

The Privacy–Utility Trade-Off of Location Tracking in Ad Personalization

Mohammad Mosaffa*
Cornell University

Omid Rafeian*
Cornell University

Abstract

Firms collect vast amounts of behavioral and geographical data on individuals. While behavioral data captures an individual’s digital footprint, geographical data reflects their physical footprint. Given the significant privacy risks associated with combining these data sources, it is crucial to understand their respective value and whether they act as complements or substitutes in achieving firms’ business objectives. In this paper, we combine economic theory, machine learning, and causal inference to quantify the value of geographical data, the extent to which behavioral data can substitute it, and the mechanisms through which it benefits firms. Using data from a leading in-app advertising platform in a large Asian country, we document that geographical data is most valuable in the early “cold-start” stage, when behavioral histories are limited. In this stage, geographical data *complements* behavioral data, improving targeting performance by almost 20%. As users accumulate richer behavioral histories, however, the role of geographical data shifts: it becomes largely *substitutable*, as behavioral data alone captures the relevant heterogeneity. These results highlight a central privacy–utility trade-off in ad personalization and inform managerial decisions about when location tracking creates value.

*We would like to thank the participants of the 2025 PhD student workshop at Cornell University and the 2026 Artificial Intelligence in Management (AIM) Conference for their feedback. We also thank Emaad Manzoor for the detailed comments, which have significantly improved the paper. Please address all correspondence to: mm3322@cornell.edu and or83@cornell.edu.

1 Introduction

Mobile devices dominate the digital economy, reshaping how consumers interact with media, commerce, and advertising. Nearly all Americans own a mobile phone, and nine in ten use a smartphone (Pew-Center, 2024), which creates an ideal environment for targeted marketing. In response, mobile advertising now accounts for more than two-thirds of all digital ad spending in the United States, surpassing \$200 billion in 2024 (eMarketer, 2024). Much of this growth stems from advances in tracking technologies that allow advertisers to monitor users' digital and physical activities and tailor messages to individual consumers. While these capabilities have made mobile advertising a central channel for marketers, they have also heightened privacy concerns, raising questions about whether the economic value of such data practices justifies the associated privacy risks.

Privacy concerns arise most directly from two types of data collected by advertising platforms: *behavioral data*, such as click histories or in-app purchases, and *geographical data*, such as GPS coordinates or IP-based locations (Ghose et al., 2019). Behavioral data captures a user's digital footprint within apps, while geographical data captures their physical footprint—their movements and proximity in the real world. Geographical information is particularly sensitive, as even a few location points can be sufficient to re-identify individuals (De Montjoye et al., 2013). When combined with behavioral data, these risks amplify: the two data types complement each other in enabling more precise inference of personal attributes and increasing the likelihood of re-identification, even when datasets are anonymized (De Montjoye et al., 2018).

Given the complementarity of geographical and behavioral information in amplifying privacy risks, it is crucial to examine whether a similar complementarity exists in the factor that motivates firms to continue collecting them—the economic value they create. In particular, it remains unclear whether behavioral and geographical data act as *substitutes* or *complements* in generating economic value for advertising platforms. Addressing this question requires first establishing how to measure the value of each data type, then comparing their relative contributions and interactions, and finally identifying the mechanisms through which geographical data may provide distinct informational value. These considerations motivate the following research questions:

1. How can we measure the value of geographical data in ad personalization relative to behavioral data?
2. To what extent do geographical data generate value once behavioral data are available, and how does this value change as behavioral data accumulate?
3. Through which mechanisms do geographical data provide distinct information?

We face several challenges in answering these questions. The first challenge is conceptual: how should we define the value of information? A common approach focuses on prediction accuracy, for example, whether adding geographical data improves click-through forecasts. However, predictive gains alone do not guarantee better advertising outcomes (Ascarza, 2018). What ultimately matters is whether available information improves decision quality. Accordingly, we draw on economic theory to define the value of geographical and behavioral data through their impact on decision outcomes. This framework allows us to assess whether the two sources act as *complements*, where their joint use delivers gains beyond the sum of

their individual effects, or as *substitutes*, where the contribution of one diminishes once the other is available. For example, if behavioral data indicate that a user is a New York Knicks fan, knowing that the user is in New York City is valuable for advertising tickets to a home game, so the two information sources may act as complements. However, for online Knicks merchandise sold online and shipped nationwide, location adds little value once behavioral interest is known, so the two sources may act as substitutes. This distinction is central, as collecting information beyond its marginal value imposes privacy costs without improving outcomes, whereas complementary information can justify broader data use.

The second challenge is empirical. Our objective is to assess whether geographical data provides incremental decision value beyond behavioral data. This requires a behavioral model that fully captures the dynamic structure of user histories; otherwise, estimated gains from geographical data may reflect behavioral misspecification rather than informational value. Because behavioral data are inherently temporal, we leverage a best-in-class Long Short-Term Memory (LSTM) network with an attention mechanism to capture temporal dependencies and heterogeneous relevance across past interactions.

The third challenge is statistical. Our objective is to test whether two sources of information act as substitutes or complements, which requires comparing policy values under alternative information sets. This, in turn, demands counterfactual outcomes that are not directly observed: we only observe user responses under the policies actually deployed, not under alternative targeting strategies. To address this limitation, we apply inverse propensity scoring (IPS) (Horvitz and Thompson, 1952), which recovers counterfactual policy values by reweighting observed outcomes using assignment probabilities. A key feature of our setting is the platform’s quasi-proportional auction mechanism, which is directly observed in the data and allocates ads in proportion to the quality-adjusted bid. This mechanism provides plausibly exogenous variation and supports reliable propensity score estimation and credible statistical inference.

Together, these challenges motivate a unified framework that combines economic theory to define information value through policy comparisons, machine learning to model user outcomes from complex behavioral and geographical data, and causal inference to enable counterfactual evaluation via IPS. We apply this framework to data from a leading mobile advertising network in a major Asian country, covering 10.3 million impressions from 439,157 users over a 10-day period. We consider four targeting scenarios: a benchmark X^\emptyset that uses only contextual information (e.g., app category or time of day); X^G , which augments context with geographical data such as exact latitude and longitude; X^B augments contextual features with behavioral information from users’ prior impression and click histories; and X^{GB} , which combines both geographical and behavioral information. Differences in data structure across scenarios motivate the use of *best-in-class* models, with gradient-boosted trees (XGBoost) for scenarios based on static and cross-sectional features (X^\emptyset and X^G) and LSTM networks with attention for scenarios that incorporate sequential behavioral histories (X^B and X^{GB}). Leveraging the observed quasi-proportional auction mechanism to support credible propensity score estimation, we apply IPS to recover counterfactual click-through rates (CTRs) and formally quantify the value of behavioral and geographical information, providing a statistical test of whether the two sources act as complements or substitutes.

At the aggregate level, we compare all targeting regimes to the baseline CTR. Targeting without user-level data (X^\emptyset) increases CTR by 20.44%. Adding geographical information (X^G) raises CTR by 28.7%,

behavioral information (X^B) by 30.3%, and combining both (X^{GB}) yields an improvement of 41.5%. Although the joint regime achieves the highest performance, aggregate results alone do not provide systematic evidence that geographical and behavioral information act as complements or substitutes. This motivates examining heterogeneity as behavioral information accumulates.

As users receive more ad impressions, advertisers observe richer behavioral histories, which change both the value of behavioral information and its interaction with geographical information. In the earliest stage (1–2 impressions), behavioral information is minimal, and performance is driven largely by geographical information, with the combined regime performing similarly to geographical data alone. In the intermediate stage (approximately 5–25 impressions), behavioral information accumulates but remains sparse. In this range, geographical and behavioral information act as complements, and using both information sources yields gains beyond either source alone. Once users exceed 25 impressions, behavioral information becomes sufficiently rich to capture preferences independently, and geographical information becomes a substitute, adding little incremental value. Overall, these results show that geographical information evolves with impression depth: complementary when it is sparse, and substitutable once it is rich.

Our findings have direct implications for managers and policymakers. Geographical information creates value primarily in the short run, acting as a temporary complement when behavioral information is sparse, but becomes substitutable as user histories accumulate. For firms, this implies that geo-targeting can be useful during the cold-start phase but delivers little incremental benefit once behavioral models are established. For regulators, the results underscore a key trade-off: location information entails substantial privacy risks while offering limited long-term value in digital advertising. Firms should therefore reassess the strategic role of geographical information, using it selectively when behavioral histories are unavailable and reducing reliance on it as richer user profiles emerge.

In summary, our paper makes several contributions to the literature. First, while prior research has emphasized the value of either behavioral or geographical data in isolation, we provide a systematic comparison of the two. Second, we develop a unified framework that combines economic theory to define information value, machine learning to model advertising responses, and causal inference to evaluate counterfactual outcomes. This framework allows managers and policymakers to assess whether different information sources act as complements or substitutes, providing a practical tool to evaluate targeting decisions while accounting for privacy costs. Substantively, we show that both geographical and behavioral data create value, but their roles evolve as behavioral information accumulates: geographical data complements behavioral data when histories are sparse and becomes substitutable once behavioral information is rich. This pattern implies that the marginal value of geo-targeting is highest when behavioral information is sparse and declines as behavioral histories become richer, raising questions about the privacy–utility trade-off.

The remainder of the paper proceeds as follows. We review related work on the value of information, advertising and privacy, and spatial economics in §2. We introduce the institutional setting of mobile advertising and describe the large-scale dataset in §3. We then formalize the problem within a decision-under-uncertainty framework to define the value of information in §4. Next, we first highlight key challenges and then present our empirical framework in §5, where we define complementarity and substitutability and combine machine-learning–based click prediction with inverse propensity scoring for counterfactual evaluation.

We report empirical findings on the value of behavioral and geographical data and their heterogeneity in §6, and we examine the mechanisms behind these patterns using a residualized spatial autocorrelation test in §7. We discuss implications for data use, privacy, and targeting efficiency in §8, and we conclude by summarizing our contributions and outlining directions for future research in §9.

2 Related Work

Our work relates to a broad literature on how the value of information is defined and measured for decision-making under uncertainty. Theoretically, [Blackwell \(1953\)](#) provides the benchmark by ordering signals according to whether they increase expected reward across all decision problems (the Blackwell order), and [Börgers et al. \(2013\)](#) extend this framework to multiple signals by characterizing when information sources act as complements or substitutes based on whether one signal’s marginal value rises or falls in the presence of another. Empirically, [Rossi et al. \(1996\)](#) quantify the value of alternative information sets for direct marketing by comparing profits from targeted couponing based on purchase histories and demographic characteristics. Complementing this perspective, [Kim et al. \(2022\)](#) show that the value extracted from rich data depends on the chosen level of data granularity and model specification, formalizing the bias–variance trade-off in information use. Subsequent work examines how data access and targeting policies affect outcomes using causal strategies. Using randomized experiments, [Ascarza \(2018\)](#) compares targeting rules based on churn risk versus treatment-effect lift, [Cui et al. \(2019\)](#) manipulate the availability of information shown to consumers, and [Wernerfelt et al. \(2025\)](#) experimentally restrict access to offsite behavioral data to estimate its incremental value. Other studies exploit natural experiments, such as [Aridor et al. \(2024\)](#), who use Apple’s App Tracking Transparency as an exogenous shock to behavioral data access to identify its impact on advertising outcomes. Finally, off-policy evaluation methods recover policy values from logged or observational data, as in [Rafieian and Yoganarasimhan \(2023\)](#) and [Rafieian \(2023\)](#). We extend this literature by proposing a unified framework that systematically compares multiple information sources and empirically tests whether they act as complements or substitutes through their marginal contributions to decision outcomes in high-dimensional settings.

Second, our study connects to the literature on privacy and personalized advertising, particularly in user tracking and engagement modeling. In display advertising, early work by [Goldfarb and Tucker \(2011\)](#) shows that ad intrusiveness and privacy sensitivity significantly affect engagement, with well-targeted yet subtle ads performing best. Subsequent research demonstrates that personalized ads can increase relevance but also intensify privacy concerns: [Tucker \(2014\)](#) show that personalization in social networks raises demand for stronger privacy controls, while [Acquisti et al. \(2016\)](#) formalize the trade-offs firms face between personalization benefits and consumer resistance to data collection. More recent studies examine the consequences of restricting data access, with [Johnson et al. \(2020\)](#) quantifying the revenue losses from consumer opt-outs and [Rafieian and Yoganarasimhan \(2021\)](#) showing that privacy restrictions reduce targeting efficiency and may affect market competition. In this regard, despite well-documented privacy risks of location data ([De Montjoye et al., 2013](#)), relatively few studies directly examine the value of geographical information. Closely related to our work, [Narang and Luco \(2025\)](#) quantify the predictive value of geo-tracking data for forecasting consumer visits. We extend this line of research by moving beyond predictive accuracy and developing a unified framework for substitutability/complementarity between geographical and behavioral

data that equips us to more precisely assess the privacy–utility trade-off.

Third, a line of empirical research in marketing and economics studies geographic dependence in outcomes using methods from spatial econometrics. Early contributions introduce diagnostics for spatial autocorrelation. [Moran \(1950\)](#) proposes Moran’s I as a global measure of spatial dependence in outcomes, while [Geary \(1954\)](#) develops Geary’s C to capture more localized spatial variation. Subsequent work applies these diagnostics to substantive economic settings. [Bronnenberg and Mahajan \(2001\)](#) model spatial dependence in market shares and promotions across neighboring markets to account for correlated, unobserved retailer actions. Focusing on longer-run demand patterns, [Bronnenberg et al. \(2009\)](#) document persistent geographic structure in brand demand, showing that brands retain higher shares near their historical origins. Related research examines diffusion and contagion by exploiting geographic or network adjacency. A body of empirical research in marketing studies how spatial structure shapes consumer behavior and diffusion outcomes ([Bradlow et al., 2005](#)). At finer levels of granularity, [Larson et al. \(2005\)](#) show that spatial layout affects consumer exposure and choice through physical shopping paths. This perspective extends to diffusion and contagion, where researchers exploit geographic or network adjacency: [Manchanda et al. \(2008\)](#) use physician networks to separate targeted communication from peer influence, [Iyengar et al. \(2011\)](#) study social networks to identify opinion leadership effects, and [Bollinger and Gillingham \(2012\)](#) use zip-code-level exposure to quantify neighborhood spillovers in solar adoption. Despite growing evidence of substantial spatial heterogeneity in advertising effectiveness, as documented by [Luo and Ranjan \(2025\)](#), existing work largely focuses on detecting and modeling spatial correlation in outcomes rather than evaluating geographical data as an input to ad personalization. We address this gap by assessing the value of geographical information for personalization and by proposing a residual spatial autocorrelation (RSA) framework that decomposes spatial correlation conditional on behavioral data to pin down the channel through which geographical information creates value.

3 Setting & Data

We start by describing the institutional setting of the mobile advertising platform (§3.1). Next, we detail the dataset, including impressions, clicks, and contextual variables (§3.2). We then explain our sampling strategy and report summary statistics (§3.3). Finally, we describe the train–test split design used to evaluate model performance (§3.4).

3.1 Setting

Our data are sourced from a leading mobile in-app advertising platform in a large Asian country, which held over 85% of the mobile advertising market during the time of our study. The platform serves as an intermediary between advertisers and mobile apps (publishers) and is responsible for delivering over 50 million ad impressions daily. This marketplace consists of four primary players:

- **Users** are mobile app consumers who generate impressions and may choose to click on displayed ads.
- **Publishers** are app developers that integrate ads into their apps and monetize based on ad clicks.
- **Advertisers** design banner ads and specify per-click bids. They can target users based on variables such as province, smartphone brand, app category, internet service provider (ISP), hour of the day,

and connectivity type. The ad network does not support detailed personalized targeting.

- **Platform or ad network** manages real-time auctions to match impressions with ads. Ads are placed as bottom banners and are refreshed every minute. Only clicks result in payment under a cost-per-click (CPC) scheme.

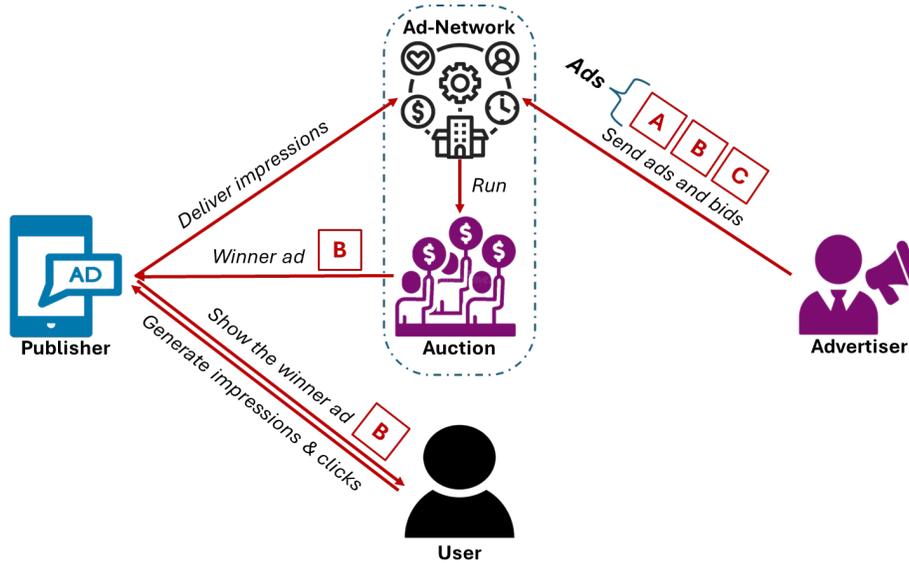


Figure 1: Schematic of the mobile in-app advertising marketplace.

Figure 1 illustrates the structure of the in-app advertising marketplace. When a user opens an app, the platform initiates a real-time auction among eligible ads, allocating impressions through a quasi-proportional rule (Mirrokni et al., 2010): ads with higher bid-quality products are assigned higher probabilities of being shown, but selection remains stochastic. For example, if two ads have bid-quality products of 2 and 1, the first ad is twice as likely to be shown as the second, but both ads retain strictly positive probabilities of exposure. As a result, propensity scores are non-zero by design, which ensures overlap across ads. Importantly, quality scores are fixed and not personalized at the user level, so allocation probabilities do not depend on unobserved user characteristics, implying unconfounded assignment by design. If the user remains active beyond one minute, a new impression is generated, and a new auction is run.

3.2 Data

Following the setting explained above, we have data on all impressions and click outcomes observed over a 30-day period from September 30, 2015, to October 30, 2015. During this period, we observe a total of 1,594,831,699 ad impressions and 14,373,293 clicks, resulting in an overall click-through rate (CTR) of approximately 0.90%. For each impression, the dataset includes detailed information across several dimensions.

For each impression, the dataset records a comprehensive set of variables. First, we observe (1) *Timestamp*, capturing the exact time of the impression, and (2) *AAID*, the Android Advertising ID, a user-resettable unique device identifier that enables anonymous tracking across applications. At the ad-delivery level, we

record (3) *AppID*, the identifier of the app displaying the ad, and (4) *CreativeID*, the identifier of the specific ad creative shown. We also observe (5) *Bid*, the advertiser’s submitted bid amount (fixed throughout the sample period), and (6) *CPC*, the cost-per-click charged in the event of a click. Geographical attributes include (7) *Latitude*, (8) *Longitude*, and (9) *Province*, which are available for most impressions. Further contextual information comprises (10) *Connectivity*, indicating whether the user was on Wi-Fi or cellular data, (11) *Brand*, the smartphone manufacturer, (12) *MSP*, the mobile service provider, and (13) *ISP*, the internet service provider. Finally, (14) *Click* is a binary indicator equal to one if the user clicked on the ad and zero otherwise.

A key aspect of this dataset is that it is sourced directly from the ad platform and includes the complete set of variables that advertisers could potentially use for targeting. As such, we observe the same information available to both the platform and advertisers at the time of impression delivery. This mitigates common concerns in observational studies around hidden targeting mechanisms or unobserved confounding. In our context, the completeness of the data supports modeling assumptions such as conditional ignorability, which are crucial for later analyses involving counterfactual analysis and policy evaluation.

3.3 Sampling and Summary Statistics

To enable user-level analysis while maintaining computational tractability, we restrict attention to impressions from the top 10 ads served between October 20 and October 30, yielding an initial sample of 16,662,783 impressions. We address missing values, particularly in variables that serve as key features in our analysis, by excluding records without geographic information, specifically latitude and longitude. After this filtering, the working dataset contains 10,336,703 impressions and has no remaining missing values in any other columns. From this dataset, we identify new users who joined the platform during the first three days of the window, October 20 to October 22, and follow their subsequent activity through October 30. This design allows us to observe user responsiveness from the start of their platform interaction, yielding a cohort of 439,157 unique users.

We now present some summary statistics for key categorical variables in the dataset. Table 1 reports, for each variable, the number of unique categories, the share of impressions associated with the top three values, and the number of non-missing observations.

Table 1: Summary Statistics for the Categorical Variables

Variable	Number of categories	Share of top categories			Number of impressions
		1st	2nd	3rd	
App	9,515	26.14%	9.46%	4.54%	10,336,703
Ad	10	22.33%	13.28%	12.67%	10,336,703
Unique User	439,157	0.02%	0.02%	0.02%	10,336,703
Smartphone brand	7	41.94%	30.59%	9.21%	10,336,703
Connectivity Type	2	54.65%	45.35%	0.00%	10,336,703
ISP	8	61.09%	26.93%	5.24%	10,336,703
Province	828	11.21%	8.78%	7.33%	10,336,703

We observe a total of 9,515 unique apps, with the top three accounting for a sizable share of total im-

pressions. Among the ten ads included by design, exposure is uneven, with the most frequently shown ad representing over 22% of impressions. The distribution of user identifiers, based on Android Advertising IDs, is highly diffuse, with no single user accounting for more than 0.02% of total impressions. Other variables, such as smartphone brand, connectivity type, ISP, and province, exhibit varying degrees of concentration, reflecting both common patterns and localized variation in usage across the population.

While Table 1 highlights variation in exposure across contextual features, it does not reveal how user responsiveness may differ across behavioral or geographical information. To explore this further, we present descriptive evidence on how CTR varies with user history and spatial location. These patterns help motivate the relevance of behavioral and geographical data for downstream modeling tasks.

3.3.1 Behavioral Heterogeneity

We present two descriptive results that illustrate how behavioral patterns shape user CTR. First, we examine how the length of a user’s exposure history relates to CTR. Second, we analyze the relationship between past click behavior and the likelihood of future clicks.

Figure 2a plots the cumulative share of impressions against the cumulative share of clicks, where users are ordered by the number of prior impressions they have seen, so that impressions are ranked from early to late in a user’s exposure history, and each point on the curve compares the share of total impressions up to that history length with the share of total clicks they generate. The curve lies well above the 45° line: a disproportionate share of clicks comes from early exposures (e.g., the first 25% of impressions, corresponding to short histories, generate nearly half of all clicks). Figure 2b shows the probability of a click on the next impression, $\Pr(\text{click}_{t+1} = 1 \mid \text{prior clicks} = k)$, as a function of the number of past clicks k . The weighted linear fit has a positive and statistically significant slope: users who have accumulated more clicks are more likely to click again on the next impression.¹

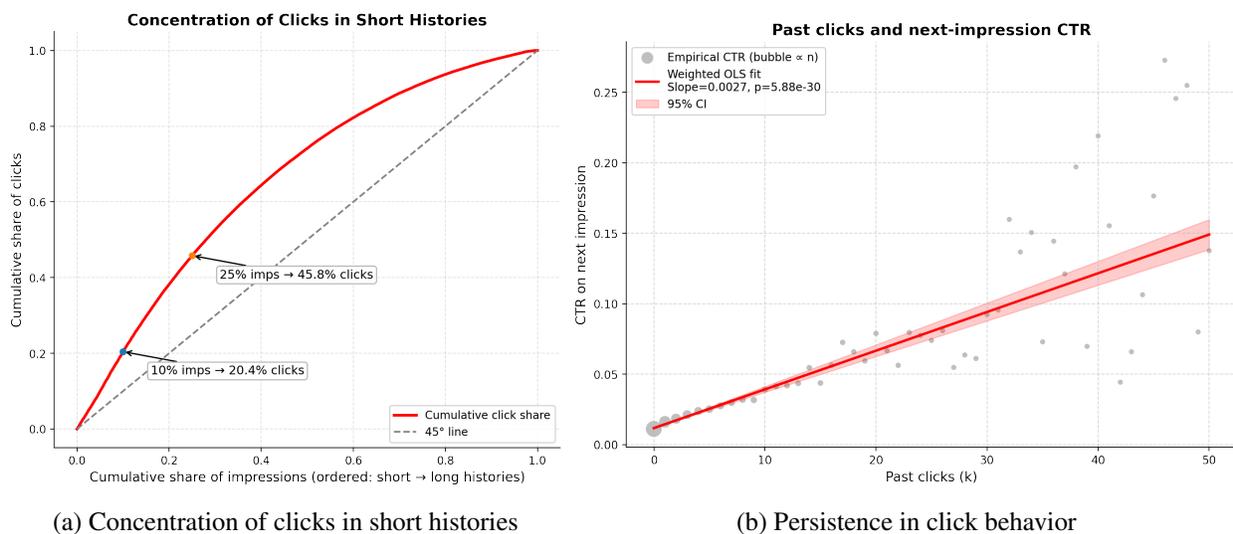


Figure 2: Behavioral heterogeneity: (a) Clicks are concentrated in early exposure histories; (b) users with prior clicks are more likely to click again.

¹Final impressions (with no $t+1$) are excluded. These patterns are descriptive and do not imply causal effects.

Taken together, these descriptive results reveal two dimensions of behavioral heterogeneity. First, Figure 2a shows that CTR declines as exposure histories lengthen, indicating the need to account for user trajectories. Second, Figure 2b shows persistence in click behavior: users with more prior clicks are more likely to click again, reflecting heterogeneity in underlying propensities and highlighting the predictive value of behavioral history.

3.3.2 Geographical Heterogeneity

We now turn to geographical patterns of responsiveness to examine whether user responses exhibit systematic spatial dependence. To do so, we group impressions by county and compute the average CTR within each spatial unit. Specifically, for a given county c , we calculate $CTR_c = \frac{1}{N_c} \sum_{i \in c} y_i$, where $y_i \in \{0, 1\}$ indicates whether impression i resulted in a click, and N_c is the number of impressions observed in county c . Figure 3 visualizes the resulting CTR values as a choropleth map.

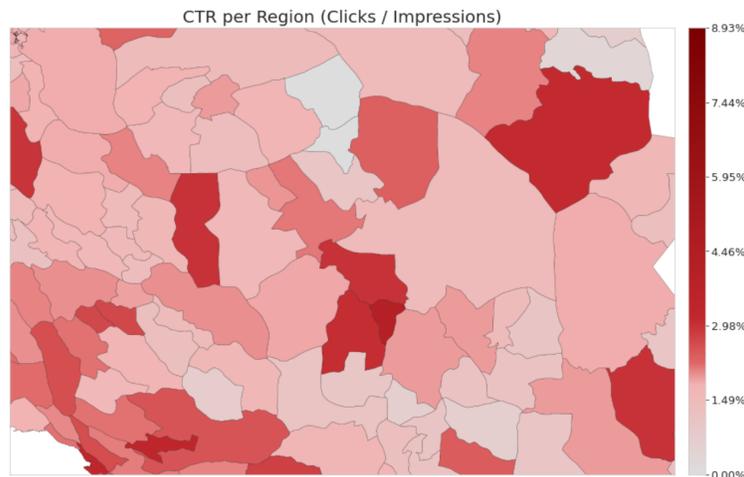


Figure 3: Average CTR by County

The figure reveals pronounced spatial correlation in responsiveness: counties with high (low) average CTRs tend to be geographically proximate to other counties with similarly high (low) CTRs, which generate visible regional clusters rather than isolated pockets of responsiveness. This spatial structure indicates that user engagement is not independently distributed across space but instead varies smoothly across neighboring locations. Such correlation suggests that geographic context captures shared latent factors, such as local demographics or socioeconomic context, that influence users’ likelihood of clicking on ads.

3.4 Train-Test Split Strategy

Finally, we split the data into training and test sets to evaluate out-of-sample predictive performance and policy effectiveness. To avoid information leakage, the split is performed at the user level rather than the impression level: each user’s full history is contained entirely within either the training or the test partition, allowing us to assess generalization to previously unseen users.

Figure 4 illustrates the procedure. A user becomes eligible once they first appear in the platform’s logs, after which we track their impression history. Each user is then randomly assigned to either the training set (60%, 263,494 users; 6,202,022 impressions) or the test set (40%, 175,663 users; 4,134,681

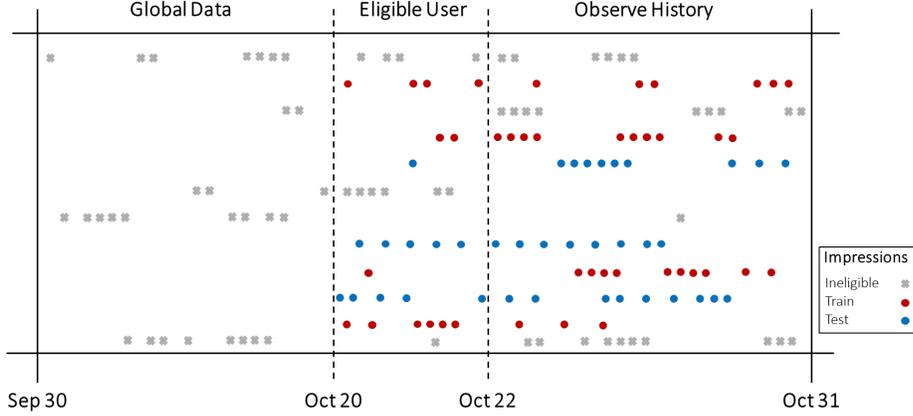


Figure 4: Schema for Data Generation and User-Based Train-Test Split

impressions). Gray dots mark ineligible impressions that precede a user’s first appearance, while red and blue dots denote impressions assigned to the training and test sets, respectively. The resulting split maintains similar distributions of activity and click behavior across both partitions.

4 Problem Definition

Advertising platforms operate in environments where advertising decisions must be made under uncertainty about user preferences. Consider the case of an ad network deciding which ad to show to a particular user. The central goal of the platform is to maximize a reward measure, such as advertising revenue or user engagement. At the same time, the platform does not fully observe the user’s type but observes a signal through observed features. As a result, the platform faces a decision-making problem under uncertainty, where the challenge lies in selecting the right action based on partial information about the user.

To analyze this decision problem, we model ad allocation as a choice among feasible actions in the presence of heterogeneous users and incomplete information. Users differ in latent preferences that determine the rewards generated by different ads, while the platform observes only an information set X that shapes its beliefs about these rewards. We now formalize the elements of this problem.

- **Environment.** The platform must choose an action from a finite set \mathcal{A} , where each $a \in \mathcal{A}$ represents a feasible option (e.g., an ad that can be shown). Users differ in unobserved characteristics captured by a latent type $\omega \in \Omega$, with Ω denoting the space of possible user types. The reward of taking action a for a user of type ω is represented by a function:

$$r : \mathcal{A} \times \Omega \rightarrow \mathbb{R}, \quad r(a, \omega) \text{ (e.g., expected CTR).}$$

This reward function encodes the payoff of each action for each user type and serves as the platform’s objective.

- **Information set and user types.** The platform does not observe the user’s type ω directly. Instead, they observe an information set $X \in \mathcal{X}$ that provides partial knowledge about ω . The information

set X may include behavioral variables, device attributes, geographical indicators, or other covariates. We assume the pair (X, ω) is drawn from a joint distribution p over $\mathcal{X} \times \Omega$, which factorizes as $p(x, \omega) = p(\omega) p(x | \omega)$, where $p(\omega) \in \Delta(\Omega)$ represents the population-level distribution of user types. Upon observing $X = x$, the platform forms a posterior belief over types:

$$q_x(\omega) := p(\omega | x) \in \Delta(\Omega).$$

For notational clarity, we also define the *random posterior* $q_X := p(\cdot | X)$, a random variable that maps each realization $X = x$ to its corresponding posterior q_x . This posterior encodes the platform’s updated belief about the user’s type after observing x .

- **Policies.** A policy specifies how the platform selects actions based on the observed information set. Formally, a policy is a measurable function:

$$\pi : \mathcal{X} \rightarrow \mathcal{A},$$

which maps each realization $x \in \mathcal{X}$ to an action $\pi(x) \in \mathcal{A}$. We focus on deterministic policies because, under linear expected reward objectives, randomization offers no additional value (by an extreme-point argument; see [Kamenica, 2019](#); [Smith et al., 2023](#))².

- **Expected reward and optimal actions.** Given a belief $q \in \Delta(\Omega)$ over user types, such as the posterior q_x induced by observing x , the expected reward of choosing action $a \in \mathcal{A}$ is

$$\mathbb{E}_{\omega \sim q_x}[r(a, \omega)].$$

Following, among all possible policies $\pi \in \Pi$ that map observed information to actions, the optimal policy $\pi^*(x)$ is the one that maximizes expected reward for each observation:

$$\pi^*(x) \in \arg \max_{a \in \mathcal{A}} \mathbb{E}_{\omega \sim q_x}[r(a, \omega)]. \quad (1)$$

- **Decision value.** Let X denote the information set observed prior to action selection, and let $q_X = p(\cdot | X)$ be the induced posterior over user types. For any deterministic policy $\pi : \mathcal{X} \rightarrow \mathcal{A}$, its ex-ante value is $V(\pi) = \mathbb{E}_{(X, \omega) \sim p}[r(\pi(X), \omega)]$. Therefore, the value of information from X is the maximum achievable value over all such policies:

$$V_X = \sup_{\pi: \mathcal{X} \rightarrow \mathcal{A}} V(\pi) = \sup_{\pi: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_X[\mathbb{E}_{\omega \sim q_X}[r(\pi(X), \omega)]]. \quad (2)$$

The optimal policy $\pi^*(x)$ defined in Equation 1 attains this supremum by maximizing expected reward for each observation.

²When \mathcal{A} is finite and the objective is linear, any mixed policy can be written as a convex combination of deterministic ones, and the maximum is achieved at an extreme point. Ties are resolved by an arbitrary but fixed rule.

In this study, we focus on click outcomes (CTR) as the primary measure of reward. Clicks are central to digital advertising because they are closely tied to both platform revenue and user engagement under prevalent cost-per-click and click-weighted pricing mechanisms. They provide an immediate and observable measure of user response that reflects the quality of the platform’s ad allocation decision. Given this setup, we operationalize rewards using click outcomes.³

Let $Y \in \{0, 1\}$ denote the binary indicator of user engagement (e.g., whether the user clicks on the displayed ad), and let $a \in \mathcal{A}$ denote the action chosen by the platform. Consistent with the model above, the realized outcome Y depends on both the selected action a and the user’s latent type $\omega \in \Omega$, which captures unobserved preferences and motivations. While ω is not directly observable, the platform may observe proxy information X that is informative about it. We categorize available user-level information into two main types:

- **Behavioral Information** ($X^B \in \mathcal{X}^B$) captures a user’s past interactions in the digital environment, which includes prior app usage patterns, ad exposure sequences, click history, and other engagement-based indicators that reflect individual preferences and interests over time.
- **Geographical Information** ($X^G \in \mathcal{X}^G$) describes the user’s spatial context, encompassing attributes such as precise location (e.g., latitude/longitude) and broader administrative regions such as city/province that may correlate with demographic or regional characteristics.

We denote by X^{GB} the information regime that combines both behavioral and geographical information. We also treat *contextual information*, denoted by X^\emptyset , as standard metadata attached to each ad impression, including device characteristics (e.g., brand and operating system) and network conditions (e.g., connectivity type). Like behavioral and geographical information, contextual variables reflect aspects of the user’s latent type $\omega \in \Omega$, but they differ in two important respects. First, they are typically available to platforms for every impression, making them a natural baseline information set. Second, consumers generally view them as substantially less privacy-sensitive than detailed behavioral or location-based data (Jerath and Miller, 2024). Accordingly, we include contextual information (X^\emptyset) in all targeting regimes. The specific variables included in each regime $X \in \{X^\emptyset, X^G, X^B, X^{GB}\}$ are described in §A.1.

While we formalize the value of information through its impact on decision outcomes in Equation 2, a common alternative approach evaluates information solely through predictive performance. This approach focuses on estimating $\hat{y} = f(X, a)$ to approximate $\mathbb{P}(Y = 1 | X, a)$ and assesses information quality using metrics such as AUC or reductions in conditional entropy $H(Y | X, a)$. However, predictive accuracy does not in general translate into better decisions or improved outcomes when targeting policies are deployed (Ascarza, 2018; Rafieian and Yoganarasimhan, 2023).

5 Empirical Framework

The central objective of this framework is to assess the value of information by quantifying and comparing how behavioral data (X^B) and geographical data (X^G) contribute to improved decision-making. The

³We focus on clicks as the primary outcome because targeting value arises from increasing the relevance and clickability of ads. However, the framework readily extends to alternative reward measures, such as revenue per click.

value function in Equation 2 provides the formal basis for this assessment by defining the maximum expected reward attainable under each information set. Implementing this framework, however, raises several challenges:

1. **Challenge 1: Information Value, Complementarity, and Substitutability:** The first challenge is to formalize how to compare the value of distinct sources of information. When multiple data sources, such as behavioral (X^B) and geographical (X^G), are available, their value is inherently *relative*: each may add incremental benefit beyond the other or overlap in what it conveys. The marginal value of one source depends on whether the other is already observed, raising the possibility of complementarity or substitutability. Without precise definitions, we cannot separate the standalone contribution of each source from the value created by using them jointly. We therefore develop a formal framework that isolates individual value, incremental value conditional on the other source, and their combined effect. We present this framework in §5.1.
2. **Challenge 2: Reward Function Estimation for Decision Policies:** The second challenge is empirical: estimating the reward function $r(a, \omega)$. Our objective is to assess whether geographical information provides incremental decision value beyond behavioral information. Doing so requires reward estimation that fully captures the dynamic structure of user histories; otherwise, estimated gains from geographical information may reflect behavioral misspecification rather than informational value. Because behavioral information is inherently temporal, reward estimation must account for this dynamic structure. We address this challenge in §5.2.
3. **Challenge 3: Information Value Estimation through Policy Evaluation:** The third challenge concerns estimating the value of information V_X . By definition, V_X is the maximum expected reward achievable under a given information set, which requires comparing alternative decision policies. Such comparisons are inherently counterfactual, since each user is observed under only one realized action and outcomes under other actions are not observed. Estimating V_X therefore requires a model-free approach to evaluating the performance of counterfactual policies using logged data, even when some actions are rarely chosen. We address this challenge in §5.3.

To address these challenges, we develop a unified framework that integrates economic theory, machine learning, and causal inference. The framework (i) formally defines information value to compare behavioral and geographical data and characterize complementarity or substitutability (§5.1), (ii) estimates reward functions that capture the dynamic structure of behavioral histories (§5.2), and (iii) evaluates information value through model-free counterfactual policy evaluation using logged data (§5.3). Together, these components quantify the decision value of behavioral and geographical information.

5.1 Substitutability and Complementarity of Information Sets

To address Challenge 1, we use formal definitions from the economic theory literature to compare the value of different information sets in a decision problem. In our context, we assess how behavioral data X^B and geographical data X^G contribute to guiding optimal actions. Intuitively, two pieces of information

complement each other if the joint incrementality in information value goes beyond a simple sum of the incrementality of each alone.

We adopt the decision-theoretic framework of substitutability and complementarity introduced by [Börgers et al. \(2013\)](#). Recall that X^B and X^G denote behavioral and geographical data, respectively. We define the combined information set as $X^{GB} := (X^B, X^G) \in \mathcal{X}^B \times \mathcal{X}^G$, which represents joint access to both sources. As a baseline, we consider X^\emptyset , corresponding to the absence of user-level information. In this case, decisions rely only on *contextual information* that informs the prior belief about user types, so the posterior reduces to the prior, $q_{X^\emptyset} \equiv p(\omega)$. The associated value, V_{X^\emptyset} , therefore represents decision-making under prior uncertainty without conditioning on individual-level signals. We define substitutability and complementarity by the marginal value of one information set conditional on access to the other, using the value function V_X (Equation 2).

Definition 1 (Substitutes). *Behavioral information X^B is a substitute for geographical information X^G if, for every decision problem ⁴ (\mathcal{A}, r) ,*

$$V_{X^G} - V_{X^\emptyset} \geq V_{X^{GB}} - V_{X^B}.$$

This inequality compares the marginal value of geographical data when used alone versus when combined with behavioral data. The left-hand side, $V_{X^G} - V_{X^\emptyset}$, reflects the value of geographical data on its own, while the right-hand side, $V_{X^{GB}} - V_{X^B}$, reflects its incremental contribution when behavioral data are already available. If the inequality holds, behavioral data reduces the marginal usefulness of geographical data, making the two substitutes.

Definition 2 (Complements). *Behavioral information X^B is a complement to geographical information X^G if, for every decision problem (\mathcal{A}, r) ,*

$$V_{X^G} - V_{X^\emptyset} \leq V_{X^{GB}} - V_{X^B}.$$

Here, the marginal value of geographical data is greater when combined with behavioral data than when used alone. The left-hand side, $V_{X^{GB}} - V_{X^B}$, measures the incremental contribution of geographical data given that behavioral data are already available, while the right-hand side, $V_{X^G} - V_{X^\emptyset}$, captures their standalone value. If the inequality holds, behavioral and geographical data enhance each other’s usefulness and are thus considered complements.

5.2 Machine Learning Framework for Reward Function Estimation

To address Challenge 2, we develop a machine learning framework for estimating expected reward when the reward function $r(a, \omega)$ is unobserved. The ad platforms cannot observe latent types ω , nor directly measure the reward associated with each action. Instead, we re-express reward in terms of observable data: user characteristics X , chosen actions A , and realized engagement outcomes $Y \in \{0, 1\}$.

⁴Following [Börgers et al. \(2013\)](#), “every” refers to the inequalities holding for *all* decision problems (A, r) . Our empirical analysis evaluates these same inequalities for a fixed problem (A, r) (CTR objective and action set), yielding complement/substitute conclusions *for this problem*.

We specify reward through the binary click outcome Y and the reward value of engagement. Suppose that showing ad $a \in \mathcal{A}$ to a user of type $\omega \in \Omega$ yields reward only when the user clicks, and that the ad platform receives a known per-click reward $v(a) \in \mathbb{R}_{\geq 0}$.⁵ In this case, the expected reward from taking action a for a user of type ω is given by

$$r(a, \omega) := \mathbb{E}[Y \mid A = a, \omega] \cdot v(a) = \Pr(Y = 1 \mid A = a, \omega) \cdot v(a).$$

This formulation captures both the probabilistic nature of engagement and the economic payoff per click, and it serves as the foundational definition of reward in our analysis. Although user types ω are unobserved, conditioning on observed characteristics $X = x$ induces a posterior belief $q_x(\omega) := p(\omega \mid x)$ about the user’s latent type. Accordingly, we focus on estimating the function:

$$f_X(a, x) := \Pr(Y = 1 \mid X = x, A = a),$$

which maps each action–characteristics pair to the probability of engagement. This function is the central object in our machine learning framework. Once f_X is estimated, the expected reward can be easily computed. This approach can be operationalized using flexible prediction models trained on observational data, provided that the model class is sufficiently expressive and the estimation procedure is properly regularized and calibrated.

To operationalize this framework, we estimate the conditional click function under several information regimes that differ in the richness of observable user data. Specifically, we consider four regimes $\{X^\emptyset, X^G, X^B, X^{GB}\}$, which correspond to progressively richer access to contextual, geographical, and behavioral information. Each regime captures a distinct informational constraint faced by the advertising platform. The specific features included in each information set are described in Appendix §A.1.

In the following, we describe learning algorithm selection f from a flexible class \mathcal{F} to approximate $\Pr(Y = 1 \mid X, A)$ in §5.2.1, and we specify the behavioral modeling architecture in §5.2.2. The estimated function \hat{f}_X delivers action-specific rewards, induces policies $\pi(x)$, and forms the basis for policy evaluation across informational regimes.

5.2.1 Learning Algorithm Selection Across Information Regimes

In regimes without behavioral histories (X^\emptyset and X^G), the information set is static across impressions. Because no new user-specific information accumulates over time, the estimation problem is inherently cross-sectional. We therefore employ flexible tree-based learners (XGBoost) to estimate the click probability as a function of contemporaneous covariates and actions. These models are well-suited to capturing nonlinearities and high-order interactions in static feature spaces without imposing strong parametric structure (Rafieian and Yoganarasimhan, 2021).

By contrast, regimes that include behavioral information (X^B and X^{GB}) generate sequential user histories whose informational content expands over time. In these settings, the relevant state variable is no longer a

⁵Since the targeting value of information arises from improved match quality and resulting changes in users’ CTR, we focus on this reward measure and set $v(a) = 1$ for simplicity. All insights extend to settings in which $v(a)$ is heterogeneous, provided that CTR and per-click valuation are independent and separable.

fixed vector but an evolving sequence of past exposures and engagement outcomes. To accommodate this structure, we develop a sequence model based on the LSTM network with an attention mechanism, which aggregates information over time and forms latent summaries of user behavior (Quadrana et al., 2018). These models map behavioral histories and current actions into predicted engagement probabilities, allowing the estimator to adapt dynamically as additional impressions are observed.

This distinction between static and sequential regimes is central to our analysis. As behavioral information accumulates over impressions, the information set X_{jt}^B becomes increasingly informative about user behavior. This richer information allows the advertiser to condition decisions on a more precise representation of past engagement, improving estimation of the engagement function $f(X_{jt}^B, a)$. As a result, targeting decisions become more accurate, and the expected reward under the induced policy increases as behavioral histories grow. We examine this empirically in §6.3.

5.2.2 Sequence-Based Behavioral Modeling with LSTM and Attention

We now describe the sequence model used for regimes that include behavioral histories in this part. We design an LSTM architecture equipped with a causal multi-head attention mechanism. The LSTM captures short- and long-range temporal dependencies in user interactions, while attention highlights the most informative parts of a user’s exposure–click history. Together, these components allow the model to represent both persistent behavioral trends and short-term recency effects, which are central to engagement modeling (Quadrana et al., 2018).

We train the model on user interaction histories using a sliding window of 150 impressions. At each time step, the input combines categorical and numerical features. We map high-cardinality categorical variables (e.g., device model, app ID, network identifiers) into dense vectors with embedding layers (Guo and Berkahn, 2016) and concatenate them with continuous covariates. To encode order, we add absolute positional embeddings (Vaswani et al., 2017). We also process the log-transformed inter-arrival time through a time-gap projection, which helps the model detect irregular timing patterns that often indicate behavioral shifts.

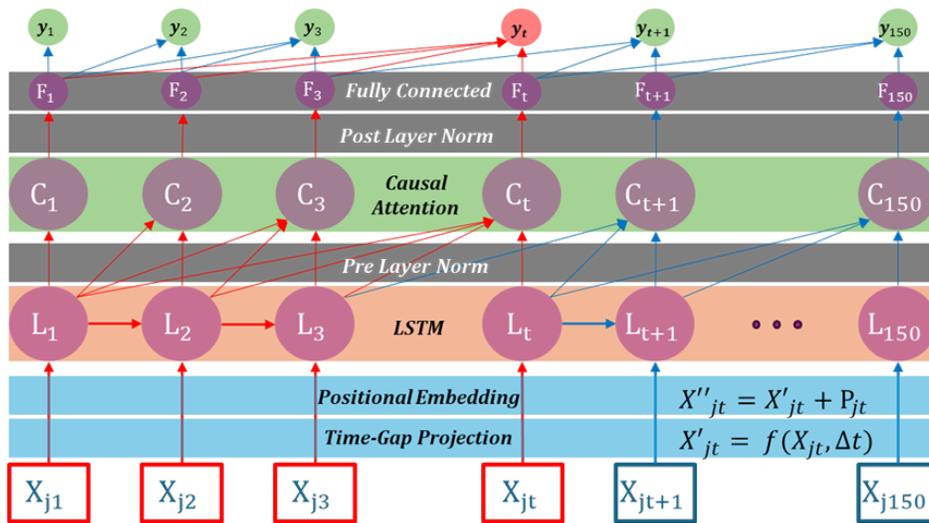


Figure 5: LSTM with causal attention architecture

We feed the enriched sequence into a four-layer LSTM with hidden size 512. On top of the LSTM outputs, we leverage causal multi-head self-attention with four heads. The attention module uses a future-masked matrix to block access to unseen impressions, which preserves temporal causality and prevents information leakage (Vaswani et al., 2017). We apply layer normalization before and after the attention block to stabilize training. Next, we add a gated projection head that runs parallel sigmoid and tanh transformations, combines them elementwise, and applies dropout for regularization. A fully connected output layer finally maps the hidden states to click probabilities at each step.

Figure 5 illustrates the designed sequence model architecture and its key components. The design directly leverages sequential order, irregular timing, and varying relevance of past impressions. As users generate longer histories, the LSTM–attention framework provides richer information about preferences and latent types. We provide additional implementation details in Web Appendix §A.2.

5.3 Estimating the Value of Information via Inverse Propensity Scoring

Challenge 3 highlights the core difficulty in estimating the value of information V_X (Equation 2):

$$V_X = \sup_{\pi: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_X [\mathbb{E}_{\omega \sim q_X} [r(\pi(X), \omega)]],$$

The data only reveal outcomes under the action chosen by the platform’s logging policy, while V_X requires utilities aggregated over all possible actions. Each impression, therefore, provides information on one action, but evaluating a policy involves counterfactual outcomes for actions that were not taken.

We address this problem with Inverse Propensity Scoring (IPS), a standard method in causal inference for adjusting observational data (Hirano et al., 2003; Rafieian and Yoganarasimhan, 2023). IPS identifies the value of a fixed target policy by reweighting observed outcomes according to how likely the logging policy was to select the same action as the target policy. By doing so, we approximate the reward that would have been realized under alternative policies, even though those policies were never deployed in practice.

Formally, for each information set X , we define a target policy π^{f_X} that deterministically selects the action maximizing the estimated expected reward based on our predictive model:

$$\pi^{f_X}(a | x) = \mathbb{I} \left\{ a = \arg \max_{a' \in \mathcal{A}} \hat{f}_X(a', x) \right\},$$

where $\hat{f}_X(a, x)$ is the estimated click probability given information set x and candidate action a . Thus, π^{f_X} prescribes, for each instance, the action expected to maximize reward under the model. The value of this policy, consistent with our decision-theoretic objective, is

$$V(\pi^{f_X}) = \mathbb{E}_X [\mathbb{E}_{\omega \sim q_X} [r(\pi^{f_X}(X), \omega)]] .$$

However, only outcomes corresponding to the logging policy’s actions are observed in the data. To estimate the value of the induced policy $V(\pi^{f_X})$, we use the IPS estimator, which reweights the observed reward outcomes for each user-impression pair by the inverse of the probability that the logging policy would have

selected the action recommended by the target policy:

$$\hat{V}(\pi^{f_X}) = \frac{1}{N} \sum_{i=1}^N \frac{\mathbb{I}[A_i = \pi^{f_X}(X_i)]}{\pi^{\mathcal{D}}(A_i | X_i)} v(A_i) Y_i,$$

where A_i is the action taken in the historical data for impression i , $\pi^{\mathcal{D}}(A_i | X_i)$ is the probability that the logging policy selected A_i given information X_i , and Y_i is the observed binary engagement outcome. Each term, therefore, reweights the realized reward to reflect how the target policy would allocate actions. For the IPS estimator to recover the true value V_X , three key conditions must hold regarding the data-generating process and the structure of the logging policy. We state these assumptions formally below.

Assumption 1 (Overlap). *For all $x \in \mathcal{X}$ and $a \in \mathcal{A}$, if $\pi^{f_X}(a | x) > 0$, then $\pi^{\mathcal{D}}(a | x) > 0$. In other words, the logging policy must assign positive probability to every action that the target policy might select.*

Assumption 2 (Unconfoundedness). *Potential outcomes are independent of the action actually taken, conditional on observed covariates X ; that is, $Y(a) \perp A | X$ for all $a \in \mathcal{A}$. This ensures that all confounding factors are captured in X .*

Assumption 3 (Policy optimality under X). *If \hat{f}_X consistently estimates $\Pr(Y = 1 | A = a, X = x)$ and the maximizer of $v(a) \Pr(Y = 1 | A = a, X = x)$ is unique almost everywhere, then $\pi^{f_X}(x) = \arg \max_a v(a) \hat{f}_X(a, x)$ is optimal under X almost everywhere.*

Proposition 1 (Standard IPS identification and recovery of V_X). *Under Assumptions 1–3, and when π^{f_X} is evaluated on an independent sample,*

$$\mathbb{E}[\hat{V}(\pi^{f_X})] = V_X, \quad \hat{V}(\pi^{f_X}) \xrightarrow{P} V_X.$$

That is, the IPS estimator is unbiased and consistent for V_X under the stated conditions.

Proof. A formal proof is provided in Web Appendix §C.1. □

The IPS approach addresses Challenge 3 by enabling estimation of the ex-ante value of any information set X through the policy it induces, even when counterfactual outcomes are unobserved. In the following, we first discuss the conditions and diagnostics required to validate propensity score estimation in §5.3.1, and then describe the empirical procedure used to estimate propensity scores in §5.3.2.

5.3.1 Estimating and Validating Propensity Scores

To implement the IPS estimator, we require estimates of the logging policy probabilities $\pi^{\mathcal{D}}(a | x)$, which govern how actions are assigned in the historical data. First, we describe the ad allocation mechanism that generates these probabilities, and then examine empirical evidence supporting the IPS identification assumptions.

Quasi-Proportional Allocation Mechanism. In our setting, impressions are allocated via a *quasi proportional auction* that induces randomized exposure across eligible ads as a function of observed bids, quality scores, and eligibility constraints (Mirrokni et al., 2010). For each impression i , let $\mathcal{A}_i \subseteq \mathcal{A}$ denote the set of ads eligible to participate in the auction. Each ad $a \in \mathcal{A}_i$ submits a bid b_a and has a platform-assigned quality score q_a , both of which are fixed and observed during our sample period. The probability that ad a wins the auction and is shown in impression i , conditional on covariates x_i , is given by:

$$\pi^{\mathcal{D}}(a | x_i) = \frac{b_a q_a}{\sum_{j \in \mathcal{A}_i} b_j q_j}.$$

This allocation rule induces a probabilistic assignment over eligible ads, with probabilities that are fully determined by observed features. In contrast to deterministic formats (e.g., second-price auctions) where only the top-ranked ad is observed, the quasi-proportional mechanism introduces randomization across all eligible ads. This variation is central to our empirical framework, as it enables estimation of counterfactual outcomes and supports off-policy evaluation.

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

This follows directly from the quasi-proportional rule: every ad with a positive $b_a q_a$ has a strictly positive probability of being selected. Hence, the overlap assumption 1 is satisfied by construction for all participating ads in each auction.

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

$$\{Y_i(a)\}_{a \in \mathcal{A}_i} \perp A_i | x_i.$$

Therefore, the unconfoundedness assumption 2 is satisfied by the transparent structure of the auction: all inputs that determine allocation, namely, bids b_a , quality scores q_a , and eligibility constraints, are fully observed and fixed over the sample period. For each impression i , we observe the complete covariate vector x_i . Advertiser bids are not dynamically adjusted, and the platform does not personalize quality scores across users. As a result, the assignment rule is fully determined conditional on x_i .

This structure implies that for any impression i , we can compute not only the probability $\pi^{\mathcal{D}}(A_i | x_i)$ of the ad that was actually shown, but also the probabilities of all counterfactual ads that were eligible in the same auction. That is, even if a particular ad a was not displayed in impression i , as long as it was eligible (i.e., $a \in \mathcal{A}_i$), we can recover the counterfactual allocation probability $\pi^{\mathcal{D}}(a | x_i)$.

Eligibility Filtering. Although the quasi-proportional rule defines probabilities $\pi^{\mathcal{D}}(a | x_i)$ for all ads in the auction, not all ads in the global set \mathcal{A} are eligible in every impression. To ensure valid counterfactual estimation, we restrict attention to ads with nonzero probability of participation in each impression. Two factors determine eligibility: (i) *Contextual Targeting*: ads may be restricted to specific provinces, times, or app categories, and are excluded when their targeting criteria are not met; (ii) *Campaign availability*: some

ads may be inactive due to budget exhaustion or campaign timing. In practice, this is rare among top ads, as we select the top 10 ads for our analysis.

We construct an eligibility matrix $E \in \{0,1\}^{N \times A}$, where $e_{i,a} = 1$ indicates that ad a was eligible to compete in impression i , based on observed targeting and availability constraints. In practice, this requires that the impression’s metadata match the ad’s targeting filters on province, hour-of-day, and app. To avoid misclassification, we drop any impressions with missing targeting variables and restrict attention to a filtered sample where eligibility can be verified.

While this filtering step identifies the support of the logging policy, it does not tell us how likely each eligible ad is to be selected. For unbiased off-policy evaluation, we must account for the non-random assignment probabilities across eligible ads. Next, we describe how we estimate these propensities and assess the validity of the unconfoundedness assumption via covariate balance.

5.3.2 Propensity Score Estimation and Covariate Balance

To correct for unequal selection probabilities inherent in the auction, we estimate the *propensity scores* $\pi^{\mathcal{D}}(a | x_i)$, defined as the probability that the logging policy assigns impression i to ad a given the observed features x_i . These propensities quantify the exposure pattern generated by the platform’s allocation mechanism and form the basis for IPS in our policy evaluation.

Although the quasi-proportional rule maps bids and quality scores into theoretical probabilities, we adopt a data-driven estimation strategy to capture the realized assignment process in practice. This approach flexibly accommodates deviations from the theoretical rule, nonlinear effects, and high-order interactions among features. We use XGBoost to estimate $\pi^{\mathcal{D}}(a | x_i)$, motivated by its empirical performance in high-dimensional classification problems (Rafieian and Yoganasimhan, 2023). To avoid overfitting and ensure out-of-sample validity, we implement a 5-fold cross-fitting procedure, so that each estimated propensity score is computed on a model trained without the corresponding observation.

While observed inputs fully determine the allocation rule, we validate our identifying assumption by testing whether inverse-propensity weighting balances the distribution of observed covariates across treatment. The variables examined include province, app context, time of day, device brand, network type, and mobile service provider, features that advertisers can directly target. For each covariate and treatment group, we compute the standardized mean difference (SMD) before and after applying the estimated propensity weights, considering absolute values below 0.2 as indicative of acceptable balance (McCaffrey et al., 2013). The diagnostics show substantial improvements in covariate alignment across ads after weighting, providing empirical support for our research design. Details of propensity score estimation and balance statistics are reported in Web Appendix §C.2.1.

6 Empirical Results

We now present the empirical results derived from our proposed framework. §6.1 reports the predictive performance of machine learning models trained on different information sets. We then turn to the core question of how behavioral and geographical information contribute to decision quality, examining it at two levels of analysis. First, at the *aggregate level* (§6.2), we quantify the overall value of each information set and test whether the two act as substitutes or complements in improving targeting performance. Second, we

extend the analysis to the *user level* (§6.3), examining how the value and interaction of these information types change with the amount of behavioral history observed by the platform.

6.1 Predictive Performance Across Information Sets

Predicting user engagement is inherently challenging due to extreme class imbalance: fewer than 2% of impressions generate a click, making accuracy an uninformative metric. We therefore evaluate predictive performance using log loss, which assesses the calibration of predicted click probabilities and aligns with the models’ training objective; relative information gain (RIG), which measures the proportional improvement in log loss relative to a baseline model that predicts the average click-through rate; and the area under the ROC curve (AUC), which captures the model’s ability to rank clicked above non-clicked impressions in a threshold-independent and imbalance-robust manner.

We evaluate four models trained under different informational regimes: (i) the contextual-only baseline (X^\emptyset), (ii) the geographical regime (X^G), (iii) the behavioral regime (X^B), and (iv) the full information set (X^{GB}). The performance of each model on the held-out test set is summarized in Table 2.

Table 2: Predictive performance of models using contextual data (X^\emptyset), geographical data (X^G), behavioral data (X^B), and both (X^{GB}).

Model	Log Loss	AUC	Relative Information Gain (%)
(X^\emptyset)	0.073	0.720	18.900
(X^G)	0.072	0.722	19.220
(X^B)	0.014	0.809	84.180
(X^{GB})	0.013	0.812	84.410

Notes: The table reports Log Loss, AUC, and RIG on the test set. RIG is computed relative to the baseline click-through rate $p = 0.017592$. Sample size is $N = 3,162,376$ impressions.

The results reveal substantial variation in predictive performance across information structures. The contextual-only baseline performs weakest, with an RIG of 18.90%. Adding geographical features yields only a modest improvement, raising RIG to 19.22% and AUC to 0.722. In contrast, incorporating behavioral information leads to a large performance gain: the behavioral model achieves an RIG of 84.18% and an AUC of 0.809, indicating both substantially richer information and a superior ability to rank click outcomes. Adding geographical information to the behavioral model produces only marginal additional improvements, suggesting that behavioral features account for most of the predictive variation.

Taken together, these results demonstrate the dominant predictive value of behavioral information in estimating click likelihood, which aligns with prior findings in the advertising (Rafeian and Yoganarasimhan, 2021). While comparable AUC levels can be achieved using XGBoost-based behavioral models, the LSTM architecture delivers substantially higher RIG, indicating improved probabilistic calibration through temporal modeling. These performance levels provide empirical support for Assumption 3. Additional comparisons between LSTM and XGBoost models are reported in Appendix §D.1.

That said, it is important to emphasize again that predictive performance does not always translate into

decision value. A model may be well-calibrated or rank instances correctly, yet offer limited benefit when used to guide actions under realistic constraints. We therefore turn next to assessing the actual targeting value generated by each model in terms of the value function.

6.2 Behavioral vs. Geographical Information: Aggregate Level

Having established model performance in §6.1, we now assess how behavioral and geographical information improve decision quality. As described in §5.3, we evaluate each targeting policy $\pi^{(fx)}$ using IPS. We proceed in two steps. First, in §6.2.1, we quantify the *aggregate value* of each information set by comparing the expected value of policies based on different data sources. Then, in §6.2.2, we test whether behavioral and geographical information act as *substitutes or complements* in shaping decision quality.

6.2.1 Aggregate Value of Behavioral and Geographical Information

We consider a set of deterministic greedy policies indexed by the information set $X \in \{X^\emptyset, X^G, X^B, X^{GB}\}$ available at the time of decision. For each impression i , the policy selects, among eligible actions \mathcal{A}_i (see §5.3.1), the ad with the highest predicted click probability, restricting attention to actions with positive estimated logging propensity $\hat{\pi}^\mathcal{D}(a | X_i) > 0$. This ensures that all selected actions lie within the support of the logging policy and satisfy the overlap condition in Assumption 1.

Each policy represents a counterfactual targeting scenario in which the advertiser optimizes using only the corresponding information set X , and its value \hat{V}_X is estimated via IPS. To assess estimator uncertainty, we report 95% confidence intervals based on cluster-robust standard errors (clustered at the user level) and the effective sample size (ESS), which captures variance inflation due to skewed importance weights (Kallus, 2019).

The results are reported in Table 3. Across all policies, we find statistically significant improvements over the historical logging baseline, with tight confidence intervals and high ESS. In particular, McCaffrey et al. (2013) advocates trimming weights until $ESS \geq 0.10N$, a threshold all policies in our setting easily exceed. This suggests that the estimated policy values are stable and inference is well-powered.

Table 3: Estimated Policy Value Using IPS

Policy	IPS Estimate	95% CI	t-stat	SE	Lift (%)	ESS
$\hat{V}_{X^{GB}}$	0.0248***	[0.0243, 0.0254]	27.611	0.000	41.450	760316
\hat{V}_{X^B}	0.0229***	[0.0223, 0.0234]	19.675	0.000	30.280	555871
\hat{V}_{X^G}	0.0226***	[0.0220, 0.0232]	16.572	0.000	28.710	592427
\hat{V}_{X^\emptyset}	0.0208***	[0.0202, 0.0214]	18.782	0.000	20.440	542747

Notes: All estimates are computed using IPS. The baseline CTR under the logging policy is 0.017592. Lift is computed relative to this baseline. SE denotes the standard error of the estimate. Effective sample size (ESS) measures the number of equally weighted observations that would yield equivalent precision. Number of observations: 3,162,376. Cluster-robust standard errors are computed using 141,595 user clusters. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Starting from the contextual information policy X^\emptyset , we observe a 20.4% improvement over the baseline CTR. Although the platform lacks access to user-level data, this policy leverages contextual variation and selects the ad with the highest average CTR. This result highlights the value of exploiting aggregate performance differences even in the absence of personalized information.

Access to richer information sets produces further gains in policy value. Geographical information (X^G) and behavioral information (X^B) both generate substantial improvements relative to the contextual policy, increasing lift to 28.7% and 30.3%, respectively. Combining both information sources (X^{GB}) yields the highest policy value, a 41.5% improvement over the logging baseline. Notably, unlike the predictive results in §6.1, where behavioral information explained most of the performance gain, geographical information contributes comparably to decision value. This pattern underscores that improvements in predictive accuracy do not necessarily translate into proportional improvements in decision quality. Additional comparisons across learning algorithms are reported in Web Appendix §D.2.

6.2.2 Complement or Substitute? Aggregate Level

We now investigate whether behavioral and geographical data act as *substitutes* or *complements* in informing ad targeting decisions, starting with the aggregate level. Formally, let:

$$\Delta = (\hat{V}_{X^{GB}} - \hat{V}_{X^B}) - (\hat{V}_{X^G} - \hat{V}_{X^\emptyset}),$$

where \hat{V}_X denotes the IPS-based value estimate for policy X . The first term in parentheses measures the incremental value of adding geographical information when behavioral data are already available, while the second term measures the incremental value of adding geographical information when no user-level data are observed.

If $\Delta > 0$, behavioral and geographical information are *complements* (Definition 2), meaning the value of combining them exceeds the sum of their stand-alone gains. If $\Delta < 0$, they are *substitutes* (Definition 1), meaning the combined gain is less than additive and the two information sets overlap in the information they provide. We estimate Δ using per-impression differences and conduct a one-sample t -test on the mean, clustering standard errors at the user level to account for correlation within users. We report the two-sided test, as one-sided results cannot be significant when the two-sided test fails to reject the null.

Table 4: Aggregate Complementarity Test for Behavioral and Geographical Information

	Estimate ($\hat{\Delta}$)	Std. Error (clustered)	t -stat	95% Confidence Interval	p -value
Value	0.000509	0.000334	1.526	[-0.000144, 0.001164]	0.127

Notes: The table reports the aggregate test of complementarity between behavioral and geographical information. Standard errors are clustered at the user level. Number of observations = 3,162,376; number of user clusters = 141,595.

As shown in Table 4, the aggregate interaction estimate is close to zero and statistically insignificant. This implies that, on average, combining behavioral and geographical data produces nearly additive gains, with no systematic evidence of complementarity or substitutability at the aggregate level. While each information set captures distinct user heterogeneity, their joint effect does not amplify or diminish targeting value in the aggregate. Given this null aggregate finding, we next examine whether the nature of the interaction varies with user impression depth to assess whether heterogeneity in complementarity or substitutability emerges across users.

6.3 Behavioral vs. Geographical Information: Heterogeneity by User Exposure

The null result at the aggregate level motivates a closer examination of how the value and interaction of behavioral and geographical information vary across users. We focus on heterogeneity by *impression depth*, the number of impressions observed from each user, which proxies how much the platform has learned about each user. We proceed in two steps. First, in §6.3.1, we assess how the decision value of each information set changes with user exposure. Then, in §6.3.2, we test whether the relationship between behavioral and geographical information shifts from substitutive to complementary as more behavioral data accumulate.

6.3.1 Heterogeneous Value of Behavioral and Geographical Information by User Exposure

Building on the aggregate results presented in §6.2.1, we now explore how the value of information varies with the user’s impression history. Specifically, we investigate whether the gain from targeting differs depending on how many impressions a user has previously seen. This analysis is motivated by the discussion in §5.2, which posits that the estimation of reward function $f(X_t^B, a)$ becomes more accurate as more behavioral data is observed. If true, the marginal value of behavioral or combined data may evolve with impression depth.

To operationalize this, we first sort all impressions for each user j by their timestamp t and assign a depth index accordingly. We then group observations into bins such that each bin contains the same number of impressions (i.e., quantile-based binning). Within each bin, we compute the absolute CTR levels for each targeting policy using the IPS-estimated click probabilities, alongside the empirical baseline CTR. These per-bin averages are then plotted against impressions have seen by user to visualize how click-through performance evolves as users receive additional exposures. Figure 6a and Figure 6b present these results for the main targeting comparisons.

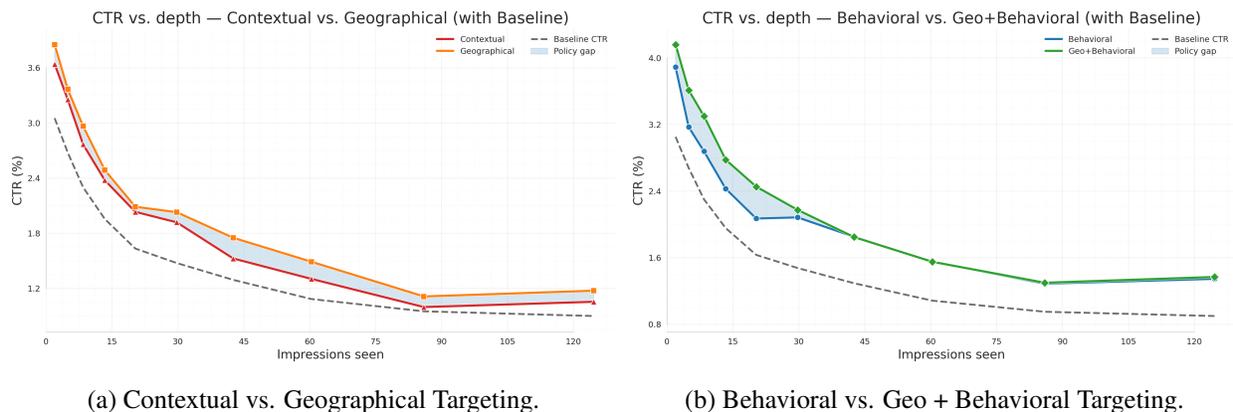


Figure 6: Heterogeneity in Policy Value by Impression Depth. Each plot shows the *absolute CTR levels* (in %) for two targeting policies and the empirical baseline. Shaded areas indicate the policy gap, while the dashed line marks the baseline CTR.

In the left panel of Figure 6a, the X^G (Geographical) policy consistently delivers higher CTR than the X^\emptyset (Contextual) policy across all impression depths. This persistent gap indicates that location-based segmentation provides a stable improvement over contextual cues alone. Although both policies rely on static

features, geographical information captures cross-regional heterogeneity that enhances targeting precision even in the absence of behavioral data.

In the right panel of Figure 6b, both the X^B (Behavioral) and X^{GB} (Combined) policies achieve substantially higher CTR than the baseline, highlighting the value of behavioral information for personalization. The X^{GB} policy consistently outperforms X^B , though the gap narrows as impression depth increases. This pattern implies that while geographical data initially enhances personalization in early exposures, its marginal contribution diminishes once rich behavioral histories accumulate.

6.3.2 Complement or Substitute? Heterogeneity by User Exposure

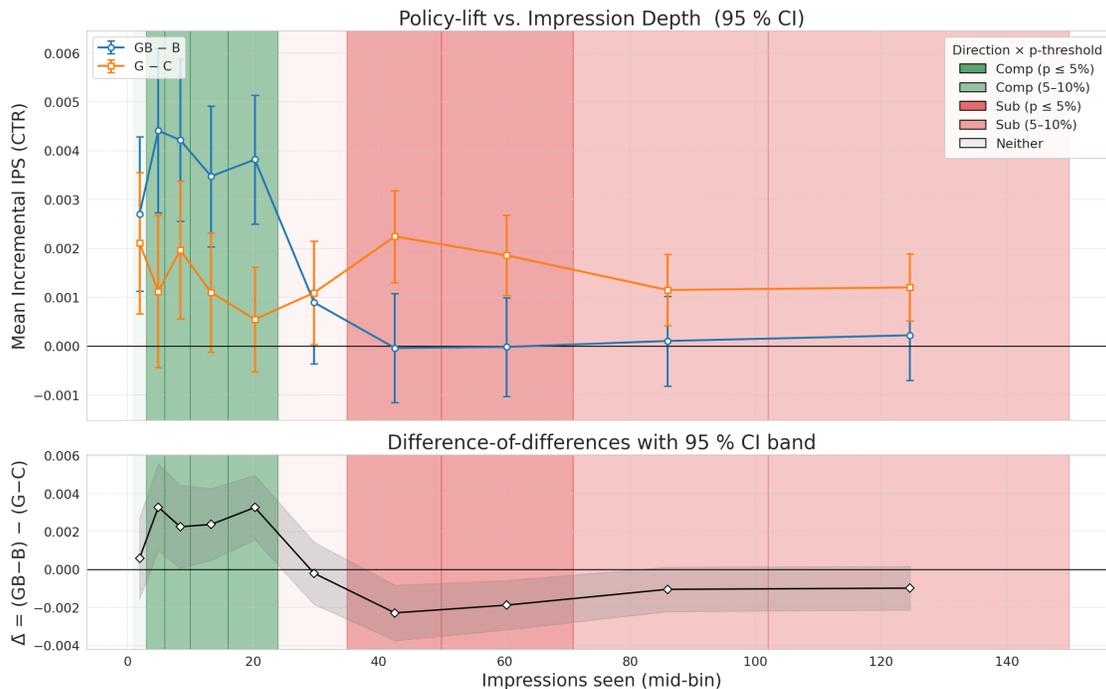


Figure 7: Complementarity and substitutability by impression depth. Top: $X^{GB} - X^B$ (blue) and $X^G - X^\emptyset$ (orange) with 95% CIs. Bottom: difference-in-differences $\bar{\Delta}$ with 95% CI; positive indicates complementarity, negative indicates substitutability.

Building on the aggregate results in Table 4, we now examine how the relationship between X^B and X^G varies with the number of impressions a user has seen. We group impressions into bins of equal size, each containing the same number of observations, defined as the number of prior ads shown to the same user. For each bin, we compute two policy-value differences: (i) the value difference between the X^{GB} and X^B policies, and (ii) the value difference between the X^G and X^\emptyset policies.

Figure 7 summarizes these patterns. The top panel plots both value differences across impression depth with 95% confidence intervals, while the bottom panel reports their difference-in-differences, $\bar{\Delta} = (X^{GB} - X^B) - (X^G - X^\emptyset)$, where positive values indicate complementarity (joint value exceeds additivity) and negative values indicate substitutability (information sets overlap in value). Shaded regions denote ranges with statistically significant effects.

The figure reveals three stages. In the *minimal behavioral history* stage (approximately 1–2 impressions), X^B contains little behavioral information. As a result, the stand-alone contribution of geographical data $X^G - X^\emptyset$ is at its highest. In this range, adding geographical data to behavioral data delivers nearly the same gain as geography alone, so $X^{GB} - X^B \approx X^G - X^\emptyset$. Consequently, X^{GB} yields little additional value relative to X^G .

In the *sparse behavioral history* stage (up to roughly 25 impressions), geographical information continues to play an important role as X^B begins to capture meaningful signals that, when combined with X^G , create synergy between the two information sources. In this range, geographical and behavioral information act as complements: the value difference $X^{GB} - X^B$ exceeds $X^G - X^\emptyset$, which indicates that adding X^G to X^B delivers incremental targeting value beyond their separate contributions.

Finally, in the *rich behavioral history* stage (from the low 20s and beyond), the marginal contribution of geographical data declines as accumulated behavioral information becomes substantially more informative. In this regime, the marginal contribution of adding X^G to X^B turns negative, indicating substitutability, as X^B already captures much of the variation that X^G would otherwise provide. For the full set of statistical test results underlying this analysis, we refer readers to Web Appendix §D.3.

7 Mechanisms Underlying the Role of Geographical Data

We now turn to the mechanisms that explain why geographical data (X^G) contributes to targeting outcomes in relation to behavioral data (X^B). Our unified framework has shown how the two information sets act as complements or substitutes in the decision value they generate. What remains unresolved is *why* geographical data improves performance in certain cases: does X^G capture an independent channel of information about user responsiveness, or does it primarily serve as a proxy for preference patterns that behavioral data eventually reveal?

To address this, we decompose spatial correlation in ad responsiveness into two distinct sources. The first is *influence*, which has a causal interpretation: users respond similarly *because* they are in the same location, which enables peer interaction and local information transmission. The second is *confounding*: users respond similarly *not because* of the same location itself, but because location is confounded with other factors, such as similar characteristics, income, demographics, or cultural norms, that independently shape preferences. This distinction mirrors the social network literature, which distinguishes *social influence* from *latent homophily* as alternative explanations for correlated behavior among connected individuals (Anagnostopoulos et al., 2008; Ma et al., 2015). We present the two sources of existing spatial correlation as follows:

1. *Spatial influence*. Under the causal interpretation, users may affect the decisions of others around them. For example, a user who sees an in-app advertisement for a product or promotion may mention it to friends or coworkers nearby, increasing their likelihood of engaging with the same ad when they encounter it. In this case, correlated responses arise from information spreading through local interactions rather than from shared preferences. Detecting such effects requires observing the spatial network that connects individuals. Geographical data (X^G) provides this information by revealing which users are located near one another and therefore more likely to interact. Behavioral data (X^B),

which records only each user’s actions, cannot capture these local interaction patterns. Consequently, if spatial influence exists, it can only be detected through geographical data that allows the underlying spatial network to be constructed.

2. *Spatial confounding.* Under the confounding interpretation, users respond similarly not because they influence one another, but because location is correlated with shared characteristics that shape preferences. For example, residents of a wealthy neighborhood may be more likely to click on luxury-product ads because they have similar income levels and consumption preferences. Geographical data (X^G) captures this pattern immediately, since location acts as a proxy for the bundle of characteristics associated with that area. However, these preferences are gradually revealed through individual behavioral data (X^B). A user who repeatedly engages with premium brands or high-end retailers reveals the same underlying preference that geography initially proxies. As behavioral histories accumulate, models can infer these preferences directly from past actions, making geographical data primarily an early but coarse proxy whose value declines as behavioral data becomes richer.

Confounding in location intuitively means that the location selection by individuals is a function of underlying characteristics that are also correlated with the outcome of interest. To formalize this, recall $\omega \in \Omega$ that denote the true latent type of a user defined in §4. The true click propensity is therefore a function of this latent type, which we denote by $g(\omega)$. Since ω is unobserved, the platform relies on observable behavioral data X^B to estimate this function. Let $f(X^B)$ denote the platform’s estimate of $g(\omega)$ based on the available behavioral history. We can therefore write click propensity as

$$y = g(\omega) + \theta_I + \varepsilon = f(X^B) + \underbrace{[g(\omega) - f(X^B)]}_{\theta_C: \text{estimation error}} + \theta_I + \varepsilon \quad (3)$$

where ε captures idiosyncratic variation in responses. In Equation (3), spatial influence is captured by θ_I , while spatial confounding appears in the estimation error $\theta_C = g(\omega) - f(X^B)$. Geographical data X^G is informative when spatial correlation remains in these terms after conditioning on behavioral data X^B . Under spatial influence, users affect the responses of nearby users through local interactions. These effects cannot be disentangled from individual preferences, regardless of how rich X^B becomes. If influence exists, it creates spatial correlation in θ_I that X^G can capture. Under spatial confounding, location is correlated with latent characteristics in ω that affect ad responsiveness. When behavioral histories are limited, $f(X^B)$ does not fully capture these characteristics, leaving spatial variation in $\theta_C = g(\omega) - f(X^B)$ that X^G can explain. As behavioral histories accumulate and $f(X^B)$ better approximates $g(\omega)$, this component shrinks. The key empirical question is whether the predictive value of X^G disappears as behavioral histories become richer or persists after conditioning on X^B .

7.1 Empirical Strategy: Residualized Spatial Autocorrelation Test

To evaluate this decomposition, we test whether spatial correlation remains after conditioning on behavioral data X^B . In Equation (3), spatial confounding enters through $\theta_C = g(\omega) - f(X^B)$ and should diminish as behavioral histories allow $f(X^B)$ to better approximate $g(\omega)$. Spatial influence, captured by θ_I , instead arises

from local interactions and therefore persists regardless of how rich X^B becomes. Our empirical strategy therefore tests whether residual spatial correlation disappears once behavioral information is incorporated. If it does, geographical data primarily reflects spatial confounding already absorbed by X^B ; if it persists, it indicates spatial influence. We implement this idea using a residualized spatial autocorrelation (RSA) test.

Our empirical strategy proceeds by aggregating outcomes at the regional level. Let Y_c denote the total number of clicks and I_c the total number of impressions in region c . Each impression can be viewed as a Bernoulli trial with region-specific click probability p_c , but because click events are rare, the Binomial distribution is well approximated by a Poisson model:

$$Y_c \mid p_c \sim \text{Poisson}(I_c p_c).$$

This specification is standard in spatial count-data analysis. The offset term $\log(I_c)$ adjusts for heterogeneous exposure across regions, while the residuals from the fitted model provide a natural basis for testing whether unexplained variation in click-through rates is spatially correlated.

- **Step 1: Baseline spatial structure in raw CTR:** As a starting point, we examine whether CTRs exhibit geographical clustering in the absence of any behavioral controls. We estimate a Poisson count model that includes only an exposure offset and define the residual:

$$Y_c \sim \text{Poisson}(\mu_c), \quad \log \mu_c = \alpha + \log(I_c), \quad \varepsilon_c^{(0)} = \log(Y_c/I_c) - \alpha,$$

Here, $\log(I_c)$ ensures that μ_c/I_c corresponds to the expected CTR in region c , while $\varepsilon_c^{(0)}$ measures deviations from the global mean CTR after adjusting for exposure volume. Evidence of spatial autocorrelation in $\varepsilon_c^{(0)}$ would indicate that raw CTRs contain location-based structure, consistent with spatial confounding or influence.

- **Step 2: Residual spatial structure after controlling for behavior:** Next, we ask whether this spatial structure remains once behavioral data are taken into account. To do so, we augment the model with the fitted behavioral component $f(X_c^B)$ and compute the adjusted residual:

$$Y_c \sim \text{Poisson}(\mu_c), \quad \log \mu_c = \alpha + f(X_c^B) + \log(I_c), \quad \varepsilon_c^{(B)} = \log(Y_c/I_c) - \alpha - f(X_c^B),$$

This residual captures variation in CTR unexplained by either exposure or behavioral features. If $f(X_c^B)$ successfully absorbs location-linked patterns, spatial dependence in $\varepsilon_c^{(B)}$ should be weaker than in $\varepsilon_c^{(0)}$.

- **Step 3: Testing for spatial dependence:** Finally, we test whether $\varepsilon_c^{(0)}$ and $\varepsilon_c^{(B)}$ exhibit spatial autocorrelation. We apply two standard measures: Moran's I (Moran, 1950), which captures global correlation across regions, and Geary's C (Geary, 1954), which is more sensitive to local clustering. For a generic set of residuals ε_c indexed by region c , these are defined as:

$$I = \frac{N}{\sum_i \sum_j w_{ij}} \cdot \frac{\sum_i \sum_j w_{ij} (\varepsilon_i - \bar{\varepsilon})(\varepsilon_j - \bar{\varepsilon})}{\sum_i (\varepsilon_i - \bar{\varepsilon})^2}, \quad C = \frac{(N-1) \sum_i \sum_j w_{ij} (\varepsilon_i - \varepsilon_j)^2}{2 \sum_i (\varepsilon_i - \bar{\varepsilon})^2},$$

where N is the number of regions, w_{ij} denotes the (i, j) element of the spatial weight matrix W , and $\bar{\varepsilon}$ is

the mean of the residuals. Significant positive values of Moran’s I (or values of Geary’s C below one) indicate that geographically proximate regions have similar residuals. By comparing the statistics for $\varepsilon_c^{(0)}$ and $\varepsilon_c^{(B)}$, we can assess whether behavioral features reduce spatial dependence, thereby clarifying whether geographical data provides independent information beyond what behavior explains.

7.2 Results: Spatial Correlation in CTRs

We present two sets of results. First, in §7.2.1, we examine spatial correlation in baseline versus behavior-adjusted CTR residuals. Second, in §7.2.2, we compare spatial dependence under sparse versus rich behavioral histories. Together, these analyses show how the role of geographical data changes once behavioral information is taken into account.

7.2.1 Residual Spatial Dependence in Baseline and Behavior-Adjusted Models

We begin by applying the RSA framework introduced in §7.1. Specifically, we compare spatial correlation in the baseline residuals $\varepsilon_c^{(0)}$, obtained without behavioral controls, to the behavior-adjusted residuals $\varepsilon_c^{(B)}$. This analysis is conducted at both the county and city levels to assess how much of the observed spatial dependence can be explained by behavioral data.

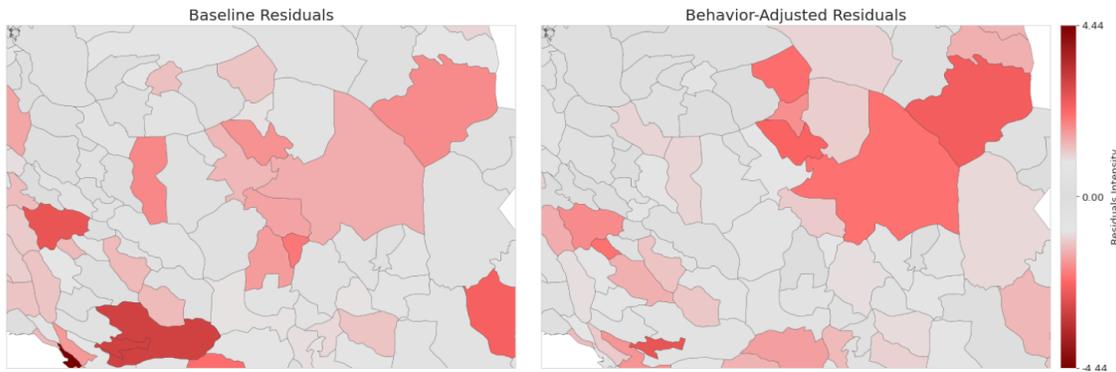


Figure 8: Spatial distribution of baseline (left) and behavior-adjusted (right) residuals at the county level. Darker red shades indicate higher residual intensity. Behavioral adjustment visibly reduces, but does not fully eliminate, spatial clustering.

Figure 8 provides an intuitive visualization for the county-level case. The left panel displays the baseline residuals $\varepsilon_c^{(0)}$ from the model without behavioral controls. Several contiguous regions show clusters of similarly high or low residual values, indicating clear spatial autocorrelation. Such clustering is consistent with the presence of spatial homophily and/or influence. In contrast, the right panel shows the behavior-adjusted residuals $\varepsilon_c^{(B)}$. Incorporating $f(X_c^B)$ visibly reduces the strength and extent of spatial clusters, particularly in the southwest, although notable pockets remain. This pattern suggests that behavioral data absorb a substantial share of the location-linked variation, but not all of it.

To quantify the patterns observed in Figure 8, Table 5 reports Moran’s I and Geary’s C at both county and city levels. At the county level, the baseline residuals $\varepsilon_c^{(0)}$ exhibit strong and significant spatial dependence: regions with unusually high (or low) CTR residuals tend to be geographically proximate. After controlling for behavioral information, both statistics decline markedly, indicating that much of this cluster-

ing is explained by differences in user behavior. However, the remaining positive spatial correlation implies that geographical data continues to capture additional variation not yet absorbed by behavior.

At the city level, the attenuation is sharper. Baseline residuals still display detectable clustering, but once behavioral features are included, both Moran’s I and Geary’s C fall to levels that are statistically indistinguishable from zero. This result suggests that at finer geographical resolution, most of the spatial variation in CTR can be accounted for by behavioral histories, leaving little independent role for geographical data. The contrast between county and city levels thus highlights how geographical data’s incremental contribution diminishes as spatial units become more granular, consistent with the idea that geographical data largely serves as a proxy for spatial confounding that behavioral data can eventually capture.

	County-level, $\varepsilon_c^{(0)}$	County-level, $\varepsilon_c^{(B)}$	City-level, $\varepsilon_c^{(0)}$	City-level, $\varepsilon_c^{(B)}$
Moran’s I	0.122	0.084	0.055	-0.001
p -value	0.00006	0.00318	0.00242	0.49156
Geary’s C	0.891	0.900	0.948	0.996
p -value	0.00094	0.00162	0.01646	0.41863

Table 5: Moran’s I and Geary’s C statistics for spatial autocorrelation of residuals before (baseline) and after (behavior-adjusted) controlling for behavioral features, at county and city levels.

Taken together, these results show that geographical data explains substantial spatial variation in ad responses when behavioral information is excluded, but much of this variation disappears once behavioral data are incorporated. Moreover, the remaining spatial structure becomes smaller at finer spatial resolutions. To examine whether this attenuation also depends on the amount of behavioral information available for each user, we next split impressions into sparse and rich behavioral-history subsets and repeat the analysis.

7.2.2 Residual Spatial Dependence under Sparse versus Rich Behavioral Histories

As discussed earlier, the ability of the behavioral model $f(X^B)$ to account for variation in CTR depends on the amount of behavioral history available for each user. To examine this effect, we split the sample into two groups: (i) a *sparse-history* group, consisting of the first 50% of impressions observed for each user, and (ii) a *rich-history* group, consisting of an equal number of later impressions for the same users. This design allows us to compare scenarios where targeting models rely on limited versus extensive behavioral histories.

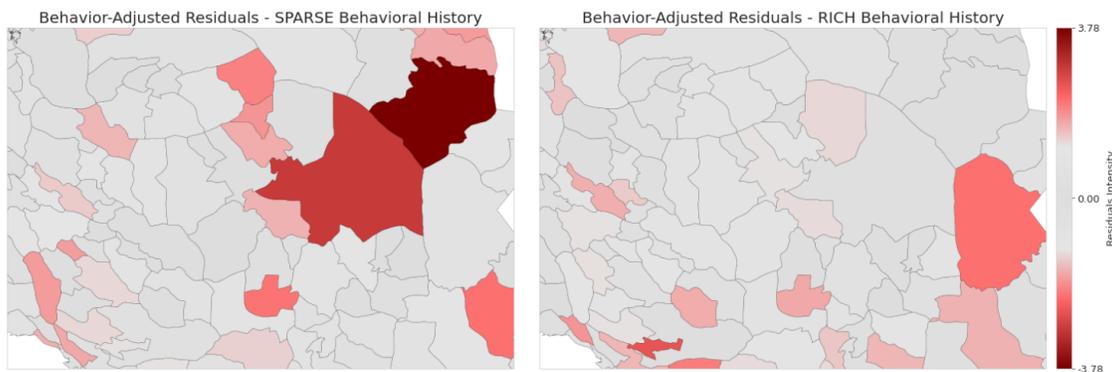


Figure 9: Behavior-adjusted residuals at the county level for sparse (left) and rich (right) behavioral histories. Darker shades indicate higher residual intensity.

Figure 9 illustrates the spatial distribution of behavior-adjusted residuals at the county level for the two groups. In the sparse-history case (left panel), several large contiguous clusters of high residuals remain even after controlling for $f(X^B)$, indicating that location continues to explain a meaningful share of CTR variation. By contrast, in the rich-history case (right panel), residual clustering is visibly weaker and more fragmented, suggesting that much of the spatial dependence disappears once users accumulate longer behavioral records.

	County-level, Sparse	County-level, Rich	City-level, Sparse	City-level, Rich
Moran's I	0.069	0.025	-0.005	0.0005
p -value	0.011	0.183	0.412	0.459
Geary's C	0.952	0.976	1.012	1.006
p -value	0.072	0.225	0.285	0.403

Table 6: Moran's I and Geary's C statistics for spatial autocorrelation of behavior-adjusted residuals under sparse and rich behavioral histories, at county and city levels.

Table 6 formalizes these visual patterns. Under sparse histories, Moran's I is positive and statistically significant at the county level, while Geary's C also indicates non-random clustering. These results confirm that when behavior is limited, geographical data captures residual variation that remains spatially structured. In the rich-history case, however, both statistics decline substantially in magnitude and lose statistical significance across county and city levels. This indicates that once behavioral data are rich, $f(X^B)$ already accounts for nearly all of the systematic spatial correlation in CTR.

The comparison between sparse and rich behavioral histories also sheds light on the mechanisms through which geographical data matters. When behavioral histories are sparse, residual spatial structure may arise from spatial influence, spatial confounding, or both. Observing how this structure evolves as behavioral data becomes richer is therefore informative about which mechanism dominates. Spatial influence is inherently relational and cannot be recovered from individual behavioral data X^B , regardless of how rich those histories become. If spatial influence were the primary driver, residual spatial autocorrelation would persist even after conditioning on X^B . Instead, we find that spatial autocorrelation largely disappears once behavioral histories are rich. This pattern rules out spatial influence and suggests that spatial confounding is the dominant channel: geographical data initially captures similarities in preferences across nearby users, but these similarities are eventually learned directly from individual behavioral histories, leaving little residual role for X^G .

8 Managerial and Policy Implications

From the perspective of advertising platforms and advertisers, our results highlight a clear trade-off in the use of geographical information for engagement modeling. Geographical data provides incremental value primarily during early cold-start phases, when behavioral histories are sparse and cannot yet support effective personalization. As behavioral data accumulates, it increasingly substitutes for the value previously provided by geographical information, making location-based targeting largely redundant. This pattern suggests that firms can restrict the use of geographical data to early cold-start stages and gradually transition to behavioral targeting as behavioral histories mature. Such an approach allows firms to reduce unnecessary collection of location information while maintaining targeting performance, and may also strengthen consumer trust through transparent, user-controlled opt-out mechanisms for geographical tracking (Martin

et al., 2017).

From a consumer perspective, the substitutability of geographical information by behavioral information implies that consumers can safely opt out of location tracking once sufficient behavioral history has accumulated. In this regime, relevant targeting can be maintained without continued use of geographical information, allowing consumers to limit exposure of sensitive location attributes without materially affecting engagement outcomes.

From a regulatory perspective, the limited incremental value of geographical information in the presence of rich behavioral histories raises questions about the justification for continued large-scale location tracking. While geographical information may serve a temporary role in early cold-start settings, its long-run substitutability suggests that stricter oversight of persistent geographical tracking is warranted, particularly in digital advertising environments where comparable targeting performance can be achieved with behavioral information alone.

9 Conclusion

In this paper, we develop a unified framework to evaluate the value of information in digital advertising, integrating economic theory, machine learning, and causal inference. Conceptually, we extend decision-based definitions of information value to test whether behavioral and geographical data act as complements or substitutes. Methodologically, we combine LSTM architectures with attention mechanisms to capture temporal and spatial dynamics and use inverse propensity scoring to recover counterfactual performance from observational data. Together, these elements provide a generalizable approach for comparing multiple information sets in high-dimensional environments, with applications extending beyond advertising.

Our results show that both geographical and behavioral data enhance targeting, but their roles change as behavioral histories become richer. At the aggregate level, complementarities and substitutions offset one another, yielding no clear overall interaction. Examining heterogeneity as users accumulate behavioral histories allows us to uncover a clear pattern. Geographical data explains most variation when behavioral information is minimal, complements behavioral data when histories are sparse, and becomes a substitute once histories are rich. This progression indicates that the value of geographical data is temporary and concentrated in the early stages before behavioral information alone suffices for accurate targeting.

Future research could extend our analysis in several ways. First, beyond in-app advertising, the value of geographical data may differ in sectors where location remains central, such as transportation, local retail, or urban planning. Second, future work should trace the privacy–utility frontier by evaluating policy value under explicit data-use constraints, including coarse geolocation, k -anonymity, differential privacy, or federated learning. Finally, our analysis focuses on online in-app advertising, where engagement is primarily shaped by behavioral information rather than physical location. We do not generalize these conclusions to industries in which location carries intrinsic value, such as ride-hailing, food delivery, local event promotions, or brick-and-mortar retail discounts. In such contexts, geographical information may remain a critical input, and further research is needed to assess whether behavioral information can fully substitute for location-based targeting.

References

- Alessandro Acquisti, Curtis Taylor, and Liad Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–492, 2016.
- Aris Anagnostopoulos, Ravi Kumar, and Mohammad Mahdian. Influence and correlation in social networks. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 7–15, 2008.
- Guy Aridor, Yeon-Koo Che, Brett Hollenbeck, Daniel McCarthy, and Maximilian Kaiser. Evaluating the impact of privacy regulation on e-commerce firms: Evidence from apple’s app tracking transparency. 2024.
- Eva Ascarza. Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, 55(1):80–98, 2018.
- David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2):265–272, 1953.
- Bryan Bollinger and Kenneth Gillingham. Peer effects in the diffusion of solar photovoltaic panels. *Marketing Science*, 31(6):900–912, 2012.
- Tilman Börgers, Angel Hernando-Veciana, and Daniel Krähmer. When are signals complements or substitutes? *Journal of Economic Theory*, 148(1):165–195, 2013.
- Eric T Bradlow, Bart Bronnenberg, Gary J Russell, Neeraj Arora, David R Bell, Sri Devi Duvvuri, Frankel Ter Hofstede, Catarina Sismeiro, Raphael Thomadsen, and Sha Yang. Spatial models in marketing. *Marketing letters*, 16(3):267–278, 2005.
- Bart J Bronnenberg and Vijay Mahajan. Unobserved retailer behavior in multimarket data: Joint spatial dependence in market shares and promotion variables. *Marketing Science*, 20(3):284–299, 2001.
- Bart J Bronnenberg, Sanjay K Dhar, and Jean-Pierre H Dubé. Brand history, geography, and the persistence of brand shares. *Journal of political Economy*, 117(1):87–115, 2009.
- Ruomeng Cui, Dennis J Zhang, and Achal Bassamboo. Learning from inventory availability information: Evidence from field experiments on amazon. *Management Science*, 65(3):1216–1235, 2019.
- Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports*, 3(1):1–5, 2013.
- Yves-Alexandre De Montjoye, Sébastien Gambs, Vincent Blondel, Geoffrey Canright, Nicolas De Cordes, Sébastien Deletaille, Kenth Engø-Monsen, Manuel Garcia-Herranz, Jake Kendall, Cameron Kerry, et al. On the privacy-conscientious use of mobile phone data. *Scientific data*, 5(1):1–6, 2018.
- eMarketer. Mobile advertising 2024: In-app ads drive growth past \$200 billion. <https://www.emarketer.com/content/mobile-advertising-2024>, 2024. Accessed August 2025.
- Robert C Geary. The contiguity ratio and statistical mapping. *The incorporated statistician*, 5(3):115–146, 1954.
- Anindya Ghose, Beibei Li, and Siyuan Liu. Mobile targeting using customer trajectory patterns. *Management Science*, 65(11):5027–5049, 2019.
- Avi Goldfarb and Catherine Tucker. Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404, 2011.
- Mihajlo Grbovic and Haibin Cheng. Real-time personalization using embeddings for search ranking at airbnb. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 311–320, 2018.
- Cheng Guo and Felix Berkhahn. Entity embeddings of categorical variables. *arXiv preprint arXiv:1604.06737*, 2016.
- Keisuke Hirano, Guido W Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.
- S Hochreiter. Long short-term memory. *Neural Computation MIT-Press*, 1997.

- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.
- Raghuram Iyengar, Christophe Van den Bulte, and Thomas W Valente. Opinion leadership and social contagion in new product diffusion. *Marketing science*, 30(2):195–212, 2011.
- Kinshuk Jerath and Klaus M Miller. Consumers’ perceived privacy violations in online advertising. *arXiv preprint arXiv:2403.03612*, 2024.
- Garrett A Johnson, Scott K Shriver, and Shaoyin Du. Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1):33–51, 2020.
- Nathan Kallus. Intrinsically efficient, stable, and bounded off-policy evaluation for reinforcement learning. In *Advances in Neural Information Processing Systems*, 2019.
- Emir Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11(1):249–272, 2019.
- Mingyung Kim, Eric T Bradlow, and Raghuram Iyengar. Selecting data granularity and model specification using the scaled power likelihood with multiple weights. *Marketing Science*, 41(4):848–866, 2022.
- Jeffrey S Larson, Eric T Bradlow, and Peter S Fader. An exploratory look at supermarket shopping paths. *International Journal of research in Marketing*, 22(4):395–414, 2005.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *International Conference on Learning Representations (ICLR)*, 2019.
- Bowen Luo and Bhoomija Ranjan. Mapping spatial heterogeneity in retail advertising effectiveness. *Journal of Marketing Research*, 62(6):1063–1080, 2025.
- Liye Ma, Ramayya Krishnan, and Alan L Montgomery. Latent homophily or social influence? an empirical analysis of purchase within a social network. *Management Science*, 61(2):454–473, 2015.
- Puneet Manchanda, Ying Xie, and Nara Youn. The role of targeted communication and contagion in product adoption. *Marketing Science*, 27(6):961–976, 2008.
- Kelly D Martin, Abhishek Borah, and Robert W Palmatier. Data privacy: Effects on customer and firm performance. *Journal of marketing*, 81(1):36–58, 2017.
- 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.
- 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.
- Patrick AP Moran. Notes on continuous stochastic phenomena. *Biometrika*, 37(1/2):17–23, 1950.
- Unnati Narang and Fernando Luco. Privacy and prediction: how useful are geo-tracking data for predicting consumer visits? *Quantitative Marketing and Economics*, 23(4):523–544, 2025.
- Pew-Center. Mobile fact sheet. <https://www.pewresearch.org/internet/fact-sheet/mobile/>, 2024. Accessed August 2025.
- Massimo Quadrana, Paolo Cremonesi, and Dietmar Jannach. Sequence-aware recommender systems. *ACM computing surveys (CSUR)*, 51(4):1–36, 2018.
- Omid Rafieian. Optimizing user engagement through adaptive ad sequencing. *Marketing Science*, 42(5): 910–933, 2023.
- Omid Rafieian and Hema Yoganasimhan. Targeting and privacy in mobile advertising. *Marketing Science*, 40(2):193–218, 2021.
- Omid Rafieian and Hema Yoganasimhan. Ai and personalization. *Artificial Intelligence in Marketing*, pages 77–102, 2023.
- Peter E Rossi, Robert E McCulloch, and Greg M Allenby. The value of purchase history data in target marketing. *Marketing Science*, 15(4):321–340, 1996.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

- Adam N Smith, Stephan Seiler, and Ishant Aggarwal. Optimal price targeting. *Marketing Science*, 42(3): 476–499, 2023.
- Catherine E Tucker. Social networks, personalized advertising, and privacy controls. *Journal of marketing research*, 51(5):546–562, 2014.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Nils Wernerfelt, Anna Tuchman, Bradley T Shapiro, and Robert Moakler. Estimating the value of offsite tracking data to advertisers: Evidence from meta. *Marketing Science*, 44(2):268–286, 2025.

Web Appendix

A Details of Machine Learning Framework

In this section, we model f_X by specifying both the information set X and the functional form $f(\cdot)$. We discuss the specific features included in each informational regime, $\{X^\emptyset, X^G, X^B, X^{GB}\}$, in §A.1. We then describe the implementation details of behavioral modeling using an LSTM with an attention mechanism, with additional details provided in §A.2.

A.1 Structure of Information Sets X

For our machine learning model, each targeting regime $X^{(\cdot)}$ corresponds to a specific set of observable features. We classify these features into three dimensions, *contextual*, *behavioral*, and *geographical*, which form the building blocks of the different regimes:

Contextual Features (X^C): Contextual features characterize the *setting* in which each ad impression occurs, independent of user identity or location. These include time of day, day of week, app, device, and network provider. Engagement often follows systematic patterns along these dimensions; for instance, click rates in gaming apps may peak in the evening when many users log in.

Behavioral Features (X^{Behave}): Behavioral features capture a user’s *online footprint*, such as prior ad exposures, cumulative clicks, time since last click, and category-specific CTRs. These variables reveal individual preferences and engagement; for example, a user who frequently clicks on sports ads is much more likely to respond to a new sports campaign.

Geographical Features (X^{Geo}): Geographical features reveal the user’s *physical footprint*, their location at the time of the impression, which may be represented at various levels of granularity, such as city or precise latitude and longitude. They reflect spatial variation in preferences and engagement, for example, ride-sharing ads see higher click rates in dense urban centers during rush hour compared to rural areas.

While we have conceptually defined the targeting regimes, here we map them to their feature representations. Each $X^{(\cdot)}$ corresponds to a specific subset of contextual, behavioral, and geographical features used in estimation:

- ($X^\emptyset = X^C$): The advertiser observes only contextual features (e.g., device, app category, time, or network conditions). No user-level information, such as behavioral or geographical data, is available, so actions are chosen based on context-dependent averages rather than latent user types ω . Such information, typically available from campaign reporting or platform analytics.
- ($X^B = X^{Behave} \cup X^C$): The advertiser observes both the context of each impression and features that summarize the user’s prior behavior. Together, these variables provide insight into the user’s likely preferences and engagement, revealing information about the latent type ω through observed patterns.
- ($X^G = X^{Geo} \cup X^C$): The advertiser observes both contextual information and the user’s geographical location. Geographical features can reveal spatial heterogeneity in preferences, such as differences due to local events, regional trends, or time zones, and thus offer indirect information about a user’s type ω .

- ($X^{GB} = X^{\text{Behave}} \cup X^{\text{Geo}} \cup X^C$): The advertiser has access to the full set of behavioral, geographic, and contextual features. This regime enables the model to capture the most complete picture of user heterogeneity to infer latent user types ω and guide targeting.

These categories mirror the types of data advertisers realistically access in practice. The no-information case serves as a benchmark for non-personalized targeting, while the other regimes reveal distinct dimensions of latent user preferences. Details on feature construction, variable definitions, and preprocessing appear in the Web Appendix (§B).

A.2 Behavioral Learning Algorithm Implementation: LSTM

To effectively model the sequential and temporal dynamics of user engagement, we employ the Long Short-Term Memory (LSTM) predictive model, a specialized class of Recurrent Neural Networks (RNNs) (Rumelhart et al., 1986). LSTMs, first introduced by Hochreiter (1997), address the limitations of traditional RNNs, such as the vanishing gradient problem, by incorporating a memory cell and gating mechanisms that regulate the flow of information. These features make LSTMs highly effective for tasks involving long-term dependencies and sequential patterns. LSTMs have been widely adopted in user engagement modeling due to their ability to handle sequential dependencies and dynamic user behavior (Grbovic and Cheng, 2018; Quadrana et al., 2018).

An LSTM unit maintains two key components at each time step t : the cell state C_t , which acts as a long-term memory reservoir, and the hidden state h_t , which encapsulates information relevant for the current time step. The evolution of these components is controlled by three gating mechanisms: the forget gate, the input gate, and the output gate, as defined:

- **Forget Gate.** The forget gate determines how much of the previous cell state C_{t-1} should be retained:

$$f_t = \sigma(W_f[h_{t-1}, x_t] + b_f),$$

where f_t is the forget gate vector, h_{t-1} is the hidden state from the previous time step, x_t is the input at the current time step, W_f and b_f are learnable parameters (weights and biases), and $\sigma(\cdot)$ is the sigmoid activation function, mapping values to the range $[0, 1]$.

- **Input Gate.** The input gate determines the extent of new information to be incorporated into the memory cell. It is defined as:

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i),$$

$$\tilde{C}_t = \tanh(W_C[h_{t-1}, x_t] + b_C),$$

where i_t is the input gate vector, \tilde{C}_t is the candidate cell state, and W_i, W_C, b_i, b_C are trainable parameters. The hyperbolic tangent (\tanh) maps the candidate state values to $[-1, 1]$.

- **Cell State Update.** The cell state C_t is updated by combining the contributions of the forget and input gates:

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t,$$

where \odot denotes element-wise multiplication.

- **Output Gate.** The output gate determines the current hidden state h_t , which is used to generate the output and influence subsequent computations:

$$o_t = \sigma(W_o[h_{t-1}, x_t] + b_o),$$

$$h_t = o_t \odot \tanh(C_t),$$

where o_t is the output gate vector and W_o, b_o are learnable parameters.

These equations collectively enable the LSTM to selectively retain, update, and output information across time steps, making it robust to long-term dependencies in sequential data.

In addition to the recurrent backbone, our predictive framework integrates several complementary components, each designed to address a specific challenge in sequential engagement modeling. First, *embedding layers* transform high-cardinality categorical inputs, such as device type, province, city, or network provider, into low-dimensional dense vectors. This approach not only reduces dimensionality but also enables the model to learn latent similarity structures among categories, improving generalization and mitigating the curse of dimensionality (Guo and Berkhahn, 2016). Second, temporal structure is enriched through two mechanisms: (i) *absolute positional embeddings*, which provide the model with awareness of an event’s position within the sequence, and (ii) *time-gap embeddings*, which capture irregular intervals between events, allowing differentiation between dense bursts of activity and prolonged inactivity. Third, a *stacked LSTM* serves as the primary sequence encoder, effectively modeling order-sensitive dependencies and evolving behavioral patterns over time. Fourth, a *causal multi-head self-attention* layer enables the model to focus adaptively on past events that are most relevant to the current prediction, capturing long-range dependencies while enforcing strict causality constraints (Vaswani et al., 2017). Finally, a *gated projection head* selectively filters and transforms the attended sequence representations before a linear readout layer maps them to click probabilities, enhancing the interpretability and robustness of the final decision stage. Table A1 reports the architecture and parameter counts.

In the following, we provide a detailed, step-by-step description of the architecture, formally defining its components and explaining how they interact to transform raw sequential inputs into engagement probability forecasts.

We model user engagement as a sequential decision prediction problem, where the state at each time step incorporates both numerical and categorical information observed up to that point. Let j index users and $t \in \{1, \dots, T\}$ denote the position in the sequence of impressions observed for user j , with a maximum sequence length $T = 150$.

Input Representation. At each time step t , we observe:

$$X_{jt}^{\text{num}} \in \mathbb{R}^n, \quad \{X_{jt}^{\text{cat},k}\}_{k=1}^c,$$

Table A1: Summary of the Sequence Model Architecture and Parameters

Layer (type)	Output Shape	Param #
Input Features	[-1, 150, 240]	0
Time-Gap Projection	[-1, 150, 240]	57,840
Absolute Positional Embedding	[1, 150, 240]	36,000
LSTM (4 layers)	[-1, 150, 512]	4,269,568
Layer Norm (pre-attention)	[-1, 150, 512]	1,024
Causal Multi-Head Attention	[-1, 150, 512]	1,049,600
Layer Norm (post-attention)	[-1, 150, 512]	1,024
Gate Projection (sigmoid)	[-1, 150, 512]	262,656
Body Projection (tanh)	[-1, 150, 512]	262,656
Dropout	[-1, 150, 512]	0
Final FC Output Layer	[-1, 150, 1]	513
Total Parameters		5,941,905

where n is the number of numerical features and c is the number of categorical features. Numerical features capture continuous signals (e.g., time-based statistics, historical engagement rates, contextual metrics). Categorical features correspond to high-cardinality identifiers such as device model, province, city, or advertisement ID. Direct one-hot encoding of categorical features would be prohibitively high-dimensional and sparse. Instead, we map each categorical variable $X_{jt}^{\text{cat},k}$ to a dense, learnable embedding vector:

$$\mathbf{e}_{jt}^{(k)} \in \mathbb{R}^{d_k},$$

where d_k is the embedding dimension for category k . Embeddings enable the model to learn latent similarity structures, e.g., provinces with similar user behavior or devices with similar performance profiles, without manual feature engineering. Concatenating all categorical embeddings yields:

$$\mathbf{e}_{jt} = [\mathbf{e}_{jt}^{(1)}, \dots, \mathbf{e}_{jt}^{(c)}] \in \mathbb{R}^{E_{\text{cat}}}, \quad E_{\text{cat}} = \sum_{k=1}^c d_k.$$

The full feature vector at time t is:

$$X_{jt} = [X_{jt}^{\text{num}}, \mathbf{e}_{jt}] \in \mathbb{R}^{d_{\text{in}}}, \quad d_{\text{in}} = n + E_{\text{cat}}.$$

Temporal Augmentation. To model sequence position and irregular arrival times, we augment X_{jt} with two temporal components:

1. *Absolute positional embedding* $\mathbf{P}_t \in \mathbb{R}^{d_{\text{in}}}$, a learnable vector for each position t , allows the model to differentiate between early- and late-stage interactions even if other features are identical.

While LSTMs can track order, they do so implicitly through hidden state transitions. Positional embeddings make temporal position an explicit feature, which is helpful in cases where position itself carries meaning (e.g., “first impressions” vs. “later impressions” behave differently). LSTM gates learn relative dependencies; positional embeddings give a direct absolute time index, which is crucial

for long sequences where LSTM memory may fade.

2. *Time-gap embedding* models the elapsed time since the previous impression:

$$\Delta t_{jt} = \text{time}(j, t) - \text{time}(j, t - 1),$$

transformed as $\log(1 + \Delta t_{jt})$ and projected to $\mathbb{R}^{d_{\text{in}}}$ through a learned linear layer. This captures behavioral intensity, distinguishing between dense bursts of activity and long idle periods.

The spacing between events often changes behavior (e.g., if a user sees an ad after 5 seconds vs. after 5 days). Neither positional embeddings nor LSTM hidden states naturally capture this irregular spacing. Positional embeddings tell us “this is the 10th impression,” but not whether it came an hour or a week after the previous one. Time-gap embeddings fill that gap.

The resulting time-aware token is:

$$\tilde{X}_{jt} = X_{jt} + \mathbf{P}_t + \text{TimeGapProj}(\log(1 + \Delta t_{jt})).$$

Sequence Encoding. The augmented sequence $\{\tilde{X}_{j1}, \dots, \tilde{X}_{jT}\}$ is processed by a stacked LSTM with hidden size $H = 512$ and $L_{\text{LSTM}} = 4$ layers:

$$h_{jt}^{(\ell)} = \text{LSTM}_{\ell}(h_{jt}^{(\ell-1)}),$$

where $h_{jt}^{(0)} = \tilde{X}_{jt}$. The LSTM captures local and medium-range temporal dependencies, preserving order-sensitive dynamics such as gradual preference shifts or decaying effects of prior interactions. Attention alone doesn’t model sequential dependencies as naturally as LSTMs, and attention without recurrence can sometimes struggle with shorter-term patterns when data is noisy. The LSTM provides a strong temporal backbone before attention refines it.

Causal Multi-Head Attention. To model longer-range dependencies and non-local feature interactions, we apply a causal multi-head self-attention layer with $H_{\text{attn}} = 4$ heads to the LSTM outputs. Attention reweights historical states based on their relevance to the current prediction. We enforce causality with a triangular mask:

$$\text{mask}(t', t) = \begin{cases} 0, & t' \leq t, \\ -\infty, & t' > t, \end{cases}$$

ensuring that predictions at time t use only past and present information. A key-padding mask derived from $m_{jt} \in \{0, 1\}$ ignores padded steps. Layer normalization before and after attention stabilizes optimization.

LSTMs compress history into a single hidden state, which can lose fine-grained details over long sequences. Attention re-opens the past and selectively retrieves important past events directly. It complements the LSTM by providing content-based memory access useful when long-range dependencies matter, e.g., a behavior 100 impressions ago is relevant now.

Gated Projection Head. The attention-enhanced representation is passed through a gated projection mechanism, which adaptively filters and transforms features:

$$g_{jt} = \sigma(W_g h_{jt}), \quad b_{jt} = \tanh(W_b h_{jt}),$$

$$z_{jt} = \text{Dropout}(g_{jt} \odot b_{jt}).$$

Here, the gate vector g_{jt} acts as a learnable feature selector, while b_{jt} generates candidate transformations. Their elementwise product modulates information flow, suppressing irrelevant history and emphasizing predictive patterns.

It lets the model selectively emphasize or suppress certain dimensions in the representation before making predictions, essentially a learned attention mechanism at the feature level. Attention looks across time, but this gating looks within the feature vector, so it’s a different form of selectivity.

Prediction Layer. The gated features are mapped to logits:

$$\tilde{y}_{jt} = w^\top z_{jt}, \quad \hat{y}_{jt} = \sigma(\tilde{y}_{jt}),$$

where \hat{y}_{jt} is the predicted click probability at time t .

Training Objective. We minimize the binary cross-entropy loss with logits:

$$\mathcal{L}(\Theta) = \sum_{j=1}^n \sum_{t=1}^T m_{jt} \ell_{\text{BCE-logits}}(y_{jt}, \tilde{y}_{jt}),$$

where m_{jt} masks out padded steps, and Θ includes all learnable parameters (embeddings, positional encodings, LSTM, attention, gating, and readout). Optimization uses AdamW (Loshchilov and Hutter, 2019), which combines adaptive learning rates with decoupled weight decay. Dropout is applied both in attention and projection layers for regularization.

This architecture integrates heterogeneous data types, numerical and categorical, into a unified, dense representation. Embeddings allow the model to exploit latent similarities among categorical entities without incurring the cost of high-dimensional sparse vectors. The positional and time-gap embeddings encode absolute sequence position and irregular temporal dynamics, both of which are crucial for interpreting user engagement behavior. The LSTM provides strong modeling of local order-sensitive dependencies, while the causal attention layer captures global relationships, enabling the model to focus selectively on distant but influential events. Finally, the gated projection head acts as a content-aware filter, passing forward only the most relevant transformed features for prediction. Together, these design elements yield a model capable of accurately forecasting engagement while respecting strict causality constraints.

B Feature Generations X and Learning Algorithm Implementation

The features X comprise **contextual** (X^C), **geographical** (X^{Geo}), and **behavioral** (X^{Behave}) components. Contextual features include static metadata (e.g., time, device, network), geographical features encode spatial attributes (latitude/longitude, province, city) using low-dimensional embeddings, and behavioral fea-

tures capture dynamic engagement patterns. We detail feature generation, preprocessing, and the variables associated with each information set in the following subsections.

B.1 Contextual features X^C

To capture the cyclical nature of time, the hour and minute of the day are encoded using sine and cosine transformations. For hour h and minute m , these features are defined as:

$$\text{Hour_SIN} = \sin\left(2\pi\frac{h}{24}\right), \quad \text{Hour_COS} = \cos\left(2\pi\frac{h}{24}\right),$$

$$\text{Minute_SIN} = \sin\left(2\pi\frac{m}{60}\right), \quad \text{Minute_COS} = \cos\left(2\pi\frac{m}{60}\right).$$

These transformations ensure that 23 : 00 and 00 : 00 are numerically close, accurately reflecting the cyclical nature of time. For each categorical variable—media package (app) name (M), device model (D), brand ID (B), operator ID (O), ISP ID (I), and advertisement ID (A)—embeddings are generated to convert them into numerical representations. While M, D, B, O , and I focus on the top 50 categories, A is restricted to 5 predefined categories:

$$f_C(c) = \begin{cases} k, & \text{if } c \in \text{Top 50 categories, } k \in \{1, \dots, 50\}, \\ 0, & \text{otherwise, for } C \in \{M, D, B, O, I\}, \\ k, & \text{if } c \in \bar{5} \text{ Categories, } k \in \{1, \dots, 5\}, \\ 0, & \text{otherwise, for } C = A. \end{cases}$$

Here, C represents each feature set, and k corresponds to the embedded category index. This encoding reduces complexity by focusing on the most common categories while ensuring advertisement ID (A) maintains its specific 5-category structure.

B.2 Geographical features X^{Geo}

Geographical information captures the user’s spatial context at varying levels of granularity. In our setting, these features include precise coordinates, *longitude* (λ) and *latitude* (ϕ), as well as administrative divisions such as *province* (P) and *city* (C). The coordinates (λ, ϕ) provide fine-grained locational information that can be used to compute spatial proximity between users or match impressions to regional market conditions. Province and city identifiers serve as coarser geographical categories that can be linked to aggregated demographic, socioeconomic, or cultural attributes. Because province and city are discrete categorical variables with many unique values, we represent them using learnable embeddings. This transformation assigns each category to a dense, low-dimensional vector in a continuous space, enabling the model to capture similarity patterns between locations without losing category-specific information. Formally:

$$f_P(p) = \mathbf{e}_p \in \mathbb{R}^{d_p}, \quad f_C(c) = \mathbf{e}_c \in \mathbb{R}^{d_c},$$

where \mathbf{e}_p and \mathbf{e}_c are embedding vectors of dimensions d_p and d_c , respectively, learned jointly with the predictive model.

Integrating X^G with contextual features X^C can enhance targeting performance in multiple ways. First, location variables may interact with device and network attributes, for example, certain smartphone brands or ISPs may dominate in specific regions, allowing the model to capture joint patterns that are not evident from either source alone. Second, combining spatial context with temporal features can improve the detection of region-specific engagement cycles, such as peak usage hours in different cities. Finally, location information, when paired with contextual metadata such as operator ID or ISP ID, can serve as a proxy for unobserved market segmentation factors, thereby enriching the representation of the user’s latent type ω without substantially increasing privacy risks. Models such as gradient-boosted decision trees (e.g., XGBoost) are particularly well-suited for this integration, as they can flexibly capture nonlinearities and high-order interactions between heterogeneous feature types, including continuous coordinates, and categorical embeddings, without requiring strong parametric assumptions.

B.3 Behavioral features X^{Behave}

The behavioral features capture user interactions and engagement dynamics at different levels. These features are categorized into three main types: *user-level features*, which track overall user exposure, clicks, and recency of interactions; *ad-level features*, which measure ad performance, frequency, and user engagement with specific ads; and *app-level features*, which focus on app usage, preferences, and their influence on user clicks. Together, these behavioral features provide a comprehensive view of user activity, enabling the model to identify patterns and relationships that drive engagement and response.

B.3.1 User Behavioral Features

Exposure Count (EC): Tracks the cumulative number of ad exposures for user i up to but not including the current impression n :

$$E_{i,n} = n - 1.$$

This feature provides a measure of user exposure history, which helps evaluate the frequency of ad interactions over time.

Click History (CH): Represents the cumulative number of clicks by user i up to but not including the current impression n :

$$C_{i,n} = \sum_{k=1}^{n-1} \text{CLICK}_{i,k}.$$

By excluding the current click, this feature highlights past user engagement and serves as a strong indicator of historical interaction behavior.

Session Click-Through Rate (SCTR): Measures the user’s CTR within the session, calculated as the ratio of past clicks to past exposures for user i at impression n :

$$S_{i,n} = \frac{C_{i,n}}{E_{i,n}}.$$

This feature captures the effectiveness of ads in driving clicks during a session and reflects user engagement intensity.

Time Since Last Exposure (TSE): Represents the time gap between the current and previous ad exposures for user i :

$$TSE_{i,n} = \text{TIME}_{i,n} - \text{TIME}_{i,n-1}.$$

This feature helps identify the recency of ad interactions, which can influence user responsiveness to repeated exposures.

Time Since Last Click (TCE): Tracks the time elapsed since the last click by user i , up to the current impression n :

$$TCE_{i,n} = \text{TIME}_{i,n} - \text{LAST_CLICK_TIME}_i.$$

By measuring click recency, this feature provides insight into how recent engagement influences future click behavior.

B.3.2 Ad-Level Behavioral Features

Ad Frequency (F): Represents the cumulative number of times user i has been exposed to ad j up to the current impression n :

$$F_{i,j,n} = n.$$

This feature helps track ad repetition for a specific user, which is critical for understanding the impact of ad frequency on engagement.

Ad Click-Through Rate (Ad CTR): Measures the click-through rate for ad j by user i , based on clicks up to but not including the current impression n :

$$\text{CTR}_{i,j,n} = \frac{\sum_{k=1}^{n-1} \text{CLICK}_{i,j,k}}{n-1}.$$

By excluding the current click, this feature captures past user engagement with a specific ad, offering insights into its historical performance.

Overall Ad Click-Through Rate (Overall Ad CTR): Represents the overall click-through rate for ad j across all users, based on clicks up to but not including the current impression n :

$$\text{CTR}_j^{\text{Overall}} = \frac{\sum_i \sum_{k=1}^{n_{i,j}-1} \text{CLICK}_{i,j,k}}{\sum_i (n_{i,j} - 1)}.$$

This feature benchmarks the performance of an ad across the entire user base, providing a global measure of its effectiveness.

B.3.3 App-Level Behavioral Features

Usage Share of App (U): Represents the proportion of exposures for app a relative to all exposures for user i up to but not including the current impression n :

$$U_{i,a,n} = \frac{\text{Exposures}_{i,a,n}}{\text{Total_Exposures}_{i,n}}.$$

This feature highlights user preferences for specific apps by tracking their exposure distribution across all apps.

Effect of App (E): Measures the contribution of app a to cumulative user clicks, excluding the current click at impression n :

$$E_{i,a,n} = \frac{\text{Clicks}_{i,a,n}}{\text{Total_Clicks}_{i,n}}.$$

By isolating historical click behavior, this feature evaluates the influence of a specific app on user engagement.

Preference for App (P): Represents the user's preference for app a within the context of ad j , based on exposures up to but not including the current impression n :

$$P_{i,a,j,n} = \frac{\text{Exposures}_{i,a,j,n}}{\text{Total_Exposures}_{i,j,n}}.$$

This feature captures how user interactions with apps vary within specific ad contexts, providing insight into app-specific ad performance.

Influence of App (I): Quantifies the impact of app a on user clicks within the context of ad j , excluding the current click at impression n :

$$I_{i,a,j,n} = \frac{\text{Clicks}_{i,a,j,n}}{\text{Total_Clicks}_{i,j,n}}.$$

This feature measures the effectiveness of an app in driving clicks for a specific ad, offering app-level insights for targeted strategies.

Overall App Usage (U^{Overall}): Represents the share of exposures for app a across all users, based on cumulative exposures up to but not including the current impression n :

$$U_{a,n}^{\text{Overall}} = \frac{\text{Exposures}_{a,n}}{\text{Total_Exposures}_n}.$$

This feature tracks the overall popularity of an app by monitoring its relative exposure share across the dataset.

Overall App Effect (E^{Overall}): Measures the contribution of app a to overall clicks across all users, excluding the current click at impression n :

$$E_{a,n}^{\text{Overall}} = \frac{\text{Clicks}_{a,n}}{\text{Total_Clicks}_n}.$$

This feature evaluates the global impact of an app on user engagement, providing a benchmark for app performance.

C IPS and Propensity Score Evaluations

C.1 Proof of Proposition 1 (deterministic π^{fx})

Define the deterministic IPS estimator

$$\hat{V}(\pi^{fx}) = \frac{1}{N} \sum_{i=1}^N \frac{\mathbf{1}\{A_i = \pi^{fx}(X_i)\}}{\pi^{\mathcal{D}}(A_i | X_i)} v(A_i) Y_i,$$

and treat π^{fx} as fixed w.r.t. the evaluation sample (by using a holdout or cross-fitting).

- *Step 1:* By i.i.d. sampling of (X_i, A_i, Y_i) from the logging DGP ($X \sim p_X, A | X \sim \pi^{\mathcal{D}}(\cdot | X), Y | X, A \sim \Pr(\cdot | X, A)$) and boundedness (Overlap Assumption 1 ensures the denominator is nonzero whenever $\mathbf{1}\{A = \pi^f(X)\} = 1; Y \in \{0, 1\}; v(\cdot)$ finite), the strong law of large numbers (SLLN) yields

$$\hat{V}(\pi^{fx}) \xrightarrow{a.s.} \mathbb{E} \left[\frac{\mathbf{1}\{A = \pi^{fx}(X)\}}{\pi^{\mathcal{D}}(A | X)} v(A) Y \right].$$

- *Step 2:* First law of iterated expectations condition on X yields:

$$\mathbb{E} \left[\frac{\mathbf{1}\{A = \pi^f(X)\}}{\pi^{\mathcal{D}}(A | X)} v(A) Y \right] = \mathbb{E}_X \left[\mathbb{E} \left[\frac{\mathbf{1}\{A = \pi^f(X)\}}{\pi^{\mathcal{D}}(A | X)} v(A) Y \mid X \right] \right].$$

- *Step 3:* Second law of iterated expectations condition on $A | X$ and expand the sum yields:

$$\mathbb{E} \left[\frac{\mathbf{1}\{A = \pi^{fx}(X)\}}{\pi^{\mathcal{D}}(A | X)} v(A) Y \mid X \right] = \sum_{a \in \mathcal{A}} \pi^{\mathcal{D}}(a | X) \frac{\mathbf{1}\{a = \pi^{fx}(X)\}}{\pi^{\mathcal{D}}(a | X)} v(a) \mathbb{E}[Y | X, a].$$

- *Step 4:* The logging propensity cancels given Overlap Assumption 1:

$$= \sum_{a \in \mathcal{A}} \mathbf{1}\{a = \pi^{fx}(X)\} v(a) \mathbb{E}[Y | X, a] = v(\pi^{fx}(X)) \mathbb{E}[Y | X, A = \pi^{fx}(X)].$$

Therefore,

$$\mathbb{E} \left[\frac{\mathbf{1}\{A = \pi^f(X)\}}{\pi^{\mathcal{D}}(A | X)} v(A) Y \right] = \mathbb{E}_X [v(\pi^f(X)) \mathbb{E}[Y | X, A = \pi^f(X)]] .$$

- *Step 5:* By Unconfoundedness Assumption 2 and the usual consistency ($Y = Y(A)$),

$$\mathbb{E}[Y | X, a] = \mathbb{E}[Y(a) | X].$$

Hence

$$\mathbb{E}_X [v(\pi^{fx}(X)) \mathbb{E}[Y | X, A = \pi^{fx}(X)]] = \mathbb{E}_X [v(\pi^{fx}(X)) \mathbb{E}[Y(\pi^{fx}(X)) | X]].$$

- *Step 6:* By the definition, $u(a, \omega) = v(a) \Pr(Y = 1 | a, \omega)$ and $q_x(\omega) = \Pr(\omega | X)$, so for any x and a ,

$$v(a) \mathbb{E}[Y(a) | X = x] = v(a) \sum_{\omega} \Pr(Y = 1 | a, \omega) q_x(\omega) = \sum_{\omega} u(a, \omega) q_x(\omega).$$

Plugging $a = \pi^{fx}(x)$ gives:

$$\mathbb{E}_X [v(\pi^{fx}(X)) \mathbb{E}[Y(\pi^{fx}(X)) | X]] = \mathbb{E}_X \left[\sum_{\omega} u(\pi^{fx}(X), \omega) q_X(\omega) \right] = V(\pi^{fx}).$$

- *Step 7:* Combining Steps 1–6,

$$\mathbb{E}[\hat{V}(\pi^{fx})] = V(\pi^{fx}), \quad \hat{V}(\pi^{fx}) \xrightarrow{a.s.} V(\pi^{fx}).$$

- *Step 8:* If Policy optimality under X Assumption 3 holds, then

$$V(\pi^{fx}) = \sup_{\pi} V(\pi) =: V_X.$$

Therefore,

$$\mathbb{E}[\hat{V}(\pi^{fx})] = V_X, \quad \hat{V}(\pi^{fx}) \xrightarrow{p} V_X.$$

□

C.2 Propensity Score Estimation and Validation

Our IPS implementation requires estimates of the platform’s logging policy probabilities (*propensity scores*) and verification of covariate balance. Appendix §C.2.1 details the estimation procedure, and Appendix §C.2.2 reports balance diagnostics confirming small post-weighting differences.

C.2.1 Propensity Score Estimation

Support and Eligibility Filtering. In our setting, not all ads are eligible to appear in every impression due to platform-level targeting constraints such as location, time-of-day, device type, or campaign-specific restrictions. To encode these constraints, we construct an *eligibility matrix* following [Rafeian and Yoganarasimhan \(2021\)](#):

$$E \in \{0, 1\}^{N \times A},$$

where $e_{i,a} = 1$ if and only if ad a is eligible to be shown in impression i . The matrix E is obtained by intersecting eligibility indicators across multiple categorical dimensions, province, hour-of-day, app category, smartphone brand, network type, and media package name, where each dimension’s support is derived from the set of observed ad–context combinations in the historical data.

The eligibility matrix E serves two roles in our estimation procedure. First, the columns $\{e_{\cdot,a}\}_{a=1}^A$ are included as binary features in the propensity model, allowing the learner to incorporate eligibility information directly into its predictions. Second, after model estimation, we enforce the *support condition* by setting $\hat{\pi}_{i,a}^{\mathcal{D}} = 0$ whenever $e_{i,a} = 0$, followed by renormalization across eligible ads. This guarantees that

$$\hat{\pi}^{\mathcal{D}}(a | x_i) > 0 \quad \text{only if} \quad e_{i,a} = 1,$$

while still permitting the model to flexibly learn variation in serving probabilities within the support.

Outcome and Covariates. The propensity score model predicts, for each impression $i \in \{1, \dots, N\}$, which ad from the candidate set $\mathcal{A} = \{a^{(1)}, a^{(2)}, \dots, a^{(K)}\}$ was displayed. The outcome variable is the multi-class indicator $\mathbf{1}\{A_i = a^{(j)}\}$, where A_i denotes the ad shown in impression i .

The covariate vector x_i is constructed to capture all observed determinants of ad allocation and includes the following components:

1. **Eligibility indicators:** For each ad a , we include the binary feature $e_{i,a} \in \{0, 1\}$ from the eligibility matrix E described above. This ensures that the model incorporates ad-specific availability and targeting rules directly into its predictions.
2. **Targeting variables:** Province, app category, hour-of-day, smartphone brand, network connection type, and mobile service provider, each factorized into integer-encoded indices to capture categorical variation.
3. **Fine-grained location:** Latitude and longitude, enabling the model to detect geographical heterogeneity beyond province-level effects.
4. **Temporal controls:** Time-of-day features measured at the minute level, absorbing short-run fluctuations in bids or quality scores due to diurnal cycles or transient platform shocks.
5. **Device and context characteristics:** Device identifiers, fraud detection codes, and media package identifiers, capturing residual variation in serving probabilities not explained by higher-level targeting variables.

All categorical covariates are transformed into numeric indices using factorization, preserving a consistent mapping across impressions. The eligibility indicators $e_{i,a}$ for all $a \in \mathcal{A}$ are appended to the base covariate set, so the final feature matrix combines both contextual features and ad-specific availability indicators.

Estimation Procedure. The platform’s quasi-proportional allocation rule implies that the probability of displaying ad $a \in \mathcal{A}$ in impression i can be expressed as

$$\pi^{\mathcal{D}}(a | x_i) = \frac{b_a(x_i) q_a(x_i)}{\sum_{a' \in \mathcal{A}_i} b_{a'}(x_i) q_{a'}(x_i)},$$

where $b_a(x_i)$ and $q_a(x_i)$ denote, respectively, the bid and quality score functions conditional on the observed covariates x_i , and $\mathcal{A}_i \subseteq \mathcal{A}$ is the set of ads eligible in context x_i .

In practice, these bid and quality components are not fully observed, and the realized allocation process may deviate from the theoretical form due to unobserved adjustments, platform-level interventions, and complex interactions among covariates. We therefore estimate the realized propensities $\pi^{\mathcal{D}}$ directly from the historical allocation data using a flexible, nonparametric method. Specifically, we adopt a binary one-vs-all approach: for each $a^{(j)} \in \mathcal{A}$, we estimate the conditional probability

$$\pi_j^{\mathcal{D}}(x_i) = \Pr(A_i = a^{(j)} | x_i),$$

By estimating a separate binary classifier for each ad $a^{(j)}$, we capture ad-specific allocation patterns without imposing a single functional form across all ads—an important flexibility given heterogeneous eligibility rules, competition, and campaign objectives. This one-vs-all formulation also mitigates the severe class imbalance of large-scale ad data: rare ads, which would exert little influence in a multiclass model’s joint decision boundary, are given dedicated models in which their appearances are positive events. Moreover, it avoids the restrictive multiclass assumption that the same covariate effects and interactions govern all alternatives equally, an implausible condition in our setting.

For each ad-specific classifier, the training set includes all N impressions, even those in which the ad was ineligible ($e_{i,a^{(j)}} = 0$), allowing the model to learn to assign zero probability in contexts where the ad cannot be served. Each classifier is implemented as an XGBoost with a logistic loss for binary outcomes. To balance predictive accuracy and computational efficiency in a high-dimensional feature space, we adopt hyperparameters for the learning algorithm XGBoost motivated by [Rafieian and Yoganarasimhan \(2021\)](#).

To ensure that estimated propensities are free from in-sample overfitting, we adopt a K -fold cross-fitting procedure with $K = 5$. Let the set of impressions be indexed by $\{1, \dots, N\}$, and let $\mathcal{I}_1, \dots, \mathcal{I}_K$ denote a partition of this set into K disjoint folds of (approximately) equal size. For each fold \mathcal{I}_k , the ad-specific binary classifiers are trained on the complement set $\bigcup_{m \neq k} \mathcal{I}_m$, and predictions $\hat{\pi}_{i,a}^{\mathcal{D}}$ are generated only for impressions $i \in \mathcal{I}_k$. By construction, every propensity score is obtained from a model that did not observe the corresponding impression during training. Concatenating these out-of-fold predictions across all $k = 1, \dots, K$ yields a complete $N \times A$ matrix of out-of-sample estimates. This procedure preserves the predictive structure learned by the models while eliminating the upward bias in inverse propensity estimators that arises from reusing training predictions.

Post Processing. Following propensity estimation, we align the predictions with the platform’s targeting constraints by imposing the *support condition*. Given $E \in \{0, 1\}^{N \times A}$ be the eligibility matrix, where $e_{i,a} = 1$ iff ad a is eligible in impression i . We set

$$\hat{\pi}_{i,a}^{\mathcal{D}} = 0 \quad \text{for all } a \notin \mathcal{A}_i,$$

where $\mathcal{A}_i = \{a \in \mathcal{A} : e_{i,a} = 1\}$ is the eligible set for impression i . The remaining probabilities are renormalized:

$$\hat{\pi}_{i,a}^{\mathcal{D}} \leftarrow \frac{\hat{\pi}_{i,a}^{\mathcal{D}}}{\sum_{a' \in \mathcal{A}_i} \hat{\pi}_{i,a'}^{\mathcal{D}}} \quad \forall a \in \mathcal{A}_i,$$

so that $\sum_{a \in \mathcal{A}_i} \hat{\pi}_{i,a}^{\mathcal{D}} = 1$ and

$$\hat{\pi}^{\mathcal{D}}(a | x_i) > 0 \iff e_{i,a} = 1.$$

This adjustment preserves the relative allocation probabilities learned by the model among eligible ads while ensuring no probability mass is assigned to ineligible alternatives.

Results from Propensity Score Estimation. The estimated propensities capture systematic variation in ad allocation without overfitting to deterministic rules. The macro-averaged AUC is 0.6335 and the micro-averaged AUC is 0.6276, indicating moderate discriminatory power above random guessing but far from perfect classification. This aligns with expectations: if the platform’s allocation contains substantial quasi-

randomization (subject to eligibility and proportional bidding), a very high AUC would be undesirable, as it would imply near-deterministic serving and reduced overlap. Likewise, the macro and micro log-loss values (0.4862 and 0.4388) are low enough to indicate well-calibrated predictions but not so low as to suggest overconfidence.

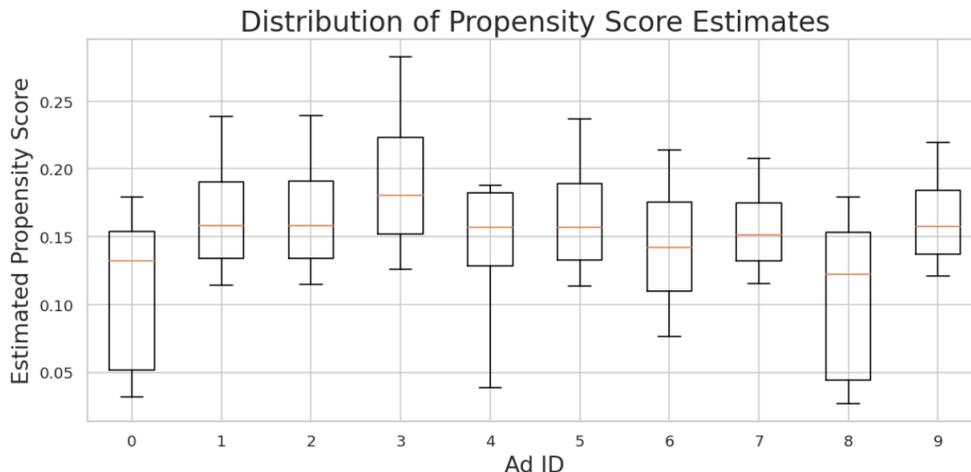


Figure A.1: Distribution of estimated propensity scores across ads. Dispersion and non-degeneracy indicate adequate common support.

Figure A.1 plots the distribution of estimated propensities for each ad. For most ads, the median serving probability lies between 0.12 and 0.18, with interquartile ranges spanning several percentage points. This dispersion indicates that the model assigns a range of probabilities across impressions rather than concentrating mass at extreme values, which is essential for maintaining overlap. Crucially, no ad exhibits degenerate propensities concentrated near zero (indicating near-never assignment) or near one (indicating near-deterministic assignment). Even the rarest ads display a nontrivial upper tail, implying that all ads have a positive probability of being shown in multiple contexts. To further assess whether reweighting achieves covariate balance across treatment arms, we next compute the standardized bias (SB) before and after applying the estimated inverse propensity weights in §C.2.2.

C.2.2 Covariate Balance Diagnostics

After estimating the logging policy propensities $\hat{\pi}_{i,a}^{\mathcal{D}}$ and enforcing support, we evaluate whether inverse-propensity weighting balances the distribution of each covariate across treatment arms using the *standardized bias* (SB) metric. The goal is to check whether, after weighting, the covariate distribution for impressions assigned to any given ad matches that of the eligible population, thereby providing evidence consistent with the unconfoundedness assumption under overlap. Let X denote a generic covariate, X_i its value for impression i , and A_i the ad shown. For ad $a \in \mathcal{A}$, the *unweighted* mean of X in impressions where $A_i = a$ is

$$\bar{X}_a = \frac{\sum_{i=1}^N \mathbf{1}(A_i = a) X_i}{\sum_{i=1}^N \mathbf{1}(A_i = a)},$$

while the mean in the *eligible* population is

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N X_i.$$

Let σ_X be the standard deviation of X in the eligible population. The *unweighted* standardized bias for X when assigned to a is then

$$SB_a(X) = \frac{|\bar{X}_a - \bar{X}|}{\sigma_X},$$

and the worst-case imbalance for X across all ads is

$$SB(X) = \max_{a \in \mathcal{A}} SB_a(X).$$

Intuitively, $SB_a(X)$ measures the relative shift (in standard deviation units) between the covariate mean for impressions shown ad a and the overall eligible population mean; $SB(X)$ takes the largest such shift across ads. Following [McCaffrey et al. \(2013\)](#), we adopt the benchmark $|SB(X)| < 0.20$ as the criterion for acceptable balance. Values above this threshold indicate potentially meaningful imbalance in X between treatment arms. After weighting by the inverse of the estimated propensities, the weighted mean of X for impressions shown ad a is

$$\bar{X}_a^{\hat{\pi}} = \frac{\sum_{i=1}^N \mathbf{1}(A_i = a) \frac{X_i}{\hat{\pi}_{i,a}}}{\sum_{i=1}^N \mathbf{1}(A_i = a) \frac{1}{\hat{\pi}_{i,a}}}.$$

Let $\bar{X}^{\text{elig}(a)} = \frac{1}{|\mathcal{S}_a|} \sum_{i \in \mathcal{S}_a} X_i$ and $\sigma_X^{\text{elig}(a)}$ denote, respectively, the mean and standard deviation of X in the ad-specific eligible set \mathcal{S}_a (i.e., rows with support for ad a). The *post-weighting* standardized bias for X when assigned to a is

$$SB_a^{\hat{\pi}}(X) = \frac{|\bar{X}_a^{\hat{\pi}} - \bar{X}^{\text{elig}(a)}|}{\sigma_X^{\text{elig}(a)}},$$

and the worst-case post-weighting imbalance is

$$SB^{\hat{\pi}}(X) = \max_{a \in \mathcal{A}} SB_a^{\hat{\pi}}(X).$$

If $SB^{\hat{\pi}}(X)$ satisfies the balance threshold for every X , then the inverse-propensity weighting has successfully aligned the covariate distributions across all treatment arms (conditional on the enforced support). Comparing $SB(X)$ (pre-weighting) and $SB^{\hat{\pi}}(X)$ (post-weighting) quantifies the extent to which our weighting procedure mitigates initial selection bias. [Figure A.2](#) reports the worst-case standardized bias $SB(X)$ and its post-weighting counterpart $SB^{\hat{\pi}}(X)$ for each covariate, as defined.

Before weighting (blue), several covariates, most notably HOUR, exhibit large imbalances, with $SB(X) > 0.20$, indicating substantial divergence between the covariate distribution for impressions assigned to certain ads and that of the eligible population. After inverse-propensity weighting (orange), all covariates satisfy the balance criterion, i.e., $SB^{\hat{\pi}}(X) < 0.20$. This demonstrates that reweighting has effectively aligned the covariate distributions across treatment arms, substantially mitigating the selection bias present in the unweighted

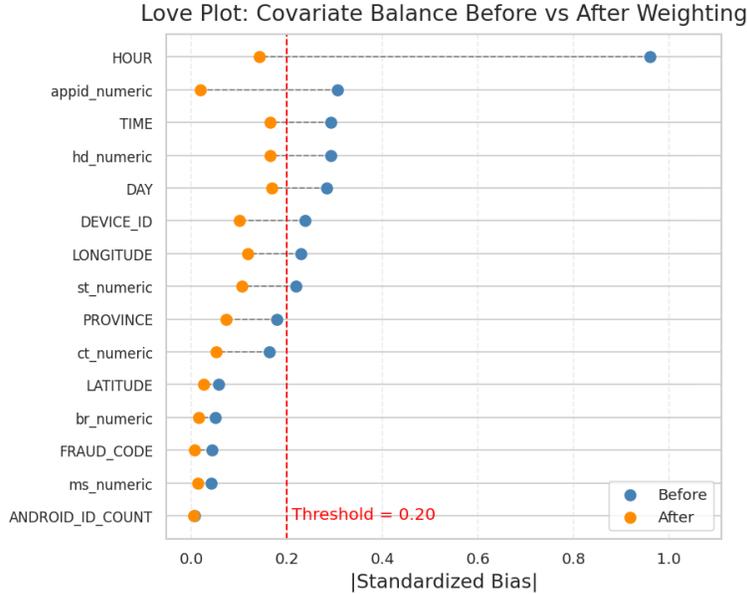


Figure A.2: Standardized bias for each covariate before and after inverse-propensity weighting. The red dashed line marks the 0.20 balance threshold from McCaffrey et al. (2013).

data. The improvement is consistent across both high-variance contextual features (e.g., PROVINCE, LONGITUDE) and stable identifiers (e.g., DEVICE_ID), providing empirical support for the overlap assumption in our setting.

D Additional Empirical Results

D.1 XGBoost vs. LSTM in Predictive Performance

To benchmark our models against approaches used in the literature, we implement an XGBoost model following the specifications in Rafeian and Yoganarasimhan (2021). This allows us to directly compare the predictive performance of our LSTM model with a widely-used tree-based method in ad-targeting research.

For both behavioral (X^B) and combined (X^{GB}) feature sets, we train the XGBoost model using the same training/test splits, hyperparameter tuning protocol, and feature preprocessing pipeline applied to our LSTM models. Hyperparameters are selected via grid search to maximize AUC on the validation set, and class weights are adjusted to account for click sparsity. Table A2 reports performance metrics for the two methods.

Model	Log Loss	AUC	Relative Information Gain
XGBoost (X^B)	0.0552	0.8235	37.64%
XGBoost (X^{GB})	0.0549	0.8333	37.98%
LSTM (X^B)	0.0140	0.8088	84.18%
LSTM (X^{GB})	0.0138	0.8115	84.41%

Table A2: Comparison of model performance for behavioral (X^B) and combined (X^{GB}) features using XGBoost and LSTM model. RIG computed as baseline CTR of test set is $p = 0.017592$.

The results reveal three key patterns. First, consistent with our earlier findings, behavioral features provide the dominant source of predictive power. Second, adding geographical information to behavioral features yields only marginal improvements, confirming that behavioral data capture most of the variation in click propensity. Third, the LSTM model substantially outperforms XGBoost in log loss and RIG for the same feature set, despite almost similar AUC values. This indicates that the LSTM’s ability to model temporal dependencies in behavioral sequences translates into better probabilistic calibration, even if its rank-ordering performance (AUC) is comparable to the tree-based benchmark.

It is important to emphasize that high predictive performance does not necessarily translate into higher decision value in targeting. A model may accurately rank or calibrate predicted click probabilities but still yield limited improvements in actual campaign outcomes under realistic budget and exposure constraints. In the next section, we therefore evaluate these models in terms of their realized *decision value* as measured by the targeting value function.

D.2 XGBoost vs. LSTM in Value of Information Performance

To complement the predictive performance comparison in Appendix §D.1, we benchmark the *decision value* delivered by LSTM and XGBoost models. Decision value is measured using the IPS-based targeting value function, which quantifies the incremental CTR lift relative to the logging-policy baseline under realistic targeting constraints. This analysis addresses the fact that models with similar predictive accuracy can differ in their ability to generate value when deployed in practice.

For both behavioral (X^B) and combined (X^{GB}) feature sets, we use the same training/test partitions and preprocessing as in the predictive exercise, ensuring that differences in decision value are attributable to the learning algorithm rather than data handling. IPS estimates are computed following the same procedure described in §5.3, and standard errors are clustered at the user level. Table A3 reports the estimated policy values, confidence intervals, t -statistics, standard errors, percentage lifts over the baseline CTR of $p = 0.017592$, and effective sample sizes.

Table A3: Estimated Policy Value Using IPS

Policy	IPS Estimate	95% CI	t-stat	SE	Lift (%)	ESS
$\hat{V}_{X^{GB}} \text{ LSTM}$	0.024883***	[0.024366, 0.025401]	27.611	0.000264	41.45%	760316
$\hat{V}_{X^{GB}} \text{ XGBoost}$	0.023531***	[0.023056, 0.024006]	24.521	0.000242	33.76%	803214
$\hat{V}_{X^B} \text{ LSTM}$	0.022918***	[0.022387, 0.023448]	19.675	0.000271	30.28%	555871
$\hat{V}_{X^B} \text{ XGBoost}$	0.022738***	[0.022261, 0.023214]	21.183	0.000243	29.25%	735365

Notes: All estimates are computed using IPS. The baseline CTR under the logging policy is 0.017592. Lift is computed relative to this baseline. SE denotes the standard error of the estimate. Effective sample size (ESS) measures the number of equally weighted observations that would yield equivalent precision. Number of observations: 3,162,376. Cluster-robust standard errors are computed using 141,595 user clusters. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The results yield two main insights. First, the LSTM model consistently achieves higher IPS estimates than XGBoost for the same feature set. This gap is particularly pronounced for the combined feature set, where LSTM delivers a 41.45% lift versus 33.76% for XGBoost. These differences align with the log-loss and RIG advantages observed in the predictive exercise, suggesting that the LSTM’s temporal-sequence

modeling translates into superior decision value in deployment. Notably, while XGBoost achieves comparable, and in the case of AUC, slightly higher rank-ordering performance in the predictive task, these apparent advantages do not translate into higher realized value. This underscores that strong predictive metrics do not guarantee superior targeting outcomes. Second, the divergence between predictive and decision-value rankings highlights the importance of evaluating targeting algorithms using realized performance metrics, not solely predictive scores. While AUC and calibration remain useful for model assessment, the ultimate measure of a targeting model’s reward is the incremental value it generates in live allocation settings.

D.3 Statistical Test of Complementarity vs. Substitutability

Here we present the detailed statistical results underlying Figure 7. The figure shows mean incremental policy lifts across impression depth, while Table A4 reports the corresponding formal test of whether geographical data complements or substitutes behavioral data at different stages of exposure. The test logic follows a difference-in-differences (DoD) design. Specifically, we define

$$\hat{\Delta} = (V^{GB} - V^B) - (V^G - V^{\emptyset}),$$

where V^{GB} is the value of the combined geographical + behavioral policy, V^B is behavior only, V^G is geographical + context, and V^{\emptyset} is context only. Intuitively, $\hat{\Delta}$ compares the incremental contribution of geographical data when behavioral information is present versus when it is absent. A positive $\hat{\Delta}$ indicates that geographical acts as a *complement* to behavior (its marginal value is higher once behavior is available), while a negative $\hat{\Delta}$ implies *substitutability* (its marginal value declines once behavior is available).

Inference is based on cluster-robust standard errors at the user level. This adjustment ensures that repeated impressions from the same individual do not bias the test statistics. For each impression-depth bin, the table reports the point estimate $\hat{\Delta}$, its 95% confidence interval, the t -statistic, a two-sided p -value, and one-sided bootstrap probabilities $p(\Delta > 0)$ and $p(\Delta < 0)$. The final two columns summarize the decision rule (Complement, Substitute, or Inconclusive) and the corresponding significance tier.

The results confirm the visual evidence from Figure 7: during early exposures (up to roughly 20–25 impressions), $\hat{\Delta}$ is positive and statistically significant, indicating that geographical data provide complementary value when behavioral histories are sparse. Beyond 50 impressions, $\hat{\Delta}$ becomes negative and significant, showing that geographical data turns into a substitute once behavioral data are sufficiently rich. Mid-range bins (around 30–40 impressions) produce inconclusive or transitional estimates. Together, the figure and table establish a dynamic pattern: geographical data is initially complementary, but ultimately substitutable as behavioral histories accumulate.

Table A4: Complement vs. Substitute by Impression Depth (Clustered DoD Test)

Last-Impr	$\hat{\Delta}$	95% CI	<i>t</i> -stat	<i>p</i> (two)	<i>p</i> ($\Delta > 0$)	<i>p</i> ($\Delta < 0$)	Decision	Signif.
3	0.0006	[-0.0017, 0.0029]	0.50	0.616	0.308	0.692	Inconclusive	none
6	0.0033	[0.0008, 0.0057]	2.63	0.009	0.004	0.996	Complement	strong
10	0.0023	[-0.0001, 0.0046]	1.86	0.063	0.032	0.968	Complement	mod
16	0.0024	[0.0003, 0.0045]	2.23	0.026	0.013	0.987	Complement	mod
24	0.0033	[0.0014, 0.0052]	3.37	0.001	0.000	1.000	Complement	strong
35	-0.0002	[-0.0021, 0.0017]	-0.20	0.839	0.581	0.419	Inconclusive	none
50	-0.0023	[-0.0039, -0.0006]	-2.69	0.007	0.996	0.004	Substitute	strong
71	-0.0019	[-0.0034, -0.0004]	-2.49	0.013	0.994	0.006	Substitute	strong
102	-0.0010	[-0.0024, 0.0003]	-1.54	0.123	0.938	0.062	Substitute	weak
150	-0.0010	[-0.0024, 0.0005]	-1.31	0.189	0.905	0.095	Substitute	weak

Notes: $\hat{\Delta} = (V^{GB} - V^B) - (V^G - V^{\emptyset})$ reports the difference-in-differences estimate of complementarity versus substitutability. *t*-statistics and confidence intervals are based on cluster-robust standard errors at the user level. *p*($\Delta > 0$) and *p*($\Delta < 0$) are one-sided bootstrap probabilities.