impressions to a single ad that has the highest average CTR. This is similar to the practice of using a fixed ad slot where the whole session is allocated to one ad. The objective in this case is the same as Equation (4.18) only for $t = 1$.[9]

  This policy provides some insight into the ad sequencing problem because it has two distinct features. First, it captures the potential gains from using a short-lived ad slot as compared to the fixed ad slot. Second, it demonstrates the value of adaptive session-level information.

- **Random sequencing policy:** In this sequencing policy, the publisher randomly selects ads from the ad inventory. While this is a naive policy, it can serve as a benchmark

[9] One could argue that the optimal single-ad that is selected for the entire session may be different from the optimal ad for the first exposure. We acknowledge this issue and check the robustness of our results by using a dynamic optimization constrained by a single ad to be shown for the entire session. In the main text, however, we use the more straightforward approach of allocating the entire session to the ad with the highest CTR in the first exposure.

showing how well we can do without any model.

4.6.2 Empirical Evaluation

We now explain how we use our empirical estimates for both click and leave outcomes to develop and evaluate all the policies presented in §4.6.1. We start with the solution concept for the baseline policies in §4.6.2. These policies are more straightforward and easier to derive. We then present the solution concept for the adaptive forward-looking sequencing policy in §4.6.2. We finally discuss how we can evaluate the performance of these policies.

Solution Concept for Baseline Policies

We now describe how we obtain the baseline policies described in §4.6.1: (1) adaptive myopic sequencing policy, (2) non-adaptive single-ad policy, and (3) random sequencing policy. Since none of these policies incorporate expected future rewards, we need to mainly rely on our estimates for the reward function with the click-maximizing objective, i.e., $\hat{y}_{i,t}(a; S_{i,t})$. We present our empirical approach to solve for these policies below:

- **Adaptive myopic sequencing policy:** As presented in Equation (4.18), this policy selects the ad with the highest CTR at any point. Therefore, we can write $a_t^m(S_{i,t}) = \text{argmax}_{a \in \mathcal{A}_i} \hat{y}_{i,t}(a; S_{i,t})$, where $a_t^m(S_{i,t})$ indicates the ad to be shown under adaptive myopic sequencing in state $S_{i,t}$. We can characterize the best policy under adaptive myopic sequencing as follows:

$$\hat{\pi}^m(a \mid S_{i,t}) = \begin{cases} 1 & a = a_t^m(S_{i,t}) \\ 0 & a \neq a_t^m(S_{i,t}) \end{cases} \tag{4.19}$$

- **Non-adaptive single-ad policy:** This case is the solution for the adaptive myopic sequencing policy at $t = 1$. As such, we have $a_t^s(S_{i,t}) = a_t^m(S_{i,1})$ for any t in session i. Thus, the non-adaptive single-ad policy for each session i can be characterized as follows:

$$\hat{\pi}^s(a \mid S_{i,t}) = \begin{cases} 1 & a = a_t^s(S_{i,t}) \\ 0 & a \neq a_t^s(S_{i,t}) \end{cases} \tag{4.20}$$

- **Random sequencing policy:** This policy basically gives the same chance to all the ads in the inventory. Thus, we can write the random sequencing policy as follows:

$$\hat{\pi}^r(a \mid S_{i,t}) = \begin{cases} \frac{1}{|\mathcal{A}_i|} & a \in \mathcal{A}_i \\ 0 & a \notin \mathcal{A}_i \end{cases}, \tag{4.21}$$

  where $\mathcal{A}_i$ is the set of ads competing in session i.

Solution Concept for the Adaptive Forward-looking Sequencing Policy

Solving for the best policy in an MDP can become a daunting task when the state space is high dimensional. Given the set of features we use for our estimation tasks, we need to store the full history within the session. For example, to update session-level features such as $Click_{i,a,t}$ or $Breadth_{i,t}$, we need to know the entire sequence of actions and outcomes within the session up until exposure number t, which can be computationally burdensome for large t.

The session-level history that we need to store for the transition function contains the sequence of all ads and the click outcomes. That is, for exposure number t in session i, this history is defined as $G_{i,t} = \langle A_{i,1}, Y_{i,1}, ..., A_{i,t-1}, Y_{i,t-1} \rangle$. This set can take $(2|\mathcal{A}|)^{t-1}$ unique values for each t, indicating that it grows exponentially in t. Thus, one way to resolve this problem is to consider a finite horizon case with a reasonable T that is sufficiently large to let us exploit the dynamic framework, while reasonably small to help us avoid computational issues. We argue that that $T = 6$ satisfies both these conditions since the user is quite likely not to get to the seventh exposure at all given the results we show in Figure 4.5a: around 75% of the sessions have shown at most six exposures. Thus, optimizing the exposure after that point has only a marginal effect on the performance of the click-maximizing policy.

To empirically derive the adaptive forward-looking sequencing policy, we need two key estimands – $\hat{y}_{i,t}(a; S_{i,t})$ and $\hat{l}_{i,t}(a; S_{i,t})$. The former affects both the reward function and state transitions, whereas the latter only affects the state transitions. For notational convenience and brevity, let $\tilde{V}_t(a, S_{i,t})$ denote our estimate for the sum of both current period reward and expected future rewards given action a being taken in state $S_{i,t}$. Namely, it is the estimated objective in Equation (4.16). Using our empirical estimates of expected click and leave probability, we can write $\tilde{V}_t(a, S_{i,t})$ as follows:

$$\begin{aligned}\tilde{V}_t(a, S_{i,t}) =& \hat{y}_{i,t}(a; S_{i,t}) + \left(1 - \hat{l}_{i,t}(a; S_{i,t})\right) \hat{y}_{i,t}(a; S_{i,t}) V_{t+1}\left(\langle S_{i,t}, a, Y_{i,t} = 1\rangle\right) \\ &+ \left(1 - \hat{l}_{i,t}(a; S_{i,t})\right) \left(1 - \hat{y}_{i,t}(a; S_{i,t})\right) V_{t+1}\left(\langle S_{i,t}, a, Y_{i,t} = 0\rangle\right),\end{aligned} \tag{4.22}$$

where $\langle S_{i,t}, a, Y_{i,t} = 1\rangle$ and $\langle S_{i,t}, a, Y_{i,t} = 0\rangle$ denote the state variables in the next period. The equation above shows that we can easily break the expected future rewards into two deterministic pieces. We now propose a backward induction solution concept for this case as described below:

1. We start from the last period, T. We assume that this is the last exposure number within the session. Hence, the problem is static in that period. We can write:

$$\hat{V}_T(S_{i,T}) = \max_{a \in \mathcal{A}_i} \hat{y}_{i,T}(a; S_{i,T}) \tag{4.23}$$

$$\hat{a}^d_T(S_{i,T}) = \operatorname*{argmax}_{a \in \mathcal{A}_i} \hat{y}_{i,T}(a; S_{i,T}) \tag{4.24}$$

   That is the maximum value the publisher can extract from this particular state variable in the last time period.

2. For any $t < T$, we can write:

$$\hat{V}_t(S_{i,t}) = \max_{a \in \mathcal{A}_i} \tilde{V}_{i,t}(a, S_{i,T}) \tag{4.25}$$

$$\hat{a}^d_t(S_{i,t}) = \operatorname*{argmax}_{a \in \mathcal{A}_i} \tilde{V}_{i,t}(a, S_{i,T}) \tag{4.26}$$

   Now, if we go backward and solve for the value functions, everything on the right-hand side of Equation (4.25) is known because we have already solved for the value functions in the next period. Therefore, we can find the value function for all the states in time period and continue this process until exposure number 1.

Once we have estimated the value function and best ad for all the states, we can simply characterize the adaptive forward-looking sequencing policy as follows:

$$\pi^d(a \mid S_{i,t}) = \begin{cases} 1 & a = a_t^d(S_{i,t}) \\ 0 & a \neq a_t^d(S_{i,t}) \end{cases} \tag{4.27}$$

We use backward induction as our main approach to determine the adaptive forward-looking sequencing policy. The only limitation of this approach is that we have to set an endpoint for the session, which does not allow us to fully exploit exposures with $t > T$. As a robustness check, we also use an infinite horizon with more memory efficient state variables, where only we consider state variables that do not require storing the entire session-level history.

4.6.3 Evaluation

One of our main goals in this chapter is to examine to what extent the dynamic framework helps publishers achieve better user engagement with ads. As such, we need to use an evaluation metric that allows us to evaluate and compare different policies. For any exposure t, we denote a t-step trajectory by g_t and define it as the sequence of states, actions, and rewards in all the steps as follows:

$$g_t = \langle s_1, a_1, r_1^c, \ldots, s_{t-1}, a_{t-1}, r_{t-1}^c, s_t \rangle,$$

where any s_k is determined by the sequence prior to that exposure k and the distribution of transitions, a_k is determined given the policy, and r_k^c is the reward for the pair of s_k and a_k with the click-maximizing objective. Let τ denote the distribution of transitions and π denote any given policy. The joint distribution (τ, π) then determines the probability of each sequence g_t. Let $\rho_T(\pi; S_{i,1})$ denote the expected number of clicks generated in session i for the horizon length T, when using policy π. We can write:

$$\rho_T(\pi; S_{i,1}) = \mathbb{E}_{g_t \sim (\tau,\pi)} \left[\sum_{t=1}^{T} \beta^{t-1} r_t^c \right] \tag{4.28}$$

Now, we can use our estimates for the distribution of transitions τ and estimate $\rho_T(\pi; S_{i,1})$ for any policy π as follows:

$$\hat{\rho}_T(\pi; S_{i,1}) = \sum_{t=1}^{T} \sum_{g_t \in \mathcal{G}_T} \sum_{a \in \mathcal{A}_i} \pi(a \mid S_{i,t}) \hat{R}_t^c(a; S_{i,t}) P(g_t \mid \tau, \pi), \tag{4.29}$$

where $\mathcal{G}_T$ is the set of all possible trajectories and g_t denotes the trajectory in the first t periods. The last component in Equation (4.29) is the probability that a specific trajectory happens given the policy and distribution of transitions.

The main reason why we employ on this approach for evaluation is that it gives us session-level performance metrics on each policy. However, for robustness, we employ other approaches such as importance sampling and doubly robust method.

4.6.4 Results from the Click-Maximizing Policy

We now present the results on the performance of different sequencing policies when the main objective is to maximize the expected number of clicks per session. We first illustrate session-level outcomes for different sequencing policies and examine the gains from adopting an adaptive forward-looking policy. Next, we explore the heterogeneity in the gains from adopting an adaptive forward-looking sequencing policy across sessions. Finally, we use the distribution of counterfactual sequencing policies to explain the differences in them.

Gains from the Adaptive Forward-looking Sequencing Policy

We start by demonstrating the gains from the adaptive forward-looking sequencing policy compared to other policies described in §4.6.1. We use the direct method presented in §4.6.3 to evaluate the performance of these sequencing policies. As such, we only focus on the first six exposures and do not evaluate sessions after the sixth exposure. We draw a random sample of 1000 users from our test data that gives us 12,136 unique sessions. We estimate the optimal policy under all the sequencing policies defined above and present the results in Table 4.5.

As expected, all model-based sequencing policies lead to major improvements in the expected number of clicks per session over random allocation of ads. Adopting an adaptive forward-looking sequencing policy results in 0.1772 clicks per session which translates to an 80.36% improvement over the random sequencing, 7.87% over the non-adaptive single-ad sequencing, and 1.50% over the adaptive myopic sequencing. These gains demonstrate the value of each piece of information and modeling paradigm. Our results indicate that both session-level and future information play an important role in ad allocation problem.

	Sequencing Policies			
	Fully Dynamic	*Adaptive Myopic*	*Single-Ad*	*Random*
Expected No. of Clicks	2150.60	2118.81	1993.75	1192.40
Expected No. of Clicks Per Session	0.1772	0.1745	0.1642	0.0983
Expected No. of Impressions	39365.30	39460.26.37	38925.97	38859.06
Expected Session Length	3.24	3.25	3.21	3.20
Expected CTR	5.46%	5.37%	5.12%	3.07%
% Click Increase over Random	80.36%	77.69%	67.21%	0.00%
No. of Users	1000	1000	1000	1000
No. of Sessions	12,136	12,136	12,136	12,136

Table 4.5: Performance of different sequencing policies for a sequence size of 6 with a click-maximizing objective

Next, we focus on the gap between the adaptive forward-looking and non-adaptive single-ad policy. This comparison relates to a broader question on whether the publisher should use a short-lived ad slot. Our results suggest that the use of dynamic ad slot leads to considerable gains and value creation compared to the fixed ad slot, justifying the current trend of using short-lived ad slots in the industry.[10]

We then examine the performance of adaptive forward-looking sequencing compared to that of adaptive myopic sequencing policy. Our results suggest that there are gains from adopting a forward-looking objective in adaptive ad sequencing. It also provides evidence regarding the interdependence of ads shown within the sequence. It is worth noting that 1.50% improvement is a lower bound for what the publisher can achieve by adopting a forward-looking objective. This is because we only focus on the first six exposures for computational reasons. The sequencing effects likely grow in later exposures. Further, we only focus on a limited ad inventory, with 15 ads. We expect the gains from the adaptive forward-looking ad sequencing to improve as a result of expanding ad inventory.

One likely explanation for the gains from an adaptive forward-looking sequencing policy is that it is the only model that takes users' leave decision into account. As such, one channel through which this policy might improve the outcome is to increase the session length, thereby enhancing the total number of clicks. To see whether this is the channel, we estimate the expected session length for all the policies. We find no significant difference between our policies in terms of the session length. Rather, the adaptive forward-looking sequencing policy generates a higher CTR, which suggests that users are more engaged with ads as a result of adaptive forward-looking ad sequencing.

Heterogeneity in Gains from the Dynamic Framework

While the results in Table 4.5 establish the average gains in clicks from adopting the adaptive forward-looking sequencing policy, they do not explain how these gains vary across sessions. In this section, we are interested in identifying the sessions for which the adaptive forward-looking sequencing is most helpful. As such, we focus on the session-level gains from the adaptive forward-looking compared to the adaptive myopic and non-adaptive single-

[10]However, the caveat here is that our approach in single-ad policy is not the equivalent of the fixed ad slot, as there are very short interruptions in every one minute this ad is being shown, whereas the fixed ad slot keeps the ad immobile throughout the session.

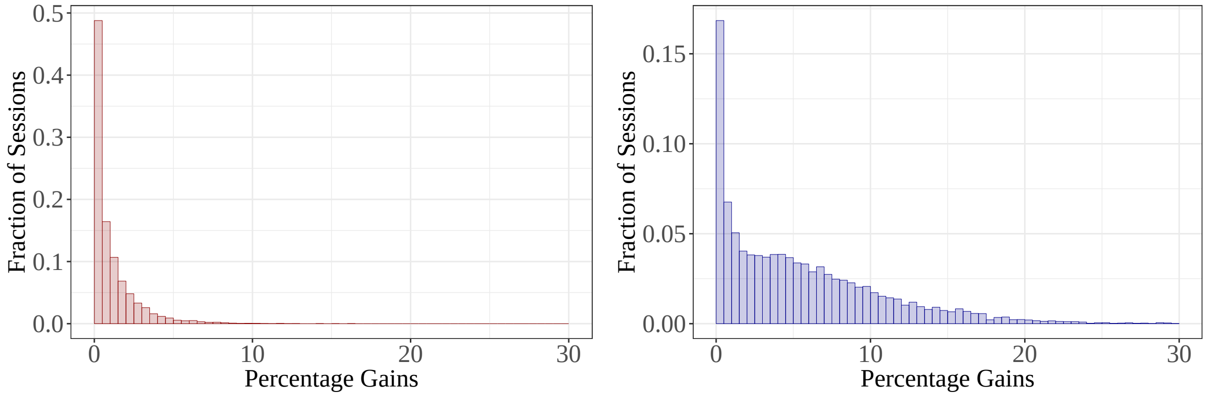

(a) Forward-Looking over Myopic

(b) Forward-Looking over Single-Ad

Figure 4.12: Distribution of session-level gains from the adaptive forward-looking over both adaptive myopic and single-ad sequencing policies

ad policies and explore the heterogeneity across sessions. Let $Gain^f_m$ and $Gain^f_s$ denote the percentage gains from the adaptive forward-looking over adaptive myopic and single-ad sequencing respectively. In Figure 4.12, we show the distribution of these two session-level gains. Both figures show significant heterogeneity in the session-level gains from the adaptive forward-looking sequencing policy over adaptive myopic and single-ad policies.

To further explore this heterogeneity across user history, we regress both $Gain^f_m$ and $Gain^f_s$ on a set of historical variables, controlling for user fixed effects. We also control for the number of ads competing in a session as it is an important factor determining the gains from different sequencing policies. The results are presented in Table 4.6. The first two columns show the results using the gains from the adaptive forward-looking over the adaptive myopic sequencing policy as the dependent variable. In column 1, we examine how the relative gains from the adaptive forward-looking over adaptive myopic sequencing changes as a user participates in more sessions. On the one hand, as the user participates in more sessions, the publisher has more history and data on him, and it may improve the gains from adopting a forward-looking objective. On the other hand, as the user becomes more experienced in the platform, the relative effectiveness of sequencing and temporal interventions may shrink [106]. As shown in column 1, we find some evidence for the latter: the relative gains from the adaptive forward-looking over adaptive myopic sequencing reduces as the user becomes more experienced.

	Dependent variable			
	$Gain_m^d$	$Gain_m^d$	$Gain_s^d$	$Gain_s^d$
$Session\ Number_{i,u}$	-0.000032***	0.000028**	-0.000447***	-0.000276***
	(-4.51)	(-1.95)	(-17.29)	(-5.41)
$No.\ of\ Competitors_i$	0.000032	-0.000253***	0.004188***	0.003070***
	(0.27)	(-2.17)	(9.90)	(7.31)
$Imp_{i,u}$		0.000019***		0.000056***
		(11.79)		(9.66)
$Breadth_{i,u}$		-0.000401***		-0.001681***
		(-15.61)		(-18.13)
User FE	✓	✓	✓	✓
No. of Obs.	12,136	12,136	12,136	12,136
R^2	0.3021	0.3223	0.2864	0.3097
Adjusted R^2	0.2469	0.2686	0.2300	0.2550
Note:				*p<0.05; **p<0.01; ***p<0.001

Table 4.6: OLS estimates of session-level gains from the adaptive forward-looking sequencing policy across user history

In column 2, we include two historical variables to further explore which parts of user's prior experience are accountable for the findings in column 1 – total number of impressions ($Imp_{i,u}$), and the total number of distinct ads ($Breadth_{i,u}$) the user has seen prior to the session. With these controls, we find that the number of prior impressions has a positive correlation with the outcome, implying that having more data on a user benefits the adaptive forward-looking sequencing more than the myopic one. However, note that the number of distinct ads the user has seen that makes adaptive forward-looking sequencing less effective. This is probably because the sequencing effects mostly come from users' information processing and differentiation of ads in a sequence. Once the user has processed many different ads in prior sessions, it is less likely that the optimal sequencing of them largely affects her decision.

In columns 3 and 4, we show the results with another dependent variable – gains from the adaptive forward-looking over non-adaptive single-ad sequencing. Again, we find that users' tenure in the platform makes adaptive forward-looking sequencing less effective, and this is because users with more experience have seen more diverse ads, and it is harder to influence their decision by dynamic sequencing of ads. It is worth noting that the results in all columns must be interpreted in the relative terms. That is, a negative coefficient does not mean that adaptive myopic or non-adaptive single-ad policies perform better than the adaptive forward-looking sequencing policy. Rather, it identifies where adaptive forward-looking sequencing performs relatively worse.

Distribution of Micro-Interventions

In the previous sections, we established the gains from the adaptive forward-looking sequencing policy over other sequencing policies and showed how these gains vary across different sessions based on user-level history. These gains mainly come from the active selection of session-level features through the selection of ads in the adaptive forward-looking sequencing policy. In this section, we first illustrate how these sequencing policies are different in micro-interventions at any point and how that results in different session-level features. We then aim to uncover the reasons behind such differences, given the goal of each sequencing policy.

To make sense of how sequences evolve under each policy, we focus on binary micro-interventions as defined in §4.3.3. There are two key points about these micro-interventions. First, these binary micro-interventions together constitute our session-level features. For example, breadth-increasing change at any point results in a higher breadth of variety in a session, whereas repeat increases the number of prior exposures of an ad. Second, as shown in the prior research, all these three micro-interventions can be effective policies from a broader perspective. Recent experimental papers document the positive effects of multiple exposures on sales of an ad [118]. While confirming the positive effects of multiple exposures of an ad within a session, [106] show that both breadth-constant and breadth-increasing changes will increase the likelihood of clicking on the next ad. The fact that these micro-interventions are mutually exclusive makes the decision very challenging from a dynamic point-of-view, as using each comes with trade-offs. Thus, comparing distributions of these micro-interventions under different sequencing policies can help in understanding the decision-making process for each policy.

We present the average values for these micro-interventions at different time periods under each sequencing policy in Figure 4.13. Since these are binary variables, the y-axis reports the percentage of making any micro-intervention. As shown in these figures, both dynamic and adaptive myopic sequencing policies employ a mix of these micro-interventions. As such, the lines for them lie somewhere between the non-adaptive single-ad and random sequencing policies.

We further explore the difference between the adaptive forward-looking and adaptive myopic sequencing to reflect how these micro-interventions change if the publisher takes expected future rewards into account as in the adaptive forward-looking sequencing. We find that the adaptive forward-looking sequencing policy tends to repeat ads in the earlier time periods, whereas change is a more popular decision in the adaptive myopic sequencing across all time periods. This is likely because repeating the same ad is a sub-optimal decision at the moment, but will help transition to a state where we can reinforce an ad with prior exposures. This is why the adaptive forward-looking sequencing policy takes advantage of repeating in the beginning, while in the last period, it tends to act similarly to the adaptive myopic sequencing.

Next, we look into the difference in the distribution of breadth-increasing and breadth-constant changes at different time periods. Since more prior exposures and changing the last ad helps at the moment, the adaptive myopic sequencing policy employs breadth-constant changes even more than the random sequencing. The adaptive forward-looking sequencing policy, however, seems to have a mix of both micro-interventions.

Overall, Figure 4.13 illustrates the differences in sequencing policies at a micro level. Our findings demonstrate the trade-offs between these micro-interventions by comparing different sequencing policies. It is generally hard to offer a unifying explanation for the optimal policy in a Markov Decision Process, as it uses an elaborate objective function. Thus, it is important to notice that these findings are descriptive and must be interpreted cautiously.

4.7 Implications

Our findings in this chapter have several implications for marketing practitioners and policy makers. The most direct set of implications is for publishers and ad-networks. We examine the opportunity cost of using a fixed ad slot, which is the current practice in many digital ad platforms. We document considerable loss as a result of running a fixed ad slot, in terms of the value created by making a better match between ads and users. This finding has implications for publishers and ad-networks that want to design their ad format. However, these results must be interpreted with caution.

More importantly, we establish the gains from the adaptive forward-looking sequencing of ads compared to adaptive myopic sequencing. While it significantly adds to the computational complexity of the problem, our findings indicate that inter-temporal trade-offs in ad allocation problem play an important role in optimal policy design. As such, publishers can create value by dynamically sequencing ads.

While advertisers are not the main target of the implications in this chapter, our findings offer them some new insights. Given the gains from the adaptive forward-looking sequencing of ads, it can be beneficial for advertisers to incorporate session-level information while targeting their ads. Further, this chapter adds to the understanding of short-lived ad formats, which, in turn, can provide some insights for advertisers with regards to their banner design.

More broadly, our general framework can be extended to any context with adaptive interventions. For example, in the context of mobile health, a growing body of work focuses Just-In-Time Adaptive Interventions (JITAI) in mobile apps and studies the impact of them in shaping consumers' health behavior including physical fitness and activity, smoking, alcohol use, and mental illness [98]. Similarly, in any context that adaptive interventions can be used to educate people, we can specify a dynamic framework that helps us achieve better outcomes [90]. These showcases can serve as motivation for the public sector to use these tools in cases where collective action is required, such as environmental protection and political participation.

4.8 Conclusion

Mobile in-app advertising has grown exponentially over the last years. The ability to exploit the time-varying information about a user to personalize ad interventions over time is a key factor in the growth of in-app advertising. Despite the dynamic nature of the information, publishers often use myopic decision-making frameworks to select ads. In this chapter, we examine whether a dynamic decision-making framework benefits the publisher in terms of the user engagement with ads, as measured by the number of clicks generated per session. Our dynamic framework has two main components: (1) a theoretical framework that specifies the domain structure such that it captures inter-temporal trade-offs in the decision to show what ad at any time period and (2) an empirical framework that breaks the policy design problem into a combination of machine learning tasks. We apply our framework to large-scale data from the leading in-app ad-network of an Asian country. Our results indicate that the adaptive forward-looking sequencing of ads results in significant gains in the expected number of clicks per session, compared to a set of benchmark policies. Next, we document heterogeneity in gains across sessions and show that adopting an adaptive forward-looking policy is most effective when users are new to the platform. Finally, we illustrate the differences in sequencing policies using a descriptive approach.

This chapter makes three contributions to the literature. First, from a methodological point-of-view, we develop a unified dynamic framework that starts with a theoretical framework that specifies the domain structure in mobile in-app advertising, and an empirical framework that breaks the problem into tasks that can be solved using a combination of machine learning methods and causal inference tools. Second, from a substantive standpoint, we are the first to document the gains from adopting an adaptive forward-looking sequencing policy as compared to the adaptive myopic sequencing policy. This comparison is of particular importance as the adaptive sequencing is the current approach in the industry. Third, we establish the gains from using a short-lived ad slot as compared to a fixed ad slot. The answer to this question informs the publisher's decision to use which type of ad slot.

Nevertheless, there are some limitations in our study that serve as excellent avenues for future research. First, our counterfactual policy evaluation is predicated on the assumption that users do not change their behavior in response to sequencing policies. While we exploit randomization to obtain our counterfactual estimate, it would be important to validate these findings in a field experiment. Further, we use the training data offline to learn counterfactual estimates for click and leave outcomes. Extension of our framework to an online setting that captures exploration/exploitation trade-off is important as offline evaluation may be costly. Finally, we use the entire within-session history to update state variables. Future research can look into more parsimonious frameworks that can be scalable to longer time horizons.

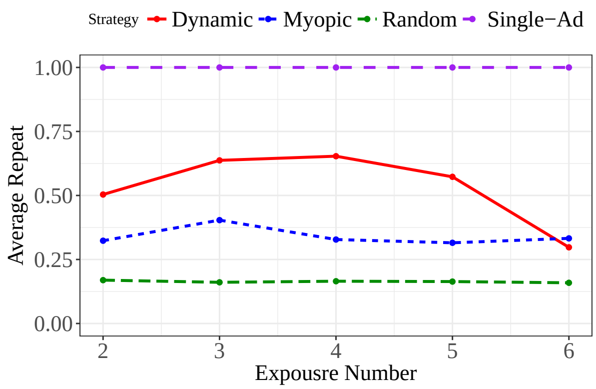

(a) Repeat

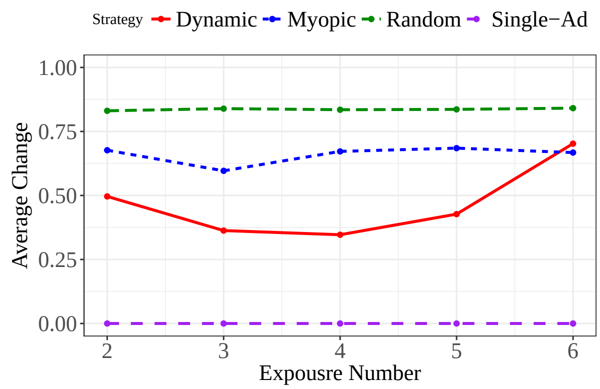

(b) Change

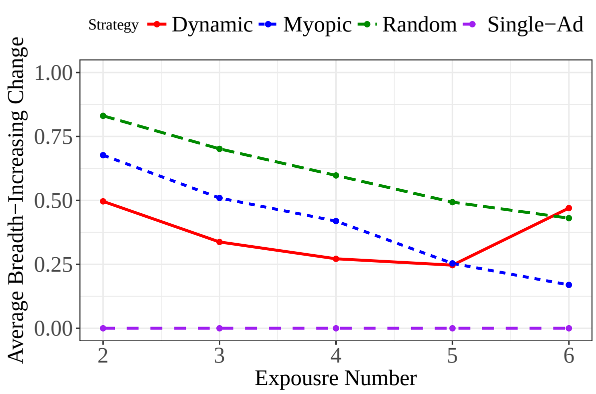

(c) Breadth-increasing change

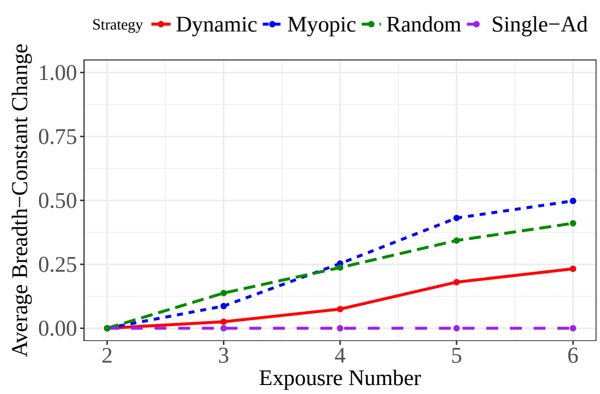

(d) Breadth-constant change

Figure 4.13: Distribution of micro-interventions at different time periods under different sequencing policies

Chapter 5

ADAPTIVE AD SEQUENCING: VALUE EXTRACTION

5.1 Introduction

Mobile in-app advertising is now a significant source of revenue for publishers and ad networks. In 2018, over 56% of the total digital ad spend came from in-app advertising [33]. Like other digital advertising environments, mobile publishers use an auction to determine which ad to show inside an app. An auction is a set of rules that characterizes how to allocate each advertising space and how much each advertiser has to pay, given advertisers' bids. Market outcomes under each auction format can accordingly be different since advertisers can strategically vary their bidding behavior. Thus, auction design plays a central role in the success of the digital advertising ecosystem.

Publishers and ad-networks often use auctions that maximize their revenues.[1] The common practice in this industry is to use a first- or second-price auction with an optimally set reserve price. This is in light of the findings from the seminal paper by [96] that has shown these auctions are revenue-optimal for a single item, under the regularity assumption. In an advertising environment, these auctions are revenue-optimal if the publisher can treat each advertising space as a single item: i.e., if advertising spaces are independent, and the auction outcome for one advertising space does not create externalities affecting other advertising spaces.

In the context of mobile in-app advertising, [104] provides empirical evidence on the interdependence of exposures within a session where a user is exposed to multiple short-lived ads. He shows that each ad exposure creates externalities that affect future exposures and

[1] In this chapter, we use the publisher, ad-network, and platform interchangeably, when we refer to the agent who designs the auction and makes the ad allocation decision.

documents the publisher's gain from dynamic sequencing of ads, i.e., the policy that captures both the immediate and future outcomes in a session and selects the ad that maximizes the expected number of clicks from that point onward. These findings rule out the independence of advertising spaces within a session, which in turn, imply that first- or second-price auction with an optimally set reserve price is not revenue-optimal in the context of mobile in-app advertising.

While the results in [104] elucidate an opportunity to create value in this market by dynamic sequencing of ads by enhancing consumer engagement and match values, the extent to which the publisher can extract this value as revenue is not clear. Notice that advertisers are strategic agents who can change their bids in response to any change in the allocation mechanism, and thereby appropriate most of this created value. The prior empirical literature on advertising auctions has highlighted cases where publishers cannot necessarily link the improvement in the match to higher revenues in a competitive environment [7, 107]. Thus, when the revenue is the primary outcome of interest, it is crucial for the publisher to incorporate advertisers' bidding behavior as well as users' ad response when designing the ad sequencing policy.

This brings us to the question of optimal auction design, wherein the publisher designs an auction that maximizes her revenues. Our main goal is to theoretically develop the revenue-optimal dynamic auction and compare its outcomes with the revenue-optimal static auction (second-price auction with optimal reserve price). Overall, we aim to answer the following three research questions in this chapter:

1. How can we design a revenue-optimal dynamic auction that captures both intertemporal trade-offs in ad sequencing and advertisers' strategic bidding behavior?

2. How can we build an empirical framework to evaluate the market outcomes such as publisher's revenues and advertisers' surplus under any auction mechanism?

3. What are the gains from using a revenue-optimal dynamic auction as compared to the

static one in mobile in-app advertising? How is the advertisers' surplus distributed across advertisers? Do all advertisers benefit when the publisher uses a dynamic auction?

We need to overcome three major challenges to answer these questions. First, to design a revenue-optimal dynamic auction, we need to specify an allocation rule that incorporates inter-temporal trade-offs in ad interventions and a payment rule that governs advertisers' strategic bidding behavior. Second, to empirically evaluate market outcomes (e.g., publisher's revenues) under counterfactual auctions, we need to obtain accurate estimates of both advertisers' and users' behavior that are valid any counterfactual auction. For the former, we need to estimate the distribution of advertisers' click valuations as it is the main structural parameter that governs their bidding behavior in any auction. For the latter, we need to estimate users' behavior: their likelihood of clicking on an ad and leaving the session after seeing an ad under any counterfactual allocation policy. Finally, to measure the gains from both dynamic and static revenue-optimal auctions, we need to first solve for the equilibrium outcome under these counterfactual auctions and then use an evaluation method that estimates the corresponding outcomes for each session.

We present an overview of our approach in Figure 5.1. This figure illustrates the general framework in the top row and the details specific to our problem in the bottom row. In the general framework, we begin with a theoretical framework that informs our empirical approach regarding how to develop and evaluate optimal auctions. More specifically, we start with designing a revenue-optimal dynamic auction that gives us a combination of allocation and payment rules. This combination captures both the inter-temporal trade-offs and advertisers' bidding strategies. Since our theoretical framework involves both advertisers' and users' behavior, our empirical framework requires an estimation procedure that cover both these components. We accordingly break our empirical framework into two separate tasks: (1) estimation of the distribution of advertisers' click valuations since click valuation is the key structural parameter governing advertisers' bidding behavior in any auction, and

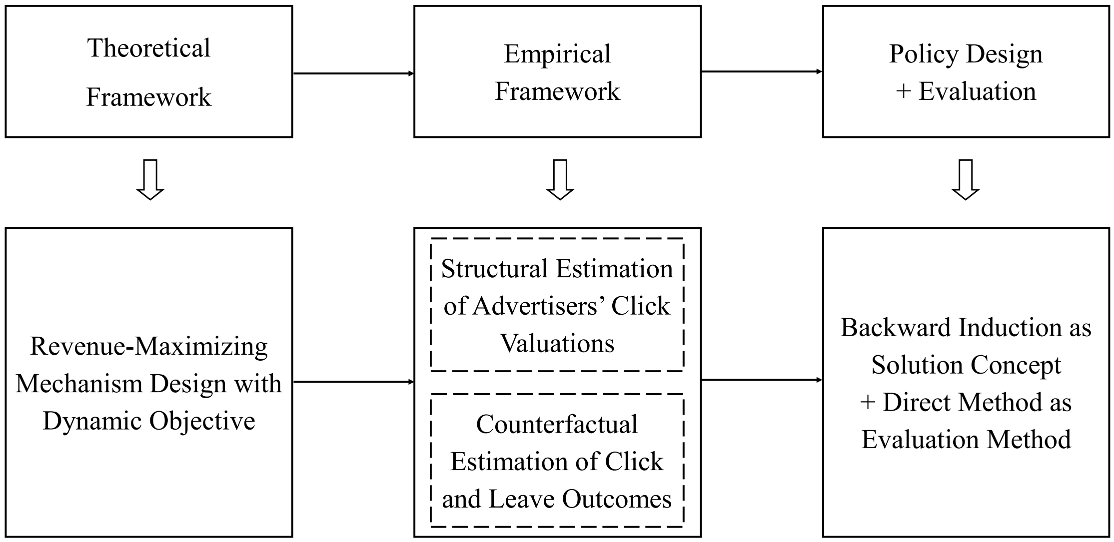

Figure 5.1: An overview of our approach. The top row presents our general framework and the bottom row shows the specific approach we take in this chapter.

(2) personalized counterfactual estimation of click and leave outcomes, which are the two user-dependent outcomes that affect the expected revenue per session. The second task is the same as the empirical task in [104]. We use the same approach in this chapter to estimate users' behavior under counterfactual auctions. Finally, we use all our estimates and numerically derive the optimal policy using backward induction and evaluate this policy using the direct method.

Our theoretical framework directly addresses the first challenge and paves the way to address other challenges. We build our theoretical model on the recent literature on dynamic mechanism design that extends the approach in [96] to a dynamic setting. The intuitive idea in this literature is to use the Revelation Principle [96] and exclusively focus on the case where all bidders report their type truthfully [73, 102]. We use the modeling framework in [73] as their separability assumption is particularly suitable in our context: the value an advertiser extracts from an impression is the product of his private click valuation and the expected probability of click on his ad. It allows us to write down the reward function in the Markov

Decision Process (MDP henceforth) in terms of virtual valuations and click probabilities (match valuations), thereby deriving the optimal allocation. Using this allocation function, we can then set the payments such that advertisers will participate in the auction and have no incentive to deviate from truthful reporting (IR and IC constraints). We then show that the auction with these allocation and payment rules is revenue-optimal.

To address our second challenge, we propose a structural framework to estimate the distribution of advertisers' click valuation from their observed bids in the data. The key challenge is that the auction format in the data is a quasi-proportional auction, where truthful bidding is not the equilibrium strategy for advertisers. We first characterize the advertisers' utility function in this setting and then derive the equilibrium properties of this auction. Using the first-order condition, we then write the advertisers' click valuation in terms of their cost and the allocation function used by the ad-network. Since both cost and allocation functions can be estimated from the distribution of observed bids and auction configurations, click valuations are identified under the assumption that advertisers are utility-maximizing. This allows us to estimate the distribution of advertisers' click valuations as well as each advertiser's click valuation.

Next, to develop the ad sequencing policy in the revenue-optimal dynamic auction, we need to solve the MDP for the allocation function. While the transition function is the same as the one in [104], the rewards have an additional multiplicative factor – each advertiser's virtual valuation. We can estimate it using the estimated click valuation for each ad as well as the distribution of click valuations. We plug these estimates into the reward function and solve the dynamic allocation policy using a backward induction solution concept. Finally, like [104], we use a direct method approach for evaluation that directly uses our estimates to simulate a session, equilibrium outcomes, and how it evolves. This method allows us to evaluate the revenue outcome for each session.

We first present the results from our auction estimation framework. We theoretically show that advertisers bid roughly half of their click valuations in the quasi-proportional auction. As such, the distribution of bids alone can approximate the distribution of click valuations. Further, it suggests that the current mechanism (quasi-proportional auction) leads to a substantial loss for the platform in terms of both revenue and efficiency. We then focus on the estimated distribution of click valuations from our structural framework, and empirically show that the regularity assumption is satisfied in our context: the virtual valuations are strictly increasing in click valuations. This is an important requirement for our counterfactual analysis, as the solution to the optimal auction is tractable given this assumption.

Next, we conduct our counterfactual analysis to examine the gains from the revenue-optimal dynamic auction. We set the benchmark as the second-price auction with an optimal reserve price, as it is the revenue-optimal static auction. Our results indicate that the expected revenue per session is 1.60% higher under the revenue-optimal dynamic auction compared to that in the revenue-optimal static auction. This is particularly important because most platforms currently use a version of the static revenue-optimal auction. Thus, our results suggest that publishers and ad-networks can significantly benefit from adopting an optimal dynamic auction. Further, we find that the expected number of clicks per session also improve by 1.80% under the dynamic case, suggesting that the gains in revenues can mostly be attributed to the improved match between users and ads as a result of dynamic sequencing, and not to the greater ability of the publisher to extract rent from advertisers.

We then focus on other market outcomes and show that the optimal dynamic auction achieves better outcomes than the optimal static auction in terms of both total surplus and average advertisers' surplus: the total surplus (efficiency) and the average advertisers' surplus increases by 1.77% and 3.00% under the optimal dynamic auction respectively. Hence, we show that the optimal dynamic auction does not achieve revenue optimality at the expense of efficiency. We then explore the surplus gains across advertisers to see whether the market will become more concentrated as a result of using the revenue-optimal dynamic auction. Using a Herfindahl-Hirschman Index (HHI), we find that the optimal dynamic auction has a lower concentration index than the optimal static auction. This suggests that the reason is that the optimal dynamic auction allocates more to the ad with the second largest surplus, thereby closing the gap between the top two advertisers in an auction.

In sum, this chapter makes three key contributions to the literature. First, from a methodological point-of-view, we propose a unified dynamic framework that captures both advertisers' and users' behavior to optimize publisher's revenue. A key contribution of our framework is in illustrating how we can use a theoretical framework to break a complex applied problem into a composite of structural estimation and machine learning tasks. To our knowledge, this is the first work to empirically examine the revenue gains from dynamic sequencing of ads using an optimal dynamic auction. Second, we present a structural estimation framework to recover the distribution of bidders' private valuations from their observed bidding behavior in a quasi-proportional auction. This is the first work to propose an estimation procedure for quasi-proportional auctions. Our framework can easily be extended to auctions with non-deterministic allocation rules. Third, from a substantive viewpoint, we establish the revenue gains from adopting a dynamic objective in allocating ads, as opposed to a static objective. This is of particular importance, as the current practice in the industry is to use a static objective. We expect our findings to be of relevance to publishers and ad-networks.

5.2 Related Literature

First, this chapter relates to the growing literature on dynamic mechanism design. While early papers in this literature start in 1980s [11, 97, 112], most of the major developments in this literature appear more recently, with generic characterizations of both efficient mechanisms [14, 8] and revenue-maximizing (optimal) mechanisms [73, 102]. The majority of applied papers on dynamic mechanism design focus on cases where the inter-temporal trade-offs arise through the dynamics of arrival, departure, or population [133, 101, 43, 120]. This chapter adds to this literature by empirically evaluating dynamic mechanisms in a digital advertising context. To our knowledge, this is the first work to provide an empirical framework to examine the performance of dynamic mechanisms.

Second, this chapter relates to the literature on the intersection of mechanism design and online advertising. Early papers in this area examines the theoretical properties of different auctions in sponsored search context [30, 131, 80]. More specific to our context, a series of work takes externalities in search advertising into account and revisits the question of mechanism design in a context where the higher position of an ad may affect the user's decision to even see lower ranked ads [44, 75, 45]. In the context of video ads, [74] adopt a cascade model similar to [75], and provide a mechanism for selection and ordering of video ads. While this stream of work proposes simple mechanisms for allocation, they are only applicable to very basic and unrealistic case where the externality is only imposed through the user's leaving decision. We extend this literature by offering a dynamic framework that captures more complex externalities, under a plausible separability assumption.

Lastly, this chapter relates to economics and marketing literature on the estimation of auctions. A significant breakthrough in this literature comes from [53] who base their identification strategy on the fact that the equilibrium outcome is achieved when all agents maximize their profits given the distribution of others' behavior. While they study the first-price auction with symmetric independent private valuations and without unobserved heterogeneity, other papers in this literature build on this work and extend it to the cases with affiliated private valuations [87], asymmetric private valuations [19], unobserved heterogeneity [54, 78], and also different auctions such as scoring auctions [10] and beauty contest auctions [136]. Related to our setting, a few papers study online advertising auctions and propose different empirical approaches for the estimation fi advertising auctions [7, 135, 23]. This chapter adds to this literature by proposing an estimation approach for quasi-proportional auctions that can easily be extended to any randomization-based auction. Further, our counterfactual analysis is the first to consider a dynamic mechanism design in an online advertising context.

5.3 Optimal Auctions

The main results in [104] establish the gains from the dynamic sequencing of ads in terms of the expected number of clicks generated per session. Intuitively, we expect the increase in the number of clicks to be linked to higher publisher revenues, as clicks are the key revenue-generating source for publishers. However, the fact that advertisers can respond to the change in the allocation by changing their bids in an auction environment makes it unclear how the publisher can extract more revenues from dynamic sequencing of ads. For example, if advertisers decrease their bid as a response to better allocation by dynamic sequencing in the equilibrium, the publisher may end up selling more clicks at a lower price. Thus, it is crucial to take advertisers' bidding behavior into account if the publisher wants to maximize her revenues through adaptive ad sequencing.

To examine revenue gains from adopting a dynamic framework, we incorporate advertisers' utility model into our framework and design revenue-optimal auctions with both dynamic and static objectives. Our framework, in turn, captures both users' behavior and advertisers' strategic bidding. The publisher's problem is then to design an auction with certain allocation and payment rules that maximizes her revenues, i.e., the total payments made by advertisers. To find the revenue-optimal auctions under each objective, we build on the seminal paper by [96] and design the allocation and payment rules in the auction to maximize publisher revenues.

This section proceeds as follows: We first define our model environment and assumptions in §5.3.1. We then introduce the case where the publisher's objective is static and discuss the optimal auction in this case in §5.3.2. In §5.3.3, we focus on the optimal auction with a dynamic objective and present the allocation and payments rules in this case.

5.3.1 Auction Environment and Assumptions

We now describe the auction environment in this problem. For each session i, there are $\mathcal{A}_i$ risk-neutral bidders competing for impressions in this session. Each advertiser a has a private click valuation x_a which is drawn from the distribution F_a with support $[\underline{x}_a, \bar{x}_a]$. This parameter is a private signal that reflects how much the advertiser values a click. We assume that each ad's private valuation is independent of other ads' private valuation and does not vary across impressions.[2] The total value generated from showing ad a at exposure t of

[2]We can extend our theoretical analysis to the cases where click valuations vary across sessions and exposures. However, in our empirical analysis, we can only identify one value for each ad. This is why we

session i will then be the product of that ad's click valuation and the probability of click on that ad, which can be written as follows:

$$w_{i,a,t}(x_a; S_{i,t}) = x_a P(Y_{i,t} \mid a, S_{i,t}), \tag{5.1}$$

where $w_{i,a,t}(x_a; S_{i,t})$ is the total value ad a receives from being shown at exposure t of session i, and $S_{i,t}$ is the state variable at that point which captures the prior history of actions and outcomes within the session, as well as the pre-session information (Please see §4.2.1 in [104] for more details on state variables).

The publisher's problem is to design a mechanism/auction that maximizes her revenues. Each mechanism is characterized by two components – (1) allocation rule, and (2) payment rule. We can define any mechanism as follows:

Definition 5 *A mechanism $M(q, e)$ is defined as a combination of an allocation rule $q(\cdot)$ and a payment rule $e(\cdot)$. Given a profile of reported bids $b = (b_1, b_2, \ldots, b_{\mathcal{A}})$, we can characterize the allocation and payment rules for each ad in the exposure number t in session i as follows:*

- *The allocation rule $q^M_{i,a,t}(b; S_{i,t})$ determines the probability that the item is allocated to a.*
- *The payment rule $e^M_{i,a,t}(b; S_{i,t})$ determines what ad a pays in expectation.*

We assume that the publisher has full commitment power, i.e., the players believe that the publisher follows the rules. Now, under any mechanism $M(q, e)$, we can characterize advertiser's utility function. The only decision variable for any advertiser is to submit a bid that reflects their willingness to pay for a click on their ad. Given advertiser a's bid, we can characterize their utility in exposure t of session i as follows:

$$u^M_{i,a,t}(b_a; x_a, S_{i,t}) = \mathbb{E}_{b_{-a}}\left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b; S_{i,t}) - e^M_{i,a,t}(b; S_{i,t})\right], \tag{5.2}$$

where b_{-a} is a bid profile of all ads except a, and b is the profile of all bids. We assume that advertisers maximize their utility.

restrict our attention to this case.

5.3.2 Warm-Up Case: Optimal Static Auction

We begin by describing the case where the publisher's objective is static. In this case, the publisher only considers the current period rewards. As such, at any point, the goal is to sell the slot to the ad that maximizes the publisher's revenues. The analysis of this case is almost identical to that of the seminal paper by [96] on optimal auctions. However, we present this case as a warm-up example for our main goal – deriving the optimal dynamic auction.

In general, the choice of the optimal auction may seem impossible as there is no bound on the set of feasible auctions. However, in light of the Revelation Principle and without loss of generality, we can only focus on direct-revelation mechanisms wherein advertisers truthfully bid their click valuations [96]. A direct revelation mechanism is feasible if it satisfies:

1. *Plausibility:* For any profile of reported click valuations x, we have $\sum_{a \in \mathcal{A}_i} q_{i,a,t}(x; S_{i,t}) \leq 1$ and $q_{i,a,t}(x; S_{i,t}) \geq 1$ for all $a \in \mathcal{A}_i$. This condition guarantees that each impression is allocated to at most one ad.

2. *Individual Rationality (IR):* Given reported click valuations, all advertisers receive non-negative utility, i.e., $u^M_{i,a,t}(x_a; x_a, S_{i,t}) \geq 0$ for all $a \in \mathcal{A}_i$. This condition guarantees that all competing ads for a session will participate in the auction.

3. *Incentive Compatibility (IC):* No advertiser has incentive to deviate from bidding truthfully, given that everyone else bids truthfully. Therefore, we have:

$$u^M_{i,a,t}(x_a; x_a, S_{i,t}) \geq u^M_{i,a,t}(b_a; x_a, S_{i,t}) \tag{5.3}$$

   for any $b_a \in [\underline{x}_a, \bar{x}_a]$. This condition guarantees that reporting truthfully is a Bayesian Nash Equilibrium for all advertisers.

The Revelation Principle helps us reduce the set of all auction to the set of feasible direct revelation mechanism denoted by $\mathcal{M}_{\text{direct}}$. As such, our search is over a more structured set with clear constraints. We can write the publisher's optimization problem as follows:

$$\max_{M \in \mathcal{M}_{\text{direct}}} \mathbb{E}_x \left[\sum_{a \in \mathcal{A}_i} e^M_{i,a,t}(x; S_{i,t}) \right] \tag{5.4}$$

In this optimization, $M \in \mathcal{M}_{\text{direct}}$ implies that the mechanism must satisfy all three constraints presented above. While focusing only on feasible direct revelation mechanisms helps constrain the problem, we still need further transformations to find the optimal solution. One key transformation is to use envelope condition instead of the IC constraint. The following lemma shows the link between these two:

Lemma 1 *If the mechanism $M(q, e)$ is IC, we have:*

$$\begin{aligned} u^M_{i,a,t}(x_a; x_a, S_{i,t}) = \;& u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) \\ &+ P(Y_{i,t} \mid a, S_{i,t}) \int_{x'_a}^{x_a} \mathbb{E}_{x_{-a}} \left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) \right] db_a \end{aligned} \tag{5.5}$$

Now, we can use this lemma to derive the publisher's expected revenue under any IC mechanism as follows:

Lemma 2 *If the mechanism $M(q, e)$ is IC, the publisher's expected revenue can be written as follows:*

$$\mathbb{E}_x \left[\sum_{a \in \mathcal{A}_i} \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) \right] - \sum_{a \in \mathcal{A}_i} u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) \tag{5.6}$$

The result of Lemma 2 is transformative in finding the optimal auction, as for any given environment, the expected revenue of an auction only depends on the allocation at the equilibrium and advertisers' expected utility of their lowest click valuation. We can now use Equation (5.6) as the objective function and maximize it subject to the *Plausibility* and *Individual Rationality (IR)* constraints and obtain the optimal auction.

Another important feature of writing the objective function in Equation (5.6) is that it additively separates allocation and payment functions: the first component is independent of the payment function. Thus, one candidate for the optimal auction is to find a plausible allocation function that maximizes the first component and a payment function that minimizes the second component. This brings us to the following proposition:

Proposition 4 *The mechanism $M(q, e)$ is optimal if q maximizes*

$$\mathbb{E}_x \left[\sum_{a \in \mathcal{A}_i} \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) \right] \tag{5.7}$$

subject to q being plausible and $\mathbb{E}_{x_{-a}} \left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t}) \right]$ increasing in x_a, and the payment function e is

$$e^M_{i,a,t}(x; S_{i,t}) = w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) - P(Y_{i,t} \mid a, S_{i,t}) \int_{\underline{x}_a}^{x_a} q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) db_a \tag{5.8}$$

As shown in Proposition 4, we can solve for the optimal allocation and payment functions. In next sections, we discuss the details of each component.

Allocation Rule

Proposition 4 offers a constrained optimization to find the allocation function under the static case. Without the constraint on $\mathbb{E}_{x_{-a}} \left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t}) \right]$ being increasing in x_a, we can simply maximize the objective function by allocating to the ad with the highest $\left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t})$. However, there is no guarantee whether the constraint is satisfied unless we impose the following assumption on the distribution F_a for all ads:

Assumption 1 *Distribution F_a is regular, i.e., the function $c_a(x_a) = x_a - \frac{1 - F_a(x_a)}{f_a(x_a)}$ is strictly increasing in x_a.*

It is worth noting that this is not an unrealistic assumption, since most familiar distributions satisfy this condition. Under this assumption, the resulting allocation rule will allocate the item to ad with the highest non-negative $c_a(x_a) P(Y_{i,t} \mid a, S_{i,t})$ for any exposure t in session i. If all values are negative, the publisher does not sell the item. It is easy to show that under this allocation, the constraint on $\mathbb{E}_{x_{-a}} \left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t}) \right]$ being increasing in x_a is satisfied: an increase in x_a will increase the expected probability of winning the item for ad a, since $c_a(x_a)$ is increasing in x_a.

Payments

Following Equation (5.8) in the second part of Proposition 4, we can now determine the optimal payment functions, consistent with the allocation function presented in §5.3.2. We can show that in this case, losing ads will not pay anything: the first component in Equation (5.8) is zero, and it is easy to show that the integral is also zero. The payment for the winning ad, however, can be calculated as follows:

$$\begin{aligned} e^M_{i,a,t}(x; S_{i,t}) &= w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) - P(Y_{i,t} \mid a, S_{i,t}) \int_{\underline{x}_a}^{x_a} q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) db_a \\ &= w_{i,a,t}(x_a; S_{i,t}) - P(Y_{i,t} \mid a, S_{i,t}) \int_{x^*_a}^{x_a} db_a \\ &= x^*_a P(Y_{i,t} \mid a, S_{i,t}), \end{aligned} \tag{5.9}$$

where x^*_a is the minimum bid that still wins the impression for the winning ad. The simple allocation rule in the static case helps us find the analytical solution to the integral in Equation (5.8). With this payment rule, it is easy to check the incentive compatibility of the proposed optimal mechanism.

Finally, if we consider the case of symmetric click valuations, we can simplify the optimal auction to a greater extent. In this case, instead of having ad-specific distributions F_a for each ad, we have one distribution F from which all click valuations are drawn independently. We can show the following corollary for this specific case:

Corollary 1 *With symmetric independent click valuations, the optimal auction is a second-price auction with a reserve price of* $c^{-1}(0)$*. The auction allocates the item to the ad with the highest* $w_{i,a,t}(x_a; S_{i,t})$ *and that ad pays the second-highest* $w_{i,a,t}(x_a; S_{i,t})$ *in expectation.*

Thus, the optimal auction either allocates the impression to the highest valuation advertiser or does not allocate it at all, meaning that it never allocates an impression to an advertiser who does not have the highest valuation for that impression.

5.3.3 Optimal Dynamic Auction

We now focus on the optimal auction with the dynamic objective. In this case, the publisher wants to incorporate the expected future revenues as well as expected revenues from the current period. This is in contrast with optimal static auction where the publisher only cares about her expected revenues in the current period. The publisher's goal in this dynamic environment is to design a mechanism $M(q, e)$ that maximizes her expected revenues from the session, i.e., $\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(\sum_{a \in \mathcal{A}_i} e^M_{i,a,t}(b; S_{i,t})\right)\right]$.

This change in the publisher's objective clearly changes critical aspects of our environment. The most important change is that the publisher no more sells each exposure t in session i, but rather auctions off the entire session i. As such, the optimal auction in the static case is no more optimal under the current objective. An important factor that helps us simplify the dynamic problem is that only one piece is private to each advertiser in the entire session: their click valuation x_a. That is, advertisers' click valuation will not change within the session, and anything that changes their overall valuation of each impression (e.g., the probability of click in different states) is common knowledge. Thus, we can treat this problem as a case of static auction where advertisers just submit only one bid and the publisher decides how to allocate the exposures within the session, using a dynamic objective.

In line with the change in the publisher's objective, we must re-define advertisers' utility function for the session as follows:

$$U^M_{i,a}(b_a; x_a, S_i) = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b; S_{i,t}) - e^M_{i,a,t}(b; S_{i,t})\right)\right], \tag{5.10}$$

where the expectation is taken over bidding strategies and the mechanism, and S_i denotes the pre-session information.

Similar to the static case, without loss of generality, we only focus on dynamic direct revelation mechanisms to find the optimal auction [97]. As such, for the dynamic case, we re-write the requirements for a feasible direct revelation mechanism as follows:

1. *Plausibility:* This condition states that at each time period, the exposure is allocated

to at most one ad, given that advertisers report their click valuations truthfully. That is, we have $\sum_{a \in \mathcal{A}} q_{i,a,t}(x; S_{i,t}) \leq 1$ and $q_{i,a,t}(x; S_{i,t}) \geq 1$ for all $a \in \mathcal{A}$ and for any t in session i.

2. *Individual Rationality (IR):* Given truthful reporting of click valuations, the advertiser's expected utility over the session is non-negative, i.e., $U^M_{i,a}(x_a; x_a, S_i) \geq 0$ for all $a \in \mathcal{A}_i$.

3. *Incentive Compatibility (IC):* No advertiser has incentive to deviate from bidding truthfully, given that everyone else reports their bid truthfully. Hence, we can write:

$$U^M_{i,a}(x_a; x_a, S_i) \geq U^M_{i,a}(x'_a; x_a, S_i) \tag{5.11}$$

where the expectation is taken over other advertisers' click valuations and the mechanism. This constraint guarantees that truth-telling is a Bayesian Nash Equilibrium for all advertisers.

While this restriction to the direct revelation mechanisms reduces the set of auctions the publisher considers, we still need some transformations in these constraints to be able to solve for the optimal mechanism. Like the static case, we show the resulting envelope condition from the IC constraint in the dynamic case as follows:

Lemma 3 *If the mechanism $M(q, e)$ is IC, we have:*

$$\begin{aligned} U^M_{i,a}(x_a; x_a, S_i) = \; & U^M_{i,a}(x'_a; x'_a, S_i) \\ & + \int_{x'_a}^{x_a} \mathbb{E}_{x_{-a}} \left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) \right] db_a \end{aligned} \tag{5.12}$$

Now, with the dynamic version of the envelope condition, we can show that the publisher's expected revenues under any IC mechanism can be written as follows:

Lemma 4 *If the mechanism $M(q,e)$ is IC, the publisher's expected revenue can be written as follows:*

$$\begin{aligned}\mathbb{E}_x\Bigg[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a;S_{i,t})\right)\\ -\sum_{a\in\mathcal{A}_i}U^M_{i,a}(\underline{x}_a;\underline{x}_a,S_i)\Bigg]\end{aligned} \tag{5.13}$$

Lemma 4 is the equivalent of Lemma 2 for the dynamic case. Now, if we optimize this new objective subject to both *Plausibility* and *Individual Rationality*, we can find the optimal auction. Further, an important finding of this lemma is that the publisher's revenues cannot exceed the first component in Equation (5.13). Thus, roughly speaking, if we find a dynamic allocation policy that maximizes the first component and payments are such that the second component is zero, the corresponding mechanism is optimal. More precisely, we can write the following proposition on the optimal auction with dynamic objective as follows:

Proposition 5 *The mechanism $M(q,e)$ is optimal if q maximizes*

$$\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a;S_{i,t})\right)\right] \tag{5.14}$$

subject to q being plausible and $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ increasing in x_a, and the payment function e is

$$\begin{aligned}e^M_{i,a}(x;S_i)=&\,\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right)\right]\\ &-\int_{\underline{x}_a}^{x_a}\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(b_a,x_{-a};S_{i,t})\right]db_a,\end{aligned} \tag{5.15}$$

where the expectation is taken over the stochasticity induced by the dynamic process and state transitions, and not over other advertisers' click valuations.

This proposition states that if an allocation mechanism that maximizes the first part of Equation (5.13) and sets the payment to make the second part zero, this mechanism is optimal if the allocation mechanism satisfies both plausibility and monotonicity conditions. We explain the details of both allocation and payment components in next sections.

Allocation Rule

We start by finding the optimal allocation rule. Again, we impose the assumption that the distribution F_a is regular for each ad a, i.e., $c_a(x_a) = x_a - \frac{1-F_a(x_a)}{f_a(x_a)}$ is increasing in x_a. In the static case, we show that under this assumption, the optimal ad at any time period is simply the one that maximizes $c_a(x_a)P(Y_{i,t} \mid a, S_{i,t})$. In the dynamic case, however, we want to design a dynamic allocation policy that maximizes the expected revenues for the entire session.

An important result of Lemma 4 is that we can re-write the reward function that is independent of payments. We present a generic definition of reward function with revenue-maximizing objective as follows:

$$R_t^r(a; S_{i,t}) = \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)}\right) P(Y_{i,t} \mid a, S_{i,t}) \tag{5.16}$$

Now, we can use the reward function in Equation (5.16) to write down the publisher's optimization problem and find the optimal allocation. The following lemma characterizes the optimal allocation function:

Lemma 5 *If the distribution F_a is regular for each ad a (i.e., $c_a(x_a) = x_a - \frac{1-F_a(x_a)}{f_a(x_a)}$ is increasing in x_a), then the optimal allocation is the solution to the following Markov Decision Process:*

$$\underset{a}{\operatorname{argmax}}\, R_t^r(a; S_{i,t}) + \beta \mathbb{E}_{G_{i,t+1}|S_{i,t},a}\left[V_{t+1}^r(G_{i,t+1})\right], \tag{5.17}$$

where the value function is defined for each exposure number as follows:

$$V_{t+1}^r(S_{i,t}) = \max_a R_t^r(a; S_{i,t}) + \beta \mathbb{E}_{G_{i,t+1}|S_{i,t},a}\left[V_{t+1}^r(G_{i,t+1})\right] \tag{5.18}$$

Depending on the context of our problem, we can use various approaches to design the optimal dynamic allocation policy. Based on this dynamic policy, we can then easily define the allocation function $q_{i,a,t}(x, S_{i,t})$. This lemma guarantees that under the assumption that click valuations come from a regular distribution, the chosen q based on Equation (5.17) and Equation (5.18) satisfies both plausibility and monotonicity constraints.

Payments

The intuition behind the payment function in the dynamic case is the same as that in the static case. Each advertiser pays the expected valuation they received from the session, minus an informational rent which is determined by integration of their allocation over the set of lower possible values. This payment function guarantees that the IR constraint is satisfied and the second component in Equation (5.13) will be zero.

While the informational rent has an analytical solution in the static case as shown in Equation (5.9), it is harder to derive it analytically in the dynamic case, since the allocation function contains more elaborate rules. The first component in Equation (5.15) is the expected valuation advertiser a receives from the session, and the second component is the informational rent. To calculate the amount of this rent, we basically need to move down from the true click valuation and see how the number of impressions allocated to advertiser a shrinks. Integrating over this function over the possible values will then give us the amount of rent advertiser a is able to extract. Intuitively, it is the total leeway advertiser a has in this optimal auction.

5.4 Empirical Strategy

In light of the theoretical results in Proposition 4 and 5, we can characterize the reward function at any point for both static and dynamic cases as follows:

$$R_t^r(a; S_{i,t}) = \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t}), \tag{5.19}$$

where $S_{i,t}$ is the set of state variables in session i at exposure t. Since the reward specification is the same at any point, we can present a unifying specification of the publisher's optimization in both static and dynamic cases as follows:

$$\operatorname*{argmax}_a R_t^r(a; S_{i,t}) + \beta \mathbb{E}_{S_{i,t+1}|S_{i,t},a} \left[V_{t+1}^r(S_{i,t+1}) \right], \tag{5.20}$$

where $\beta = 0$ when the objective is static and the value function is defined as the maximum future rewards at a given state.

Given the unified optimization problem in Equation (5.20), we need to know three key elements to determine the outcomes: (1) distribution of click valuations, (2) match values or click probabilities, and (3) distribution of transitions. The first two elements are required for both static and dynamic objectives, whereas the distribution of transitions is only required for the dynamic case. To conduct an empirical analysis of both static and dynamic optimal auctions presented in §5.3, we need to obtain estimates for all three unknown elements in Equation (5.20). Thus, we can identify three empirical tasks as follows:

Task 1: For any ad a in the data, we want to estimate the click valuations x_a based on their observed bidding behavior in the data, under the quasi-proportional auction.

Task 2: For any set of state variables observed in the data, we want to accurately estimate the click probability for all ads if shown in that impression. That is:

$$\hat{y}_{i,t}(a; S_{i,t}) = P(Y_{i,t} \mid a, S_{i,t}), \forall a \tag{5.21}$$

Task 3: For any set of state variables observed in the data, we want to accurately estimate the leave probability for all ads if shown in that impression. That is:

$$\hat{l}_{i,t}(a; S_{i,t}) = P(L_{i,t} \mid a, S_{i,t}), \forall a \tag{5.22}$$

The empirical strategy for Tasks 2 and 3 is presented in [104] in great details. We use the same approach for these tasks. We present our empirical strategy to estimate the click valuations in this section.

5.4.1 Setting

The setting of this problem is the same as that of [104]. However, since we are interested in estimation of click valuations, we present more details on the auction environment and advertisers' decisions in this section.

Auction Mechanism

For any exposure that is recognized, the ad-network runs an auction to serve an ad. Unlike the common practice in this industry, the ad-network runs a quasi-proportional auction to select the ad for each exposure. The most notable feature of this auction is in the probabilistic allocation rule that is in contrast with the commonly used mechanisms such as first- or second-price auctions.

1. **Reserve Price:** There is a reserve price b_0 that advertisers' bid must exceed to participate in the auction.

2. **Allocation Function:** For any exposure i and any set of participating ads $\mathcal{A}_i$ with bidding profile $(b_1, b_2, \ldots, b_{|\mathcal{A}_i|})$, ad a has the following probability to win each exposure t within session i:

$$q^p_{i,a,t}(b; z) = \frac{b_a z_a}{\sum_{j \in \mathcal{A}_i} b_j z_j}, \tag{5.23}$$

   where z_a is ad a's quality score which is a measure reflecting the profitability of ad a. The rationale behind using such quality score adjustments is to run the auction based on the expected revenue the ad-network can extract from the ad, rather than their willingness to pay per click. For example, there may be an ad with a very high bid but no chance of getting a click. So despite its high willingness to pay per click, the ad-network cannot actually extract much from it as it will not get many clicks.

   While quality scores can technically vary across auctions, the ad-network does not update them regularly. They only take value zero when the ad is not available for the auction due to budget exhaustion or targeting decisions. Likewise, bids do not change across t, since advertisers cannot choose their bids per auction[3]. Further, changing bids is unlikely as bids reflect advertisers' structural parameters that are stable across auctions.

3. **Payment Scheme:** The ad-network employs a cost-per-click (CPC henceforth) payment scheme wherein advertisers are only when charged when a user clicks on their ad. The amount an ad is charged per click is determined by a next-price rule similar

[3]There are over 500 auctions run in a second and the information shared with advertisers is minimal about them, disabling them from distinguishing auctions.

to that of Google's sponsored search auctions. That is, ads are first ranked based on their product of bid and quality score, and each ad pays the minimum amount that guarantees their rank, if a click happens on their ad:

$$e^{p}_{i,a,t}(b;z) = \inf\left\{ b' \mid \sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b' z_a \leq b_j z_j) = \sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b_a z_a \leq b_j z_j) \right\}, \tag{5.24}$$

where $\sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b_a q_a \leq b_j q_j)$ indicates the rank of advertiser a, and the payment b' is the minimum amount of bid that guarantees the same rank for ad a. For example, if there are three bidders with bids 1, 2, and 3, and quality scores 0.1, 0.2, and 0.3, the scores will be 0.1, 0.4, and 0.9. Now, if the second-ranked bidder gets a click, she will pay the price that would have guaranteed her second rank. That is, she only needs to pay $\frac{1\times 0.1}{0.2} = 0.5$, as it guarantees her score to be higher than the third-ranked bidder. Since the item being sold is a click, advertisers' bids reflect their willingness to pay per click. Our goal is to use advertisers' observed bid to estimate their click valuations.

Advertisers' Decisions

Advertisers can make the following decisions:

- **Bid:** Advertisers can set their bid indicating how much they are willing to pay per click.
- **Targeting:** Advertisers can specify their targeting decisions on the following variables: (1) app category, (2) province, (3) hour of the day, (4) smartphone brand, (5) connectivity type, and (6) mobile service provider (MSP). As such, they can exclude some categories from the variables listed above. It will guarantee that their ad will not be shown in the excluded categories.
- **Budget:** Advertisers become unavailable if they do not have enough budget in their account. They can refill their budget and make their campaigns available again.
- **Design of the Banner:** They can design a small banner for their ad.

While we mainly focus on their bidding behavior, it is important to take into account what other decisions they can make. Further, it is important to notice what information they receive from the auction. Each advertiser has a profile in which she can track some key performance metrics on an hourly basis such as the number of impressions, number of clicks, and the total cost of clicks. As such, they do not have granular access to exposure-level information and can only incorporate this information at an aggregate level.

5.4.2 Data

We use data from a leading in-app ad-network in a large Asian country for over a one week time period from October 22 to 30 in 2015. The analysis sample is the same as the one in [104]. However, for estimation of auction, we use all the impressions as it adds to the precision of our estimates. The original data are at the impression-level, indicating the characteristics of each impression, the ad shown, and the final clicking decision by the user. Please see § 3.2.1 in [104] for the description of impression-level data.

The data required for the estimation of auction are not readily available, but we can construct that from the impression-level data. Auction information in the impression-level data is limited to the winning ad and the potential CPC for that particular ad. However, for estimation of auction, we need more information on each auction (impression) regarding the losing ads as well their bids. We recover these two pieces of information using the following approach:

- **Set of Competing Ads:** This task is equivalent to the task of identifying ads that could have been shown in a session. Thus, we use the filtering strategy presented in §5.1.1 in [104].
- **Inferring Bid Amounts:** Actual bids are not directly observed from the original data. The amount reported in the data is the one presented in Equation (5.24). Since the CPC is always lower than or equal to the actual bid, the maximum observed CPC can be the best approximation for the actual bid. We can write:

$$b_a \approx \sup_i \inf \left\{ b' \mid \sum_{j \in \mathcal{A}_i, j \neq a} \mathbb{1}(b' z_a \leq b_a z_a) = \sum_{j \in \mathcal{A}_i, j \neq a} \mathbb{1}(b_a z_a \leq b_a z_a) \right\} \tag{5.25}$$

Again, being in a data-rich environment with extensive variation in the next-price enormously helps. One intuitive validation for this method would be the ending digits of approximated bids as bidders are more likely to round up their bids. We find that this is the case for most approximated bids.

Overall, it provides us with 547,626,756 impressions and 398 distinct ads competing to win the impressions.

Summary Statistics

Here we present some summary statistics of the data. Overall, there are 398 ads participating in the timeline of the study. Table 5.1 shows mean, standard deviation, minimum and maximum of key ad-level variables. All these variables are defined by ads and reflect an important aspect of advertiser's decision, including bidding, targeting, and budget.

Variable	**Mean**	**Std. Dev**	**Min**	**Max**
Bid	411.13	208.25	300.00	2976.00
Avg. CPC	363.30	103.98	300.00	1375.84
Total Hours of Availability	48	70	1	217
No. of Impressions	1,379,411	5,402,346	7	66,228,977
No. of Clicks	12,619	55,208	0	656,680
Click-Through Rate	0.0109	0.0112	0.0000	0.1429
Total Expenditure	4,880,013	21,032,815	0	216,434,237

Table 5.1: Summary statistics of key ad-level variables

As shown in this table, there is substantial variation in ads participating in this auction, ranging from smaller ads with a very short lifetime and few clicks, to larger ads with permanent availability and significant expenditure. Further, Table 5.1 presents number of categories targeted out of all 86 targeting categories, indicating that most ads are not targeting at all.

Figure 5.2 shows the the empirical CDF of ads' bids and their average CPC. The figure on the left (Figure 5.2b) shows the distribution for all values of bids. As shown in this figure, there is a reserve price 300 that censors the left side of the graph. With no reserve price, there could have been bids lower than 300. This raises an important identification challenge that we address in the estimation approach, especially because a substantial portion of ads are reserve bidders.

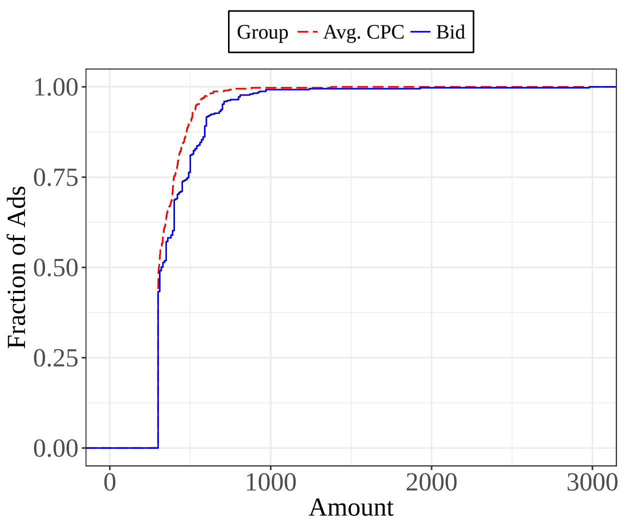

(a) Zoomed Out

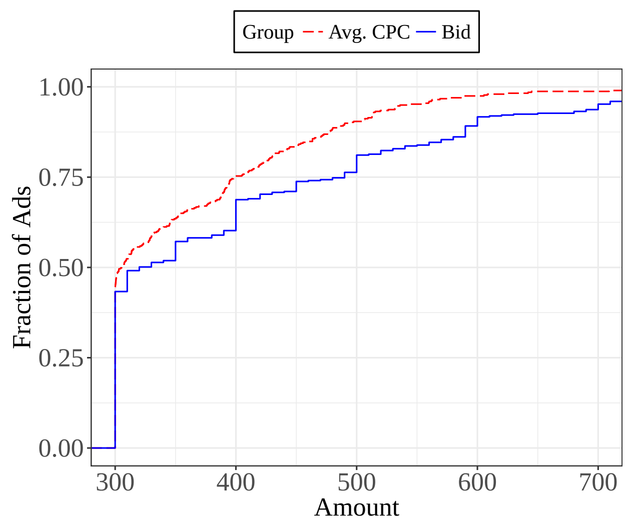

(b) Zoomed In

Figure 5.2: Empirical CDF of ads' bids and average CPCs.

Figure 5.2b zooms into the bids for a shorter interval of amounts that capture over 90% of all ads. While there are discontinuities in bids especially at round numbers, average CPCs show a smoother pattern. Overall, both figures demonstrate a first-order stochastic dominance relationship between the empirical CDF of average CPC and bids which emphasizes the fact an ad's CPC cannot exceed her bid.

5.4.3 Estimation of Auction

We now present our approach to estimate the distribution of click valuations from the observed auction data. We first develop advertisers' utility model and provide and equilibrium analysis. We then state the set of assumptions required for the identification of the distribution of click valuations. Finally, we propose our estimation approach.

Advertisers' Model

We begin by characterizing advertisers' utility model in the context of our data. Given the allocation and payment functions defined in §5.4.1, we know that advertiser a receives the following utility from exposure t in session i:

$$\mathbb{1}(A_{i,t} = a, Y_{i,t} = 1) \left(x_a - e^p_{i,a,t}(b)\right),$$

where the indicator function takes value one if ad a is selected through the allocation mechanism and got clicked in that impression. Here we explicitly assume that advertisers only receive utility from clicks and there is no utility from an impression alone.[4]

[94] study this auction with the case of identity cost function, i.e., $e^p_{i,a,t}(b) = b_a$. They show that advertisers' optimal bidding strategy depends on their expected probability of winning. While this expected probability varies across auctions in our case, we do not observe bid-changing behavior. This is possibly because the effect of the expected probability of winning on their bid is very small in a competitive market and it does not exceed the bid-changing cost.[5]

The observation that advertisers do not change their bids informs us in modeling advertisers' utility function. We can characterize the expected utility of advertiser a from

[4]We believe this assumption is valid in our context, because almost all ads are mobile apps whose objective is more app installs. This assumption may be violated in the presence of brand ads whose objective is more reach.

[5]Please notice that this is just a behavioral cost and the platform does not charge them for bid-changing. As shown in [107], the marginal effect of the expected winning probability α is $\frac{1}{1-\alpha}$. Hence, in a market with many homogeneous competitors, α has a very small effect on advertisers' equilibrium bidding strategy.

quasi-proportional auction p as follows:

$$u_a^p(b_a; x_a) = m_a \left(x_a \tilde{q}_a^p(b_a) - \tilde{e}_a^p(b_a) \right), \tag{5.26}$$

where m_a is the expected probability of click on ad a conditional on winning. Since we treat this probability as independent to other parts of the expected utility, we separate that in the specification. Further, functions $\tilde{q}_a^p$ and $\tilde{e}_a^p$ are advertiser a's expected allocation and payment functions. Before defining these two functions, we define an important distribution that we use in our estimation. Let $\mathcal{C}_a$ denote the joint distribution of auctions that advertiser a participates in. Each draw C_a from this distribution contains the information about the distribution of bids and quality scores, the impression characteristics (session i and exposure t), and the number of ads competing in that auction. Using this distribution, we can now characterize the allocation function as follows:

$$\tilde{q}_a^p(b_a) = \mathbb{E}_{\mathcal{C}_a} \left[q_{i,a,t}^p(b_a, b_{-a}^{C_a}) \right], \tag{5.27}$$

where the expectation is taken over all the configurations C_a. This function essentially returns the expected probability of winning by advertiser a in a random impression. Similarly, we can define the payment function as follows:

$$\tilde{e}_a^p(b_a) = \mathbb{E}_{\mathcal{C}_a} \left[\mathbb{1} \left(A_{i,t} = a; C_a \right) e_{i,a,t}^p \left(b_a, b_{-a}^{C_a}; z_a, z_{-a}^{C_a} \right) \right] \tag{5.28}$$

The term $\mathbb{1}\left(A_{i,t} = a; C_a \right)$ in Equation (5.28) shows whether ad a wins the impressions and the second term basically computes the cost-per-click using Equation (5.24). Now, using the utility specification in Equation (5.26), we can write the first-order condition as follows:

$$x_a = \frac{\partial \tilde{e}_a^p(b_a)}{\partial b_a} \left(\frac{\partial \tilde{q}_a^p(b_a)}{\partial b_a} \right)^{-1} \tag{5.29}$$

This equation lays out our estimation approach, as we need to empirically estimate both $\tilde{q}$ and $\tilde{e}$ using the data at hand. In the next section, we discuss the assumptions required for identification and overall identification strategy.

Assumptions and Identification

We make a series of assumptions required for our estimation task. Some of these are commonly used in the context of auction estimation, whereas some other assumptions are more specific to the context of quasi-proportional auction in our data. While the former is necessary for identification, we impose the latter mostly for the ease of estimation. For robustness check, we relax those specific assumptions and show the results will not change.

Our first assumption characterizes how advertisers make decisions regarding their own bidding strategy:

Assumption 2 *[**Profit-Maximizing Advertisers**] Advertisers are profit-maximizing, that is, they choose the bid that maximizes their profit.*

In light of Assumption 2, we can treat observed bids in the data as equilibrium bids and use Lemma 6 to estimate click valuations. As shown in Equation (5.26), advertisers choose their optimal bids given their own private click valuations and their belief about other advertisers. The following assumption characterizes other advertisers' private click valuations:

Assumption 3 *[**Independent Private Values (IPV)**] Advertisers' private click valuations are drawn independently from a distribution $F(\cdot)$ with a continuous density.*

The assumption of Independent Private Values (IPV) is an assumption used in most auction settings [53, 6]. This assumption hints a straightforward approach to estimate both $\tilde{q}$ and $\tilde{e}$ in Equation (5.26), by simulating the distribution of $\mathcal{C}_a$ for each ad.

Now, we make assumptions more specific to the case of quasi-proportional auctions. Some of these assumptions are necessary to use the first-order condition in Lemma 6. However, we impose some of these assumptions to provide an analytically simpler solution. We start with the following assumption:

Assumption 4 *[**Zero Impression Value**] Advertisers' valuation from an unclicked impression is zero.*

This assumption is widely used in both theoretical and empirical literature on cost-per-click auctions [30, 131, 7]. In our case, we believe the value from impressions is negligible as most ads are mobile apps whose objective is to get more app installs and can be considered as performance ads. In the next assumption, we make an assumption about the role of budget in advertisers bidding behavior:

Assumption 5 *[**Budget Independent Bidding**] Advertisers' bidding behavior is independent of their budget.*

In principle the equilibrium bidding behavior may change for budget-constrained advertisers [16, 5]. We make this assumption in our case for two reasons. First, we observe quite a few advertisers who have run out of budget during our study and refilled it later. However, there is no difference in their bidding behavior. The second reason is more empirical, as we do not observe the exact budget in our data. As such, we need to approximate this budget which can be quite noisy.

The assumptions so far are enough to establish the identification of the distribution of click valuations for bidders who bid above the reserve price. However, we make two additional assumption that helps us derive a simpler analytical solution for this auction.

Assumption 6 *[**Separability of Allocation**] We can separate the allocation function from the payment function as follows:*

$$\tilde{e}^p_a(b_a) = \tilde{q}^p_a(b_a)\mathbb{E}_{\mathcal{C}_a}\left[e^p_{i,a,t}\left(b_a, b^{C_a}_{-a}; z_a, z^{C_a}_{-a}\right)\right]$$

Assumption 7 *[**Proportional Functional Form of Allocation**] We can write the functional form for the function $\tilde{q}$ as follows:*

$$\tilde{q}^p_a(b_a) = \frac{b_a z_a}{b_a z_a + \zeta_a}, \tag{5.30}$$

where ζ_a captures the overall competition ad a faces.[6]

[6]It is worth mentioning that the functional form of the expected probability of winning is not necessarily quasi-proportional, but it is well-approximated by a quasi-proportional functional form.

These assumptions allow us to simplify the first-order condition presented in Equation (5.29). However, it is worth noting that we can actually relax both assumptions and neither of them is crucial for identification.

For brevity, let $\epsilon_a^p(b_a) = \mathbb{E}_{\mathcal{C}_a}\left[e_{i,a,t}^p\left(b_a, b_{-a}^{C_a}; z_a, z_{-a}^{C_a}\right)\right]$. It is easy to show that the function ϵ_a^p is monotonic for each advertiser a. However, we need to make the following assumption about this function:

Assumption 8 *[**Twice Differentiability of Payment**] The expenditure function ϵ_a^p is twice differentiable in b_a.*

Now, we can use all these assumptions and show the following lemma which is the key idea behind our identification:

Lemma 6 *If advertiser a's equilibrium bid is greater than the reserve price ($b_a > b_0$) and the function $b^2 \frac{\partial \epsilon_a^p(b)}{\partial b}$ is increasing in the local neighbourhood around b_a, her click valuation x_a can be written in terms of equilibrium bids as follows:*

$$x_a = \epsilon_a^p(b_a) + \frac{b_a \frac{\partial \epsilon_a^p(b)}{\partial b}\,|_{b=b_a}}{1 - \tilde{q}_a(b_a)} \tag{5.31}$$

Lemma 6 shows the first-order condition for the case where the advertiser's equilibrium bid is greater than the reserve price. We need $b^2 \frac{\partial \epsilon_a^p(b)}{\partial b}$ to be increasing at b_a to satisfy the second-order condition. Intuitively, if the function e is too concave, this assumption may not hold. However, it is a testable assumption as we can empirically test it for all the bidders.

Now, we focus on the advertisers whose bid is equal to the reserve price. For these advertisers, we cannot use the inverse bidding equation in Equation (5.31), as their optimal bid could have been lower than the reserve price, if they had been allowed to bid lower. Instead of the first-order condition, we have other conditions for reserve price bidders. First, we know that their participation constraint is satisfied, i.e., $x_a \geq b_0$. Second, we know that their first-derivative at b_0 is lower than or equal to zero, as the utility must be decreasing at the truncation point. Together, we can write the following lemma to characterize the link between the click valuation and reserve bidding behavior:

Lemma 7 *If advertiser a's equilibrium bid is equal the reserve price ($b_a = b_0$), we can obtain lower and upper bounds for her click valuation x_a as follows:*

$$b_0 \leq x_a \leq \epsilon_a^p(b_0) + \frac{b_0 \frac{\partial \epsilon_a^p(b)}{\partial b} |_{b=b_0}}{1 - \tilde{q}_a(b_0)} \tag{5.32}$$

In light of Lemma 5.32, we cannot point-identify click valuations for the reserve bidders. However, we can use lower and upper bounds in Equation (5.32). To complete our identification for the distribution of valuations for all participating advertisers, we need one more assumption on the click valuations of the reserve bidders:

Assumption 9 *[**Uniformity of Reserve Bidders' Valuations**] The click valuation x_a for any reserve price bidder is drawn from a uniform distribution with the following bounds:*

$$x_a \sim \mathcal{U}\left(b_0, \epsilon_a^p(b_0) + \frac{b_0 \frac{\partial \epsilon_a^p(b)}{\partial b} |_{b=b_0}}{1 - \tilde{q}_a(b_0)}\right) \tag{5.33}$$

With this assumption, we can now state our identification proposition:

Proposition 6 *If all Assumptions 2 to 9 hold, then the distribution of advertisers' private click valuations are non-parametrically identified.*

The proof for this part is similar to the identification proof for most auction models with independent private values [53, 6]. We can directly observe all the elements in Lemma 6 and 7. Thus, we can estimate advertisers' click valuations and form the distribution F.

Estimation Method

Our estimation approach relies on the findings in §5.4.3. We first estimate the joint distribution of configurations $\mathcal{C}_a$ for all ads. This involves the estimation of quality scores, distribution of observed bids, and observed impressions. We then use this distribution to form both allocation and payment functions, which in turn, allows us to estimate the distribution of click valuations.

Before describing the estimation procedure, we need to define a time period η at which advertisers update their bids. While advertisers do not change their bid in our data, we set this time period for two main reasons. First, advertisers can technically change their bids if they want. Therefore, it is more reasonable to assume that they revise their decision every once in a while. Second, from an empirical point-of-view, it allows us to capture the variance in the set of advertisers competing. Thus, following the arguments in [106], we set an hourly time period such that each η distinguish a different hour-day combination. The denote the last time period by L.

We provide a step-by-step procedure for our estimation as follows:

- **Step 1:** We estimate all the quality scores $z_{a,\eta}$ for all ads and time periods. We use the proportional nature of the allocation to identify quality scores. If ads a and a' both participate in the auction for exposure t in session i, the denominator of their winning probability is the same. Thus, the odds ratio of these two ads can be written as follows:

$$\frac{\Pr(A_{i,t} = a)}{\Pr(A_{i,t} = a')} = \frac{b_a z_a}{b_{a'} z_{a'}}$$

Since we observe their bids, we can easily estimate the ratio $\frac{z_a}{z_{a'}}$ by calculating the number of impressions awarded to each ad over the set of auctions where they both have participated. Let $\mathcal{I}^{\eta}_{a,a'}$ denote the set of impressions wherein both ads a and a' participate in time period η. We can write:

$$\frac{\hat{z}_{a,\eta}}{\hat{z}_{a',\eta}} = \frac{b_{a'} \sum_{(i,t)\in\mathcal{I}^{\eta}_{a,a'}} \mathbb{1}(A_{i,t} = a)}{b_a \sum_{(i,t)\in\mathcal{I}^{\eta}_{a,a'}} \mathbb{1}(A_{i,t} = a')} \tag{5.34}$$

As such, we can estimate all the ratios. For ad a^* that has participated in all the impressions, we set $\hat{z}_{a^*,\eta} = 1$ for all time periods. This allows us to estimate all quality scores $\hat{z}_{a,\eta}$.

- **Step 2:** We empirically estimate the joint distribution $\mathcal{C}_{a,\eta}$ for all ads in all time periods. This distribution contains the information about all the impressions that ad a has participated in, number of bidders in corresponding impressions, and the joint distribution of bids and quality scores. We call the estimated distribution $\hat{\mathcal{C}}_{a,\eta}$.

- **Step 3:** Given all estimated configurations $\hat{\mathcal{C}}_{a,\eta}$, we can estimate the allocation probability for all ads over all time periods as follows:

$$\hat{q}_{a,\eta}(b_a) = \frac{1}{N_a} \sum_{(i,t)\sim\hat{\mathcal{C}}_{a,\eta}}^{N_a} \mathbb{1}(A_{i,t} = a), \tag{5.35}$$

where N_a is the number of draws we get from the distribution $\hat{\mathcal{C}}_{a,\eta}$. It is worth noting that we do not need to fully estimate the allocation function as the relationship in Equation (5.29) only depends on the final allocation probability given the observed bid.

- **Step 4:** Again, we use the distribution estimate configurations $\hat{\mathcal{C}}_{a,\eta}$ to estimate our cost-per-click function $\hat{\epsilon}_{a,\eta}(b)$. For any value b, we can estimate the cost-per-click function as follows:

$$\hat{\epsilon}_{a,\eta}(b) = \frac{1}{N_a} \sum_{(i,t)\sim\hat{\mathcal{C}}_{a,\eta}}^{N_a} \inf\left\{ b' \mid \sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b' z_a \leq b_j z_j) - \mathbb{1}(b z_a \leq b_j z_j) = 0 \right\} \tag{5.36}$$

This gives us the estimate for the cost-per-click function. We can then take the numerical derivatives of $\hat{\epsilon}_{a,\eta}(b)$ and estimate $\hat{\epsilon'}_{a,\eta}(b)$ for all values of b.

- **Step 5:** Using the estimates for the allocation and cost-per-click functions, we can now identify click valuations $x_{a,\eta}$ for any combination of a and η as follows:

$$\hat{x}_{a,\eta} = \begin{cases} \hat{\epsilon}_{a,\eta}(b_a) + \frac{b_a \hat{\epsilon'}_{a,\eta}(b_a)}{1-\hat{q}_{a,\eta}(b_a)} & b_a > b_0 \\ x_0 \sim \mathcal{U}\left(b_0, \hat{\epsilon}_{a,\eta}(b_a) + \frac{b_a \hat{\epsilon'}_{a,\eta}(b_a)}{1-\hat{q}_{a,\eta}(b_a)}\right) & b_a = b_0 \end{cases} \tag{5.37}$$

We also obtain single estimates for each x_a without using the time periods. The procedure is the same as what described above. We only need to use $\hat{\mathcal{C}}_a$ instead of this distribution for all time periods. We then use these estimates throughout for click valuations in our analysis.

- **Step 6:** Finally, we use the estimates in Equation (5.37) to form the distribution of click valuations. We can write:

$$\hat{F}(x) = \frac{1}{L}\sum_{\eta=1}^{L}\frac{1}{|\mathcal{A}_\eta|}\sum_{a\in\mathcal{A}_\eta}\mathbb{1}(x_{a,\eta}\leq x) \tag{5.38}$$

$$\hat{f}(x) = \frac{1}{L}\sum_{\eta=1}^{L}\frac{1}{|\mathcal{A}_\eta|}\sum_{a\in\mathcal{A}_\eta}\frac{1}{h}\mathcal{K}\left(\frac{x-\hat{x}_{a,\eta}}{h}\right), \tag{5.39}$$

where $\mathcal{K}$ is the kernel function. In this case, we use Epanechnikov kernel, i.e., $\mathcal{K}(u) = \frac{3}{4}(1-u^2)\mathbb{1}(|u|\leq 1)$.

5.4.4 Results

We now present our results on the estimation of click valuations. We first show the main findings on the distribution of click valuations. Next, we discuss the validity of our assumption that the distribution of click valuations is regular.

Estimated Distribution of Click Valuations

We first show the estimated distribution of click valuations for top ads. We only focus on the top ads, as it is the set of ads that we use for the counterfactual evaluation of the optimal auctions. However, we can technically recover the distribution for any set of ads. Figure 5.3 show the empirical CDF and estimated density of the distribution. As shown in Figure 5.3a, the range between 700 to 1000 constitutes the vast majority of click valuations. Over 80% of click valuations lie in this range.

Figure 5.3b shows the estimated density for the distribution of click valuations. We use Epanechnikov kernel with bandwidth size of 75. The shape of density is similar to a bell curve with a low variance. The mean and standard deviation for this distribution are 854.91 and 164.73 respectively.

On the Regularity of the Distribution of Click Valuations

Now, we use the results shown in Figure 5.3 and check whether the estimated distribution is regular. As described in Assumption 1, the estimated distribution $\hat{F}$ is regular if $x - \frac{1-\hat{F}(x)}{\hat{f}(x)}$

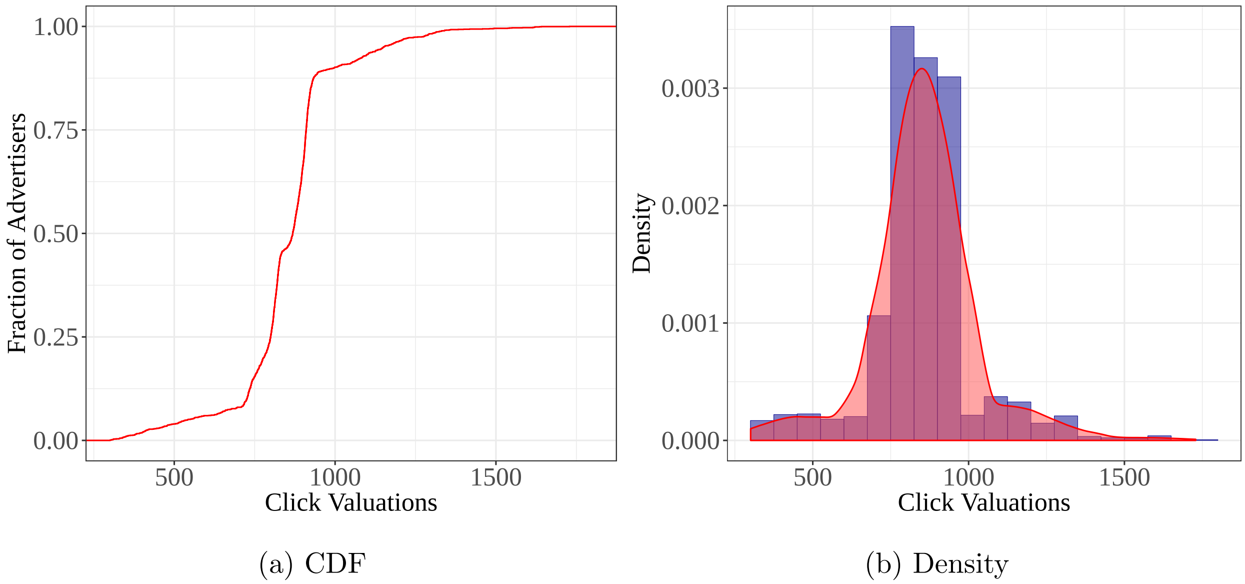

Figure 5.3: Estimated distribution of click valuations for top 15 ads.

is strictly increasing in x. Using our estimates for $\hat{F}$ and $\hat{f}$, we plot the virtual valuations in Figure 5.4. As demonstrated in this figure, virtual values are monotonic in click valuations.

5.5 Counterfactual Evaluation of Optimal Auctions

In this section, we take our theoretical models in §5.3.2 and §5.3.3 to the data and present our approach to evaluate the counterfactual outcomes under optimal auctions. In both cases, we need to use the reward function that is presented Equation (5.16) in a generic way. More specific to our context, we can write it as follows:

$$R_t^r(a; S_{i,t}) = \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t}), \tag{5.40}$$

Now, we want to estimate the reward function in this case, given our data. We use estimates for both click probability at any state (i.e., $P(Y_{i,t} \mid a, S_{i,t})$) as well as click valuations x_a for all ads from [104]. We also impose symmetry assumption on the distribution of click valuations and estimate $\hat{F}$ and $\hat{f}$, based on the estimated click valuations. Given these estimates, we

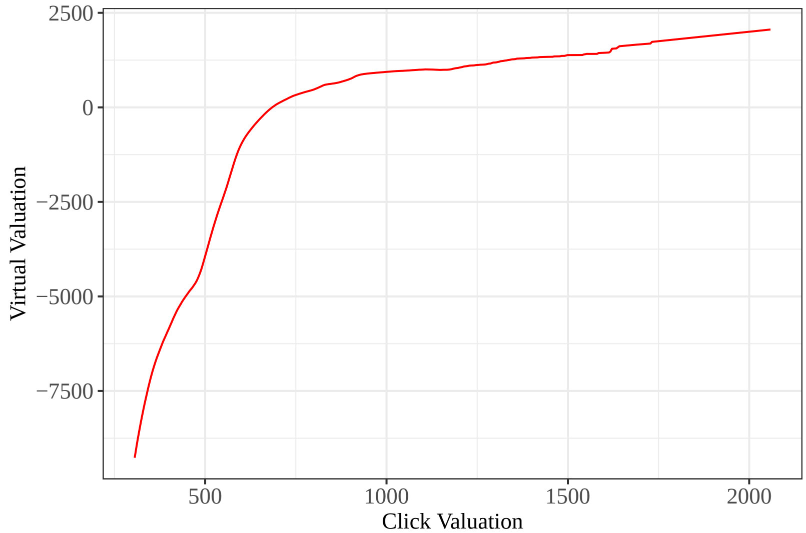

Figure 5.4: Virtual valuations against click valuations

can now estimate the reward function in Equation (5.41) and re-write it as follows:

$$\hat{R}_t^r(a; S_{i,t}) = \left(\hat{x}_a - \frac{1 - \hat{F}(\hat{x}_a)}{\hat{f}(\hat{x}_a)} \right) \hat{y}_{i,t}(a; S_{i,t}) \tag{5.41}$$

We use our estimates for the reward function in this case to obtain optimal auctions in both static and dynamic cases and then present our how we evaluate the resulting auctions.

5.5.1 Solution Concept for the Optimal Static Auction

We start with the simpler case where the publisher's objective is static. In §5.3.2, we show how we can derive the optimal mechanism $M^m(q^m, e^m)$. More specifically, we show that when $\hat{F}$ is symmetric, the optimal auction will be a second-price auction with an optimal reserve price. We can estimate the reserve price $\hat{e}_0$ as follows:

$$\hat{e}_0 = \hat{c}^{-1}(0),$$

where $\hat{c}(x) = \frac{1-\hat{F}(x)}{\hat{f}(x)}$ and we can find $\hat{e}_0$ by solving for $\hat{e}_0 = \frac{1-\hat{F}(\hat{e}_0)}{\hat{f}(\hat{e}_0)}$. Using this reserve price, we can then derive the optimal allocation q^m as follows:

$$q^m_{i,a,t}(\hat{x}; S_{i,t}) = \begin{cases} 1 & \hat{R}^r_t(a; S_{i,t}) > \max_{a' \in \mathcal{A}_i \setminus a} \left(\hat{R}^r_t(a'; S_{i,t}), \hat{e}_0 \right) \\ 0 & \text{otherwise} \end{cases} \tag{5.42}$$

This allocation rule indicates that the highest reward ad (which is the ad with the highest expected valuation) will win the impression.[7] As specified in §5.3.2, the payments are determined as follows:

$$e^m_{i,a,t}(\hat{x}; S_{i,t}) = \begin{cases} \max_{a' \in \mathcal{A}_i \setminus a} \left(\hat{R}^r_t(a'; S_{i,t}), \hat{e}_0 \right) & q^m_{i,a,t}(x \mid S_{i,t}) = 1 \\ 0 & \text{otherwise} \end{cases} \tag{5.43}$$

We can now use both Equation (5.42) and Equation (5.43) to find the optimal mechanism and apply it to each session.[8]

5.5.2 Solution Concept for the Optimal Dynamic Auction

Now, we focus on the optimal auction with dynamic objective and show how we can use our estimates and determine the optimal mechanism $M^d(q^d, e^d)$. As shown in §5.3.3, we can find the optimal allocation by solving a Markov Decision Process that incorporates expected future rewards as well as the current period reward. Like [104], we consider a finite case with and solve for the optimal allocation using backward induction. For notational convenience,

---

[7]It is worth noting that in a case where there are multiple highest reward ads, we randomly allocate the item to one of the highest reward ads. However, since this a very rare event in an empirical setting, for brevity, we exclude that from Equation (5.42).

[8]It is important to notice that we can recover the distribution of click valuations specific to each session. For simplicity, we focus on the case one global distribution for all sessions. As a robustness check, we also consider the case where distributions are session-specific.

we first define the function $\tilde{V}_t^r(a, S_{i,t})$ for a pair of action and state as follows:

$$\begin{aligned} \tilde{V}_t^r(a, S_{i,t}) = & \left(\hat{x}_a - \frac{1-\hat{F}(\hat{x}_a)}{\hat{f}(\hat{x}_a)}\right)\hat{y}_{i,t}(a; S_{i,t}) \\ & + \left(1-\hat{l}_{i,t}(a; S_{i,t})\right)\hat{y}_{i,t}(a; S_{i,t}) V_{t+1}^r\left(\langle S_{i,t}, a, Y_{i,t}=1\rangle\right) \\ & + \left(1-\hat{l}_{i,t}(a; S_{i,t})\right)\left(1-\hat{y}_{i,t}(a; S_{i,t})\right) V_{t+1}^r\left(\langle S_{i,t}, a, Y_{i,t}=0\rangle\right), \end{aligned} \tag{5.44}$$

where $\hat{y}_{i,t}(a; S_{i,t})$ and $\hat{l}_{i,t}(a; S_{i,t})$ are estimated leave and click probabilities respectively. We can use Equation (5.44) and describe our backward induction solution concept as follows:

1. We begin from the last period T. Since there is no expected future at that point, the value function can be estimated as follows:

$$\hat{V}_T^c(S_{i,T}) = \max_{a\in\mathcal{A}_i} \hat{R}_t^r(a, S_{i,t}) \tag{5.45}$$

2. For any $t < T$, we can determine the value function as follows:

$$\hat{V}_t^c(S_{i,t}) = \max_{a\in\mathcal{A}_i} \tilde{V}_t^r(a, S_{i,T}) \tag{5.46}$$

   We can easily estimate the value function for any $t < T$ if we have all the value functions for the next periods. We can satisfy that by going backward and find the value function for all states at any t and continue this process until exposure number 1.

Once we have identified the value function for all the states, we can find the optimal allocation with dynamic objective as follows:

$$q_{i,a,t}^d(\hat{x}; S_{i,t}) = \begin{cases} 1 & a = \text{argmax}_{a\in\mathcal{A}_i} \tilde{V}_t^r(a, S_{i,t}), t < T \\ 1 & a = \text{argmax}_{a\in\mathcal{A}_i} \hat{R}_t^r(a, S_{i,t}), t = T \\ 0 & \text{otherwise} \end{cases} \tag{5.47}$$

Now, given the allocation, we can determine the payments using Equation (5.15). It is important to notice that the payment is determined in expectation over the entire session for each ad. For any ad a, the first component in Equation (5.15) is the average value ad a would get given the allocation, and the second component is the integral of the expected value ad a would get if she reduces her bid, taken over all possible bids that she could submit. The first component is easier to calculate as it is an expected sum of the total value each ad receives given each sequence. The second part, however, is more computationally intensive as it involves a numerical integration. To do that, we first need to estimate the function inside the integral. Let $\hat{Q}_{i,a}(b_a; S_{i,1})$ denote the expectation inside the integral in Equation (5.15). This will be the expected number of clicks ad a would get for any bid, given other players report their click valuations. We interpolate this function by finding its values for a set of points on the interval $[\underline{x}_a, x_a]$. Since the maximum number of expected number of clicks an ad could get is 6 when $T = 6$, we split this interval into 6 equally length intervals and find the function values for the points splitting the interval. We operationalize that with $h_a = \frac{\hat{x}_a - \hat{\underline{x}}_a}{T}$, that indicates the length of each interval for ad a. As such, if $b_a \in (\hat{\underline{x}}_a + (i-1)h_a, \hat{\underline{x}}_a + ih_a]$, we can estimate this function as follows:

$$\hat{Q}_{i,a}(b_a, \hat{x}_{-a}; S_{i,1}) = \sum_{t=1}^{T} \sum_{g_t \in \mathcal{G}_T} q^d_{i,a,t}(\hat{x}_a + ih_a, \hat{x}_{-a}; S_{i,t}) P(g_t \mid \tau, \pi), \tag{5.48}$$

where g_t contains the states and ads prior to period t. In Equation (5.48), we approximate the function $\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right]$ with a step function, which makes the integral computation significantly easier. Now, we can determine the payments as follows:

$$\begin{aligned} e^d_{i,a}(x; S_{i,1}) = &\sum_{t=1}^{T} \sum_{g_t \in \mathcal{G}_T} \hat{x}_a \hat{y}_{i,t}(a; S_{i,t}) q^d_{i,a,t}(x; S_{i,t}) P(g_t \mid \tau, \pi) \\ &- h_a \sum_{t=1}^{T} \hat{Q}_{i,a}(\hat{x}_a + th_a, \hat{x}_{-a}; S_{i,1}), \end{aligned} \tag{5.49}$$

where the first component is the total value ad a receives from the session given the allocation, and the second component is the estimate of advertiser's surplus.

5.5.3 Evaluation

To evaluate the performance of each auction, we implement the direct method with the reward function with the revenue-maximizing objective. As such, we can define the expected revenue from session i with the horizon length T as $\rho^r_T(M; S_{i,1})$ for any mechanism M as follows:

$$\rho_T(M; S_{i,1}) = \mathbb{E}_{g_t \sim (\tau, M)} \left[\sum_{t=1}^{T} \beta^{t-1} r^r_t \right], \tag{5.50}$$

where the sequence g_t is determined by the joint distribution of transitions τ, and the allocation rule in mechanism M. Further, r^r_t denotes the reward function with the revenue-maximizing objective for the pair of state and action shown in the corresponding g_t. We can use our estimates for the distribution of transitions and evaluate $\rho_T(M; S_{i,1})$ for any mechanism M as follows:

$$\hat{\rho}_T(M; S_{i,1}) = \sum_{t=1}^{T} \sum_{g_t \in \mathcal{G}_T} \sum_{a \in \mathcal{A}_i} q^M_{i,a,t}(\hat{x}; S_{i,t}) \hat{R}^r_t(a; S_{i,t}) P(g_t \mid \tau, M) \tag{5.51}$$

This gives us the expected revenue that the publisher can extract from session i in the first T periods, when using the mechanism M.

5.6 Results

5.6.1 Gains from the Optimal Dynamic Auction

We start by presenting different session-level outcomes under both dynamic and static optimal auctions as well as the actual outcomes under the current auction – quasi-proportional auction. We evaluate the counterfactual outcomes under optimal auctions using our approach in §5.5. The sample of sessions that we use is the same as the one in [104]. In our evaluation, we only focus on the first six exposures in any session, as we set $T = 6$ as the length of the horizon. We present our results in Table 5.2.

As expected, both optimal auctions generate substantial gains over the current auction in terms of all session-level outcomes, except the expected advertisers' surplus. Three key components of the current auction explain its performance relative to the optimal auctions.

First, as discussed before, advertisers bid roughly half of their click valuation in this auction, lowering the revenues the publisher is able to extract. This, in turn, explains higher advertisers' surplus in the current auction as advertisers can extract huge informational rent by bidding half of their valuations. Second, the allocation mechanism is probabilistic in the current auction, which enables advertisers with low valuations to win impressions and clicks. This significantly reduces the efficiency or total surplus generated in the current auction compared to optimal auctions. Finally, as mentioned before, the ad-network does not incorporate sophisticated customization tools to show more relevant ads to users. This also greatly contributes to the gap between the current auction and other optimal auctions.

	Auctions		
	Dynamic	*Static*	*Current*
Expected Publisher's Revenue Per Session	126.49	124.49	31.63
Expected Total Surplus Per Session	143.45	140.96	66.60
Expected Advertisers' Surplus Per Session	16.96	16.47	34.97
Expected No. of Clicks Per Session	0.1577	0.1549	0.0823
Expected Session Length	3.24	3.25	3.12
Expected CTR	4.86%	4.76%	2.64%
No. of Users	1000	1000	1000
No. of Sessions	12,136	12,136	12,136

Table 5.2: Market outcomes under different auctions for a sequence size of 6

Next, we focus on our main goal in this chapter, and compare the session-level outcomes under dynamic and static optimal auctions. We find that using an optimal dynamic auction leads to 1.60% increase in the expected revenue per session, compared to the optimal static auction. This is of particular importance to the publishers and ad-networks as they usually use optimal static auctions (e.g., second-price auction).

While revenue is the main outcome most publishers and ad-networks care about, they do not usually want to achieve better revenue outcomes at the expense of efficiency and advertisers' surplus. Our results show that the optimal dynamic auction yields higher total and advertisers' surplus than the optimal static auction: the expected total surplus and advertisers' surplus grow by 1.77% and 3.00% respectively. Thus, advertisers also benefit from the revenue-optimal dynamic auction, as compared to the revenue-optimal static auction. Finally, we focus on the user-level outcomes – expected number of clicks per session and the session length. While the difference in the session length is very small, we find 1.83% increase in the expected number of clicks. This finding is in line with the findings in [104]. More importantly, it suggest that the revenue gains from the optimal dynamic auction likely come from the improvement in the match between ads and users, and not from the greater ability to extract informational rent from advertisers.

5.6.2 Distribution of Advertisers' Surplus

Our results in Table 5.2 indicate that the average advertisers' surplus increases by 3.00% in the optimal dynamic auction, as compared to the optimal static auction. We now explore how this surplus is distributed across advertisers. We use two main approaches to examine the distribution of advertisers' surplus: (1) number of advertisers who benefit from the optimal dynamic auction compared to the optimal static auction, and (2) the Herfindahl-Hirschman Index (HHI) which is a well-known measure to study market concentration.

We first focus on the distribution of advertisers' surplus under both auctions. We compute the log of each advertiser's surplus over all 12,136 sessions and present it in Figure 5.5. Interestingly, we find that only 3 out of 15 advertisers benefit from the optimal dynamic auction compared to the one with the optimal static auction. In contrast, 9 out of 15 advertisers prefer the optimal static auction. However, it is worth noting that these advertisers have a very small surplus in both cases and their gains are small in magnitude. Therefore, the average advertisers' surplus is higher under the optimal dynamic auction.

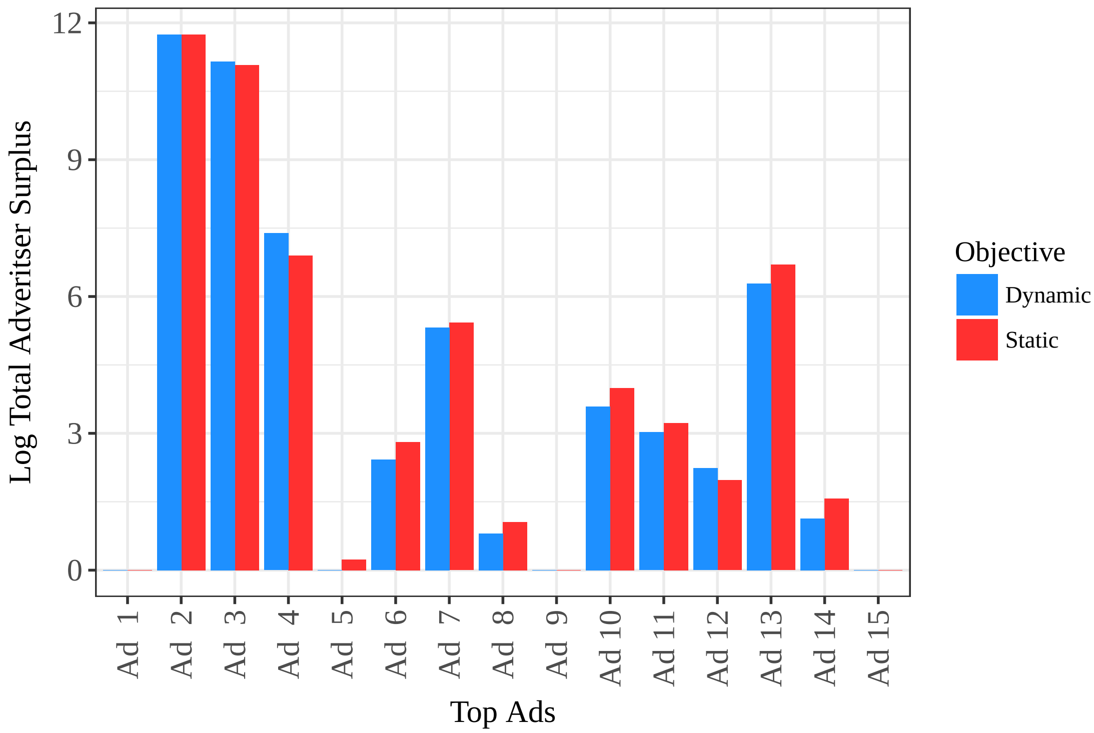

Figure 5.5: Log advertiser surplus under auctions with both dynamic and static objectives. The values are computed over all 12,136 session.

Next, we focus on the concentration of advertiser surplus in the market and calculate the Herfindahl-Hirschman Index (HHI) under both auctions. While more advertisers prefer the optimal static auction, the concentration in the optimal dynamic auction is lower: the Herfindahl-Hirschman Index (HHI) for the optimal dynamic auction is 0.5285, whereas it is 0.5398 for the optimal static auction. Although both auctions seem quite concentrated as they both allocate most impressions to one ad, optimal dynamic auction achieves a lower concentration by allocating more to the second-largest ad, thereby closing the gap between the two largest advertisers.

5.7 Conclusion

Mobile in-app advertising has grown exponentially over the last years. One important reason contributing to this growth is the publishers' ability to make adaptive interventions – using the time-varying information about the users within the session to personalize ad interventions for these users. [104] focuses on the match between ads and users as the main outcome of interest and shows that dynamic sequencing of ads improves the match outcome per session, compared static sequencing of ads. While dynamic sequencing of ads leads to improvements in the match outcome, it is not clear whether it can be linked to higher revenues, as advertisers can change their bids in response to the change in the allocation. In this chapter, we investigate the revenue gains from adopting a dynamic sequencing framework in a competitive environment, as opposed to a static sequencing framework. We propose a unified framework that contains two key components: (1) a theoretical framework that solves for the revenue-optimal auction with the dynamic objective, and (2) an empirical framework that estimates the counterfactual market outcomes under this auction. Our empirical framework comprises of structural estimation of advertisers' private valuations as well as personalized estimation of the click outcome given any pair of user-ad using machine learning methods. We apply our framework to large-scale data from the leading in-app ad-network of an Asian country. We demonstrate that adopting a dynamic sequencing framework increases the expected revenue by 1.60% compared to the static sequencing framework. We then show that the improved match outcome is the key factor in achieving these gains. Further, we explore other outcomes such as the total surplus (efficiency) and advertisers' surplus and document gains from the dynamic framework over the static framework. Thus, our optimal dynamic auction leads to improvement in all primary market outcomes.

This chapter makes three key contributions to the literature. First, from a methodological standpoint, we present a unified dynamic framework that incorporates both advertisers' and users' behavior and examines the market outcomes under the optimal dynamic auction. To our knowledge, this is the first work to empirically study the revenue gains from adopting an optimal dynamic auction. Second, we propose a methodological framework for structural estimation of quasi-proportional auction. Our framework adds to the literature on the non-parametric estimation of auctions and can easily be extended to any auction that employs a randomized allocation rule. Finally, from a substantive viewpoint, we show that dynamic sequencing of ads can lead to considerable gains in the publisher's revenue, over the existing auctions. This is particularly important, because the current practice in marketing ignores the gains from using a dynamic framework.

Our findings in this chapter have several implications for marketing practitioners. First, we propose a framework that helps publishers evaluate market outcomes under counterfactual auctions. Second, our substantive finding shows that adopting an optimal dynamic auction results in considerable gains in terms of three key market outcomes – revenue, total surplus, and advertisers' surplus. Thus, we expect this finding to inform the publishers' decision as to whether use an optimal dynamic auction. Further, we highlight that adaptive interventions require more careful consideration of players' incentives, when there are strategic players whose actions can affect market outcomes. Thus, our framework has implications for the practitioners and policy makers that increasingly use adaptive interventions in the presence of strategic players.

Nevertheless, there remains some limitations in our study that serve as excellent avenues for future research. First, our optimal dynamic auction involves a bit complex allocation and payment rules. As a result, it may be possible that advertisers will not behave in the expected way. Finding auctions with easier rules that achieve approximately the same outcomes would be an interesting avenue for future study. Second, our structural estimation framework assumes that advertisers are risk-neutral and aware of how their bid affects their probability of winning. However, advertisers' behavior would be different if they are risk-averse. Future research can study risk aversion in probabilistic auctions and empirically identify the risk elements in advertisers' utility function. Finally, while this chapter provides counterfactual estimates optimal auctions, we do not run these auctions in the field. An interesting future research is to test these auction in a field setting and examine market outcomes.

BIBLIOGRAPHY

[1] Alessandro Acquisti, Curtis Taylor, and Liad Wagman. The Economics of Privacy. *Journal of Economic Literature*, 54(2):442–92, 2016.

[2] Wilfred Amaldoss, Kinshuk Jerath, and Amin Sayedi. Keyword Management Costs and "Broad Match" in Sponsored Search Advertising. *Marketing Science*, 35(2):259–274, 2015.

[3] Nick Arnosti, Marissa Beck, and Paul Milgrom. Adverse selection and auction design for internet display advertising. *American Economic Review*, 106(10):2852–66, 2016.

[4] Neeraj Arora and Ty Henderson. Embedded premium promotion: Why it works and how to make it more effective. *Marketing Science*, 26(4):514–531, 2007.

[5] Arash Asadpour, Mohammad Hossein Bateni, Kshipra Bhawalkar, and Vahab Mirrokni. Concise bid optimization strategies with multiple budget constraints. In *International Conference on Web and Internet Economics*, pages 263–276. Springer, 2014.

[6] Susan Athey and Philip A Haile. Nonparametric approaches to auctions. *Handbook of econometrics*, 6:3847–3965, 2007.

[7] Susan Athey and Denis Nekipelov. A structural model of sponsored search advertising auctions. In *Sixth Ad Auctions Workshop*, volume 15, 2010.

[8] Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

[9] Automated. Ad Refresh - A Beginner's Guide For Publishers, 2019.

[10] Patrick Bajari, Stephanie Houghton, and Steven Tadelis. Bidding for incomplete contracts: An empirical analysis of adaptation costs. *American Economic Review*, 104(4):1288–1319, 2014.

[11] David P Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and policy*, 1(3):267–302, 1984.

[12] Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966.

[13] Dirk Bergemann and Alessandro Bonatti. Targeting in Advertising Markets: Implications for Offline Versus Online Media. *The RAND Journal of Economics*, 42(3):417–443, 2011.

[14] Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.

[15] Thomas Blake, Chris Nosko, and Steven Tadelis. Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment. *Econometrica*, 83(1):155–174, 2015.

[16] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*, pages 531–540. ACM, 2007.

[17] Leo Breiman. Arcing Classifier. *The Annals of Statistics*, 26(3):801–849, 1998.

[18] Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. *Classification and Regression Trees*. CRC press, 1984.

[19] Sandra Campo, Isabelle Perrigne, and Quang Vuong. Asymmetry in first-price auctions with affiliated private values. *Journal of Applied Econometrics*, 18(2):179–207, 2003.

[20] L Elisa Celis, Gregory Lewis, Markus Mobius, and Hamid Nazerzadeh. Buy-It-Now or Take-a-Chance: Price Discrimination Through Randomized Auctions. *Management Science*, 60(12):2927–2948, 2014.

[21] Olivier Chapelle, Eren Manavoglu, and Romer Rosales. Simple and Scalable Response Prediction for Display Advertising. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4):61, 2015.

[22] Tianqi Chen and Carlos Guestrin. Xgboost: A Scalable Tree Boosting System. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794. ACM, 2016.

[23] Hana Choi and Carl F Mela. Online marketplace advertising. *Available at SSRN*, 2016.

[24] David Roxbee Cox. Planning of experiments. 1958.

[25] Hannes Datta, George Knox, and Bart J Bronnenberg. Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Science*, 37(1):5–21, 2017.

[26] Alexandre De Corniere and Romain De Nijs. Online Advertising and Privacy. *The RAND Journal of Economics*, 47(1):48–72, 2016.

[27] Jean-Pierre Dubé and Sanjog Misra. Scalable price targeting. Technical report, National Bureau of Economic Research, 2017.

[28] Miroslav Dudík, Dumitru Erhan, John Langford, Lihong Li, et al. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):485–511, 2014.

[29] Daria Dzyabura and Hema Yoganarasimhan. Machine Learning and Marketing. In *Handbook of Marketing Analytics*. Edward Elgar Publishing, 2018.

[30] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.

[31] Jim Edwards. Apple Has Quietly Started Tracking iPhone Users Again, And It's Tricky To Opt Out, 2012.

[32] Ariel Edwards-Levy and Dana Liebelson. Even Trump Voters Hate This Bill He Just Signed, 2017.

[33] eMarketer. Mobile In-App Ad Spending , 2018.

[34] eMarketer. US Ad Spending , 2018.

[35] eMarketer. Digital Ad Spending 2019 , 2019.

[36] eMarketer. Digital Ad Spending 2019 , 2019.

[37] eMarketer. Time Spent with Mobile, US , 2019.

[38] eMarketer. Smartphone Users and Penetration in Worldwide, 2020.

[39] eMarketer. US Total Media Ad Spending, by Media, 2020.

[40] Jerome Friedman, Trevor Hastie, Robert Tibshirani, et al. Additive Logistic Regression: A Statistical View of Boosting (with Discussion and a Rejoinder by the Authors). *The Annals of Statistics*, 28(2):337–407, 2000.

[41] Jerome H Friedman. Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics*, pages 1189–1232, 2001.

[42] Jerome H Friedman. Stochastic Gradient Boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, 2002.

[43] Jérémie Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.

[44] Arpita Ghosh and Mohammad Mahdian. Externalities in online advertising. In *Proceedings of the 17th international conference on World Wide Web*, pages 161–168. ACM, 2008.

[45] Arpita Ghosh and Amin Sayedi. Expressive auctions for externalities in online advertising. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pages 371–380, New York, NY, USA, 2010. ACM.

[46] David Godes and Dina Mayzlin. Using online conversations to study word-of-mouth communication. *Marketing science*, 23(4):545–560, 2004.

[47] Avi Goldfarb. What is Different About Online Advertising? *Review of Industrial Organization*, 44(2):115–129, 2014.

[48] Avi Goldfarb and Catherine Tucker. Advertising Bans and the Substitutability of Online and Offline Advertising. *Journal of Marketing Research*, 48(2):207–227, 2011.

[49] Avi Goldfarb and Catherine Tucker. Online Display Advertising: Targeting and Obtrusiveness. *Marketing Science*, 30(3):389–404, 2011.

[50] Avi Goldfarb and Catherine Tucker. Search Engine Advertising: Channel Substitution When Pricing Ads to Context. *Management Science*, 57(3):458–470, 2011.

[51] Avi Goldfarb and Catherine E Tucker. Privacy Regulation and Online Advertising. *Management science*, 57(1):57–71, 2011.

[52] Brett R Gordon, Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky. A comparison of approaches to advertising measurement: Evidence from big field experiments at facebook. *Marketing Science*, 38(2):193–225, 2019.

[53] Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.

[54] Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. *Econometrica*, 77(4):1193–1227, 2009.

[55] Seungyeop Han, Jaeyeon Jung, and David Wetherall. A study of third-party tracking by mobile apps in the wild. *Univ. Washington, Tech. Rep. UW-CSE-12-03-01*, 2012.

[56] David A Harrison and Katherine J Klein. What's the difference? diversity constructs as separation, variety, or disparity in organizations. *Academy of management review*, 32(4):1199–1228, 2007.

[57] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. The Elements of Statistical Learning. *NY Springer*, 2001.

[58] John R Hauser, Guilherme Liberali, and Glen L Urban. Website morphing 2.0: Switching costs, partial exposure, random exit, and when to morph. *Management science*, 60(6):1594–1616, 2014.

[59] John R Hauser, Olivier Toubia, Theodoros Evgeniou, Rene Befurt, and Daria Dzyabura. Disjunctions of Conjunctions, Cognitive simplicity, and Consideration Sets. *Journal of Marketing Research*, 47(3):485–496, 2010.

[60] John R Hauser, Glen L Urban, Guilherme Liberali, and Michael Braun. Website morphing. *Marketing Science*, 28(2):202–223, 2009.

[61] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, et al. Practical Lessons from Predicting Clicks on Ads at Facebook. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, pages 1–9. ACM, 2014.

[62] Günter J Hitsch and Sanjog Misra. Heterogeneous treatment effects and optimal targeting policy evaluation. *Available at SSRN 3111957*, 2018.

[63] Stephen J Hoch, Eric T Bradlow, and Brian Wansink. The variety of an assortment. *Marketing Science*, 18(4):527–546, 1999.

[64] David W Hosmer Jr, Stanley Lemeshow, and Rodney X Sturdivant. *Applied Logistic Regression*, volume 398. John Wiley & Sons, 2013.

[65] Patrick Hummel and R Preston McAfee. When Does Improved Targeting Increase Revenue? *ACM Transactions on Economics and Computation (TEAC)*, 5(1):4, 2016.

[66] Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.

[67] Jacob Jacoby. Information load and decision quality: Some contested issues. *Journal of Marketing Research*, pages 569–573, 1977.

[68] Jacob Jacoby. Perspectives on information overload. *Journal of consumer research*, 10(4):432–435, 1984.

[69] Przemysław Jeziorski and Sridhar Moorthy. Advertiser prominence effects in search advertising. *Management Science*, 2017.

[70] Garrett A Johnson, Randall A Lewis, and Elmar I Nubbemeyer. The online display ad effectiveness funnel & carryover: A meta-study of predicted ghost ad experiments. 2016.

[71] Garrett A Johnson, Randall A Lewis, and Elmar I Nubbemeyer. Ghost ads: Improving the economics of measuring online ad effectiveness. *Journal of Marketing Research*, 54(6):867–884, 2017.

[72] Barbara E Kahn and Brian Wansink. The influence of assortment structure on perceived variety and consumption quantities. *Journal of consumer research*, 30(4):519–533, 2004.

[73] Sham M Kakade, Ilan Lobel, and Hamid Nazerzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013.

[74] Wreetabrata Kar, Viswanathan Swaminathan, and Paulo Albuquerque. Selection and ordering of linear online video ads. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, pages 203–210, New York, NY, USA, 2015. ACM.

[75] David Kempe and Mohammad Mahdian. A cascade model for externalities in sponsored search. In *International Workshop on Internet and Network Economics*, pages 585–596. Springer, 2008.

[76] Jaehwan Kim, Greg M Allenby, and Peter E Rossi. Modeling consumer demand for variety. *Marketing Science*, 21(3):229–250, 2002.

[77] Jason Kint. Opinion: Europe's Strict New Privacy Rules Are Scary but Right, 2017.

[78] Elena Krasnokutskaya. Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies*, 78(1):293–327, 2011.

[79] Vijay Krishna. *Auction theory*. Academic press, 2009.

[80] Sébastien Lahaie, David M Pennock, Amin Saberi, and Rakesh V Vohra. Sponsored search auctions. *Algorithmic game theory*, pages 699–716, 2007.

[81] T. Lattimore and Cs. Szepesvári. *Bandit Algorithms*. Cambridge University Press (to be printed soon), August 2020.

[82] Jonathan Levin and Paul Milgrom. Online Advertising: Heterogeneity and Conflation in Market Design. *The American Economic Review*, 100(2):603–607, 2010.

[83] Randall A Lewis and Justin M Rao. The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics*, 130(4):1941–1973, 2015.

[84] Hairong Li, Steven M Edwards, and Joo-Hyun Lee. Measuring the intrusiveness of advertisements: Scale development and validation. *Journal of advertising*, 31(2):37–47, 2002.

[85] Hongshuang Li and PK Kannan. Attributing Conversions in a Multichannel Online Marketing Environment: An Empirical Model and a Field Experiment. *Journal of Marketing Research*, 51(1):40–56, 2014.

[86] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670. ACM, 2010.

[87] Tong Li, Isabelle Perrigne, and Quang Vuong. Structural estimation of the affiliated private value auction model. *RAND Journal of Economics*, pages 171–193, 2002.

[88] Xiaoliang Ling, Weiwei Deng, Chen Gu, Hucheng Zhou, Cui Li, and Feng Sun. Model ensemble for click prediction in bing search ads. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 689–698, 2017.

[89] Tyler Lu, Dávid Pál, and Martin Pál. Contextual multi-armed bandits. In *Proceedings of the Thirteenth international conference on Artificial Intelligence and Statistics*, pages 485–492, 2010.

[90] Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. Offline policy evaluation across representations with applications to educational games. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1077–1084. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[91] Leigh McAlister. A dynamic attribute satiation model of variety-seeking behavior. *Journal of Consumer Research*, 9(2):141–150, 1982.

[92] Daniel F McCaffrey, Beth Ann Griffin, Daniel Almirall, Mary Ellen Slaughter, Rajeev Ramchand, and Lane F Burgette. A tutorial on propensity score estimation for multiple treatments using generalized boosted models. *Statistics in medicine*, 32(19):3388–3414, 2013.

[93] H Brendan McMahan, Gary Holt, David Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, et al. Ad Click Prediction: A View from the Trenches. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1222–1230. ACM, 2013.

[94] Vahab Mirrokni, S Muthukrishnan, and Uri Nadav. Quasi-Proportional Mechanisms: Prior-Free Revenue Maximization. In *Latin American Symposium on Theoretical Informatics*, pages 565–576. Springer, 2010.

[95] Sendhil Mullainathan and Jann Spiess. Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106, 2017.

[96] Roger B Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[97] Roger B Myerson. Multistage games with communication. *Econometrica: Journal of the Econometric Society*, pages 323–358, 1986.

[98] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. Just-in-time adaptive interventions (jitais) in mobile health: key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine*, 52(6):446–462, 2017.

[99] Sridhar Narayanan and Kirthi Kalyanam. Position effects in search advertising and their moderators: A regression discontinuity approach. *Marketing Science*, 34(3):388–407, 2015.

[100] Michael Ostrovsky and Michael Schwarz. Reserve Prices in Internet Advertising Auctions: A Field Experiment. Working Paper, 2016.

[101] David C Parkes and Satinder P Singh. An mdp-based approach to online mechanism design. In *Advances in neural information processing systems*, pages 791–798, 2004.

[102] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.

[103] Rik Pieters, Edward Rosbergen, and Michel Wedel. Visual Attention to Repeated Print Advertising: A Test of Scanpath Theory. *Journal of Marketing Research*, pages 424–438, 1999.

[104] Omid Rafieian. Optimizing user engagement through adaptive ad sequencing. 2020.

[105] Omid Rafieian. Revenue-optimal dynamic auctions for adaptive ad sequencing. 2020.

[106] Omid Rafieian and Hema Yoganarasimhan. How does variety of previous ads influence consumer's ad response? 2020.

[107] Omid Rafieian and Hema Yoganarasimhan. Targeting and Privacy in Mobile Advertising. *Marketing Science*, 2020.

[108] Rebecca K Ratner and Barbara E Kahn. The impact of private versus public consumption on variety-seeking behavior. *Journal of Consumer Research*, 29(2):246–257, 2002.

[109] Rebecca K Ratner, Barbara E Kahn, and Daniel Kahneman. Choosing less-preferred experiences for the sake of variety. *Journal of consumer research*, 26(1):1–15, 1999.

[110] Joseph P Redden. Reducing Satiation: The Role of Categorization Level. *Journal of Consumer Research*, 34(5):624–634, 2008.

[111] John G Riley and William F Samuelson. Optimal Auctions. *The American Economic Review*, 71(3):381–392, 1981.

[112] Michael H Riordan and David EM Sappington. Information, incentives, and organizational mode. *The Quarterly Journal of Economics*, 102(2):243–263, 1987.

[113] Lorenzo Rosasco, Ernesto De Vito, Andrea Caponnetto, Michele Piana, and Alessandro Verri. Are loss functions all the same? *Neural Computation*, 16(5):1063–1076, 2004.

[114] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

[115] Matt Rosoff. The research firm that once thought Microsoft would beat the iPhone has given up on Windows Phone , 2015.

[116] Donald B Rubin. Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American Statistical Association*, 75(371):591–593, 1980.

[117] Oliver J Rutz and Randolph E Bucklin. From generic to branded: A model of spillover in paid search advertising. *Journal of Marketing Research*, 48(1):87–102, 2011.

[118] Navdeep S Sahni. Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quantitative Marketing and Economics*, 13(3):203–247, 2015.

[119] Navdeep S Sahni. Advertising spillovers: Evidence from online field experiments and implications for returns on advertising. *Journal of Marketing Research*, 53(4):459–478, 2016.

[120] Maher Said. Auctions with dynamic populations: Efficiency and revenue maximization. *Journal of Economic Theory*, 147(6):2419–2438, 2012.

[121] Amin Sayedi. Real-time bidding in online display advertising. *Marketing Science*, 37(4):553–568, 2018.

[122] David W Schumann, Richard E Petty, and D Scott Clemons. Predicting the effectiveness of different strategies of advertising variation: A test of the repetition-variation hypotheses. *Journal of Consumer Research*, 17(2):192–202, 1990.

[123] Eric M Schwartz, Eric T Bradlow, and Peter S Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017.

[124] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, July 1948.

[125] Bradley Shapiro. Advertising in health insurance markets. *Browser Download This Paper*, 2016.

[126] Zhen Sun, Milind Dawande, Ganesh Janakiraman, and Vijay Mookerjee. Not just a fad: Optimal sequencing in mobile in-app advertising. *Information Systems Research*, 28(3):511–528, 2017.

[127] Georgios Theocharous, Philip S Thomas, and Mohammad Ghavamzadeh. Personalized ad recommendation systems for life-time value optimization with guarantees. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[128] Olivier Toubia, Theodoros Evgeniou, and John Hauser. Optimization-based and Machine-learning Methods for Conjoint Analysis: Estimation and Question Design. In Anders Gustafsso, Andreas Herrmann, and Frank Huber, editors, *Conjoint Measurement: Methods and Applications*, pages 231–258. Springer, 4 edition, 2007.

[129] Catherine E Tucker. Social Networks, Personalized Advertising, and Privacy Controls. *Journal of Marketing Research*, 51(5):546–562, 2014.

[130] Glen L Urban, Guilherme Liberali, Erin MacDonald, Robert Bordley, and John R Hauser. Morphing banner advertising. *Marketing Science*, 33(1):27–46, 2013.

[131] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.

[132] William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of finance*, 16(1):8–37, 1961.

[133] Gustavo Vulcano, Garrett Van Ryzin, and Costis Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 48(11):1388–1407, 2002.

[134] Jeffrey M Wooldridge. *Econometric analysis of cross section and panel data*. MIT press, 2010.

[135] Song Yao and Carl F Mela. A Dynamic Model of Sponsored Search Advertising. *Marketing Science*, 30(3):447–468, 2011.

[136] Hema Yoganarasimhan. Estimation of beauty contest auctions. *Marketing Science*, 35(1):27–54, 2015.

[137] Hema Yoganarasimhan. Search personalization using machine learning. *Management Science*, 66(3):1045–1070, 2020.

[138] Tong Zhang, Bin Yu, et al. Boosting with Early Stopping: Convergence and Consistency. *The Annals of Statistics*, 33(4):1538–1579, 2005.

Appendix A
APPENDIX FOR CHAPTER 2

A.1 Propensity Scores and Unconfoundedness Assumption

Our empirical strategy relies on randomization in the allocation of ads, which gives us the unconfoundedness assumption and allows us to perform targeting counterfactuals. As outlined in Chapter 2, one diagnostic test for unconfoundedness assumption is to assess covariate balance. For this, we first need to estimate propensity scores, and then examine whether adjusting for these propensity scores achieves covariate balance. We do the former in §A.1.1 and the latter in §A.1.2.

A.1.1 Estimation of Propensity Scores

We now propose our approach to estimate propensity scores. First, recall the allocation rule presented in Equation (3.1):

$$\pi_{ia} = \frac{b_a q_a}{\sum_{j \in \mathcal{A}_i} b_j q_j}$$

Our goal is to estimate π_{ia} for all top 37 ads in each impression i. In our filtering procedure, we identify ads with zero propensity of being shown and filter them out of the competition. However, our filtering procedure does not estimate the magnitude of the propensity score for ads with non-zero propensity scores. We therefore discuss our procedure to estimate propensity scores below.

The outcome and covariates for our propensity score model are listed below:

1. *Outcome*: The main outcome of interest is the ad shown. This is a categorical variable that has 38 classes, i.e., $a_i \in A = \{a^{(1)}, a^{(2)}, \ldots, a^{(37)}, a^{(s)}\}$, where $a^{(1)}, a^{(2)}, \ldots, a^{(37)}$ refer to the identities of the top 37 ads, and all the smaller ads are grouped into one

category denoted by $a^{(s)}$.

2. *Covariates*: The set of covariates that we control for are:

   - Filtering indicators for each ad $e_{i,a}$. This controls for whether the ad a is in $\mathcal{A}_i$.
   - All the targeting variables, including province, app category, hour of the day, smartphone brand, connectivity type, and MSP.
   - Exact location (latitude and longitude), to control for any geographical targeting beyond province-level targeting.
   - Minute-level time of the day, to control for any shocks or changes in the quality scores and bids in the system. Note that we do not expect bids or quality scores to change within the observation window. Nevertheless, this control simply ensures that there are no biases even if they did due to system bugs etc.

We use a multi-class XGBoost model to regress our outcome of interest on the set of covariates. We use multi-class logarithmic loss as the evaluation metric, which is defined as follows:

$$\mathcal{L}^{mlog\ loss}(\hat{\pi}, \mathbf{a}) = -\frac{1}{N}\sum_{i=1}^{N}\sum_{j \in A} \mathbb{1}(a_i = a^{(j)}) \log(\hat{\pi}_{i,j}), \tag{A.1}$$

where the summation over j refers to all the outcome classes that we have. Using a softmax objective, we can then estimate the propensity scores. Figure A.1 shows the distribution of propensity scores for top eight ads, using box plots. As shown in this figure, the average propensity of winning for each ad is below 0.25. Further, there is no disproportionately high propensity of winning: the highest is slightly above 0.50. We also notice that Ad 5 has higher median propensity score than Ad 2, even though the total number of impressions is higher for Ad 2. This is because Ad 5 is targeting more narrowly, and is not available for all the impressions.

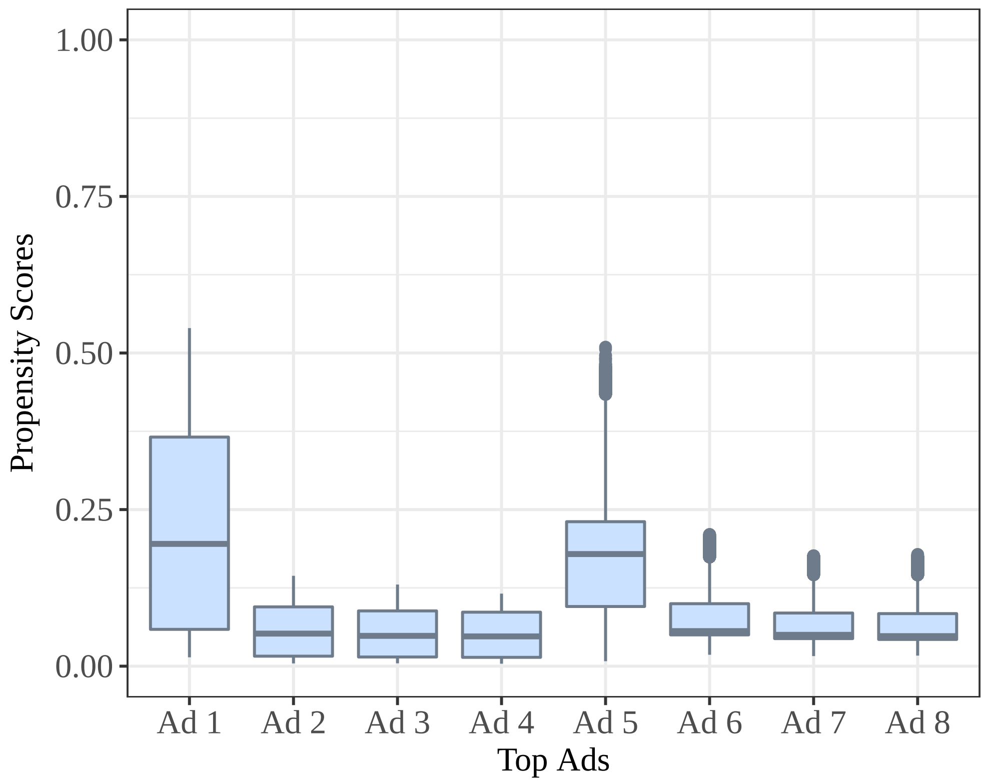

Figure A.1: Distribution of propensity scores for top 8 ads

A.1.2 Assessing Covariate Balance Using Inverse Probability Weighting

Now, we use our estimated propensity scores to assess covariate balance across different ads. We begin by describing the balance for a simple case in which we have a fully randomized experiment with one treatment and one control group. We then extend the same notion to the case where we want to assess balance using propensity scores across multiple treatment arms.

Covariate Balance in a Fully Randomized Experiment with a Binary Treatment

To assess the covariate balance in a fully randomized experiment with two treatment arms, we simply need to show that the distribution of covariate Z is the same across treated and control samples. As such, we rely on standardized bias, which is the absolute mean difference in covariate Z across the treatment and control groups, divided by the standard deviation of the Z in the pooled sample:

$$SB(Z) = \frac{|\bar{Z}_{\text{treated}} - \bar{Z}_{\text{control}}|}{\sigma_Z}, \tag{A.2}$$

where $\bar{Z}_{\text{treated}}$ and $\bar{Z}_{\text{control}}$ are the mean value of covariate Z for the treated and control groups, and σ_Z is the standard deviation of Z for the pooled sample.

Clearly, the lower the standardized bias is, we have greater balance in covariate Z across treatment and control groups. In practice, different cutoffs are considered acceptable by the researchers. A common cutoff is 0.20 – that is, a standardized bias greater than 0.20 is taken as evidence for imbalance [92]. In this work, we follow this norm and set our cutoff to be 0.20.

Covariate Balance Using Inverse Propensity Weights for Multiple Treatments

Our empirical setting is more complicated than a fully randomized experiment with a binary treatment because of two reasons: (1) ad assignment is not fully exogenous, but a function of observed covariates, so covariate balance may not be satisfied for the unweighted distribution of covariates, and (2) unlike a binary treatment, we have multiple treatment arms (ads), so the balance is not simply the difference between two groups. To address the first issue, we need to use our estimated propensity scores to weight our original sample such that an observation with a lower propensity score gets a larger weight to represent what the sample would have been if the platform had run a fully randomized experiment. For the second issue, we follow the approach in [92] and compare the weighted mean of covariate Z when assigned to ad with the population mean of that covariate and check for the balance by measuring the standardized bias. In principle, if weights are accurate, the weighted mean of covariate Z when assigned to ad a must be the same as the population mean. As such, we define the weight-adjusted standardized bias for covariate Z when assigned to ad a as follows:

$$SB_a^{\hat{\pi}}(Z) = \frac{|\bar{Z}_a^{\hat{\pi}} - \bar{Z}|}{\sigma_Z}, \tag{A.3}$$

where superscript $\hat{\pi}$ indicates the propensity score estimates used to weight the samples, and $\bar{Z}_a^{\hat{\pi}}$ is defined as follows:

$$\bar{Z}_a^{\hat{\pi}} = \frac{\sum_{i=1}^{N} \frac{\mathbb{1}(a_i=a)}{\hat{\pi}_{ia}} Z_i}{\sum_{i=1}^{N} \frac{\mathbb{1}(a_i=a)}{\hat{\pi}_{ia}}}, \tag{A.4}$$

where Z_i is the value of covariate Z in impression i, the term $\mathbb{1}(a_i = a)$ indicates whether ad a is actually shown in that impression, and the weights are simply the inverse propensity scores.

We can then calculate the weight-adjusted standardized bias measure for each ad a and define the standardized bias for variable Z as follows:

$$SB^{\hat{\pi}}(Z) = \max_a SB_a^{\hat{\pi}}(Z), \tag{A.5}$$

where by construction, if $SB^{\hat{\pi}}(Z)$ satisfies the balance criteria, it means that each $SB_a^{\hat{\pi}}(Z)$ satisfies this criteria. To have a baseline of how imbalanced covariates were before adjusting for inverse propensity weights, we define the unweighted measure of standardized bias for covariate Z when assigned to ad a as follows:

$$SB(Z) = \max_a SB_a(Z) = \max_a \frac{|\bar{Z}_a - \bar{Z}|}{\sigma_Z}, \tag{A.6}$$

Comparing $SB(Z)$ (unweighted standardized bias) and $SB^{\hat{\pi}}(Z)$ (weight-adjusted standardized bias) allows us to see how adjusting for inverse propensity scores helped balance covariates.

Now, we focus on all the targeting variables in our sample and calculate both unweighted and weight-adjusted measures of standardized bias. Targeting variables are the main candidates that could potentially fail the balance test because advertisers can target these variables. Since all these variables are categorical, for each targeting variable (Z), we need to assess balance for the dummy variable for each separate category within that targeting variable.

For a given targeting variable (e.g., province), let $\mathcal{Z}$ denote the set of dummy variables corresponding to each sub-category within that targeting variable. We define the standardized bias as the maximum standardized bias across all its sub-categories. As such, for both unweighted and weight-adjusted methods, we have:

$$SB^{\hat{\pi}}(\mathcal{Z}) = \max_{Z \in \mathcal{Z}} SB^{\hat{\pi}}(Z) = \max_{Z \in \mathcal{Z}} \max_{a} SB^{\hat{\pi}}_a(Z)$$
$$SB(\mathcal{Z}) = \max_{Z \in \mathcal{Z}} SB(Z) = \max_{Z \in \mathcal{Z}} \max_{a} SB_a(Z)$$

Figure A.2 shows both unweighted and weight-adjusted measures of standardized bias for all of our targeting variables (on the filtered test data sample, which is used for the counterfactual analysis). Notice that all six targeting variables are initially imbalanced in the actual data. In particular, province, hour of the day, and MSP are the variables with the largest magnitude of imbalance. However, after adjusting for inverse propensity scores, we find that all measures of standardized bias are below the 0.2 cutoff. Thus, we have covariate balance.

What is striking about Figure A.2 is the very large measures of unweighted standardized bias for some targeting variables. This is due to the presence of some ads with very low propensity scores and sparse categories within certain targeting variables. For example, there are some provinces in our data that constitute less than 1% of our observations. Likewise, there are some ads with less than 1% share of observations. As a result, there are very few observations where these ads are shown in these provinces. This may result in substantial imbalance in the original data.

Next, we focus only on ads that are shown at least in 1% of our sample and make the same balance plot in Figure A.3. In this case, the measures of standardized bias are significantly lower compared to Figure A.2, especially the unweighted measures. Interestingly, we find that four out of six covariates are balanced in the unweighted sample: app category, connectivity type, MSP, and smartphone brand. This is because these variables are not widely used for targeting and their sub-categories have a sufficiently large number of observations. The only two variables that are imbalanced in the original sample are state and hour of the day, most likely because these are the variables are most used for targeting among advertisers. We also

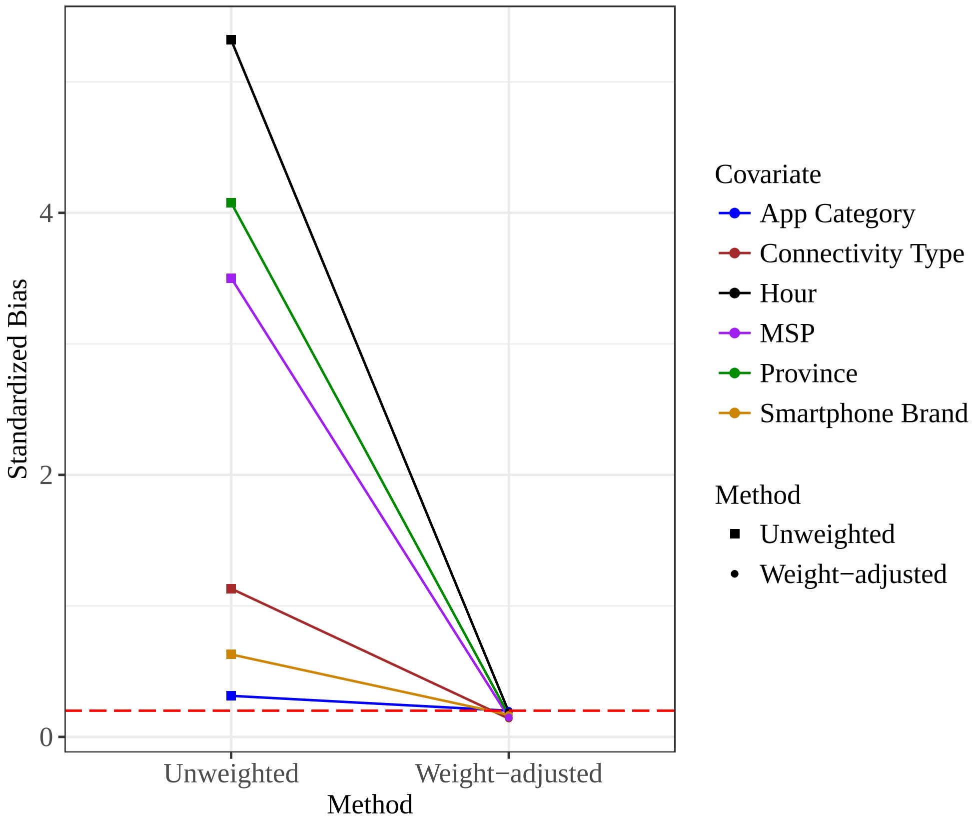

Figure A.2: Weight-adjusted and unweighted measures of standardized bias for all targeting variables (balance plot). The dotted red line shows the 0.2 cut-off.

find that in both cases, hour of the day has the highest standardized bias (even though we achieve balance at the 0.2 cutoff in both cases). This may also be due to the unavailability of some ads because of reasons other than targeting such as budget exhaustion.

In sum, the above analysis presents strong empirical evidence for covariate balance in our data after adjusting for propensity scores. This suggests the unconfoundedness assumption is valid in our setting.

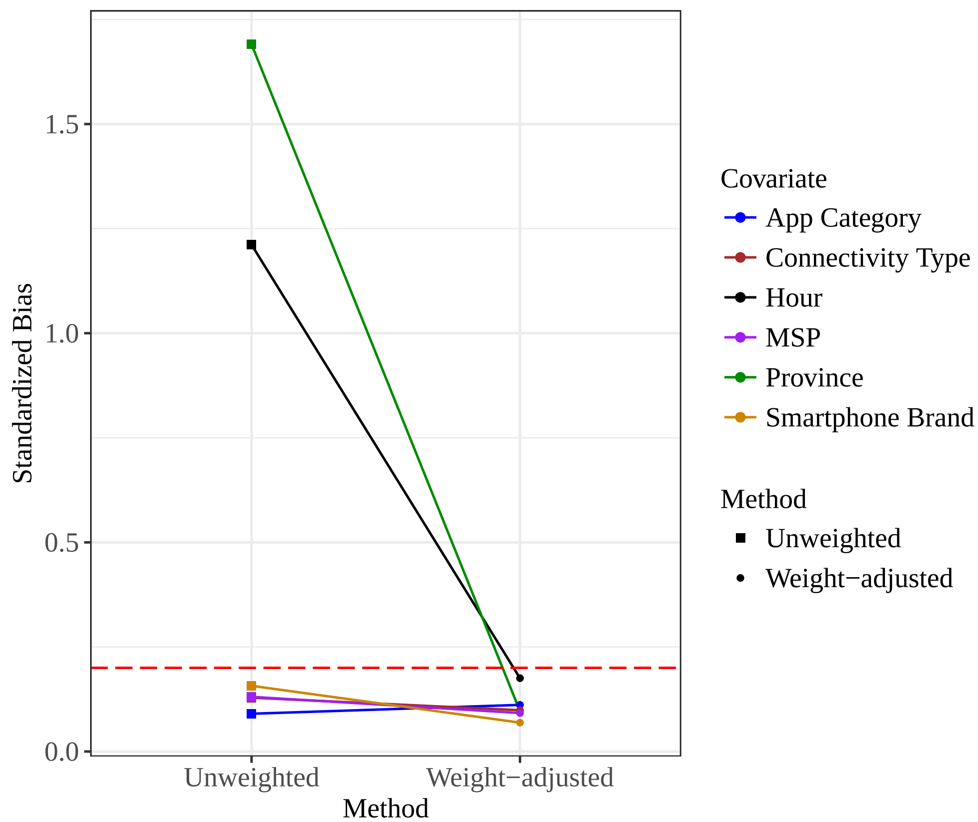

Figure A.3: Weight-adjusted and unweighted measures of standardized bias for all targeting variables, when we only consider ads with over 1% share of total number of observations in the filtered sample in test data (balance plot). The dotted red line shows the 0.2 cut-off.

A.2 Feature Generation Framework

We need a set of informative features that can help accurately predict the probability of click for a given impression. Given our research agenda, our features should be able to capture the contextual and behavioral information associated with an impression over different lengths of history preceding the impression (long-term, short-term, and session-level). Therefore, we adopt the main ideas from the functional feature generation framework proposed by [137]. In §A.2.1, we describe the inputs to the feature functions, in §A.2.2, we describe the functions, and finally in §A.2.3, we present the full table of features and their classification.

A.2.1 Inputs for Feature Functions

Each impression i in our training, validation, and test data can be uniquely characterized by the following four variables:

- the ad a_i which was shown in the impression, where $a_i \in A = \{a^{(1)}, a^{(2)}, \dots, a^{(37)}, a^{(s)}\}$. The elements $a^{(1)}, a^{(2)}, \dots, a^{(37)}$ refer to the identities of the top 37 ads, whereas all the smaller ads are grouped into one category, which is denoted by as $a^{(s)}$. A denotes the set of all ads.
- the app p_i within which it occurred, where $p_i \in P = \{p^{(1)}, p^{(2)}, \dots, p^{(50)}, p^{(s)}\}$. $p^{(1)}$ through $p^{(50)}$ refer to the top 50 apps, and $p^{(s)}$ refers to the category of all smaller apps (which we do not track individually).
- the hour of the day during which the impression occurred (denoted by t_i). t_i can take 24 possible values ranging from 1 through 24 and the set of these values is: $T = \{1, 2, \dots, 24\}$.
- the user u_i who generated the impression, where $u_i \in U = \{u^{(1)}, \dots, u^{(728{,}340)}\}$. U is thus the full sample of users in the training, validation, and test data.

We also have an indicator for whether the impression was clicked or not (denoted by y_i). In addition, we have a history of impressions and clicks observed in the system prior to this impression.

To generate a set of features for a given impression i, we use feature functions that take some inputs at the impression level and output a corresponding feature for that impression. Our feature functions are typically of the form $F(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie})$ (though some functions may have fewer inputs). We now define each of these inputs:

1. *Ad related information*: The first input θ_{ia} captures the ad related information for impression i. It can take two possible values: $\theta_{ia} \in \Theta_{ia} = \{\{a_i\}, A\}$. If $\theta_{ia} = \{a_i\}$, then the feature is specific to the ad shown in the impression. Instead, if $\theta_{ia} = A$, then it means that the feature is not ad-specific and is aggregated over all possible ads.

2. *Contextual information*: The second and third inputs θ_{ip} and θ_{it} capture the context (where and when) for impression i.

   (a) $\theta_{ip} \in \Theta_{ip} = \{\{p_i\}, P\}$ can take two possible values. If $\theta_{ip} = \{p_i\}$, then the feature is specific to the app where the impression was shown. Instead, if $\theta_{ip} = P$, the feature is aggregated over all apps.

   (b) $\theta_{it} \in \Theta_{it} = \{\{t_i\}, T\}$ characterizes the time-related information for impression i. If $\theta_{it} = \{t_i\}$, then the feature is specific to the hour during which the impression occurred. Instead, if $\theta_{it} = T$, the feature is aggregated over all hours of the day.

3. *Behavioral information*: The fourth input, $\theta_{iu} \in \Theta_{iu} = \{\{u_i\}, U\}$, captures the behavioral information for impression i. If $\theta_{iu} = \{u_i\}$, the feature is specific to user u_i who generated the impression. Instead, if $\theta_{iu} = U$, the feature is aggregated over all users.

4. *History*: The last two inputs, η_{ib} and η_{ie}, capture the history over which the function is aggregated. η_{ib} denotes the **b**eginning or starting point of the history and η_{ie} denotes the **e**nd-point of the history.

   (a) $\eta_{ib} \in \mathcal{H}_{ib} = \{l, s, o_i\}$ can take three possible values which we discuss below:

   - $\eta_{ib} = l$ denotes **l**ong-term history, i.e., the starting point of the history is September 30, 2015, the date from which data are available.
   - $\eta_{ib} = s$ denotes **s**hort-term history, i.e., only data on or after October 25, 2015, are considered.
   - $\eta_{ib} = o_i$ denotes **o**ngoing session-level history, i.e., the history starts from the beginning of the session within which the impression occurs.[1] This history is the most accessible in the user's short-term memory. Note that by definition, $\eta_{ib} = o_i$ implies that the function is calculated at the user level.

[1]A session ends when there is a five-minute interruption in a user's exposure to the ads. So if the time difference between two consecutive impressions shown to a user-app combination is more than five minutes, we assume that the two impressions belong to different sessions.

(b) $\eta_{ie} \in \mathcal{H}_{ie} = \{g, r_i\}$ can take two possible values which we define them as follows:

- $\eta_{ie} = g$ implies that the feature is calculated over the **g**lobal data, i.e., data after October 27, 2015, are not considered. These features are calculated without using any impression from the training, validation and test data-sets.
- $\eta_{ie} = r_i$ refers to **r**eal-time aggregation, i.e., the feature is calculated over all the information up till the focal impression i.

In Figure A.4, we present a picture with five example users to illustrate different types of history.

A.2.2 Feature Functions

We now use the nomenclature described above to define the following functions.

1. *Impressions*$(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie})$: This function counts the number of impressions with the characteristics given as inputs over the specified history.

$$\textit{Impressions}(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = \sum_{j \in [\eta_{ib}, \eta_{ie}]} \mathbb{1}(a_j \in \theta_{ia})\mathbb{1}(p_j \in \theta_{ip})\mathbb{1}(t_j \in \theta_{it})\mathbb{1}(u_j \in \theta_{iu}), \tag{A.7}$$

where j denotes an impression whose time-stamp lies between the starting point of the history η_{ib} and end point of the history η_{ie}.

For example, *Impressions*$(\{a_i\}, \{p_i\}, \{t_i\}, \{u_i\}; l, r_i)$ returns the number of times ad a_i is shown to user u_i while using app a_i at the hour of day t_i in the time period starting September 30, 2015, and ending with the last impression before i. Instead, if we are interested in the number of times that user u_i has seen ad a_i over all apps and all hours of the days from October 25, 2015, ($\eta_{ib} = s$) till the end of the global data ($\eta_{ie} = g$),

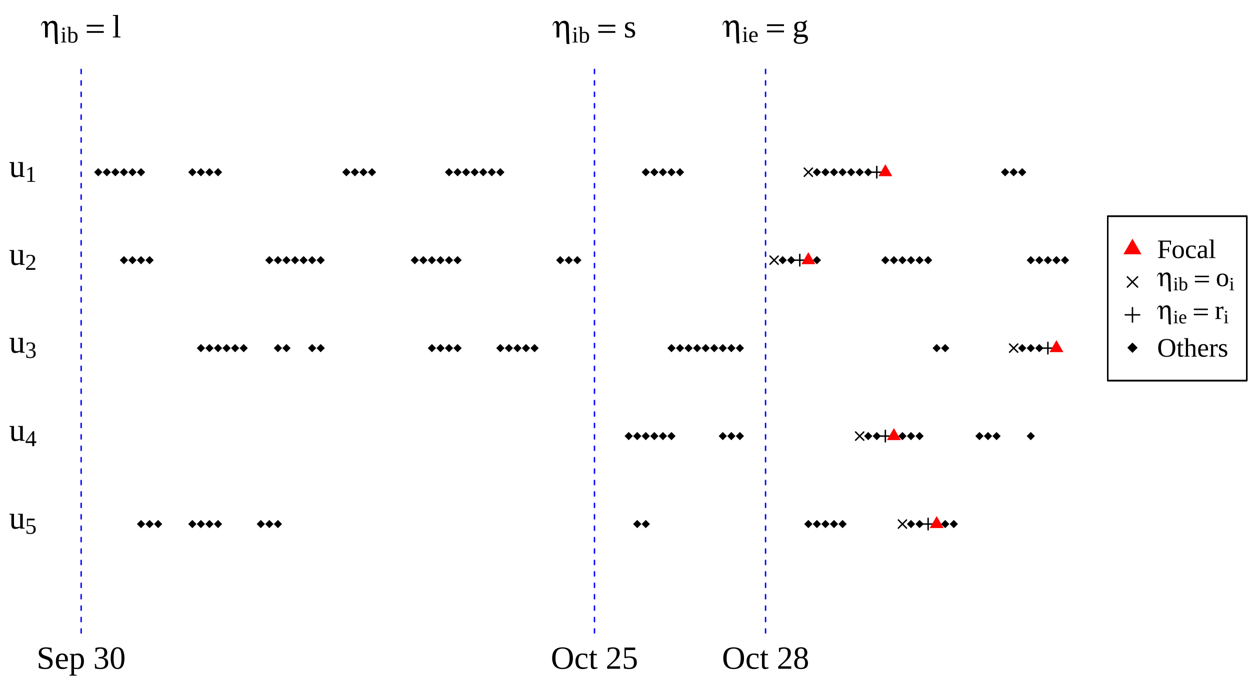

Figure A.4: Depiction of history for five example users.The red ▲ refers to the focal impression i for which the features are being generated. $+$ denotes the last impression just before the focal impression, and $\times$ refers to the first impression in the session in which the focal impression occurs.

then we would have:

$$\begin{aligned} \mathit{Impressions}(\{a_i\}, P, T, \{u_i\}; s, g) &= \sum_{j \in [s,g]} \mathbb{1}(a_j \in \{a_i\}) \mathbb{1}(p_j \in P) \mathbb{1}(t_j \in T) \mathbb{1}(u_j \in \{u_i\}) \\ &= \sum_{j \in [s,g]} \mathbb{1}(a_j \in \{a_i\}) \mathbb{1}(u_j \in \{u_i\}) \end{aligned}$$

Intuitively, the *Impressions* function captures the effects of repeated ad exposure on user behavior, and ad exposure has been shown to yield higher ad effectiveness in the recent literature [118, 70]. In some cases, it may also capture some unobserved ad-specific effects, e.g., an advertiser may not have enough impressions simply because he does not have a large budget.

2. $Clicks(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie})$: This function counts the number of clicks with the characteristics given as inputs over the history specified. It is similar to the *Impressions* function, but the difference is that *Clicks* only counts the impressions that led to a click.

$$Clicks(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = \sum_{j \in [\eta_{ib}, \eta_{ie}]} \mathbb{1}(a_j \in \theta_{ia}) \mathbb{1}(p_j \in \theta_{ip}) \mathbb{1}(t_j \in \theta_{it}) \mathbb{1}(u_j \in \theta_{iu}) \mathbb{1}(y_j = 1) \tag{A.8}$$

Clicks is a good indicator of both ad and app performance. Moreover, at the user-level, the number of clicks captures an individual user's propensity to click, as well as her propensity to click within a specific ad and/or app. Thus, this is a very informative metric.

3. $CTR(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie})$: This function calculates the click-through rate (CTR) for a given set of inputs and history, *i.e.*, the ratio of clicks to impressions. If $Impressions(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie}) > 0$, then:

$$CTR(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = \frac{Clicks(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie})}{Impressions(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie})}, \tag{A.9}$$

else if $Impressions(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = 0$, $CTR(\theta_{ia}, \theta_{ip}, \theta_{it}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = 0$. The *CTR* function is a combination of *Impressions* and *Clicks*. It captures the click-propensity at the user, ad, app, and time-levels as well their interactions.[2]

4. $AdCount(\theta_{ip}, \theta_{iu}; \eta_{ib}, \eta_{ie})$: This function returns the number of distinct ads shown for a given set of inputs.

$$AdCount(\theta_{ip}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = \sum_{a \in A^F} \mathbb{1}(Impressions(\{a\}, \theta_{ip}, T, \theta_{iu}; \eta_{ib}, \eta_{ie}) > 0), \tag{A.10}$$

[2] In principle, we do not need to *CTR* over and above *Clicks* and *Impressions* if our learning model can automatically generate new features based on non-linear combinations of basic features. Machine learning models like Boosted Trees do this naturally and it is one of their big advantages. However, other methods like OLS and Logistic regressions cannot do automatic feature combination and selection. So we include this feature to help improve the performance of these simpler models. This ensures that we are not handicapping them too much in our model comparisons.

We use the full set of ads seen in our data (denoted by A^F) and not just the top 37 ads in this function.[3] This ensures that we are capturing the full extent of variety in ads seen by users. Features derived using this function capture the variety in ads for a given set of inputs. In a recent paper, [106] show that higher variety of previous ads increases the CTR on the next ad.

An example of a feature based on this function is $AdCount(\{p_i\}, U; l, r_i)$, which counts the number of unique ads that were shown in app p_i across all users in the time period starting September 30 2015 and ending with the last impression before i.

5. $Entropy(\theta_{ip}, \theta_{iu}; \eta_{ib}, \eta_{ie})$: This function captures the entropy or dispersion in the ads seen by a user. We use [**?**]'s measure of diversity as our entropy metric.

$$Entropy(\theta_{ip}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = \frac{1}{\sum_{a \in A^F} Impressions(\{a\}, \theta_{ip}, T, \theta_{iu}; \eta_{ib}, \eta_{ie})^2} \tag{A.11}$$

Entropy contains information over and above just the count of the number of unique ads seen in the past since it captures the extent to which ads were dispersed across impressions. Consider two users who have both seen five impressions in the past, but with the following difference – (a) the first user has seen five distinct ads, or (b) the second has seen same ad five times. Naturally, the dispersion of ads is higher in case (a) than in case (b). The *Entropy* function reflects this dispersion – in case (a) the entropy is $\frac{1}{1^2+1^2+1^2+1^2+1^2} = 0.2$, whereas in case (b) it is $\frac{1}{5^2} = 0.04$. Thus, entropy is higher when ads are more dispersed.

The literature on eye-tracking has shown that individuals' attention during advertisements reduces as ad-repetition increases [103]. We therefore expect the dispersion of previous ads to influence a user's attention span in our setting. Since attention is a prerequisite to clicking, we expect entropy-based measures to have explanatory power in our click prediction model.

---

[3] We observe 263 unique ads in our data. Thus, the range of *AdCount* in our data goes from 0 to 263.

6. $AppCount(\theta_{ia}, \theta_{iu}; \eta_{ib}, \eta_{ie})$: This function calculates the number of distinct apps in which a given ad is shown to a specific user. Similar to *AdCount*, it can be defined as follows:

$$AppCount(\theta_{ia}, \theta_{iu}; \eta_{ib}, \eta_{ie}) = \sum_{p \in P} \mathbb{1}(Impressions(\theta_{ia}, \{p\}, T, \theta_{iu}; \eta_{ib}, \eta_{ie}) > 0), \quad \text{(A.12)}$$

where P is the set of apps, as defined in §A.2.1.[4] Previous studies have documented the spillover effects in multichannel online advertising [85]. We therefore expect click probabilities to vary if a user saw an ad in just one app vs. saw it in many apps. We also expect the click behavior of users to vary based on the number of apps they use. *AppCount*-based features help us capture these differences.

7. $TimeVariability(\theta_{iu}; \eta_{ib}, \eta_{ie})$: This function measures the variance in a user's CTR over different hours of the day. We can define this function as follows:

$$TimeVariability(\theta_{iu}; \eta_{ib}, \eta_{ie}) = \mathrm{Var}_t[CTR(A, P, \{t\}, \theta_{iu}; \eta_{ib}, \eta_{ie})] \quad \text{(A.13)}$$

Features based on this function capture variations in the temporal patterns in a user's clicking behavior, and we expect this information to help predict clicks.

8. $AppVariability(\theta_{iu}; \eta_{ib}, \eta_{ie})$: This function measures the variance in a user's CTR across different apps and is analogous to *TimeVariability*. This function is defined as follows:

$$AppVariability(\theta_{iu}; \eta_{ib}, \eta_{ie}) = \mathrm{Var}_p[CTR(A, \{p\}, T, \theta_{iu}; \eta_{ib}, \eta_{ie})] \quad \text{(A.14)}$$

Note that both *TimeVariability* and *AppVariability* are defined at the user level.

We use the functions defined above to generate 98 features for each impression. In addition, we include a few standalone features such as a dummy for each of the top ads, the mobile and Internet service providers, latitude, longitude, and connectivity type as described below:

[4] P has 51 elements, the top 1–50 apps and the 51^{st} element being the group of all smaller apps. In principle, we could use all the apps observed in the data and not just the Top 50 apps (like we did in the case of *AdCount*). However, we found that doing so does not really improve model performance. So we stick with this simpler definition.

9. $Bid(a_i)$: The bid submitted for ad a_i shown in impression i.

10. $Latitude(u_i)$: The latitude of user u_i when impression i occurs.

11. $Longitude(u_i)$: The longitude of user u_i when impression i occurs.

12. $WiFi(u_i)$: Denotes whether the user is connected via Wi-Fi or mobile data plan during impression i. It can take two possible values, $\{0, 1\}$.

13. $Brand(u_i)$: Denotes the brand of the smartphone that user u_i is using. We observe eight smartphone brands in our data and therefore capture each brand using a dummy variable. So the range for this function is $\{0, 1\}^8$.

14. $MSP(u_i)$: Denotes the Mobile Service Provider used by user u_i during impression i. We observe four MSPs in our data. So the range for this function is $\{0, 1\}^4$.

15. $ISP(u_i)$: Denotes the Internet service providers (ISP) of user u_i. We observe nine unique ISPs in our data. So the range of this function is $\{0, 1\}^9$.

16. $AdDummy(a_i)$: Generates dummy variables for each of the top ads in the data. Although we expect our main features to capture many of the ad-fixed effects, we include ad dummies to capture any residual ad-fixed effects, e.g., banner design. The range of this function is $\{0, 1\}^{38}$, since there are 37 top ads and a 38^{th} category comprising all the smaller ads.

A.2.3 Table of Features

In Table A.1, we now present our full list of 160 features derived from the feature-functions described above as well as their classification as behavioral, contextual, and/or ad-specific based on the categorization scheme described in Chapter 2.

Table A.1: List of Features for impression i.

No.	Feature Name	Feature Classification			Contextual Features	
		Behavioral	Contextual	Ad-specific	App-specific	Time-specific
1	*Impressions* $(A, P, T, u_i; l, r_i)$	✓				
2	*Impressions* $(a_i, P, T, U; l, g)$			✓		
3	*Impressions* $(A, p_i, T, U; l, g)$		✓		✓	
4	*Impressions* $(A, P, t_i, U; l, g)$		✓			✓
5	*Impressions* $(a_i, P, T, u_i; l, r_i)$	✓		✓		
6	*Impressions* $(A, p_i, T, u_i; l, r_i)$	✓	✓		✓	
7	*Impressions* $(A, P, t_i, u_i; l, r_i)$	✓	✓			✓
8	*Impressions* $(a_i, p_i, T, U; l, g)$		✓	✓	✓	
9	*Impressions* $(a_i, P, t_i, U; l, g)$		✓	✓		✓
10	*Impressions* $(A, p_i, t_i, U; l, g)$		✓		✓	✓
11	*Impressions* $(a_i, p_i, T, u_i; l, r_i)$	✓	✓	✓	✓	
12	*Impressions* $(a_i, p_i, t_i, U; l, g)$		✓	✓	✓	✓
13	*Impressions* $(A, P, T, u_i; s, r_i)$	✓				
14	*Impressions* $(a_i, P, T, U; s, g)$			✓		
15	*Impressions* $(A, p_i, T, U; s, g)$		✓		✓	
16	*Impressions* $(A, P, t_i, U; s, g)$		✓			✓
17	*Impressions* $(a_i, P, T, u_i; s, r_i)$	✓		✓		
18	*Impressions* $(A, p_i, T, u_i; s, r_i)$	✓	✓		✓	
19	*Impressions* $(A, P, t_i, u_i; s, r_i)$	✓	✓			✓
20	*Impressions* $(a_i, p_i, T, U; s, g)$		✓	✓	✓	
21	*Impressions* $(a_i, P, t_i, U; s, g)$		✓	✓		✓
22	*Impressions* $(A, p_i, t_i, U; s, g)$		✓		✓	✓
23	*Impressions* $(a_i, p_i, T, u_i; s, r_i)$	✓	✓	✓	✓	
24	*Impressions* $(a_i, p_i, t_i, U; s, g)$		✓	✓	✓	✓
25	*Impressions* $(A, P, T, u_i; o_i, r_i)$	✓				
26	*Impressions* $(a_i, P, T, u_i; o_i, r_i)$	✓		✓		
27	*Clicks* $(A, P, T, u_i; l, r_i)$	✓				

Continued on next page

Table A.1 – continued from previous page

No.	Feature Name	Feature Classification			Contextual Features	
		Behavioral	Contextual	Ad-specific	App-specific	Time-specific
28	*Clicks* $(a_i, P, T, U; l, g)$			✓		
29	*Clicks* $(A, p_i, T, U; l, g)$		✓		✓	
30	*Clicks* $(A, P, t_i, U; l, g)$		✓			✓
31	*Clicks* $(a_i, P, T, u_i; l, r_i)$	✓		✓		
32	*Clicks* $(A, p_i, T, u_i; l, r_i)$	✓	✓		✓	
33	*Clicks* $(A, P, t_i, u_i; l, r_i)$	✓	✓			✓
34	*Clicks* $(a_i, p_i, T, U; l, g)$		✓	✓	✓	
35	*Clicks* $(a_i, P, t_i, U; l, g)$		✓	✓		✓
36	*Clicks* $(A, p_i, t_i, U; l, g)$		✓		✓	✓
37	*Clicks* $(a_i, p_i, T, u_i; l, r_i)$	✓	✓	✓	✓	
38	*Clicks* $(a_i, p_i, t_i, U; l, g)$		✓	✓	✓	✓
39	*Clicks* $(A, P, T, u_i; s, r_i)$	✓				
40	*Clicks* $(a_i, P, T, U; s, g)$			✓		
41	*Clicks* $(A, p_i, T, U; s, g)$		✓		✓	
42	*Clicks* $(A, P, t_i, U; s, g)$		✓			✓
43	*Clicks* $(a_i, P, T, u_i; s, r_i)$	✓		✓		
44	*Clicks* $(A, p_i, T, u_i; s, r_i)$	✓	✓		✓	
45	*Clicks* $(A, P, t_i, u_i; s, r_i)$	✓	✓			✓
46	*Clicks* $(a_i, p_i, T, U; s, g)$		✓	✓	✓	
47	*Clicks* $(a_i, P, t_i, U; s, g)$		✓	✓		✓
48	*Clicks* $(A, p_i, t_i, U; s, g)$		✓		✓	✓
49	*Clicks* $(a_i, p_i, T, u_i; s, r_i)$	✓	✓	✓	✓	
50	*Clicks* $(a_i, p_i, t_i, U; s, g)$		✓	✓	✓	✓
51	*CTR* $(A, P, T, u_i; l, r_i)$	✓				
52	*CTR* $(a_i, P, T, U; l, g)$			✓		
53	*CTR* $(A, p_i, T, U; l, g)$		✓		✓	
54	*CTR* $(A, P, t_i, U; l, g)$		✓			✓
55	*CTR* $(a_i, P, T, u_i; l, r_i)$	✓		✓		
56	*CTR* $(A, p_i, T, u_i; l, r_i)$	✓	✓		✓	

Continued on next page

Table A.1 – continued from previous page

No.	Feature Name	Feature Classification			Contextual Features	
		Behavioral	Contextual	Ad-specific	App-specific	Time-specific
57	$CTR\ (A, P, t_i, u_i; l, r_i)$	✓	✓			✓
58	$CTR\ (a_i, p_i, T, U; l, g)$		✓	✓	✓	
59	$CTR\ (a_i, P, t_i, U; l, g)$		✓	✓		✓
60	$CTR\ (A, p_i, t_i, U; l, g)$		✓		✓	✓
61	$CTR\ (a_i, p_i, T, u_i; l, r_i)$	✓	✓	✓	✓	
62	$CTR\ (a_i, p_i, t_i, U; l, g)$		✓	✓	✓	✓
63	$CTR\ (A, P, T, u_i; s, r_i)$	✓				
64	$CTR\ (a_i, P, T, U; s, g)$			✓		
65	$CTR\ (A, p_i, T, U; s, g)$		✓		✓	
66	$CTR\ (A, P, t_i, U; s, g)$		✓			✓
67	$CTR\ (a_i, P, T, u_i; s, r_i)$	✓		✓		
68	$CTR\ (A, p_i, T, u_i; s, r_i)$	✓	✓		✓	
69	$CTR\ (A, P, t_i, u_i; s, r_i)$	✓	✓			✓
70	$CTR\ (a_i, p_i, T, U; s, g)$		✓	✓	✓	
71	$CTR\ (a_i, P, t_i, U; s, g)$		✓	✓		✓
72	$CTR\ (A, p_i, t_i, U; s, g)$		✓		✓	✓
73	$CTR\ (a_i, p_i, T, u_i; s, r_i)$	✓	✓	✓	✓	
74	$CTR\ (a_i, p_i, t_i, U; s, g)$		✓	✓	✓	✓
75	$AdCount\ (A, u_i; l, g)$	✓		✓		
76	$AdCount\ (p_i, U; l, g)$		✓	✓	✓	
77	$AdCount\ (p_i, u_i; l, g)$	✓	✓	✓	✓	
78	$AdCount\ (A, u_i; s, g)$	✓		✓		
79	$AdCount\ (p_i, U; s, g)$		✓	✓	✓	
80	$AdCount\ (p_i, u_i; s, g)$	✓	✓	✓	✓	
81	$AdCount\ (A, u_i; o_i, r_i)$	✓		✓		
82	$AppCount\ (A, u_i; l, g)$	✓	✓		✓	
83	$AppCount\ (a_i, U; l, g)$		✓	✓	✓	
84	$AppCount\ (a_i, u_i; l, g)$	✓	✓	✓	✓	
85	$AppCount\ (A, u_i; s, g)$	✓	✓		✓	

Continued on next page

Table A.1 – continued from previous page

No.	Feature Name	Feature Classification			Contextual Features	
		Behavioral	Contextual	Ad-specific	App-specific	Time-specific
86	*AppCount* $(a_i, U; s, g)$		✓	✓	✓	
87	*AppCount* $(a_i, u_i; s, g)$	✓	✓	✓	✓	
88	*Entropy* $(A, u_i; l, g)$	✓		✓		
89	*Entropy* $(p_i, U; l, g)$		✓	✓	✓	
90	*Entropy* $(p_i, u_i; l, g)$	✓	✓	✓	✓	
91	*Entropy* $(A, u_i; s, g)$	✓		✓		
92	*Entropy* $(p_i, U; s, g)$		✓	✓	✓	
93	*Entropy* $(p_i, u_i; s, g)$	✓	✓	✓	✓	
94	*Entropy* $(A, u_i; o_i, r_i)$	✓		✓		
95	*TimeVariability* $(u_i; l, g)$	✓	✓			✓
96	*TimeVariability* $(u_i; s, g)$	✓	✓			✓
97	*AppVariability* $(u_i; l, g)$	✓	✓		✓	
98	*AppVariability* $(u_i; s, g)$	✓	✓		✓	
99	*Latitude* (u_i)	✓				
100	*Longitude* (u_i)	✓				
101	*WiFi* (u_i)	✓				
102-109	*Brand* (u_i)	✓				
110-113	*Operator* (u_i)	✓				
114-122	*ISP* (u_i)	✓				
123-160	*AdDummy* (a_i)			✓		

A.3 XGBoost: Overview and Implementation

A.3.1 Overview of XGBoost

We start by considering a generic tree ensemble method as follows: let y_i and $\mathbf{x}_i$ denote the click indicator and the vector of features for impression i such that $y_i \in \{0, 1\}$ and $\mathbf{x}_i \in \mathbb{R}^k$, where k is the number of features. Then, a tree ensemble method is defined as follows:

$$\hat{y}_i = \phi(\mathbf{x}_i) = \sum_{j=1}^{J} \tau_j(\mathbf{x}_i) = \sum_{j=1}^{J} w^{(j)}_{q_j(\mathbf{x}_i)}, \tag{A.15}$$

where $q_j : \mathbb{R}^k \to \{1, 2, \ldots, L_j\}$ and $w^{(j)} \in \mathbb{R}^{L_j}$ constitute the j^{th} regression tree τ_j with L_j leaves. Here q_j maps an impression to the leaf index and $w^{(j)}$ represents the weight on leaves. Hence, $w^{(j)}_{q_j(\mathbf{x}_i)}$ indicates the weight on the leaf that $\mathbf{x}_i$ belongs to. The tree ensemble method uses J additive functions to predict the output. In order to estimate the set of functions, [22] minimize the following regularized objective:

$$\mathcal{L}(\phi) = \sum_i l(\hat{y}_i, y_i) + \gamma \sum_j L_j + \frac{1}{2}\lambda \sum_j \left\| w^{(j)} \right\|^2, \tag{A.16}$$

where γ and λ are the regularization parameters to penalize the model complexity, and l is a differentiable convex loss function. Here, in contrast to MART, XGBoost penalizes not just tree depth but leaf weights as well.

Since the regularized objective in Equation (A.16) cannot be minimized using traditional optimization methods in Euclidean space, we employ Newton boosting in function space to train the model in an additive manner. Formally, if we define $\hat{y}_i^{(j)}$ as the prediction of the i^{th} impression at the j^{th} iteration, we will add τ_j to minimize the following objective:

$$\mathcal{L}(\phi) = \sum_i l\big(\hat{y}_i^{(j-1)} + \tau_j(\mathbf{x}_i), y_i\big) + \gamma L_j + \frac{1}{2}\lambda \left\| w^{(j)} \right\|^2 \tag{A.17}$$

In each iteration, we greedily add the tree that most improves our model according to the objective function in Equation (A.16). Since the model uses a greedy algorithm to find the best split at each iteration, it is impossible to search over all tree structures. Thus, we restrict the set of trees by specifying the maximum depth of a tree. To optimize this objective function, [40] propose a second-order approximation:

$$\mathcal{L}(\phi) \simeq \sum_i \left[l(\hat{y}_i^{(j-1)}, y_i) + \mathcal{G}_i \tau_j(\mathbf{x}_i) + \frac{1}{2}\mathcal{H}_i \tau_j^2(\mathbf{x}_i) \right] + \gamma L_j + \frac{1}{2}\lambda \left\| w^{(j)} \right\|^2, \tag{A.18}$$

where $\mathcal{G}_i = \partial_{\hat{y}_i^{(j-1)}} l(\hat{y}_i^{(j-1)}, y_i)$ and $\mathcal{H}_i = \partial^2_{\hat{y}_i^{(j-1)}} l(\hat{y}_i^{(j-1)}, y_i)$ are first and second order gradient statistics on the loss function. This approximation is used to derive the optimal tree at the step.

To implement this optimization routine, the researcher needs to provide a set of hyper-parameters. These include the regularization parameters γ and λ defined in Equation (A.16). Further, XGBoost uses two additional parameters to prevent over-fitting. The first is called shrinkage parameter (ν) introduced by [42], which functions like learning rate in stochastic optimization. The second is the column sub-sampling parameter (χ_s), which is used to pick the fraction of features supplied to a tree and ranges from 0 to 1. In addition, we also need to specify parameters that structure the optimization problem and control the exit conditions for the algorithm. The first is d_{max}, the maximum depth of trees. As mentioned earlier, this ensures that the model searches over a finite set of trees. Second, we set the number of maximum number of iterations. Finally, the last parameter defines the early stopping rule, $\kappa \in \mathbb{Z}^+$, which stops the algorithm if the loss does not change in κ consecutive iterations. This parameter also helps prevent over-fitting.

Note all that these parameters cannot be inferred from the training data alone and should be set based on a scientific validation procedure. In the next section of the Appendix, §A.3.2, we provide a step by step explanation and the final values used for these hyper-parameters in our analysis.

A.3.2 Validation

The goal of validation is to pick the optimal tuning parameters. They cannot be inferred from the training data alone because they are hyper-parameters. The validation procedure uses two separate data-sets – training and validation data (from October 28 and 29, as shown in Figure 2.1) to pin down the hyper-parameters. It is worth noting that at this stage, the test data is kept separate and is brought out only at the end after the model has been finalized to evaluate the final model performance.

As discussed earlier in §A.3.1, XGBoost uses five hyper-parameters that need tuning. Let $\mathcal{W} = \{\gamma, \lambda, \nu, d_{\max}, \chi_s\}$ denote the set of hyper-parameters, where γ and λ are the regularization parameters, ν is the shrinkage parameter or learning rate, $d_{\max}$ is maximum depth of trees, and χ_s is the column sub-sampling parameter, which refers to the percentage of features randomly selected in each round. For each of these parameters, we consider the following sets of values:

- $\gamma \in \{7, 9, 11\}$
- $\lambda \in \{0, 1, 2\}$
- $\nu \in \{0.05, 0.1, 0.5, 1\}$
- $d_{\max} \in \{5, 6, 7\}$
- $\chi_s \in \{0.5, 0.75\}$

Overall, $\mathcal{W}$ contains 216 elements. We now describe the validation and testing procedure in detail.

- Step 1: For each element of $\mathcal{W}$, train the model on the first two-thirds of the training and validation data and evaluate the model's performance on the remaining one-third. This is the typical hold-out procedure [57]. Boosted trees typically over-fit when they are stopped at some point. So when training the model, we use the early stopping rule to avoid this problem [138].
- Step 2: Choose the set of hyper-parameters that gives the best model performance on the validation set (as measured by RIG in our case). Denote this as $\mathcal{W}^*$.
- Step 3: Using $\mathcal{W}^*$, train the final model on the full training and validation data, and here too we use the early stopping rule.
- Step 4: Evaluate the final model's performance on the test data.

Based on the above procedure, we derive the set of optimal hyper-parameters for our empirical setting as $\mathcal{W}^* = \{\gamma = 9, \lambda = 1, \nu = 0.1, d_{\max} = 6, \chi_s = 0.5\}$. Further, when training on the full training and validation data (Step 3), we do not see any improvement in model fit after 276 steps, so at this iteration number (following the early stopping rule).

Note that our validation procedure addresses two potential complications with our data and problem setting. First, our data and features have a time component. Thus, if we are not careful in how we split the validation and training data-sets, we can end up in situations where we use the future to predict the past (e.g., train on data from period t and validate on data from $t-1$). To avoid this problem, the splits should be chosen based on time. Our validation procedure does this by using the first two-thirds of the validation+training data for training and the latter one-third for validation. However, doing so gives rise to a second problem – by choosing the most recent data for validation (instead of training), we forgo the information in the most recent impressions. In CTR prediction, it is important not to waste the most recent impressions while fitting the model because these impressions are more likely to be predictive of the future [93]. To address this, in Step 3, after choosing the optimal hyper-parameters, we train a final model on the full training and validation data. This model is then used as the final model for results and counterfactuals.

A.4 Evaluating Efficient Targeting Policy Using Importance Sampling

As discussed in §2.5.2, the overlap between our data and the efficient targeting policy provides a unique opportunity to use model-free approaches to evaluate the efficient targeting policy. Intuitively, if we weight the overlapping sample such that it represents the full sample, we can estimate the outcome of interest under the efficient targeting policy. In particular, we use *importance sampling* whereby we weight the impressions in the overlapping area by the inverse propensity score of those impressions. We can now define the average match value under the efficient targeting as follows:

$$\hat{m}_{IS}^{\tau^*} = \frac{1}{N_F} \sum_{i=1}^{N_F} \sum_{a \in \mathcal{A}_i} \frac{\mathbb{1}(a_i = a)\mathbb{1}(\tau^*(i) = a)}{\hat{\pi}_{ia}} y_{i,a}, \tag{A.19}$$

where $y_{i,a}$ is the click outcome for ad a in impression i, and the indicator functions both take value one only if the efficient targeting policy coincides the actual data. As such, for non-zero elements in the summation above, we actually observe the actual outcome of interest $y_{i,a}$. Now, we can even show that under unconfoundedness assumption, $\hat{m}_{IS}^{\tau^*}$ is a consistent estimator for the match value under the efficient targeting policy, conditional on observed covariates. We can write:

$$
\begin{aligned}
\mathbb{E}\left[\hat{m}_{IS}^{\tau^*} \mid X\right] &= \mathbb{E}\left[\frac{1}{N_F}\sum_{i=1}^{N_F}\sum_{a\in\mathcal{A}_i}\frac{\mathbb{1}(a_i = a)\mathbb{1}(\tau^*(i) = a)}{\hat{\pi}_{ia}}y_{i,a} \mid X\right] \\
&= \frac{1}{N_F}\sum_{i=1}^{N_F}\sum_{a\in\mathcal{A}_i}\mathbb{E}\left[\frac{\mathbb{1}(a_i = a)\mathbb{1}(\tau^*(i) = a)}{\hat{\pi}_{ia}}y_{i,a} \mid X\right] \\
&= \frac{1}{N_F}\sum_{i=1}^{N_F}\sum_{a\in\mathcal{A}_i}\frac{\hat{\pi}_{ia}\mathbb{1}(\tau^*(i) = a)}{\hat{\pi}_{ia}}\mathbb{E}\left[y_{i,a} \mid X\right] \\
&= \frac{1}{N_F}\sum_{i=1}^{N_F}\sum_{a\in\mathcal{A}_i}\mathbb{1}(\tau^*(i) = a)\mathbb{E}\left[y_{i,a} \mid X\right] \\
&= \mathbb{E}\left[\frac{1}{N_F}\sum_{i=1}^{N_F}\sum_{a\in\mathcal{A}_i}\mathbb{1}(\tau^*(i) = a)y_{i,a} \mid X\right],
\end{aligned}
\tag{A.20}
$$

where the last term is the true match value under the efficient targeting policy. In the first two lines in Equation (A.20), we just use the linearity of the expectation function. In the third line, we use the unconfoundedness assumption. In the last two lines, we again use the linearity of the expectation function.

Now, we use the specification in Equation (A.20) to estimate the match value under the efficient targeting policy. We find that the efficient targeting policy coincides with the ads shown in the Filtered Sample in 67,152 impressions. We use inverse propensity weights to adjust for this fraction of the sample and find that the efficient targeting policy improves the average CTR in the current regime by 65.53%. This is very similar to our original finding using a direct method based on our match value estimates.

A.5 Appendix for Robustness Checks

A.5.1 Other Evaluation Metrics

We consider three alternative evaluation metrics.

- Area Under the Curve (AUC): It calculates the area under the ROC curve, which is a graphical depiction of *true positive rate (TPR)* as a function of *false positive rate (FPR)*. This metric is often used in classification problems, but is less appropriate for prediction tasks such as ours because of two reasons. First, it is insensitive to the transformation of the predicted probabilities that preserve their rank. Thus, a poorly fitted model might have higher AUC than a well-fitted model [64]. Second, it puts the same weight on *false positive rate (FPR)* and *false negative rate (FNR)*. However, in CTR prediction, the penalty of FNR is usually higher than FPR [**?**].

- 0/1 Loss: This is a simple metric used to evaluate correctness in classification tasks. It is simply the percentage of incorrectly classified impressions. As with AUC, it is not very useful when accuracy of the prediction matters since it is not good at evaluating the predictive accuracy of rare events (e.g., clicks). For example, in our case, the loss is lower than 1% loss even if we blindly predict that none of the impressions will lead to click.
- Mean Squared Error (MSE): This is one of the most widely used metrics for measuring the goodness of fit. It is similar to LogLoss, which we use to calculate *RIG*. Both LogLoss and SquareLoss are often used for probability estimation and boosting in the machine learning literature. Let d_i be the Euclidean distance between the predicted value and actual outcome for impression i. This can be interpreted as the misprediction for the corresponding impression. SquareLoss and LogLoss for this impression will then be d_i^2 and $-\log(1-d_i)$ respectively. Since both functions are convex with respect to d_i, they penalize larger mispredictions more than smaller ones. However, a big difference is that SquareLoss is finite, whereas LogLoss is not. In fact, LogLoss evaluates $d_i = 1$ as infinitely bad. In our problem, this translates to either predicting 1 for non-clicks or predicting 0 for clicks. Therefore, the model optimized by LogLoss will not predict 0 or 1. Given that we do not know how users interact with the app at the moment, it is quite unrealistic to predict 1 for an impression, especially because each impression only lasts a short time. Thus, we choose LogLoss as our main metric, which is also the most common choice in the literature on CTR prediction [**?**].
- Confusion Matrix: We also present the Confusion Matrix for our main models. Like AUC and 0/1 Loss, Confusion Matrix is mostly used in classification problems. However, this gives us an intuitive metric to evaluate the performance of our model.

We now take the optimized models presented in §3.6 and evaluate their performance on these alternative metrics. The results from this exercise are presented in Table A.2. Overall, all our substantive results remain the same when we use alternative evaluation metrics.

A.5.2 Other Learning Methods

We now compare the performance of XGBoost with other five other learning algorithms and present the results in Table A.3. Note that XGBoost outperforms all of them.

For each learning model, we describe the hyper-parameters associated with it and our approach to optimizing these hyper-parameters. Note that in all the cases, we use the same high-level validation procedure described in §A.3.2.

- Least squares does not use any hyper-parameters and hence does not require validation. In this case, we simply train the model on the full training and validation data to infer the model parameters, and report the model's performance on the test data.
- For LASSO, the validation procedure is straightforward. The only hyper-parameter to set is the L1 regularization parameter, λ. We search over 100 values of λ and pick the one which gives us the best performance on the validation set ($\lambda = 6.3 \times 10^{-4}$). We then use this parameter and train the model on the entire training and validation set and test the performance on the test set.
- For CART, we use the package *rpart* in R, which implements a single tree proposed by [18]. We use recursive partitioning algorithm with a complexity parameter (c_{tree}) as the only hyper-parameter that we need to select through validation. The main role of this parameter is to avoid over-fitting and save computing time by pruning splits that are not worthwhile. As such, any split that does not improve the fit by a factor of complexity parameter is not attempted. We search for the optimal complexity parameter over the grid $c_{tree} \in \{0.1, 0.05, 0.01, 0.005, 0.001, 0.0005, 0.0001, 0.00005, 0.00001\}$, and based on the validation procedure, we derive the optimal complexity parameter as 5^{-5}. That is, the model adds another additional split only when the R-squared increments by at least

Data	Evaluation Metric	Behavioral	Contextual	Full
Full Sample	AUC	0.7910	0.7014	0.8230
Top Ads/Apps	AUC	0.8082	0.7192	0.8410
Filtered Sample	AUC	0.8073	0.7410	0.8410
Full Sample	0/1 Improvement (%)	0.63%	0.00%	4.74%
Top Ads/Apps	0/1 Improvement (%)	1.07%	0.00%	8.23%
Filtered Sample	0/1 Improvement (%)	1.34%	0.00%	8.47%
Full Sample	MSE Improvement (%)	3.41%	0.55%	8.59%
Top Ads/Apps	MSE Improvement (%)	4.86%	0.67%	13.33%
Filtered Sample	MSE Improvement (%)	4.89%	0.67%	12.93%
Full Sample	True Positives (#)	905	0	5934
Full Sample	True Negatives (#)	9533282	9533605	9531972
Full Sample	False Positives (#)	323	0	1633
Full Sample	False Negatives (#)	91325	92230	86296
Top Ads/Apps	True Positives (#)	757	0	5621
Top Ads/Apps	True Negatives (#)	6057514	6057744	6056235
Top Ads/Apps	False Positives (#)	230	0	1509
Top Ads/Apps	False Negatives (#)	50010	50767	45146
Filtered Sample	True Positives (#)	517	0	2928
Filtered Sample	True Negatives (#)	4424624	4424738	4424343
Filtered Sample	False Positives (#)	114	0	395
Filtered Sample	False Negatives (#)	29379	29896	26968
Full Sample	RIG	12.14%	5.25%	17.95%
Top Ads/Apps	RIG	14.82%	5.98%	22.85%
Filtered Sample	RIG	14.74%	6.77%	22.45%

Table A.2: Model performance for the two samples (full and top ads/apps) when evaluated on the alternative metrics.

Method	*RIG* **over Baseline**
Least Squares	7.72%
LASSO	7.92%
Logistic Regression	11.17%
Regression Tree	15.03%
Random Forest	15.75%
XGBoost	17.95%

Table A.3: *RIG* of different learning methods for the test data.

5^{-5}.

- For Random Forest, we use the package *sklearn* in Python. There are three hyper-parameters in this case – (1) n_{tree}, the number of trees over which we build our ensemble forest, (2) χ_s, the column sub-sampling parameter, which indicates the percentage of features that should be randomly considered in each round when looking for the best split, and (3) $n_{\min}$, the minimum number of samples required to split an internal node. We search for the optimal set of hyper-parameters over the following grid:
  - $n_{tree} \in \{100, 500, 1000\}$
  - $\chi_s \in \{0.33, 0.5, 0.75\}$
  - $n_{\min} \in \{100, 500, 1000\}$

  Based on our validation procedure, we find the optimal set of hyper-parameters to be: $\{n_{tree} = 1000, \chi_s = 0.33, n_{\min} = 500\}$.

A.5.3 Robustness Checks on Feature Generation

We also ran the following robustness checks on the feature generation framework.

- We start by considering different ways of aggregating over the history. One possibility is to use $\eta_{ie} = r_i$ for all the features, i.e., update all the features in real-time. We find no difference in terms of the prediction accuracy when we adopt this approach, though it increases the computational costs significantly. Therefore, we stick to our current approach where we use a combination of global and real-time features.

- Next, we examine the model's performance under an alternative definition of long- and short-term history for features that are updated in real-time. The idea is to fix the length of history, instead of fixing η_{ib} for each impression. In other words, instead of aggregating over $[l, r_i]$ and $[s, r_i]$ where l and s are fixed, we aggregate over $[l_i, r_i]$ and $[s_i, r_i]$ where l_i and s_i are no more fixed, but the length of $[l_i, r_i]$ and $[s_i, r_i]$ are fixed. For example, l_i for impression i on October 28 is the same time on September 30, while l_i for impression i on October 30 is the same time on October 2. Under this new approach, we find a slight decrease in the performance: the *RIG* drops to 17.69% improvement over the baseline.
- We also consider a model with dummy variables for the top apps in our feature set (similar to what we now do for top ads). Again, we find no discernible differences in the results without app dummies: the *RIG* is 17.97% over the baseline. This may be due to the fact that our feature generation framework captures the fixed effects of apps well.

Overall, we find that our feature set works well and any additional features or more complex feature generation mechanisms do not provide any significant benefits in *RIG*.

A.5.4 Sampling and Data Adequacy

We conduct our analyses using a relatively large sample consisting of 727,354 users in the train, validation, and test data-sets. This corresponds to 17,856,610 impressions in the training and validation data, and 9,625,835 impressions in the test data. We now examine the adequacy the rate the adequacy of our sample by calculating the RIG for different (lower) sample sizes. That is, we quantify how much our model gains by using more data, and at what point the marginal value of additional data is minimal.

To calculate the *RIG* for a given sample size of N_u, we do the following: 1) We take ten random samples of N_u users, and generate two data sets – the training data and the test data. 2) For each sample, we train the model using the training data and then test the model's performance on the test data.[5] 3) We then calculate the mean and standard deviation of the RIG for each sample. We perform this exercise for nine sample sizes starting with $N_u = 1000$ and going up till $N_u = 600,000$. The results from this exercise are shown in Table A.4. We also report the average sample size of train and test data respectively as $\overline{N}_{train}$ and $\overline{N}_{test}$.

In principle, we can perform the above exercise for each sample size with only one sample instead of ten. However, such an approach is likely to be error-prone, especially at smaller sample sizes, since there is heterogeneity among users and each sample is random. So we may randomly find a smaller sample to have a higher RIG than a larger sample in one particular instance. To avoid making incorrect inferences due to the particularities of one specific sample and to minimize the noise in our results, we employ the bootstrap procedure described above.

Table A.4 suggests that after about 100,000 users, increasing the sample size improves the prediction only slightly. However, increasing sample sizes also increase the training time and computational costs. Given the cost-benefit trade-off, our sample of 727,354 users is more than sufficient for our purposes.

A.5.5 Other Validation Techniques

The validation procedure outlined in Appendix §A.3.2 is the first-best validation procedure in data-rich situations such as ours. Nevertheless, we examine two other commonly used techniques:

- Hold-out validation – very similar to our current approach, except that at Step 3, instead of training the final model on the combination of training and validation data, we simply use the best model (using the optimal $\mathcal{W}^*$) trained based on the training data (from Step

[5]We use the hyper-parameters obtained from the validation exercise that we performed in the main model for training. This is likely to help the performance of the models trained on smaller samples because if we were to tune the model using smaller data, the estimated hyper-parameters are likely to be worse. Thus, the gains reported here more favorable than what we would obtain if we also validated/tuned the model using the smaller data samples.

User Sample Size (N_u)	$\overline{N}_{train}$	$\overline{N}_{test}$	RIG over Baseline CTR	
			Coefficient	Std. error
1,000	24,500	13,274	13.76%	3.17%
5,000	124,521	66,820	14.25%	1.76%
10,000	249,749	139,123	15.26%	1.48%
20,000	486,007	266,497	16.14%	0.40%
50,000	1,220,394	663,569	16.84%	0.28%
100,000	2,436,037	1,332,894	17.27%	0.23%
200,000	4,875,586	2,654,110	17.58%	0.20%
400,000	9,749,402	5,327,471	17.84%	0.18%
600,000	14,699,589	7,928,275	17.91%	0.15%

Table A.4: RIG for different sample sizes. $\overline{N}_{train}$ and $\overline{N}_{test}$ are respectively the average size of train and test data after sampling users.

2) as the final model. Thus, we do not use the validation data to train the final model. This can lead to some information loss (especially from the recent impressions). We find that the model performance on the test set drops when we use hold-out validation: our *RIG* is 17.21% which is lower than that of our validation procedure.

- k-fold cross-validation – we find no improvement in the performance of the model selected by 5-fold cross-validation (*RIG* is 17.33%). Please see Footnote 7 in [137] and [57] for a detailed discussion on the pros and cons of k-fold cross validation.

A.6 Detailed Analysis of the Example Presented in Figure 2.8

In an important paper, [82] argue that sharing too much targeting information with advertisers can thin auction markets which in turn would soften competition and make the platform worse. We now present a simple example to highlight this idea. Consider a platform with two advertisers ($a^{(1)}$ and $a^{(2)}$) competing for the impressions of two users ($u^{(1)}$ and $u^{(2)}$). Assume that the platform uses second price auctions with Cost per Impression (CPI) pricing, where the highest bidder wins the impression and pays the bid of the second-highest bidder. These auctions have the useful property of truthful bidding [132]. Let the advertisers be symmetric in their valuation of a click (normalized to 1 hereafter). Further, let the match values between advertisers and users be as shown in Equation (A.21). Match values can be interpreted as the eCTR of an impression for the advertiser-user pair. Notice that advertiser $a^{(1)}$ has a better match with user $u^{(1)}$ and advertiser $a^{(2)}$ with user $u^{(2)}$.

$$eCTR = \begin{array}{c} \\ a^{(1)} \\ \\ a^{(2)} \end{array} \begin{array}{c} \begin{array}{cc} u^{(1)} & u^{(2)} \end{array} \\ \begin{pmatrix} 0.5 & 0.1 \\ & \\ 0.1 & 0.3 \end{pmatrix} \end{array} \longrightarrow \begin{array}{c} \bar{u} \\ \begin{pmatrix} 0.3 \\ \\ 0.2 \end{pmatrix} \end{array} \tag{A.21}$$

We can formally show that, for a given targeting strategy, both CPI and Cost per Click (CPC mechanisms generate the same revenues for the platform and advertisers. So we focus on the CPI case throughout the text and refer readers to Appendix A.8 for the CPC example and analysis.

Consider the advertiser's bidding strategy and outcomes under two regimes – 1) No data disclosure by the platform, and 2) Full disclosure of match values by the platform. The results from these two regimes are laid out in Table A.5 and discussed below:

- No data disclosure – Here advertisers only know their aggregate match value for both users. So $a^{(1)}$ and $a^{(2)}$'s expected match values for the two users are 0.3 and 0.2. In a second price auction, advertisers bid their expected valuations. So $a^{(1)}$ wins both impressions and pays the next highest bid, $b_1^{(2)} = b_2^{(2)} = 0.2$, for each impression. Therefore, platform's revenue is $R = 0.4$, advertisers' surplus is $W^{(1)} = 0.2$ and $W^{(2)} = 0$, and the total surplus is $S = 0.6$.
- Full data disclosure – Since advertisers now have information on their match for each impression, they submit targeted bids that reflect their valuations as shown in Table A.5. Therefore, the advertiser who values the impression more wins it. However, because of the asymmetry in advertisers' valuation over impressions, the competition over each

No Data Disclosure (or No Targeting)	Full Data Disclosure (or Perfect Targeting)
For both impressions:	For User $u^{(1)}$'s impression:
Bids:	Bids:
Advertiser $a^{(1)}$: $b_1^{(1)} = b_2^{(1)} = 0.3$	Advertiser $a^{(1)}$: $b_1^{(1)} = 0.5$; and $a^{(2)}$: $b_1^{(2)} = 0.1$
Advertiser $a^{(2)}$: $b_1^{(2)} = b_2^{(2)} = 0.2$	Outcome:
	Advertiser $a^{(1)}$ wins $u^{(1)}$'s impression and pays 0.1
Outcome:	
Advertiser $a^{(1)}$ wins both impressions and pays 0.2 per impression	For User $u^{(2)}$'s impression:
	Bids:
	Advertiser $a^{(1)}$: $b_2^{(1)} = 0.1$; and $a^{(2)}$: $b_2^{(2)} = 0.3$
	Outcome:
	Advertiser $a^{(2)}$ wins $u^{(2)}$'s impression and pays 0.1
Platform's expected revenue:	Platform's expected revenue:
$R = 2 \times 0.2 = 0.4$	$R = 0.1 + 0.1 = 0.2$
Advertiser's expected surplus:	Advertiser's expected surplus:
$W^{(1)} = 2 \times (0.3 - 0.2) = 0.2$	$W^{(1)} = 0.5 - 0.1 = 0.4$
$W^{(2)} = 0$	$W^{(2)} = 0.3 - 0.1 = 0.2$
Total expected surplus:	Total expected surplus:
$S = 0.6$	$S = 0.8$

Table A.5: Example depicting two regimes: 1) No data disclosure and 2) Full disclosure.

impression is softer. This ensures higher advertiser revenues, with $W^{(1)} = 0.4$ and $W^{(2)} = 0.2$. However, the platform's revenue is now lower, at $R = 0.2$. Thus, even though ads are matched more efficiently and the total surplus generated is higher, the platform extracts less revenue.

This example illustrates the platform's trade-off between value creation and value appropriation, and highlights the platform's incentives to withhold targeting information from advertisers.

A.7 Proofs

A.7.1 Proof of Proposition 1

We first prove the proposition for the general case of two sets of bundles $\mathcal{I}^{(1)}$ and $\mathcal{I}^{(2)}$, where $\mathcal{I}^{(1)}$ is at least as granular as $\mathcal{I}^{(2)}$. Under a given targeting regime, risk-neutral advertisers' valuation for an impression in a bundle is their aggregate valuation for that bundle, as they cannot distinguish between different impressions within a bundle. Thus, all impressions within a bundle have the same expected valuation for an advertiser. As such, for any $i \in I_j$, incentive compatibility constraint in the second-price auction induce advertiser a to places the bid $b_{ia} = \frac{1}{|I_j|} \sum_{k \in I_j} v_{ka}$.

According to Definition 4, for any $I_k^{(2)} \in \mathcal{I}^{(2)}$, there exist $I_{j_1}^{(1)}, \ldots, I_{j_l}^{(1)}$ such that $\bigcup_{s=1}^{l} I_{j_s}^{(1)} = I_k^{(2)}$. We can write:

$$\max_a \sum_{i \in I_k^{(2)}} v_{ia} \leq \sum_{s=1}^{l} \max_a \sum_{i \in I_{j_s}^{(1)}} v_{ia} \tag{A.22}$$

where the LHS is the surplus generated from the impressions in bundle $I_k^{(2)}$ for targeting regime $\mathcal{I}^{(2)}$, and RHS is the surplus generated from the same impressions for targeting regime $\mathcal{I}^{(1)}$. Since inequality (A.22) holds for any $I_k^{(2)} \in \mathcal{I}^{(2)}$, we can show that $S^{(1)} \geq S^{(2)}$, *i.e.*, the surplus is increasing with more granular targeting.

However, platform revenues can go either way. If the second-highest valuation is exactly the same as the highest valuation for all impressions, the revenue and surplus are identical, i.e., $R^{(1)} \geq R^{(2)}$. On the other hand, consider a case where, for any given impression, one advertiser has a high valuation and the rest of advertisers have the same (lower) valuation. In this case, the second-highest valuation is the minimum valuation for each impression. Thus, we can write (A.22) with min function instead of max and change the sign of inequality. This gives us $R^{(1)} \leq R^{(2)}$. Therefore, the platform revenues can be non-monotonic with more granular targeting.

Finally, it is worth noting that this result can be applied to any efficient auction in light of revenue-equivalence theorem.

A.7.2 Proof of Proposition 2

To write the FOC, we first need to specify advertisers' expected utility function. If ad a gets a click in impression i, the utility he extracts from this impression is $v_a^{(c)} - CPC_{ia}$. As such, ad a's utility from impression i is

$$\left(v_a^{(c)} - CPC_{ia}\right) \mathbb{1}(a_i = a) \mathbb{1}(y_{i,a} = 1),$$

where $\mathbb{1}(a_i = a)$ indicates whether ad a is shown in impression i and $\mathbb{1}(y_{i,a} = 1)$ is an indicator for whether this impression is clicked. Clearly, ad a only extracts utility when his ad is being shown and clicked in an impression. We further assume that $\mathbb{1}(a_i = a)$ is independent of CPC_{ia} with respect $\mathcal{G}_a$. Now, we define the expected utility function EU_a as the expectation over impressions as follows:

$$\begin{aligned} EU_a(b_a; \mathcal{G}_a) &= \mathbb{E}_{\mathcal{G}_a}\left[\left(v_a^{(c)} - CPC_{ia}\right) \mathbb{1}(a_i = a) \mathbb{1}(y_i = 1)\right] \\ &= \left(v_a^{(c)} - c_a(b_a)\right) \pi_a(b_a) \mu_a \\ &= \left(v_a^{(c)} - c_a(b_a)\right) \frac{b_a q_a}{b_a q_a + Q_{-a}} \mu_a \end{aligned} \tag{A.23}$$

where $\mu_a = \mathbb{E}_{\mathcal{G}_a}\left[\mathbb{1}(y_i = 1)\right]$ is the average probability of click that ad a expects to receive. The second line of Equation (A.23) is resulted because of the independence of $\mathbb{1}(a_i = a)$ and CPC_{ia} with respect $\mathcal{G}_a$.

Now, to write the FOC, we need to take the first derivative of the expected utility function for ad a. We can write:

$$\frac{\partial EU_a(b_a;\mathcal{G}_a)}{\partial b_a} = \left(v_a^{(c)}\pi_a'(b_a) - c_a(b_a)\pi_a'(b_a) - c_a'(b_a)\pi_a(b_a)\right)\mu_a \tag{A.24}$$

Now, to satisfy the FOC, we need to have $\frac{\partial EU_a(b_a^*;\mathcal{G}_a)}{\partial b_a^*} = 0$, where b_a^* is the equilibrium bid. Using Equation (A.24), we can write the FOC as follows:

$$\begin{aligned} v_a^{(c)} &= c_a(b_a^*) + \frac{c_a'(b_a^*)\pi_a(b_a^*)}{\pi_a'(b_a^*)} \\ &= c_a(b_a^*) + \frac{c_a'(b_a^*)b_a^*}{1-\pi_a(b_a^*)}, \end{aligned} \tag{A.25}$$

where the second line is because we have $\frac{\pi_a(b_a)}{\pi_a'(b_a)} = \frac{b_a}{1-\pi_a(b_a)}$ which is easy to show:

$$\frac{\pi_a(b_a)}{\pi_a'(b_a)} = \frac{\frac{b_a q_a}{b_a q_a + Q_{-a}}}{\frac{q_a Q_{-a}}{(b_a q_a + Q_{-a})^2}} = \frac{b_a(b_a q_a + Q_{-a})}{Q_{-a}} = \frac{b_a}{1-\pi_a(b_a)} \tag{A.26}$$

Now, to complete the proof, we need to show that the second-order condition (SOC) is also satisfied. We start by writing a useful property in the relationship between the first and second derivative of the allocation function:

$$\frac{\pi_a'(b_a)}{\pi_a''(b_a)} = \frac{\frac{q_a Q_{-a}}{(b_a q_a + Q_{-a})^2}}{\frac{-2q_a^2 Q_{-a}}{(b_a q_a + Q_{-a})^3}} = \frac{(b_a q_a + Q_{-a})}{2q_a} = -\frac{b_a}{2\pi_a(b_a)} \tag{A.27}$$

We can now write the second derivative of the expected utility function as follows:

$$\begin{aligned} \frac{\partial^2 EU_a(b_a^*;\mathcal{G}_a)}{\partial {b_a^*}^2} &= \mu_a\left(\left(v_a^{(c)} - c_a(b_a^*)\right)\pi_a''(b_a^*) - c_a'(b_a^*)\pi_a'(b_a^*) - c_a'(b_a^*)\pi_a'(b_a^*) - c_a''(b_a^*)\pi_a(b_a^*)\right) \\ &= \mu_a\left(\frac{c_a'(b_a^*)\pi_a(b_a^*)}{\pi_a'(b_a^*)}\pi_a''(b_a^*) - 2c_a'(b_a^*)\pi_a'(b_a^*) - c_a''(b_a^*)\pi_a(b_a^*)\right) \\ &= \mu_a\left(c_a'(b_a^*)\pi_a(b_a^*)\left(\frac{\pi_a''(b_a^*)}{\pi_a'(b_a^*)} - 2\frac{\pi_a'(b_a^*)}{\pi_a(b_a^*)}\right) - c_a''(b_a^*)\pi_a(b_a^*)\right) \\ &= \mu_a\left(c_a'(b_a^*)\pi_a(b_a^*)\left(\frac{2\pi_a(b_a^*)}{b_a^*} - 2\frac{1-\pi_a(b_a^*)}{b_a^*}\right) - c_a''(b_a^*)\pi_a(b_a^*)\right) \\ &= \mu_a\left(-\frac{2}{b_a^*}c_a'(b_a^*)\pi_a(b_a^*) - c_a''(b_a^*)\pi_a(b_a^*)\right) \\ &= \mu_a\pi_a(b_a^*)\left(-\frac{2}{b_a^*}c_a'(b_a^*) - c_a''(b_a^*)\right) \end{aligned} \tag{A.28}$$

In the equation above, we know that both μ_a and $\pi_a(b_a^*)$ are positive, so it is sufficient to show that $-\frac{2}{b_a^*}c_a'(b_a^*) - c_a''(b_a^*) \leq 0$, which is resulted from the assumption in Proposition 2. Thus, the SOC is satisfied and this completes the proof.

A.8 Analysis of Cost-per-Click Payment Mechanism

We now present an analysis of targeting under the Cost-per-Click (CPC) mechanism, where an advertiser's bid indicates the maximum price he is willing to pay per click. In this case, having a more accurate estimation of match values for impressions does not change advertisers' bidding behavior because they do not pay per impression. Theoretically, if advertisers have infinite budget, they should bid their valuation for click, even if they know they have higher CTR for some impressions.

However, the ad-network has an incentive to make more efficient matches in order to generate more clicks since clicks are their main source of revenue. For example, if there is an advertiser with a very high bid but very low CTR, the ad-network cannot make money by selling the slot to this advertisers. Let v_{ia} and m_{ia} respectively denote the valuation for click and the match value for advertiser a for impression i. The maximum revenue that platform could get is then $\max_a v_{ia}m_{ia}$. Thus, defining b_{ia} as advertiser a's bid on impression i, the platform should sell the ad slot to $\text{argmax}_a\, b_{ia}m_{ia}$ and charge her the minimum bid with which she still wins. It generates the expected revenue of the second-highest $b_{ia}m_{ia}$. (In fact, this is how Google's sponsored search auctions work.)

Table A.6 shows how the example in §2.6.1 generalizes to the CPC case. Although CPC and CPI have different properties, we can easily prove that their revenue is the same under different levels of targeting. This stems from the fact that under second-price auction, bidders bid their valuation. Thus, the platform's expected revenue from impression i under both mechanisms is the second highest $v_{ia}m_{ia}$, as long as the match-values (or targeting strategies) are the same under both mechanisms. Conceptually, there exists a one-to-one mapping between CPC and CPI mechanisms if and only if there is a one-to-one mapping between the targeting regime and resulting match-value matrix. In the CPI mechanism, m_{ia}

No Targeting	Perfect Targeting
For both impressions: Bids: Advertiser 1: $b_{11} = b_{21} = 1$ Advertiser 2: $b_{12} = b_{22} = 1$ Match values: Advertiser 1: $m_{11} = m_{21} = 0.3$ Advertiser 2: $m_{12} = m_{22} = 0.2$ Outcomes: Advertiser 1 wins both impressions and pays $\frac{0.2 \times 1}{0.3} = 0.66$ per click	For User 1's impression: Bids: Advertiser 1: $b_{11} = 1$, Advertiser 2: $b_{12} = 1$ Match values: Advertiser 1: $m_{11} = 0.5$, Advertiser 2: $m_{12} = 0.1$ Outcome: Advertiser 1 wins User 1's impression and pays $\frac{1 \times 0.1}{0.5} = 0.2$ per click For User 2's impression: Bids: Advertiser 1: $b_{11} = 1$, Advertiser 2: $b_{12} = 1$ Match values: Advertiser 1: $m_{11} = 0.1$, Advertiser 2: $m_{12} = 0.3$ Outcome: Advertiser 2 wins User 2's impression and pays $\frac{1 \times 0.1}{0.3} = 0.33$ per click
Platform's expected revenue: $R = 0.66 \times (0.5 + 0.1) = 0.4$	Platform's expected revenue: $R = 0.2 \times 0.5 + 0.33 \times 0.3 = 0.2$
Advertiser's expected surplus: $W_1 = (1 - 0.66) \times (0.5 + 0.1) = 0.2$ $W_2 = 0$	Advertiser's expected surplus: $W_1 = (1 - 0.2) \times 0.5 = 0.4$ $W_2 = (1 - 0.33) \times 0.3 = 0.2$
Total expected surplus: $S = 0.6$	Total expected surplus: $S = 0.8$

Table A.6: Example depicting no targeting and perfect targeting under CPC pricing mechanism.

enters the advertisers' bidding strategy, i.e., the targeting decision is made by the advertisers. The ad-network's decision consists of whether to share targeting information with advertisers or not. In contrast, under the CPC mechanism, the ad-network directly decides the extent of targeting to engage in and the advertisers always bid their valuation.

Our empirical results suggest that under the CPI mechanism, ad-networks may have incentive to limit behavioral targeting. In the CPC context, this translates to the following result: the ad-network has an incentive to not use behavioral information for targeting, as compared to contextual information. In both cases, the platform has an incentive to protect users' privacy by ensuring that behavioral information is not used for targeting purposes. In sum, the empirical estimates of platform's revenue, advertisers' revenues, and the implications for user-privacy are similar in both settings.

A.9 Robustness Checks for Analysis of Revenue-Efficiency Trade-off

A.9.1 Alternative Methods for Estimation of Click Valuations

As discussed in §2.6.3, we make three simplifications to derive Equation (2.15).

1. **Simplification 1:** Advertisers' probability of winning an impression is approximately zero, i.e., $\pi_a(b_a) \approx 0$, for all a.

2. **Simplification 2:** Advertisers' cost-per-click is approximately their bid, i.e., $c_a(b_a) \approx b_a$, for all a.

3. **Simplification 3:** The first-order condition shown in Equation (2.14) is satisfied for all advertisers, including reserve price bidders.

If one or more of these simplifications fail in our data, our estimates of click valuations may be biased. This, in turn, can affect our main qualitative findings. Therefore, we now present six alternative methods to estimate click valuations that progressively relax these simplifications and show that our main findings are robust to these simplifications.

Simplification	Main	AM1	AM2	AM3	AM4	AM5	AM6
Zero winning probability ($\pi_a \approx 0$)	✓	×	×	×	×	×	×
Pay-your-bid cost function ($c_a(b_a) \approx b_a$)	✓	✓	×	×	×	×	×
FOC satisfied for reserve price bidders	✓	✓	✓	✓	×	×	×
Parametric cost function	✓	✓	✓	✓	✓	✓	×

Table A.7: Table of simplifications imposed by each estimation method

Table A.7 gives an overview of the set of simplifications that each alternative method relaxes. We now describe each of these methods in detail.

1. **Alternative Method 1 (AM1):** In this method, we relax Simplification 1, but retain Simplifications 2 and 3.

   Specifically, instead of assuming that $\pi_a(b_a) \approx 0$, we estimate the advertiser's expected probability of ad allocation directly from the data as:

   $$\hat{\pi}_a(b_a^*) = \frac{\sum_{i=1}^{N_F} \mathbb{1}(a_i = a)}{\sum_{i=1}^{N_F} e_{i,a}}. \tag{A.29}$$

   This is the total number of impressions showing ad a divided by the total number of impressions ad a bid on (i.e., participated in the auction for). This is a consistent estimator of the proportion of impressions that ad a will win in the platform.

   Since we still assume that advertisers pay their own bid per click (Simplification 2), we have $c_a(b_a) = b_a$. This gives us the click value estimates as:

   $$\hat{v}_a^{(AM1)} = b_a^* + \frac{b_a^*}{1 - \hat{\pi}_a(b_a^*)}, \tag{A.30}$$

2. **Alternative Method 2 (AM2):** In this method, we relax both Simplifications 1 and 2. Like AM1, here also we estimate the probability of winning (proportion of impressions allocated) from the data using Equation (A.29). In addition, we no longer

assume that advertisers pay their bid. Instead, we model their payment upon winning as a linear function of their bid such that $c_a(b_a) = \gamma_a b_a$. Then, using Equation (2.14), we can write:

$$\hat{v}_a^{(AM2)} = \hat{c}_a(b_a^*) + \frac{\hat{c}_a(b_a^*)}{1 - \hat{\pi}_a(b_a^*)} \tag{A.31}$$

This is because if $c_a(b_a)$ is linear, we have $c_a(b_a) = c_a'(b_a)b_a$. Hence, we only need to estimate $\hat{c}_a(b_a^*)$, which is the average cost-per-click for ad a.

3. **Alternative Method 3 (AM3):** This method is similar to AM2. The main difference is that, here we model the cost function more flexibly by allowing it to be a quadratic function of bids. We can characterize our cost function as $c_a(b_a) = \gamma_a b_a + \gamma b_a^2$. Note that the coefficient for b_a is ad specific, but the one for b_a^2 is the same for all ads. The is because we cannot identify advertiser-specific curvature since advertisers do not change their bids.

   We can estimate this quadratic cost function from the data using the following regression model:

$$CPC_i = \gamma_a b_a + \gamma b_a^2 + \epsilon_i, \tag{A.32}$$

   where CPC_i is the actual cost-per-click for impression i.

   Once we have the estimates, $\hat{\gamma}_a$ and $\hat{\gamma}$, we can estimate click valuations using Equation (2.14) as follows:

$$\hat{v}_a^{(AM3)} = \hat{c}_a(b_a^*) + \frac{b_a^*(\hat{\gamma}_a + 2\hat{\gamma}b_a^*)}{1 - \hat{\pi}_a(b_a^*)} = \hat{c}_a(b_a^*) + \frac{\hat{c}_a(b_a^*)}{1 - \hat{\pi}_a(b_a^*)} + \frac{\hat{\gamma}b_a^*}{1 - \hat{\pi}_a(b_a^*)} \tag{A.33}$$

   Equation (A.33) is very similar to Equation (A.31), except for the last term $\frac{\hat{\gamma}b_a^*}{1-\hat{\pi}_a(b_a^*)}$, which is added to incorporate the curvature in the relationship between bid and cost-per-click. Of course, if the relationship is linear, (A.33) will be identical Equation (A.31).

   While both AM2 and AM3 allow the cost-per-click to differ from the actual bid, they formulate an advertiser's cost-per-click only as a function of their own bids, and not that of its competitors. Given Proposition 2, our click valuation estimates are valid to the extent that our parametric cost functions estimate $c'_a(b^*_a)$ well. Nevertheless, in AM6, we estimate the advertisers' cost function without imposing any functional form assumptions, and allow it to depend on other advertisers' bids as inputs.

4. **Alternative Method 4 (AM4):** In this method, we relax all three simplifications. We build on AM3 and incorporate the fact that there is a reserve price r_0 for all impressions (i.e., advertisers need to bid higher than or equal to this amount).

   In the presence of a reserve price, equilibrium bids may not satisfy the first-order condition since they may be right-censored. That is, there could have been bids lower than r_0 without the reserve price. For reserve-price bidders, we know that the participation constraint $v_a^{(c)} \geq r_0$ is satisfied. On the other hand, we know that $b^*_a \leq r_0$ which implies that $v_a^{(c)} \leq r_0 \left(1 + \frac{1}{1-\pi_a(b_a)}\right)$. Thus, for reserve bidders, we have $v_a^c \in [r_0, r_0(1+\frac{1}{1-\pi_a(b_a)})]$. As such, if such valuations come from a uniform distribution in this range, our estimates will be:

$$\hat{v}_a^{(AM4)} = \begin{cases} r_0 \left(1 + \frac{1}{2(1-\hat{\pi}_a(b^*_a))}\right) & \text{if } a \text{ is a reserve bidder} \\ \hat{v}_a^{(AM3)} & \text{otherwise} \end{cases}, \tag{A.34}$$

   where the estimated click valuations for a reserve bidder a is the mean of uniform distribution $\mathcal{U}\left(r_0, r_0\left(1 + \frac{1}{1-\hat{\pi}_a(b^*_a)}\right)\right)$.

5. **Alternative Method 5 (AM5):** In this method, we follow all the steps in AM4, but draw a value from distribution $\mathcal{U}(r_0, r_0(1 + \frac{1}{1-\hat{\pi}_a(b^*_a)}))$ for any ad a who bids the reserve price. As such, our estimates will be:

$$\hat{v}_a^{(AM5)} = \begin{cases} x \sim \mathcal{U}(r_0, r_0(1 + \frac{1}{1-\hat{\pi}_a(b^*_a)})) & \text{if } a \text{ is a reserve bidder} \\ \hat{v}_a^{(AM3)} & \text{otherwise} \end{cases} \tag{A.35}$$

6. **Alternative Method 6 (AM6):** Here, we implement the the non-parametric estimator proposed by [105] for quasi-proportional auctions. This is a structural approach

to estimate click valuations that relaxes all the simplifications made earlier and incorporates the full structure of the setting. The main difference between his approach and the alternative methods AM1–AM5 is in the operationalization of the cost function. Unlike the earlier methods, which parameterized the cost function, he employs a fully non-parametric method for modeling the cost distribution. As such, advertisers' expected cost function is obtained by taking the expectation over the distribution of auctions they participate in. We refer readers to [105] for more details.

Results from Alternative Methods for Click Value Estimation

We first illustrate the empirical CDFs of all these alternative methods as well as the main method used in Chapter 2. As we can see from Figure A.5, the estimated value distributions are not very different. This is mostly due to the fact that Simplification 1 is mostly valid (i.e., there are many bidders and no one has disproportionately high chance of winning), which in turn, makes the approximation method used in (2.15) quite reasonable.

Next, we use the click value estimates from each of these alternative methods to estimate the market outcomes – total surplus, platform revenues, and advertisers' surplus. The results are shown in Table A.8. While these estimates are quantitatively different from those presented in Table 2.5, notice that our main qualitative results remain the same: we find a monotonic relationship between total surplus and the granularity of targeting, but platform revenues are maximized when the platform limits the targeting level to the contextual targeting. In particular, we expect that results from AM6 to be the most accurate since it is the most assumption-free method, and even here the main qualitative finding holds.

Further, we find that the results from AM1 are the closest to the results in the main text (Table 2.5). This is because AM1 still uses the approximation that $CPC_a \approx b_a$. On the other hand, the results from AM2 to AM6 are very close quantitatively. This is because they all relax Simplification 2 and estimate the cost function from the data. As a result, the magnitude of total surplus and platform revenues is smaller (because advertisers in a quasi-proportional auction submit higher bids when they know they are not exactly charged their

Targeting	Total Surplus	Platform Revenue	Advertisers' Surplus
	A. Alternative Method 1 (AM1)		
No Targeting	8.54	8.36	0.18
Contextual Targeting	9.09	8.49	0.60
Behavioral Targeting	9.32	8.47	0.85
Full Targeting	9.60	8.48	1.12
	B. Alternative Method 2 (AM2)		
No Targeting	7.96	7.93	0.03
Contextual Targeting	8.62	8.09	0.53
Behavioral Targeting	8.83	8.05	0.78
Full Targeting	9.09	8.05	1.04
	C. Alternative Method 3 (AM3)		
No Targeting	7.91	7.89	0.02
Contextual Targeting	8.61	8.07	0.53
Behavioral Targeting	8.81	8.04	0.77
Full Targeting	9.08	8.04	1.04
	E. Alternative Method 4 (AM4)		
No Targeting	8.02	7.91	0.11
Contextual Targeting	8.61	8.08	0.53
Behavioral Targeting	8.82	8.03	0.79
Full Targeting	9.08	8.04	1.04
	D. Alternative Method 5 (AM5)		
No Targeting	8.02	7.91	0.11
Contextual Targeting	8.61	8.08	0.53
Behavioral Targeting	8.82	8.03	0.79
Full Targeting	9.08	8.03	1.05
	F. Alternative Method 6 (AM6)		
No Targeting	8.03	7.91	0.12
Contextual Targeting	8.61	8.08	0.53
Behavioral Targeting	8.82	8.03	0.78
Full Targeting	9.08	8.04	1.04

Table A.8: Platform revenues, advertisers' surplus, and total surplus for different levels of targeting using different methods to estimate click valuations. The numbers are reported in terms of the average monetary unit per impression.

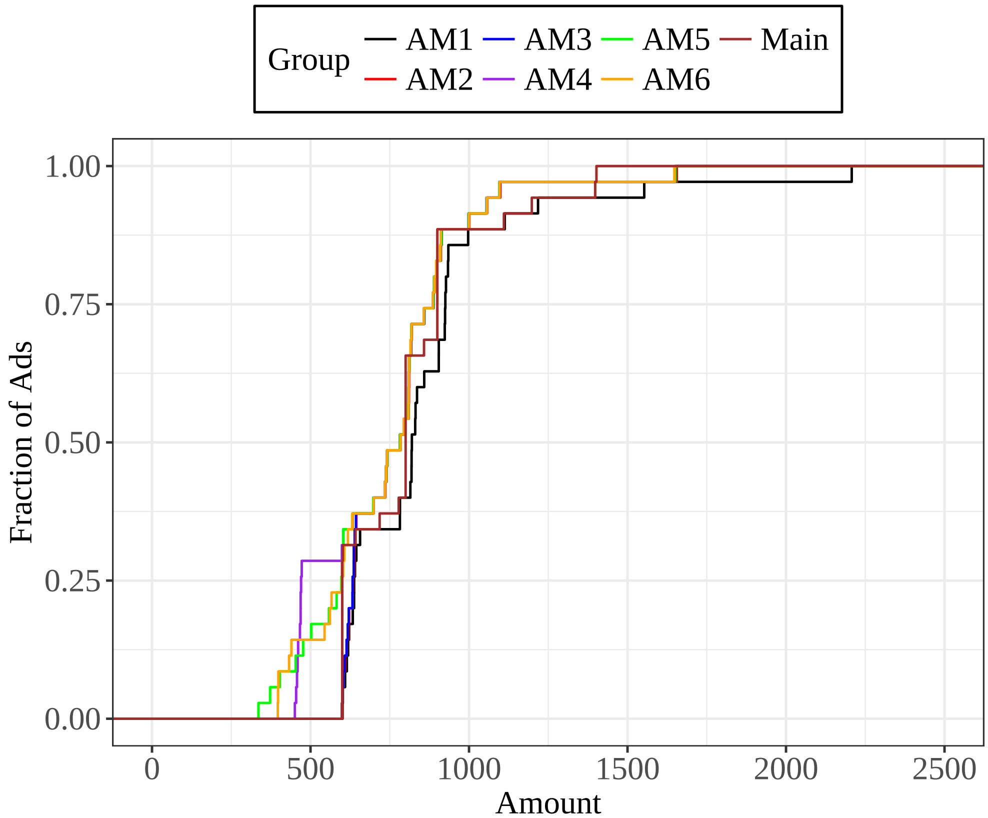

Figure A.5: Empirical CDF of estimated value distributions using alternative methods

submitted bids). However, the main qualitative findings remains the same under all these methods, since the main source for these findings is the heterogeneity in match valuations induced by more granular targeting.

A.9.2 Adding Noise to Match Valuations

As discussed in §2.6.3, our assumption that advertisers can obtain match values estimated in our machine learning framework may not be realistic. As such, we consider various scenarios here to reflect cases in which advertisers can only obtain a noisy version of our estimates.

Identically Distributed Noise Across Ads

First, we consider the case that an identically distributed noise is added to any element in the match value matrix. We operationalize this noise as follows:

$$\hat{m}_{i,a}^{(\nu)} = \hat{m}_{i,a}\epsilon_{i,a}, \tag{A.36}$$

where $\epsilon_{i,a} \sim \mathcal{U}(1-\nu, 1+\nu)$. As such, advertisers obtain estimates $\nu \times 100$ percent lower or higher than the estimate from our ML framework. For any noise distribution, we make the entire matrix and follow the procedure presented in §2.6.3 to estimate the average surplus, revenue, and advertisers' surplus. The results are shown in Table A.9. We find that noise-adding can increase the revenue under full targeting by distorting efficiency. This is in line with [7] who find that the platform's optimal decision is to use imperfect CTR estimates in a cost-per-click setting. Interestingly, we find that full and contextual targeting yield almost the same results when $\nu = 0.3$. Thus, we can interpret contextual targeting as a noise-adding strategy.

Differentially Distributed Noise Across Ads

Here we consider a more realistic case wherein ads with more impressions (more data) have better estimates (less noisy). Similar to the case with identical noise, we we operationalize the differential noise as follows for any ad a:

$$\hat{m}_{i,a}^{(\nu_a)} = \hat{m}_{i,a}\epsilon_{i,a}, \tag{A.37}$$

where $\epsilon_{i,a} \sim \mathcal{U}(1-\nu_a, 1+\nu_a)$. As such, the noise is different for each ad. Since the error rate is proportional to the square root of number of observations, we write:

$$\nu_a = \frac{\nu}{\sqrt{N_a}},$$

where N_a is the number of impressions ad a has in our Filtered sample. Now, for any ν, we can make the entire match value matrix and estimate the market outcomes. The results

Targeting	Noise (ν)	Total Surplus	Platform Revenue	Advertisers' Surplus
No Targeting	0.05	8.36	8.30	0.06
Contextual Targeting	0.05	8.99	8.44	0.55
Behavioral Targeting	0.05	9.18	8.35	0.82
Full Targeting	0.05	9.43	8.35	1.08
No Targeting	0.1	8.36	8.30	0.06
Contextual Targeting	0.1	8.99	8.44	0.55
Behavioral Targeting	0.1	9.17	8.36	0.82
Full Targeting	0.1	9.37	8.35	1.02
No Targeting	0.15	8.36	8.30	0.06
Contextual Targeting	0.15	8.99	8.44	0.55
Behavioral Targeting	0.15	9.16	8.36	0.80
Full Targeting	0.15	9.29	8.35	0.93
No Targeting	0.2	8.36	8.30	0.06
Contextual Targeting	0.2	8.99	8.44	0.55
Behavioral Targeting	0.2	9.14	8.36	0.78
Full Targeting	0.2	9.20	8.37	0.82
No Targeting	0.25	8.36	8.30	0.06
Contextual Targeting	0.25	8.99	8.44	0.55
Behavioral Targeting	0.25	9.13	8.37	0.76
Full Targeting	0.25	9.10	8.41	0.70
No Targeting	0.3	8.36	8.30	0.06
Contextual Targeting	0.3	8.99	8.44	0.54
Behavioral Targeting	0.3	9.12	8.38	0.74
Full Targeting	0.3	9.01	8.47	0.54

Table A.9: Platform revenues, advertisers' surplus, and total surplus for different levels of targeting when adding an identically distributed noise to match values. The numbers are reported in terms of the average monetary unit per impression.

are shown in Table A.10. Again, we find that the platform benefits when advertisers have imperfect estimates of their match valuations. Comparing the results in Table A.9 and A.10, we find that the total surplus declines faster with a differentially distributed noise: when the platform's revenue is the same under full and contextual targeting, the total surplus is higher under contextual targeting. Given the absence of privacy costs under contextual targeting, our results in both tables indicate that platforms benefit most when allowing only contextual targeting.

Targeting	Noise (ν)	Total Surplus	Platform Revenue	Advertisers' Surplus
No Targeting	10	8.36	8.30	0.06
Contextual Targeting	10	8.99	8.44	0.55
Behavioral Targeting	10	9.18	8.35	0.82
Full Targeting	10	9.42	8.34	1.08
No Targeting	20	8.36	8.30	0.06
Contextual Targeting	20	8.99	8.44	0.55
Behavioral Targeting	20	9.17	8.35	0.82
Full Targeting	20	9.35	8.32	1.03
No Targeting	30	8.36	8.30	0.06
Contextual Targeting	30	8.99	8.44	0.55
Behavioral Targeting	30	9.16	8.35	0.80
Full Targeting	30	9.24	8.31	0.93
No Targeting	40	8.36	8.30	0.06
Contextual Targeting	40	8.99	8.44	0.55
Behavioral Targeting	40	9.14	8.35	0.79
Full Targeting	40	9.12	8.32	0.80
No Targeting	50	8.36	8.30	0.06
Contextual Targeting	50	8.99	8.44	0.55
Behavioral Targeting	50	9.12	8.35	0.77
Full Targeting	50	8.99	8.35	0.65
No Targeting	60	8.36	8.30	0.06
Contextual Targeting	60	8.99	8.44	0.55
Behavioral Targeting	60	9.10	8.35	0.75
Full Targeting	60	8.93	8.36	0.57
No Targeting	70	8.36	8.30	0.06
Contextual Targeting	70	8.99	8.43	0.56
Behavioral Targeting	70	9.09	8.34	0.75
Full Targeting	70	8.86	8.39	0.47
No Targeting	80	8.36	8.30	0.06
Contextual Targeting	80	8.99	8.43	0.55
Behavioral Targeting	80	9.07	8.34	0.73
Full Targeting	80	8.80	8.44	0.36

Table A.10: Platform revenues, advertisers' surplus, and total surplus for different levels of targeting when adding an differentially distributed noise to match values. The numbers are reported in terms of the average monetary unit per impression.

Appendix B
APPENDIX FOR CHAPTER 3

B.1 Proof of Proposition 3

The distribution of propensity scores for ads is fully determined by the allocation rule in a quasi-proportional auction. That is, for any ad a in session i, the probability that this ad wins an impression is denoted by $\pi_{i,a}$, characterized as follows:

$$\pi_{i,a} = \mathbb{1}(a \in \mathcal{C}_i)\frac{b_{i,a}m_{i,a}}{\sum_{k\in\mathcal{C}_i} b_{i,k}m_{i,k}}, \tag{B.1}$$

where $\mathcal{C}_i$ is the set of ads competing in the auction for session i, and $b_{i,a}m_{i,a}$ is ad a's quality-adjusted bid which is the product of the ad's bid ($b_{i,a}$) and quality score ($m_{i,a}$). If we observe all the components of this equation, the distribution of propensity scores for ads is a function of observed variables by default. However, we do not observe the quality scores in our data. So we need to show that the distribution of propensity scores is still fully identified even without having the quality scores. We first use the following lemma that helps us establish the identifiability of the distribution of propensity scores in a specific sub-sample of our sessions:

Lemma 8 *Let G denote a set of sessions where the auction is identical for any two sessions i and j that $i, j \in G$. The distribution of propensity scores for ads is identified if we observe the actual ad assignment in our data.*

Proof. *If we know the actual ad assignments for impressions in G, the proportion of impressions in G that show a in is an accurate estimate of the propensity score for that ad because the distribution of ad assignment is identical across impressions in G.* ■

In light of this lemma, if we know the actual ad assignments in the data, we do not need to observe bids or quality scores to identify the distribution of propensity scores in any group G

where the auction is identical across impressions within that group. Now, if we show that such partitioning or stratification of our sessions is feasible, our proof is complete. We use the fact that the ad-network does not update quality scores throughout the day. Hence, to make sure that the quality scores remain constant in all the impressions within a partition, it is sufficient to partition our sessions such that sessions in different days are not in the same partition. Finally, since we directly observe bids, we only need to show that we can identify $\mathcal{C}_i$. This set can only vary if sessions are different in two dimensions: (1) targeting characteristics, because some ads may decide to exclude some sessions based on their targeting characteristics, and/or (2) time because some ad campaigns may be unavailable at some points in time. Since we observe all the targeting characteristics as well as the exact timestamp of each impression, we can identify the groups where all the sessions have the same $\mathcal{C}_i$, and this completes our proof.

B.2 Results for Robustness Checks

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
$Breadth_{i,t}$	0.00062	0.00168***	0.00085***	0.00046**	0.00066***	0.00052***	0.00047***
	(1.23)	(5.53)	(4.24)	(2.96)	(5.38)	(4.88)	(5.28)
Previous	0.00056	0.00170***	0.00040*	0.00045**	0.00043***	0.00023*	0.00025**
Exposure ($I_{i,t}$)	(1.06)	(5.39)	(1.98)	(2.95)	(3.65)	(2.28)	(3.13)
$Space_{i,t}$	0.00115***	0.00080***	0.00004	0.00020*	0.00030***	0.00006	-0.00006
	(4.79)	(4.98)	(0.42)	(2.54)	(4.80)	(1.14)	(-1.55)
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	449,667	519,690	573,096	616,887	652,311	681,160	711,391
R^2	0.157	0.129	0.120	0.104	0.093	0.076	0.075
Adjusted R^2	0.039	0.033	0.037	0.032	0.028	0.017	0.022
Note:					*p<0.05; **p<0.01; ***p<0.001		

Table B.1: OLS estimates of the effects of variety after controlling for temporal spacing using the intended sample with at least one prior exposure of the focal ad

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
$Breadth_{i,t}$	0.00062**	0.00070***	0.00037**	0.00018	0.00024**	0.00023**	0.00022***
	(3.01)	(4.45)	(3.26)	(1.90)	(2.98)	(3.21)	(3.70)
Previous	0.00090**	0.00098***	0.00028	0.00031*	0.00028**	0.00014	0.00013
Exposure ($I_{i,t}$)	(2.69)	(4.25)	(1.80)	(2.49)	(2.86)	(1.61)	(1.92)
$Space_{i,t}$	0.00101***	0.00080***	0.00014	0.00021**	0.00035***	0.00008	-0.00004
	(4.59)	(5.44)	(1.44)	(2.77)	(6.05)	(1.72)	(-1.11)
First Exposure	0.00240**	0.00191**	-0.00002	0.00033	0.00111**	-0.00022	-0.00049
Dummy	(3.27)	(3.24)	(-0.05)	(0.83)	(3.27)	(-0.69)	(-1.76)
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
R^2	0.099	0.088	0.087	0.077	0.071	0.063	0.061
Adjusted R^2	0.028	0.024	0.029	0.024	0.022	0.017	0.019
Note:					*p<0.05; **p<0.01; ***p<0.001		

Table B.2: OLS estimates of the effects of variety after controlling for temporal spacing using the intended sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
$Shannon_{i,t}$	0.00123***	0.00184***	0.00135***	0.00073*	0.00108***	0.00123***	0.00116***
	(3.51)	(5.47)	(4.59)	(2.56)	(4.08)	(4.72)	(4.73)
Previous	0.00042**	0.00074***	0.00040***	0.00033***	0.00017*	0.00024***	0.00023***
Exposure ($I_{i,t}$)	(2.72)	(5.80)	(4.11)	(3.91)	(2.47)	(3.79)	(4.35)
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
R^2	0.099	0.088	0.087	0.077	0.071	0.063	0.061
Adjusted R^2	0.028	0.024	0.029	0.024	0.022	0.017	0.019
Note:					*p<0.05; **p<0.01; ***p<0.001		

Table B.3: OLS estimates of the effects of variety as measure by Shannon entropy using the intended sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
			A. Breadth of Variety				
$Breadth_{i,t}$	0.00051**	0.00069***	0.00039***	0.00023**	0.00023***	0.00024***	0.00025***
	(3.29)	(5.83)	(4.56)	(3.19)	(3.91)	(4.72)	(5.71)
Previous	0.00014	0.00045***	0.00010	0.00010	0.00006	0.00010*	0.00009*
Exposure ($I_{i,t}$)	(1.05)	(4.28)	(1.33)	(1.51)	(1.12)	(2.23)	(2.51)
R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.002
			B. Consecutive Changes				
$Change_{i,t}$	0.00042**	0.00071***	0.00037***	0.00028***	0.00019***	0.00025***	0.00021***
	(2.78)	(6.30)	(4.54)	(4.16)	(3.53)	(5.37)	(5.57)
Previous	0.00012	0.00044***	0.00009	0.00011	0.00005	0.00010*	0.00009*
Exposure ($I_{i,t}$)	(0.87)	(4.23)	(1.19)	(1.67)	(0.87)	(2.26)	(2.30)
R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.002
			C. Gini-Simpson Index for Diversity				
$GiniSimpson_{i,t}$	0.00098**	0.00222***	0.00178***	0.00112*	0.00135**	0.00215***	0.00201***
	(2.80)	(5.26)	(4.20)	(2.53)	(3.11)	(4.93)	(4.79)
Previous	0.00013	0.00045***	0.00011	0.00009	0.00006	0.00012*	0.00010*
Exposure ($I_{i,t}$)	(0.96)	(4.23)	(1.38)	(1.43)	(1.04)	(2.54)	(2.53)
R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.002
Ad FE	✓	✓	✓	✓	✓	✓	✓
Province FE	✓	✓	✓	✓	✓	✓	✓
Connectivity Type FE	✓	✓	✓	✓	✓	✓	✓
MSP FE	✓	✓	✓	✓	✓	✓	✓
Smartphone Brand FE	✓	✓	✓	✓	✓	✓	✓
Hour FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426

Note: *p<0.05; **p<0.01; ***p<0.001

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
				A. Breadth of Variety			
$Breadth_{i,t}$	0.00054***	0.00072***	0.00041***	0.00028***	0.00024***	0.00028***	0.00026***
	(3.41)	(5.99)	(4.61)	(3.75)	(3.85)	(5.15)	(5.59)
Previous	0.00025	0.00057***	0.00023**	0.00021**	0.00011*	0.00017***	0.00015***
Exposure ($I_{i,t}$)	(1.85)	(5.25)	(2.83)	(3.05)	(1.97)	(3.49)	(3.65)
R^2	0.012	0.013	0.011	0.010	0.010	0.009	0.009
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
				B. Consecutive Changes			
$Change_{i,t}$	0.00046**	0.00077***	0.00039***	0.00031***	0.00022***	0.00030***	0.00023***
	(2.99)	(6.65)	(4.67)	(4.55)	(3.99)	(6.27)	(5.70)
Previous	0.00023	0.00056***	0.00022**	0.00021**	0.00010	0.00017***	0.00014***
Exposure ($I_{i,t}$)	(1.70)	(5.25)	(2.73)	(3.15)	(1.88)	(3.63)	(3.53)
R^2	0.012	0.013	0.011	0.010	0.010	0.009	0.009
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
				C. Gini-Simpson Index for Diversity			
$GiniSimpson_{i,t}$	0.00107**	0.00245***	0.00199***	0.00154***	0.00153***	0.00276***	0.00213***
	(2.98)	(5.65)	(4.54)	(3.33)	(3.35)	(5.95)	(4.73)
Previous	0.00024	0.00057***	0.00024**	0.00021**	0.00011*	0.00019***	0.00015***
Exposure ($I_{i,t}$)	(1.78)	(5.27)	(2.96)	(3.05)	(1.99)	(4.00)	(3.67)
R^2	0.012	0.013	0.011	0.010	0.010	0.009	0.009
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
Ad FE	✓	✓	✓	✓	✓	✓	✓
Province-Hour FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
Note:							*p<0.05; **p<0.01; ***p<0.001

Table B.5: OLS estimates of the effects of variety on clicks with province-hour fixed effects using the intended sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
			A. Breadth of Variety				
$Breadth_{i,t}$	0.00052***	0.00070***	0.00039***	0.00024**	0.00024***	0.00025***	0.00025***
	(3.33)	(5.91)	(4.46)	(3.27)	(3.97)	(4.82)	(5.75)
Previous	0.00015	0.00046***	0.00012	0.00012	0.00007	0.00011*	0.00009*
Exposure ($I_{i,t}$)	(1.13)	(4.35)	(1.47)	(1.85)	(1.27)	(2.36)	(2.41)
R^2	0.007	0.008	0.006	0.006	0.004	0.004	0.004
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
			B. Consecutive Changes				
$Change_{i,t}$	0.00041**	0.00072***	0.00037***	0.00028***	0.00019***	0.00025***	0.00022***
	(2.71)	(6.36)	(4.53)	(4.16)	(3.45)	(5.46)	(5.76)
Previous	0.00012	0.00045***	0.00011	0.00013*	0.00005	0.00011*	0.00008*
Exposure ($I_{i,t}$)	(0.93)	(4.30)	(1.35)	(1.98)	(1.00)	(2.38)	(2.25)
R^2	0.007	0.008	0.006	0.006	0.004	0.004	0.004
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.003
			C. Gini-Simpson Index for Diversity				
$GiniSimpson_{i,t}$	0.00099**	0.00225***	0.00173***	0.00117**	0.00134**	0.00218***	0.00211***
	(2.81)	(5.31)	(4.08)	(2.64)	(3.06)	(4.96)	(4.97)
Previous	0.00014	0.00046***	0.00012	0.00012	0.00006	0.00012**	0.00010*
Exposure ($I_{i,t}$)	(1.04)	(4.30)	(1.51)	(1.77)	(1.16)	(2.65)	(2.49)
R^2	0.007	0.008	0.006	0.006	0.004	0.004	0.004
Adjusted R^2	0.005	0.006	0.004	0.004	0.003	0.003	0.002
Ad FE	✓	✓	✓	✓	✓	✓	✓
Targeting Area FE	✓	✓	✓	✓	✓	✓	✓
Hour FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
Note:							*p<0.05; **p<0.01; ***p<0.001

Table B.6: OLS estimates of the effects of variety on clicks with targeting area and hour fixed effects using the intended sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
				A. Breadth of Variety			
$Breadth_{i,t}$	0.00083***	0.00080***	0.00042***	0.00029**	0.00032***	0.00029***	0.00021**
	(4.24)	(5.18)	(3.64)	(2.89)	(3.82)	(3.88)	(3.21)
Previous	0.00054**	0.00078***	0.00047***	0.00034***	0.00016*	0.00024***	0.00018**
Exposure ($I_{i,t}$)	(3.22)	(5.57)	(4.36)	(3.66)	(2.00)	(3.41)	(2.98)
R^2	0.147	0.130	0.124	0.109	0.101	0.090	0.085
Adjusted R^2	0.042	0.037	0.041	0.035	0.033	0.028	0.028
				B. Consecutive Changes			
$Change_{i,t}$	0.00074***	0.00079***	0.00043***	0.00034***	0.00031***	0.00039***	0.00023***
	(3.86)	(5.22)	(3.79)	(3.58)	(3.82)	(5.50)	(3.74)
Previous	0.00052**	0.00076***	0.00047***	0.00035***	0.00015	0.00026***	0.00018**
Exposure ($I_{i,t}$)	(3.08)	(5.46)	(4.33)	(3.77)	(1.88)	(3.73)	(3.05)
R^2	0.147	0.130	0.124	0.109	0.101	0.090	0.085
Adjusted R^2	0.042	0.037	0.041	0.035	0.033	0.028	0.028
				C. Gini-Simpson Index for Diversity			
$GiniSimpson_{i,t}$	0.00176***	0.00272***	0.00248***	0.00164*	0.00255***	0.00281***	0.00218**
	(3.93)	(4.78)	(4.18)	(2.53)	(3.82)	(4.00)	(3.11)
Previous	0.00054**	0.00078***	0.00051***	0.00034***	0.00018*	0.00026***	0.00019**
Exposure ($I_{i,t}$)	(3.19)	(5.56)	(4.63)	(3.64)	(2.21)	(3.64)	(3.13)
R^2	0.147	0.130	0.124	0.109	0.101	0.090	0.085
Adjusted R^2	0.042	0.037	0.041	0.035	0.033	0.028	0.028
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata ($\frac{1}{2}$ hour) FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
Note:							*p<0.05; **p<0.01; ***p<0.001

Table B.7: OLS estimates of the effects of variety on clicks with strata constructed by half an hour time intervals using the intended sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
			A. Breadth of Variety				
$Breadth_{i,t}$	0.00088***	0.00117***	0.00034*	0.00028*	0.00047***	0.00021	0.00005
	(3.61)	(5.81)	(2.17)	(2.01)	(3.81)	(1.83)	(0.47)
Previous	0.00053**	0.00063***	0.00010	0.00028*	0.00016	0.00010	0.00028**
Exposure ($I_{i,t}$)	(2.65)	(3.60)	(0.70)	(2.27)	(1.56)	(1.05)	(3.28)
R^2	0.285	0.248	0.221	0.193	0.174	0.150	0.138
Adjusted R^2	0.109	0.098	0.090	0.078	0.072	0.058	0.055
			B. Consecutive Changes				
$Change_{i,t}$	0.00077**	0.00100***	0.00039**	0.00030*	0.00030**	0.00036***	0.00007
	(3.27)	(5.28)	(2.68)	(2.35)	(2.77)	(3.69)	(0.86)
Previous	0.00049*	0.00056**	0.00011	0.00029*	0.00011	0.00017	0.00029***
Exposure ($I_{i,t}$)	(2.47)	(3.25)	(0.82)	(2.35)	(1.09)	(1.76)	(3.44)
R^2	0.285	0.248	0.221	0.193	0.174	0.150	0.138
Adjusted R^2	0.109	0.098	0.090	0.078	0.072	0.058	0.055
			C. Gini-Simpson Index for Diversity				
$GiniSimpson_{i,t}$	0.00182***	0.00339***	0.00184*	0.00095	0.00254**	0.00094	0.00162
	(3.31)	(4.70)	(2.37)	(1.09)	(2.76)	(0.95)	(1.62)
Previous	0.00053**	0.00060***	0.00013	0.00025	0.00016	0.00008	0.00034***
Exposure ($I_{i,t}$)	(2.60)	(3.38)	(0.93)	(1.93)	(1.44)	(0.79)	(3.72)
R^2	0.285	0.248	0.221	0.193	0.174	0.150	0.138
Adjusted R^2	0.109	0.098	0.090	0.078	0.072	0.058	0.055
Ad-Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426	1,054,426
Note:					*p<0.05; **p<0.01; ***p<0.001		

Table B.8: OLS estimates of the effects of variety on clicks with ad-strata fixed effects using the intended sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
				A. Breadth of Variety			
$Breadth_{i,t}$	0.00188**	0.00235***	0.00189***	0.00068	0.00135***	0.00067*	0.00069*
	(3.04)	(4.33)	(4.35)	(1.73)	(3.98)	(2.07)	(2.42)
Previous	0.00118*	0.00193***	0.00098**	0.00108***	0.00091***	0.00031	0.00079***
Exposure ($I_{i,t}$)	(2.56)	(4.63)	(2.97)	(3.73)	(3.70)	(1.36)	(4.01)
R^2	0.212	0.186	0.178	0.147	0.143	0.118	0.113
Adjusted R^2	0.075	0.067	0.074	0.053	0.058	0.040	0.041
				B. Consecutive Changes			
$Change_{i,t}$	0.00131*	0.00158***	0.00156***	0.00059	0.00060*	0.00076**	0.00052*
	(2.32)	(3.40)	(4.39)	(1.89)	(2.28)	(3.11)	(2.46)
Previous	0.00093*	0.00158***	0.00086**	0.00107***	0.00064**	0.00043	0.00079***
Exposure ($I_{i,t}$)	(2.09)	(3.98)	(2.74)	(3.84)	(2.68)	(1.89)	(4.03)
R^2	0.212	0.186	0.178	0.147	0.143	0.118	0.113
Adjusted R^2	0.075	0.067	0.074	0.053	0.058	0.040	0.041
				C. Gini-Simpson Index for Diversity			
$GiniSimpson_{i,t}$	0.00409**	0.00564**	0.00912***	0.00409	0.00858***	0.00537	0.01056***
	(3.00)	(3.01)	(4.47)	(1.76)	(3.50)	(1.96)	(3.83)
Previous	0.00120**	0.00170***	0.00118***	0.00118***	0.00105***	0.00043	0.00120***
Exposure ($I_{i,t}$)	(2.58)	(3.94)	(3.33)	(3.66)	(3.70)	(1.56)	(4.99)
R^2	0.212	0.186	0.178	0.147	0.143	0.118	0.113
Adjusted R^2	0.075	0.067	0.074	0.053	0.058	0.040	0.042
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	188,892	189,232	187,562	188,462	187,077	187,895	187,188
Note:							*p<0.05; **p<0.01; ***p<0.001

Table B.9: OLS estimates of the effects of variety on clicks for the top ad using the intended sample

	Dependent variable: Click ($Y_{i,t}$)						
	$(t=4)$	$(t=5)$	$(t=6)$	$(t=7)$	$(t=8)$	$(t=9)$	$(t=10)$
				A. Breadth of Variety			
$Breadth_{i,t}$	0.01371***	0.01700***	0.00846***	0.00888***	0.00890***	0.00807***	0.00699***
	(4.40)	(5.95)	(3.56)	(3.85)	(4.24)	(3.96)	(3.63)
Previous	0.00471	0.00411	0.00030	-0.00144	-0.00019	0.00303	0.00262
Exposure ($I_{i,t}$)	(1.68)	(1.54)	(0.13)	(-0.65)	(-0.09)	(1.53)	(1.41)
R^2	0.565	0.596	0.599	0.614	0.633	0.644	0.644
Adjusted R^2	0.035	0.059	0.031	0.028	0.036	0.026	-0.007
				B. Consecutive Changes			
$Change_{i,t}$	0.00843**	0.01242***	0.00308	0.00631**	0.00917***	0.00914***	0.00765***
	(2.67)	(4.34)	(1.31)	(2.80)	(4.47)	(4.61)	(4.08)
Previous	0.00343	0.00281	-0.00113	-0.00222	-0.00045	0.00306	0.00268
Exposure ($I_{i,t}$)	(1.23)	(1.05)	(-0.50)	(-1.01)	(-0.23)	(1.56)	(1.45)
R^2	0.564	0.596	0.599	0.614	0.633	0.644	0.644
Adjusted R^2	0.034	0.059	0.030	0.027	0.036	0.027	-0.007
				C. Gini-Simpson Index for Diversity			
$GiniSimpson_{i,t}$	0.02760***	0.05400***	0.04075***	0.04251**	0.05747***	0.05679**	0.06266**
	(3.82)	(5.07)	(3.34)	(2.83)	(3.49)	(2.99)	(3.06)
Previous	0.00447	0.00396	0.00054	-0.00168	-0.00027	0.00282	0.00270
Exposure ($I_{i,t}$)	(1.59)	(1.47)	(0.23)	(-0.75)	(-0.13)	(1.39)	(1.42)
R^2	0.565	0.596	0.599	0.614	0.633	0.643	0.644
Adjusted R^2	0.035	0.059	0.031	0.027	0.036	0.026	-0.008
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
No. of Obs.	57,363	50,684	45,882	41,332	37,773	34,717	32,241
Note:							*p<0.05; **p<0.01; ***p<0.001

Table B.10: OLS estimates of the effects of variety on clicks for the first time users are exposed to the experimental condition using the survived sample

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
Variety ($V_{i,t}$)	0.0062**	0.0116***	0.0070***	0.0061***	0.0057***	0.0070***	0.0059***
	(2.73)	(5.86)	(4.25)	(3.92)	(3.91)	(5.10)	(4.65)
breadth-constant ($Q_{i,t}$)	-0.0002	0.0042	-0.0004	0.0031	-0.0000	0.0040**	0.0040**
Changes	(-0.07)	(1.55)	(-0.20)	(1.78)	(-0.02)	(2.95)	(3.24)
Previous Exposure ($I_{i,t}$)	0.0001	0.0033	-0.0003	0.0004	0.0006	0.0024*	0.0020
	(0.05)	(1.91)	(-0.22)	(0.26)	(0.48)	(2.16)	(1.87)
Ad FE	✓	✓	✓	✓	✓	✓	✓
Province-Hour-Day FE	✓	✓	✓	✓	✓	✓	✓
N	58981	54271	50844	47273	44060	41077	38528
R^2	0.156	0.183	0.172	0.177	0.183	0.191	0.184
Adjusted R^2	0.057	0.080	0.060	0.059	0.058	0.059	0.042
Note:					*p<0.05; **p<0.01; ***p<0.001		

Table B.11: Linear regressions estimates of click model with for the first time users are exposed to the experimental condition controlling for province-hour-day fixed effects

	Dependent variable: Click ($Y_{i,t}$)						
	($t = 4$)	($t = 5$)	($t = 6$)	($t = 7$)	($t = 8$)	($t = 9$)	($t = 10$)
Variety ($V_{i,t}$)	0.0105**	0.0149***	0.0083**	0.0072*	0.0070*	0.0085**	0.0060*
	(2.58)	(4.13)	(2.69)	(2.49)	(2.49)	(3.23)	(2.39)
Breadth-Constant ($Q_{i,t}$)	-0.0013	-0.0007	-0.0047	0.0008	0.0025	0.0051	0.0037
Changes	(-0.20)	(-0.14)	(-1.23)	(0.23)	(0.81)	(1.85)	(1.46)
Previous Exposure ($I_{i,t}$)	0.0017	0.0017	0.0006	-0.0008	-0.0019	0.0025	0.0023
	(0.49)	(0.53)	(0.20)	(-0.31)	(-0.79)	(1.11)	(1.06)
Ad FE	✓	✓	✓	✓	✓	✓	✓
Strata FE	✓	✓	✓	✓	✓	✓	✓
N	58981	54271	50844	47273	44060	41077	38528
R^2	0.573	0.596	0.599	0.612	0.629	0.641	0.635
Adjusted R^2	0.070	0.092	0.078	0.083	0.092	0.096	0.055
Note:	*p<0.05; **p<0.01; ***p<0.001						

Table B.12: Linear regressions estimates of click model with for the first time users are exposed to the experimental condition controlling for auction-invariant strata fixed effects

Appendix C

APPENDIX FOR CHAPTER 5

C.1 Proofs

Proof of Lemma 1. In a mechanism M, we can write the maximum utility advertiser a can receive as follows:

$$\max_{b_a} u^M_{i,a,t}(b_a; x_a, S_{i,t}) = \max_{b_a} \mathbb{E}_{x_{-a}}\left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) - e^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right], \tag{C.1}$$

where the expectation is over other advertisers' click valuations, as we have IC. The first derivative of advertiser a with respect to her click valuation x_a as follows:

$$\frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} = \mathbb{E}_{x_{-a}}\left[P(Y_{i,a} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] \tag{C.2}$$

Given IC constraint, we know that $u^M_{i,a,t}(x_a; x_a, S_{i,t}) = \max_{b_a} u^M_{i,a,t}(b_a; x_a, S_{i,t})$. Therefore, based on envelope theorem, we have:

$$\frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} \Big|_{b_a = x_a} = \mathbb{E}_{x_{-a}}\left[P(Y_{i,a} \mid a, S_{i,t}) q^M_{i,a,t}(x; S_{i,t})\right] \tag{C.3}$$

Now, since u is differentiable, we can write:

$$\begin{aligned} u^M_{i,a,t}(x_a; x_a, S_{i,t}) - u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) &= \int_{x'_a}^{x_a} \frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} db_a \\ &= P(Y_{i,a} \mid a, S_{i,t}) \int_{x'_a}^{x_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a \end{aligned} \tag{C.4}$$

This directly implies Equation (5.5) and completes the proof for Lemma 1. ■

Proof of Lemma 2. We can write the publisher's revenues as follows:

$$\begin{aligned} \mathbb{E}_x\left[\sum_{a \in \mathcal{A}_i} e^M_{i,a,t}(x; S_{i,t})\right] &= \sum_{a \in \mathcal{A}_i} \mathbb{E}_x\left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t})\right] \\ &- \sum_{a \in \mathcal{A}_i} \mathbb{E}_x\left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) - e^M_{i,a,t}(x; S_{i,t})\right], \end{aligned} \tag{C.5}$$

where each component of the second term is advertiser a's surplus and allows us to write it as follows:

$$\mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i} e^M_{i,a,t}(x;S_{i,t})\right] = \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right] - \sum_{a\in\mathcal{A}_i}\mathbb{E}_{x_a}\left[u^M_{i,a,t}(x_a;x_a,S_{i,t})\right] \tag{C.6}$$

Using calculus theorems and Lemma 1, we can transform each element of the second term. For brevity and with slight abuse of notation, we drop $S_{i,t}$ from the function specification, and define $p_{i,a} = P(Y_{i,a} \mid a, S_{i,t})$. We can write:

$$\begin{aligned}
\mathbb{E}_{x_a}\left[u^M_{i,a,t}(x_a;x_a)\right] &= \int_{\underline{x}_a}^{\bar{x}_a} u^M_{i,a,t}(x_a)f_a(x_a)dx_a \\
&= \int_{\underline{x}_a}^{\bar{x}_a}\left(u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{\underline{x}_a}^{x_a}\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a,x_{-a})\right]db_a\right)f_a(x_a)dx_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{\underline{x}_a}^{\bar{x}_a}\int_{\underline{x}_a}^{x_a}\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a,x_{-a})\right]f_a(x_a)db_a dx_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{\underline{x}_a}^{\bar{x}_a}\int_{b_a}^{\bar{x}_a} f_a(x_a)\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a,x_{-a})\right]dx_a db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{\underline{x}_a}^{\bar{x}_a}\left(1-F_a(b_a)\right)\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a,x_{-a})\right]db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{\underline{x}_a}^{\bar{x}_a}\left(1-F_a(b_a)\right)\int_{x_{-a}} q^M_{i,a,t}(b_a,x_{-a})f_{-a}(x_{-a})dx_{-a}db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{\underline{x}_a}^{\bar{x}_a}\frac{1-F_a(b_a)}{f_a(b_a)}\int_{x_{-a}} q^M_{i,a,t}(b_a,x_{-a})f_{-a}(x_{-a})dx_{-a}f_a(b_a)db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{x}\frac{1-F_a(b_a)}{f_a(b_a)}q^M_{i,a,t}(b_a,x_{-a})f(b_a,x_{-a})dx_{-a}db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{x}\frac{1-F_a(x_a)}{f_a(x_a)}q^M_{i,a,t}(x)f(x)dx \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a}\mathbb{E}_x\left[\frac{1-F_a(x_a)}{f_a(x_a)}q^M_{i,a,t}(x)\right]
\end{aligned} \tag{C.7}$$

Now we can transform the publisher's revenues in Equation (C.6) using Equation (C.7) to complete the proof:

$$\begin{aligned}\mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i} e^M_{i,a,t}(x;S_{i,t})\right] &= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right]\\ &\quad-\sum_{a\in\mathcal{A}_i}u^M_{i,a,t}(\underline{x}_a;\underline{x}_a)-\sum_{a\in\mathcal{A}_i}p_{i,a}\mathbb{E}_x\left[\frac{1-F_a(x_a)}{f_a(x_a)}q^M_{i,a,t}(x)\right]\\ &=\mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)p_{i,a}q^M_{i,a,t}(x;S_{i,t})\right]-\sum_{a\in\mathcal{A}_i}u^M_{i,a,t}(\underline{x}_a;\underline{x}_a,S_{i,t})\end{aligned}$$

■

Proof of Proposition 4. Given the payments in Equation (5.8), it is easy to check that $u^M_{i,a,t}(\underline{x}_a;\underline{x}_a,S_{i,t})=0$. We can write:

$$\begin{aligned}u^M_{i,a,t}(\underline{x}_a;\underline{x}_a,S_{i,t}) &= \mathbb{E}_{x_{-a}}\left[w_{i,a,t}(\underline{x}_a;S_{i,t})q^M_{i,a,t}(\underline{x}_a,x_{-a};S_{i,t})-e^M_{i,a,t}(\underline{x}_a,x_{-a};S_{i,t})\right]\\ &=\mathbb{E}_{x_{-a}}\left[P(Y_{i,t}\mid a,S_{i,t})\int_{\underline{x}_a}^{\underline{x}_a}q^M_{i,a,t}(b_a,x_{-a};S_{i,t})db_a\right]\\ &=0\end{aligned}\tag{C.8}$$

This implies that the second part of the publisher's revenues in Equation (5.6) is zero, i.e., $\sum_{a\in\mathcal{A}_i}u^M_{i,a,t}(\underline{x}_a;\underline{x}_a,S_{i,t})=0$. Thus, since the q is chosen to maximize the first part of Equation (5.6) given some constraints, we know that mechanism M is optimal given those constraint. It is now sufficient to show the following two statements: 1) mechanism M is a direct revelation mechanism, and 2) any direct mechanism satisfies the constraints: q is plausible and $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ is increasing in x_a.

We start by showing that M is a direct revelation mechanism. The plausibility is satisfied by definition. We only need to show both IR and IC. Given the payment function, we can write the utility function for advertiser a as follows:

$$\begin{aligned}u^M_{i,a,t}(x_a;x_a,S_{i,t}) &= \mathbb{E}_{x_{-a}}\left[w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})-e^M_{i,a,t}(x;S_{i,t})\right]\\ &=\mathbb{E}_{x_{-a}}\left[P(Y_{i,t}\mid a,S_{i,t})\int_{\underline{x}_a}^{x_a}q^M_{i,a,t}(b_a,x_{-a};S_{i,t})db_a\right]\\ &=P(Y_{i,t}\mid a,S_{i,t})\int_{\underline{x}_a}^{x_a}\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a,x_{-a};S_{i,t})\right]db_a\end{aligned}\tag{C.9}$$

Given $u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) = 0$, this is equivalent to the envelope condition. Since the integral on the RHS is always non-negative, $u_{i,a,t}$ is increasing, which combined with $u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) = 0$ imply IR constraint. We now need to show IC constraint is also satisfied. We prove that by contradiction. Suppose that there is an x'_a that gives a higher utility to advertiser a than truthful reporting, given everyone else bidding their true click valuations. Let γ denote the gains advertiser a receives by reporting x'_a instead of x_a. We can write:

$$\begin{aligned}
\gamma &= u^M_{i,a,t}(x'_a; x_a, S_{i,t}) - u^M_{i,a,t}(x_a; x_a, S_{i,t}) \\
&= u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) - \left(u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) - u^M_{i,a,t}(x'_a; x_a, S_{i,t})\right) - u^M_{i,a,t}(x_a; x_a, S_{i,t}) \\
&= u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) - u^M_{i,a,t}(x_a; x_a, S_{i,t}) - (x'_a - x_a)P(Y_{i,t} \mid a, S_{i,t})\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] \\
&= P(Y_{i,t} \mid a, S_{i,t})\left(\int_{x_a}^{x'_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a - (x'_a - x_a)\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right]\right) \\
&= P(Y_{i,t} \mid a, S_{i,t}) \int_{x_a}^{x'_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) - q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] db_a
\end{aligned} \tag{C.10}$$

Now, given that $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a, it is easy to check $\gamma \leq 0$ regardless of whether $x'_a > x_a$ or not. This completes the proof of part 1: mechanism M is a direct revelation mechanism.

Now, we show the second part: any direct revelation mechanism satisfies the constraints: q is plausible and $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a. A direct revelation mechanism M satisfies IC constraints. Hence, for $x'_a > x_a$ we can write:

$$u^M_{i,a,t}(x_a; x_a, S_{i,t}) \geq u^M_{i,a,t}(x'_a; x_a, S_{i,t}) \tag{C.11}$$

$$u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) \geq u^M_{i,a,t}(x_a; x'_a, S_{i,t}) \tag{C.12}$$

Subtracting these two equations, we have:

$$u^M_{i,a,t}(x_a; x_a, S_{i,t}) - u^M_{i,a,t}(x_a; x'_a, S_{i,t}) \geq u^M_{i,a,t}(x'_a; x_a, S_{i,t}) - u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) \tag{C.13}$$

Simplifying Equation (C.23) gives us:

$$(x'_a - x_a)\left(\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] - \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]\right) \geq 0 \tag{C.14}$$

Since $x'_a > x_a$, we can show that $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a. This completes the proof for the second part. ■

Proof of Lemma 3. The proof for this lemma is almost identical to the proof for Lemma 1. We start by writing down the maximization problem advertiser a faces in mechanism M:

$$\max_{b_a} U^M_{i,a}(b_a; x_a, S_i) = \max_{b_a} \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}\left(w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) - e^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right)\right], \tag{C.15}$$

where the expectation is over other advertisers' click valuations as well as the stochasticity induce by the dynamic process. The reason we take the expectation over other advertisers' click valuation is the main condition in the lemma: IC constraint. We can write the first derivative of advertiser a with respect to her click valuation x_a as follows:

$$\frac{\partial U^M_{i,a}(b_a; x_a, S_i)}{\partial x_a} = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right], \tag{C.16}$$

where, again, the expectation is over other advertisers' click valuations as well as the stochasticity induce by the dynamic process. IC constraint implies that $U^M_{i,a}(x_a; x_a, S_i) = \max_{b_a} U^M_{i,a}(b_a; x_a, S_i)$. We can now apply the envelope theorem as follows:

$$\frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} |_{b_a = x_a} = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] \tag{C.17}$$

Now, due to the first-differentiability of U, we have:

$$\begin{aligned} U^M_{i,a}(x_a; x_a, S_i) - U^M_{i,a}(x'_a; x'_a, S_{i,t}) &= \int_{x'_a}^{x_a} \frac{\partial U^M_{i,a}(b_a; x_a, S_i)}{\partial x_a} db_a \\ &= \int_{x'_a}^{x_a} \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a, \end{aligned} \tag{C.18}$$

which is equivalent to Equation (5.12) and completes the proof for Lemma 3. ■

Proof of Lemma 4. The steps for this proof is almost identical to the proof for Lemma

2. We start by writing the publisher's objective function:

$$
\begin{aligned}
\mathbb{E}\left[\sum_{a\in\mathcal{A}_i} e^M_{i,a}(x;S_i)\right] &= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right] \\
&\sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t}) - e^M_{i,a}(x;S_i)\right] \\
&= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right] - \sum_{a\in\mathcal{A}_i}\mathbb{E}_{x_a}\left[U^M_{i,a}(x_a;x_a,S_i)\right],
\end{aligned}
\tag{C.19}
$$

where all the expectations are over the specified click valuations and the stochasticity induced by the dynamic process. Now, we can transform each element of the second term. For brevity and with slight abuse of notation, we drop S_i and $S_{i,t}$ from the function specification, and define $p_{i,a,t} = P(Y_{i,a} \mid a, S_{i,t})$. We can write:

$$
\begin{aligned}
\mathbb{E}_{x_a}\left[U_{i,a}^M(x_a;x_a)\right] &= \int_{\underline{x}_a}^{\bar{x}_a} U_{i,a}^M(x_a) f_a(x_a) dx_a \\
&= \int_{\underline{x}_a}^{\bar{x}_a} \left(U_{i,a}^M(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{x_a} \mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(b_a, x_{-a})\right] db_a \right) f_a(x_a) dx_a \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a} \int_{\underline{x}_a}^{x_a} \mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(b_a, x_{-a})\right] f_a(x_a) db_a dx_a \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a} \int_{b_a}^{\bar{x}_a} f_a(x_a) \mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(b_a, x_{-a})\right] dx_a db_a \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a} \left(1 - F_a(b_a)\right) \mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(b_a, x_{-a})\right] db_a \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{\underline{x}_a}^{\bar{x}_a} \left(1 - F_a(b_a)\right) \int_{x_{-a}} \sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(b_a, x_{-a}) f_{-a}(x_{-a}) dx_{-a} db_a \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a} \frac{1 - F_a(b_a)}{f_a(b_a)} \int_{x_{-a}} \sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(b_a, x_{-a}) f_{-a}(x_{-a}) dx_{-a} f_a(b_a) db_a \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + \int_{x} \frac{1 - F_a(b_a)}{f_a(b_a)} \sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(b_a, x_{-a}) f(b_a, x_{-a}) dx_{-a} db_a \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + \int_{x} \frac{1 - F_a(x_a)}{f_a(x_a)} \sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(x) f(x) dx \\
&= U_{i,a}^M(\underline{x}_a;\underline{x}_a) + p_{i,a} \mathbb{E}_x\left[\frac{1 - F_a(x_a)}{f_a(x_a)} \sum_{t=1}^{\infty} \beta^{t-1} p_{i,a,t} q_{i,a,t}^M(x)\right]
\end{aligned}
\tag{C.20}
$$

Now we can re-write Equation (C.19) using Equation (C.20) to complete the proof:

$$\begin{aligned}
\mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i} e^M_{i,a}(x;S_i)\right] &= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right] \\
&\quad - \sum_{a\in\mathcal{A}_i} U^M_{i,a}(\underline{x}_a;\underline{x}_a) - \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\frac{1-F_a(x_a)}{f_a(x_a)}\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(x)\right] \\
&= \mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i}\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)p_{i,a,t}q^M_{i,a,t}(x;S_{i,t})\right] \\
&\quad - \sum_{a\in\mathcal{A}_i} U^M_{i,a}(\underline{x}_a;\underline{x}_a,S_i) \\
&= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)p_{i,a,t}q^M_{i,a,t}(x;S_{i,t})\right] \\
&\quad - \sum_{a\in\mathcal{A}_i} U^M_{i,a}(\underline{x}_a;\underline{x}_a,S_i)
\end{aligned}$$

■

Proof of Proposition 5. The proof for this proposition is very similar to the one for Proposition 4. We begin by providing the necessary and sufficient conditions for the IC constraint. We first show that a mechanism M is IC if and only if $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ is increasing in x_a and we have the envelope condition as presented in Equation (5.12).

We start our proof by the only if part. We want to show that if the mechanism M is IC, then $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ is increasing in x_a and we have the envelope condition as presented in Equation (5.12). The latter is the result of Lemma 3. So we only need to show the that IC implies monotonicity. Given the IC, for any $x'_a > x_a$ we can write:

$$U^M_{i,a}(x_a;x_a,S_i) \geq U^M_{i,a}(x'_a;x_a,S_i) \tag{C.21}$$

$$U^M_{i,a}(x'_a;x'_a,S_i) \geq U^M_{i,a}(x_a;x'_a,S_i) \tag{C.22}$$

Subtracting these two equations, we have:

$$U^M_{i,a}(x_a;x_a,S_i) - U^M_{i,a}(x_a;x'_a,S_i) \geq U^M_{i,a}(x'_a;x_a,S_i) - U^M_{i,a}(x'_a;x'_a,S_i) \tag{C.23}$$

Simplifying Equation (C.23) gives us:

$$(x'_a - x_a)\left(\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]\right) \geq 0 \tag{C.24}$$

Since $x'_a > x_a$, we can show that $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a. This completes the *only if* part of the statement.

Now, we show the *if* part of the statement. If we have plausibility and monotonicity conditions as above, then we IC constraint is satisfied. We assume that the IC is not satisfied and then show contradiction. If IC is not satisfied, there exists an x'_a that gives a higher utility to advertiser a than truthful reporting, given IC for other advertisers. We denote the gains from deviating by γ. We can write:

$$\begin{aligned}
\gamma &= U^M_{i,a}(x'_a; x_a, S_i) - U^M_{i,a}(x_a; x_a, S_i) \\
&= U^M_{i,a}(x'_a; x'_a, S_i) - \left(U^M_{i,a}(x'_a; x'_a, S_i) - U^M_{i,a}(x'_a; x_a, S_i)\right) - U^M_{i,a}(x_a; x_a, S_i) \\
&= U^M_{i,a}(x'_a; x'_a, S_i) - U^M_{i,a}(x_a; x_a, S_i) \\
&\quad - (x'_a - x_a)\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] \\
&= \int_{x_a}^{x'_a}\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right]db_a \\
&\quad - (x'_a - x_a)\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] \\
&= \int_{x_a}^{x'_a}\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right. \\
&\quad \left. - \sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right]db_a
\end{aligned} \tag{C.25}$$

Now, given that $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a, S_{i,t})q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a, it is easy to check $\gamma \leq 0$ independent of the relationship between x'_a and x_a. This contradicts the assumption that $\gamma > 0$ and shows that mechanism M is IC.

Now we show that the proposed mechanism is optimal. This mechanism maximizes the first component in Equation (5.13) subject to the plausibility and monotonicity conditions, while setting the payment such that the second component is zero. Given the payments, we can write:

$$\begin{aligned} U^M_{i,a}(x_a; x_a, S_i) &= \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x_a, x_{-a}; S_{i,t}) - e^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right)\right] \\ &= \int_{\underline{x}_a}^{x_a} \mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a \end{aligned} \tag{C.26}$$

Equation Equation (C.26) shows that $U^M_{i,a}(\underline{x}_a; \underline{x}_a, S_i) = 0$ for all a, which implies that we have the envelope condition as in Equation (5.12). Together, these imply IR constraint. Further, the envelope condition and the monotonicity constraint are equivalent to the IC constraint. Therefore, IC is also satisfied in our case.

Now we only need to show the optimality of the mechanism M. Since $\sum_{a\in\mathcal{A}_i} U^M_{i,a}(\underline{x}_a; \underline{x}_a, S_i) = 0$, the choice of q maximizes the publisher's objective given the plausibility and monotonicity constraint. Since these two constraints are necessary for any direct revelation mechanism, the mechanism M is optimal.

This equivalence proves two statements: 1) IC is satisfied, and 2) any direct revelation mechanism must satisfy monotonicity and plausibility. ■

Proof of Lemma 5. We only need to show that

$$\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$$

is increasing in x_a. Let q^M and $q^{M'}$ denote the optimal allocation functions derived by Equation (5.17) and Equation (5.18) for click valuation profiles x and x' respectively. Since these mechanisms are optimal for corresponding cases, we can write the following two inequalities:

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a - \frac{1-F_a(x_a)}{f_a(x_a)}\right) P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x_a; S_{i,t})\right)\right] \\ \geq\ &\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a - \frac{1-F_a(x_a)}{f_a(x_a)}\right) P(Y_{i,t} \mid a, S_{i,t}) q^{M'}_{i,a,t}(x'_a; S_{i,t})\right)\right], \end{aligned} \tag{C.27}$$

and

$$\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x'_a-\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right)\right]$$
$$\geq\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x'_a-\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M}_{i,a,t}(x_a;S_{i,t})\right)\right] \tag{C.28}$$

Subtracting these two inequalities will give us the following inequality:

$$\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M}_{i,a,t}(x_a;S_{i,t})\right)\right]$$
$$\geq\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right)\right] \tag{C.29}$$

Now, suppose that x and x' are the same at each element except the a-th element, i.e., $x_a\neq x'_a$ and $x_j=x'_j$ for all $j\neq a$. Further, suppose that $x_a>x'_a$. We can then write:

$$\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M}_{i,a,t}(x_a;S_{i,t})\right]$$
$$\geq\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right] \tag{C.30}$$

Now, since the distribution F_a is regular, we have:

$$\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^{M}_{i,a,t}(x_a;S_{i,t})\right]\geq\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right] \tag{C.31}$$

The last inequality directly implies $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^{M}_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ increasing in x_a and completes the proof. ∎