

# Advertiser Learning in Direct Advertising Markets\*

Carl F. Mela<sup>†</sup>

Jason M.T. Roos<sup>‡</sup>

Tulio Sousa<sup>§</sup>

May 2, 2024

## Abstract

Direct buy advertisers procure advertising inventory at fixed rates from publishers and ad networks. Such advertisers face the complex task of choosing ads amongst myriad new publisher sites. We offer evidence that advertisers do not excel at making these choices. Instead, they try many sites before settling on a favored set, consistent with advertiser learning. We subsequently model advertiser demand for publisher inventory wherein advertisers learn about advertising efficacy across publishers' sites. Results suggest that advertisers spend considerable resources advertising on sites they eventually abandon—in part because their prior beliefs about advertising efficacy on those sites are too optimistic. The median advertiser's expected CTR at a new site is 0.23%, five times higher than the true median CTR of 0.045%.

We consider how an ad network's pooling of advertiser information remediates this problem. As ads with similar visual elements garner similar CTRs, the network's pooling of information enables advertisers to better predict ad performance at new sites. Counterfactual analyses indicate that gains from pooling advertiser information are substantial: over six months, we estimate a median advertiser welfare gain of \$2,756 (a 15.5% increase) and a median publisher revenue gain of \$9,618 (a 63.9% increase).

**Keywords:** Display advertising, Learning models, Bayesian estimation

---

\*The authors would like to thank seminar participants at the 9th Invitational Choice Symposium (Noordwijk), the 2023 Conference on Artificial Intelligence, Machine Learning, and Business Analytics, the 2024 Workshop on Platform Analytics, Bocconi University, Cologne University, Duke University, Erasmus University, Georgetown University, INSEAD, Rice University, Santa Clara University, Stanford University, Tilburg University, the University of Groningen, and the University of North Carolina, Charlotte as well as Santiago Balseiro, Siddarth Prusty, Srinivas Tunuguntla, Nils Wernerfelt, Boya Xu, and Levin Zhu for their comments and insights.

<sup>†</sup>T. Austin Finch Foundation Professor of Business Administration at the Fuqua School of Business, Duke University (email: mela@duke.edu).

<sup>‡</sup>Associate Professor of Marketing at the Rotterdam School of Management, Erasmus University (email: jroos@rsm.nl).

<sup>§</sup>Ph.D Candidate, Duke University Department of Economics (email: tuliosalvio.sousa@duke.edu).

# 1 Introduction

Advertisers seek to place ads with publishers whose readers are potential customers of the advertised goods. In the context of direct buy display advertising, a myriad of sites from which to choose makes the publisher selection problem daunting. Because sites differ in their readership and the editorial context in which display ads are served, advertisers lacking prior experience with a given site are typically uncertain about the value of placing ads there (Perlich et al. 2012, Tunuguntla and Hoban 2020). This uncertainty creates inefficiencies in ad spending that are amplified by the costly nature of direct buy ad purchases. In direct markets, advertisers often pay a fixed price (as posted on a rate card) to present an ad to all readers who visit the publisher’s site during a specified period lasting days or even weeks. Initial purchases are characterized by uncertainty in ad performance outcomes, and a commitment to delivering a potentially large number of impressions. A key goal of this paper is to explore the welfare implications of advertiser learning in the face of uncertain ad response in direct buy advertising markets. Although our emphasis is on direct buy display markets, similar contexts include video, television, retail media, print, radio, and other channels. In each of these settings, advertisers often have little information about the efficacy of their advertising in a particular channel prior to committing to a large ad buy. Consequently, our objective is to measure the degree of advertiser learning, and to ascertain the attendant welfare implications (for advertisers and publishers) of potentially miscalibrated initial beliefs about the value or efficacy of advertising across different sites.

In this regard, this research makes several advances. First, much of the canonical economic theory behind advertiser behavior assumes that advertisers know a priori the value of ad exposures (Varian 2007, Athey and Ellison 2011, Choi et al. 2020, Balseiro et al. 2015). In contrast, we do not presume that valuations are known before buying ads on a site, but rather that they must be learned. If advertiser prior beliefs are too optimistic (pessimistic), they will spend too much (too little) on ads. With experience, however, their media buys should become more efficient. A growing literature pertaining to automated advertiser learning approaches has recently appeared (Cai et al. 2017, Choi and Sayedi 2019, Ren et al. 2018, Scott 2010, Schwartz et al. 2017, Tunuguntla and Hoban 2020, Waisman et al. 2019, Balseiro and Gur 2019). Although these approaches imply advertisers should try to learn, these studies do not seek to i) demonstrate whether advertisers do indeed learn, ii) measure advertisers’ initial beliefs about advertising efficacy, iii) assess the welfare implications of those initial beliefs, nor iv) explore the potential for an intermediary, such as an ad network, to pool information across advertisers. Recently, Tadelis et al. (2023) address point i) cross-sectionally in the context of an exchange network. They find that more experienced advertisers enjoy greater lift from advertising, a result

that is consistent with advertiser learning. In contrast to their approach, our research explicitly models the learning process via a structural learning model, explores the welfare implications of learning, and suggests an approach to improve the efficiency of the media buy. In general, there is a paucity of empirical evidence in marketing and economics regarding advertisers learning by doing. As noted by Tadelis et al. (2023), the learning by doing literature has focused on the production of goods or services (e.g., Benkard 2000, Hendel and Spiegel 2014, Levitt et al. 2013). Few, if any, papers explicitly measure the degree to which advertisers learn and the attendant costs of learning.

Our second advance is to consider the setting of a direct buy ad network, as opposed to exchanges (Allouah and Besbes 2017, Balseiro and Candogan 2017, Wu 2015). According to eMarketer, direct buying represents 75% of the annual \$130B display ad market, with exchanges representing the balance. In spite of this, most prior research in marketing has focused on exchanges (Choi et al. 2020).<sup>1</sup> In exchange markets, individual impressions are purchased in real time and priced via an auction. In direct settings, impressions are bundled and sold in bulk, often with a guarantee on the number of days the ad will run, but not the precise number of impressions that will eventually be served.<sup>2</sup> Furthermore, direct inventory is often listed and sold at a fixed price.<sup>3</sup> Figure 1 shows an example of the 2024 rate card from the Star News Group, a small, local newspaper publisher headquartered in Manasquan, New Jersey. An advertiser can purchase a headline banner on the site for \$100 per month, and the Star News Group estimates that this purchase would yield roughly 30,000 impressions (larger papers, like The Guardian, charge upwards of \$100,000 for a one-day take over of its homepage). A new advertiser at the Star News Group must choose between not advertising with them, or spending a relatively large sum of money (compared to a single impression sold via exchange) in return for a relatively large number of impressions of an uncertain value. In these environments, advertisers face an overwhelming set of publishers of news, blogs, and other sites from which to choose. After choosing to run an ad with a particular publisher, advertisers effectively receive feedback on ad response in batches, that is, only after large numbers of impressions are served. In our data, the median ad buy is for 8 days, and the median number of impressions from a single ad buy is about 820,000. The soonest this information could impact subsequent choices is after the initial ad buy is over.

The differences between fixed price contexts and ad exchanges are consequential for how advertiser

---

<sup>1</sup><https://tinuiti.com/blog/performance-display/ana-programmatic-buying-guide-pov/> (Archive.org), <https://www.insiderintelligence.com/content/programmatic-ad-spending-forecast-h1-2024> (Archive.org)

<sup>2</sup>See <https://www.socialchimp.com/blog/direct-vs-programmatic-breakdown-media-buying/> (Archive.org) and <https://newormedia.com/blog/ad-exchange-vs-ad-networks/> (Archive.org).

<sup>3</sup>See [https://media-index.kochava.com/ad\\_partners?country=&channel%5B%5D=desktop-display&channel%5B%5D=mobile&category%5B%5D=ad-network&pricing\\_models%5B%5D=flat-rate](https://media-index.kochava.com/ad_partners?country=&channel%5B%5D=desktop-display&channel%5B%5D=mobile&category%5B%5D=ad-network&pricing_models%5B%5D=flat-rate) (Archive.org) for an example list of flat rate exchanges.



**Figure 1:** 2024 Star News Digital Advertising Rate Card. Source: Star News Group, [https://cdn.starnewsgroup.com/pdf/advertising/2024\\_SNG\\_Web\\_Advertising\\_Rates.pdf](https://cdn.starnewsgroup.com/pdf/advertising/2024_SNG_Web_Advertising_Rates.pdf) (Archive.org).

learning is modeled. In exchange markets, feedback (in the form of KPIs such as clicks) arrives incrementally by impression. In direct markets, a considerable amount of uncertainty can be resolved between subsequent purchases, because each purchase leads to many impressions, and the information about ad efficacy from so many impressions overwhelms prior beliefs. Under these circumstances, obtaining better priors is more pressing than test and learn, and there is little opportunity (if any) to explore and exploit using exchange algorithms, such as those discussed in Tunuguntla and Hoban (2020).<sup>4</sup>

Hence, our third contribution is to propose a scalable approach to enhance advertiser priors in order to improve their outcomes in the context of direct ad networks. We propose a procedure whereby the ad network pools its advertisers' information to improve advertisers' prior beliefs about ad efficacy at new sites.<sup>5</sup>

<sup>4</sup>Because direct channels, in contrast to exchange markets, do not involve individual targeting of ads, there are no "test and learn" algorithms used for targeting impressions in direct buy markets. Therefore, advertiser learning behavior is not algorithmically driven by direct exchanges.

<sup>5</sup>An interesting example of information pooling in practice is Meta's training of models to forecast advertiser's campaign outcomes such as reach and clicks. These tools pool information across advertisers to generate forecasts to better inform advertiser purchases. See <https://ai.meta.com/blog/ai-ads-performance-efficiency-meta-lattice/> (Archive.org).

This raises the question of what information to use, as well as how it should be shared with advertisers. Conceptually, similar advertisers should evidence similar ad performance (e.g., click-through rate, CTR) at the same site. However, a further consideration emerges about how to measure similarity between advertisers (especially those who are new to the ad network). One approach to measuring similarity between advertisers is to consider advertising copy or design, with the idea that any two advertisers producing similar ad content should also generate similar ad responses (Yao et al. 2023). Because our advertisements are images, it is possible to measure the image similarity of the display advertisements. Specifically, we create a set of concept tags for each ad in our data, and then compute a similarity score for each pair of advertisers based on their associated tags. Then, to improve advertisers’ prior beliefs, we impute advertiser CTR at each new site using a weighted average of *other* advertisers’ CTRs at that site (with the weights determined by the advertiser similarity scores). The approach thus builds on the cold start literature by using machine learning image recognition techniques in the context of an ad exchange (Gope and Jain 2017, Lika et al. 2014, Lam et al. 2008, Schein et al. 2002, Xu et al. 2022). A benefit of the approach is that it requires only first-party data to implement; as such, it would not be impacted by Google’s recent decision to phase out third-party cookies.

Owing to our approach to sharing information, a fourth contribution pertains to the literature on the role of network intermediaries in advertising (Choi et al. 2020). Typically, ad networks are construed to reduce search frictions for advertisers, making it easier for them to find publisher inventory. Yet ad networks can also improve match by helping advertisers learn which sites are more effective. An open question is whether publishers and the ad network have an incentive to inform advertisers of prior match. On the one hand, if advertisers are too optimistic in their initial beliefs about the efficacy of sites they selected for advertising, better information would attenuate advertiser spending and the network would have little incentive to inform advertisers that they were overly optimistic. On the other hand, to the extent substitute sites can be found with better match, advertisers might increase spend. Little evidence exists to suggest which of the two effects dominates, and whether networks have an incentive to share information on match values with advertisers. Accordingly, we consider the role of information sharing on advertiser, publisher, and ad network welfare.

We collect data from a direct sales ad network that consolidates direct sales display ad inventory across multiple publishers (in this case, the sites are blogs). As is common in the direct sales display advertising channel, advertisers procure the publishers’ ad inventory in advance at a fixed rate. The data from this ad network are ideal for our empirical context, because ad sales and ad prices are observed from the network’s inception and over a long duration (3 years). As a result, the behavior of all advertisers is observed from the platform’s infancy, providing an ideal context in which to observe advertisers learning. Our empirical

strategy relies upon longitudinal changes in the advertisers' propensity to choose particular sites, and this type of variation is common in the data. The structural learning model we develop presumes that advertisers choose the expected number of ad impressions across sites to maximize their profits, conditioned on their (possibly incorrect) beliefs about the efficacy of advertising on those sites. After an ad has run at a site, the advertiser observes its CTR for that site. These CTRs provide a noisy signal about the site's match with the advertiser's ads. Sites with higher (lower) CTR's are more (less) likely to be indicative of a good match, and thus sites with high CTRs are more likely to be used again by the advertiser. When advertisers select a site upon which to advertise, fail to obtain clicks, and then cease to advertise, we reason they were too optimistic in their initial beliefs about the efficacy of advertising on that site. We then use the demand-side model estimates to gain insights about advertiser conduct, and to simulate demand under counterfactual scenarios that manipulate what advertisers know.

The data evidence patterns that are broadly consistent with learning, such as advertisers initially trying many sites before settling on a smaller number, and preferring to place additional ads at sites that previously generated relatively higher CTRs. Results from a structural model of advertiser demand suggest that advertisers are overly optimistic in their initial beliefs about the efficacy of advertising. The median advertiser's choices are consistent with an expected click through rate of 0.23% when, in practice, the median CTR is .045%. The finding suggests that advertisers often choose the wrong publisher sites initially and overspend on them, while also potentially overlooking sites that would have been a better match. With pooled information provided by the platform, the median advertiser in the estimation sample increases expected total spend by about \$671 (52.8%) over six months, and generates an expected \$2,756 (15.5%) in incremental value: all advertisers see welfare gains. These welfare gains largely arise because pooling allows advertisers to sort themselves into better matching publisher sites. Owing to this better match, the median publisher also obtains an expected \$9,618 (63.9%) increase in revenue over six months (across the top 20 publishers in our data). Hence, publishers and the ad network have an incentive to better inform advertisers about match. Projecting these median effects across all advertisers suggests an overall six-month welfare gain of approximately \$5,400,000. Presumably, similar gains could accrue to other direct ad networks, as well in related contexts such as retail media and online video.

In what follows, we first overview the data and provide descriptive evidence of advertiser learning. We then outline an advertiser learning model, report the estimation results, detail the counterfactual analyses related to information sharing, and conclude with a summary of our findings and ideas for further research.

## 2 Data


### 2.1 Data Overview

To assess the effect of advertiser learning, it is necessary to observe the advertising decisions of firms over time. To this end, we have collected data provided by a fixed-price, direct-sales Internet ad aggregator. The data span 3 years, starting with the aggregator’s inception in late July 2006 and ending in early December 2009 and covers 8,000 advertisers and 3,200 publishers. These data record the transactions between advertisers and publishers. Each transaction corresponds with an individual advertiser’s ad buy (commonly called a *subscription*), and specifies the number of consecutive days (typically a week) the ad is to be shown to all visitors to the publisher’s website, and the price paid by the advertiser for that purchase. Notably, the price paid depends on the length of the subscription, and not on the eventual number of impressions served. Further, for each day the ad is active, the data record the total number of *impressions*—the number of times the ad was served to a site visitor—and the number of *clicks*—the number of times a site visitor clicked on the ad. Over the duration of the data, we are aware of no algorithmic changes that might influence the nature of advertiser learning or the allocation of advertisers to sites.

The estimation sample focuses on a subset of these data. Using k-means clustering with a Jaccard similarity metric computed from the overlap of firms advertising on sites, we identify a subset of 165 politically liberal blogs and news sites serving a similar set of advertisers that do not generally advertise on other sites.<sup>6</sup> This subset of publishers comprises 15.6% of subscriptions (ad buys) in the data, with 1801 (22.7%) of the advertisers in the data placing an ad at one or more of these sites. From this coherent grouping of publishers, the top 20 are selected in terms of total ad revenue using transactions conducted in the first half of 2007 (January 1–June 30). We then choose a random subset of 100 advertisers who placed an ad at one or more of the top 20 sites, and aggregate choices to the weekly level to comport with advertisers’ typical purchase frequency (rarely does interpurchase time fall short of a week). The unit of analysis for our empirical model is thus at the advertiser-site-week level. Over the 27 weeks in the first half of 2007, the estimation sample comprises 547 ad subscriptions. We assume that advertisers include the top 20 sites in their choice sets, unless a previously purchased subscription is already running in a given week. Some advertisers joined the ad network during the estimation window; for these we only consider choices starting with the week they first appear in the data. The estimation sample thus comprises 36,911 choices leading to

---

<sup>6</sup>The similarity matrix is computed using i) sites with at least 9 advertisements (5th percentile) and ii) advertisers with at least 20 advertisements placed at 10 or more sites. Smaller advertisers and publishers are largely inconsequential economically, and make the task of finding a unique sets of adjacent sites impracticable.



**Figure 2:** Distribution Of Realized Cost Per Click, Grouped By Action After First Ad Run Is Completed. Based on 1,784 advertisers who ran ads with any publisher in the focal cluster of 165 sites, showing only ads placed by these advertisers at these sites. Cost per click is cumulative through last observed ad placement. Medians are indicated.

547 purchases.

The daily number of impressions for each ad (in the period prior to July 2007, including the second half of 2006) are used to impute the daily average number of ad impressions at each site. This serves as an approximation to sites' expected daily traffic. The site in the estimation sample with the greatest number of daily visitors had a peak audience of 910K daily visits in the first half of 2007, and average daily traffic of 420K; the site with the least traffic peaked at 43K, with average daily traffic of 20K. Prices vary accordingly, with the most expensive placement garnering over \$755 per day (\$2265 for 3 days) and the least just \$5.50 per day (\$500 for 3 months). A typical subscription in the estimation sample lasts one week, costs \$800, and yields 821K impressions; the median click-through rate among ads in the estimation sample is 0.045%.


## 2.2 Descriptive Analyses

This subsection first presents preliminary data analyses revealing patterns in the data that are consistent with advertiser learning. Such an analysis affords evidence of the types of behaviors we seek to model. Subsequently this section characterizes how advertising outcomes (CTRs) vary across sites and advertisers. To the extent CTR variation across sites (or site-advertisers) is large, relative to CTR variation across advertisers, there is value in sharing information about CTRs among advertisers.

### 2.2.1 Evidence of Advertiser Learning

Figure 2 shows the distribution of *realized* cost per click (CPC) for all ads placed at the subset of 165 liberal blogs and news sites by 1,784 advertisers. Because ads are placed on a fixed-price basis, CPC is an outcome that is only observed *after* the ad has run. Hence, after comparing the number of clicks an ad received with






**Figure 3:** Average Daily Number of Sites Running Ads. The first 26 weeks of ads served by the focal cluster of 165 sites are shown for 939 advertisers who placed ads with two or more of these publishers. The line (and smoothed average) depicts the average daily number of publishers serving ads for advertisers actively running ads. The timing of ads is normalized relative to when the advertiser joined the ad network, truncated to show weeks 1–26, and logarithmically scaled to emphasize changes during the initial few weeks of advertising. The number of advertisers with ads running, depicted by the bars at the top of the chart, decreases from 939 in week 1 to 64 in week 26.

the amount of money spent, the advertiser can calculate the effective CPC for that ad. If the advertiser never used that particular site again, then the realized cost per click is depicted in the distribution on the right of Figure 2 (labeled “Did not advertise again with these publishers”). If the advertiser *did* place additional ads at that site, then the realized CPC is included in the distribution to the left. The difference between these two distributions is consistent with advertisers abandoning sites where the returns to advertising (as reflected in the ex-post number of clicks received) are insufficiently high to justify the cost.<sup>7</sup>

Figure 3 provides another view of the data that is consistent with advertiser learning. For each advertiser that placed at least one ad with one of the 165 liberal blogs and news sites, we calculate the number of sites on which the advertiser actively ran ads each day. To account for different advertiser cohorts commencing their advertising on different dates, start dates are normalized to the first ad placement for these time-series. For each day, we calculate the average number of sites actively used by each advertiser, conditional on that advertiser running at least one ad (because subscriptions are typically 7 days or more, observations prior to 7 days are truncated). At the end of the first week, advertisers have ads running, on average, at about 6 sites. But over the next month, the number of sites used drops by almost half. Moreover, many advertisers cease advertising entirely during this period. At the end of the first week, 939 advertisers are still running ads. By week 26, the number of advertisers still running ads at the subset of 165 sites drops to 64. This pattern is predicted by learning, as advertisers start with a larger set of sites, but eventually stop placing ads at the least

<sup>7</sup>A similar analysis is conducted at the advertiser level to control for unobserved heterogeneity. Specifically, we compute, for each advertiser the mean difference in realized cost per click between sites chosen again and sites not chosen again (for advertisers where both types of choices are observed). The analysis yields a median difference across advertisers of \$2.20 (versus an implied difference of \$2.18 in Figure 2).



**Figure 4:** Persistence in Set of Publishers Chosen. The first 26 weeks of ads purchased from the focal cluster of 165 sites are shown for 542 advertisers who placed ads with two or more publishers of these publishers in two or more two-week periods. The timing of ad buys is normalized relative to when the advertiser joined the ad network and truncated to the first 26 weeks. Similarity is calculated as the Jaccard coefficient between i) the set of publishers used in a two-week period, and ii) the set of publishers from the most recent, previous two-week period with an ad buy. Points (vertical lines) indicates the average (bootstrap 95% CI's) similarity in the set of sites used among advertisers purchasing ads. The number of advertisers purchasing ads, depicted by the bars at the top of the chart, decreases from 407 in weeks 1–2 to 45 in weeks 25–26.

effective sites.<sup>8</sup>

A third piece of evidence consistent with advertiser learning is depicted in Figure 4, which portrays changes over time in the set of sites at which advertisers place ads. Similarity is calculated as the Jaccard coefficient between two sets of publishers, over rolling, bi-weekly periods. The first set contains all publishers used by an advertiser in the focal bi-week. The second set contains all publishers used by an advertiser in the previous bi-week (exclusive of periods with no purchases). This coefficient, defined as the cardinality of the intersection of the two sets divided by the union of the two sets, ranges from 0 to 1, with a higher number meaning that the set of sites in the current and preceding buy show greater overlap. Over time, the set of sites an advertiser places ads at is more likely to resemble the previous set, as the Jaccard coefficient triples from 0.1 to 0.3. Note that it does not necessarily follow that the Jaccard index must decrease as set sizes decrease, meaning that Figure 3 and Figure 4 are not isomorphic.

Thus, not only is the number of sites chosen decreasing, but the set of publishers within those decreasing sets exhibits greater consistency over time, and advertisers are more likely to advertise again with sites that yield better CTR outcomes. Collectively, these behaviors would be consistent with advertisers finding a set of sites that works well, and sticking with that set.

<sup>8</sup>Shrinking budgets are another potential explanation for a reduction in the number of sites used by advertisers over time, but unlike learning, shrinking budgets do not explain the broader patterns in the data. First, shrinking budgets alone cannot explain why advertisers are more (less) likely to repeat buy at publisher sites with lower (higher) costs per click. Second, lower ad budgets cannot explain the tendency of advertisers in the data to explore new sites over time, as opposed to simply cutting sites (by week 26 since joining the ad network, nearly 1/3 of all ad purchases are on new sites). Third, diminished budgets cannot explain why many advertisers spend more over time, as 35% of the 853 advertisers buying ads on two or more occasions spent more in the second half of their tenure using the network.

Term	DF	Sum of Squares	Variation Explained (%)
Advertiser	1783	.0248	50.8
Publisher	164	.0044	9.1
Residuals	8205	.0196	40.1

**Table 1:** Variance Decomposition of Advertiser-Site Long-Run Click Through Rates. Based on data from 1784 advertisers who placed one or more ads at the focal set of 165 sites. The 40.1% residual variance reflects differences in advertisers’ CTRs when placing ads at different sites.

### 2.2.2 Evidence of Returns to Learning

Our conjecture is that there exists something for the advertiser to learn. If CTRs are homogeneous across sites, then ad effectiveness might not vary across publisher sites, and there should be little value to learning. Toward this end, Table 1 reports a variance decomposition of CTRs across advertisers and sites. This analysis suggests that 9.1% of the variation in CTRs can be apportioned to publishers, 50.8% to advertisers, with the remainder idiosyncratic to the advertiser-site pair. The 40.1% of variation apportioned to the advertiser-publisher residual suggests there is a benefit to advertisers in finding a match with those sites that yield better outcomes.

## 3 Model and Estimation

This section details the model of advertisers’ site choice, and how learning affects these choices. It concludes by detailing our estimation approach.

### 3.1 Advertising Payoffs

Each week  $w$ , advertiser  $a$  considers placing an ad at each of the top 20 sites,  $s$ , where the advertiser does not already have an ad running. When choosing to place an ad at site  $s$ , the advertiser considers a menu of potential subscription lengths (i.e., how many days the ad is to run),  $x$ , and prices,  $p$ , that are charged by sites. For example, an advertiser considering site  $s$  might have the following options: {7 days for \$100, 14 days for \$180, 30 days for \$350, not advertising at site  $s$ }. The option not to advertise is always available, and is represented as a subscription of length  $x = 0$  days, obtained at a price of  $p = 0$ .

The expected payoff generated by the ad running for  $x$  days at site  $s$  starting at week  $w$  depends on the expected number of site visitors who will be shown the ad. The average daily number of visitors exposed to ads at site  $s$  is denoted  $t_s$  and is common knowledge to all sites and advertisers. Because all site visitors are served ads for all active subscriptions, an ad running for  $x$  days at site  $s$  yields  $t_s x$  total expected impressions.

Advertiser  $a$ 's expected payoff from an ad running for  $x$  days at site  $s$  is given as follows:

$$\pi_{asw}(x) = \delta_{asw} \zeta_a^{-1} \log(1 + t_s x \zeta_a) - p_{sw}(x) \quad (1)$$

where the first additive term represents the advertiser's valuation from the expected advertising outcome, and the second term ( $p_{sw}(x)$ ) represents the price the publisher charges to the advertiser.<sup>9</sup> Choosing not to advertise yields a payoff of  $\pi_{asw}(0) = 0$ . This function for advertiser valuation is a special case of the generalized, translated, constant elasticity of substitution utility function discussed in Bhat (2008) and Lee and Allenby (2014). The parameter  $\zeta_a > 0$  determines the rate at which the advertiser satiates on additional impressions (Dubé et al. 2005). The term  $\delta_{asw} > 0$  determines the overall scale of payoffs, and in particular, the change in marginal payoff at the point of zero impressions.

The term  $\delta_{asw}$  is subscripted by both advertiser and site, reflecting the empirical regularity that the same number of impressions served by two similar sites can generate different payoffs for the same advertiser (Perlich et al. 2012). The term represents the *match* between the advertiser's ad content and the site's audience. A higher match value leads to higher expected returns, and thus a higher likelihood of advertising at site  $s$ . A priori, the value of  $\delta_{asw}$  is uncertain to the advertiser, but it can be learned over time, as described in Section 3.2. Were the advertiser to be overly optimistic about its match with site  $s$ , it would initially advertise too much at that site. Were the advertiser to be overly pessimistic, it would advertise too little. Hence, there is value in learning efficiently about match.

## 3.2 Advertiser-Site Match and Learning

We next discuss how advertisers learn about their match value with each site. Clicks are assumed to provide an unbiased signal about match, under the presumption that more clicks reflect greater interest in the advertised good. Hence, we first link match ( $\delta_{asw}$ ) to CTRs ( $c_{as}$ ), and then show how changes in beliefs about clicks (i.e., learning) translate into changes in beliefs about match.

### 3.2.1 Linking Clicks to Match

The proportion of site  $s$ 's audience who click on advertiser  $a$ 's ads—that is, the true CTR—is a noisy measure of the general effectiveness of advertiser  $a$ 's ads when shown to site  $s$ 's audience. We denote the true CTR

---

<sup>9</sup>We follow Berry (1992) and Bajari et al. (2007) and incorporate ad costs directly in the payoff function without a price coefficient. The corresponding interpretation of the first term in Equation (1) as a dollar valuation is consistent with the display ad literature in exchange markets, where bid payments are expressed in monetary terms, and bid valuations are interpreted in a dollar metric (Ahmadi et al. 2023, Alcobendas and Zeithammer 2021, Bompaire et al. 2021, Tunuguntla and Hoban 2020, Waisman et al. 2019). To the degree a price coefficient were to exist and be normalized, it would be absorbed into the advertiser fixed-effect contained within  $\delta_{asw}$  (described in Equation (2)), which is multiplicatively separable in the advertiser value function in the first part of Equation (1).

as  $c_{as}$ , and note that this rate is unknown to both the advertiser and the site. Conditional on this unknown quantity, the true match between advertiser and site is a function of their true CTR, as given by

$$\delta_{asw}|c_{as} = \frac{c_{as}}{\gamma_a} \exp(\xi_a + \eta_s + \phi_{\tau[a,s,w]} + \psi_{m[w]} + \varepsilon_{asw}). \quad (2)$$

As elaborated in the next subsection, the term  $c_{as}/\gamma_a$  implies that match values are proportional to the true, but a priori unknown, advertiser-site CTRs.

Although actual CTRs are initially unknown to the advertiser, the component of advertiser match given by  $\exp(\xi_a + \eta_s + \phi_{\tau[a,s,w]} + \psi_{m[w]} + \varepsilon_{asw})$  is presumed known to the advertiser (but not the econometrician) when choosing where to place ads. This fixed component of advertiser match is decomposed into the following parts. First, the term  $\xi_a$  reflects the general efficacy of a firm's advertising across sites. It accommodates the possibility that some advertisers purvey goods that are more popular than others or are endowed with better advertising, and thus have higher expected returns per ad impression. Analogously,  $\eta_s$  reflects differences in sites' efficacy that are common to all advertisers, perhaps due to the demographics of the site (as sometimes claimed in its description on the platform) or some observable (to the advertiser) aspect of its content or construction.<sup>10</sup> Although, more generally, any ad-site-week observable covariates could be incorporated into the conditional match expression for  $\delta_{asw}$ , we only observe  $\phi_{\tau[a,s,w]}$ , a fixed effect that is selected according to how long ago advertiser  $a$  first placed an ad at site  $s$ .<sup>11</sup> The purpose of this fixed effect is to reflect dynamics in the value or efficacy of advertising that are unrelated to learning about match (e.g., wear-in or wear-out; Little and Lodish 1969). The term  $\psi_{m[w]}$  is a month fixed effect, included to account for any seasonal dynamics affecting the entire ad network. Finally,  $\varepsilon_{asw}$  is an idiosyncratic demand shifter at the advertiser-site-week level that is observed by the advertiser and not the econometrician.

### 3.2.2 Learning About Clicks

As mentioned, the expression  $c_{as}/\gamma_a$  indicates the ratio of advertiser  $a$ 's true (but a priori unknown) CTR at site  $s$  to the parameter  $\gamma_a$ . The parameter  $\gamma_a$ , which is known to the advertiser but not the econometrician,

<sup>10</sup>Equations (1) and (2) accommodate a restricted degree of heterogeneity in satiation from additional impressions at different sites. This is because the term in Equation (1) representing the value from advertising is proportional to  $c_{as} \exp(\xi_a + \eta_s) \zeta_a^{-1} \log(1 + t_s x \zeta_a)$ . At the point of no advertising, the marginal value of the first impression is proportional to  $c_{as} \exp(\xi_a + \eta_s)$ ; and at the point of an expected  $t_s x$  impressions, the marginal value of an additional impression is proportional to  $c_{as} \exp(\xi_a + \eta_s) (1 + t_s x \zeta_a)^{-1}$ . Moreover, as advertisers learn about their true CTRs ( $c_{as}$ ), the extent of satiation is expected to change. Following the broader literature on learning models in marketing and economics, the satiation parameter is not assumed to be fully unrestricted over advertisers and time,  $\zeta_{asw}$ . Relaxing this restriction could prove a useful extension to the learning literature.

<sup>11</sup>There are separate fixed effects for  $\tau = 1, \dots, 4$  weeks since first placing an ad at a given site (the fixed effect for  $\tau = 0$  is normalized to 0). Weeks 5–6, 7–8, 9–12, 13–16, 17–32, 33–52, and 52+ are grouped and represented by common fixed effects,  $\tau \in \{5, \dots, 11\}$ . As  $\phi_2$  captures the second week since an ad first appeared, it could be correlated with lagged advertising and thus endogenous (all other  $\phi$  are invariant to lag purchase). The correlation with lag advertising is 0.37 and the correlation with clicks and impressions are  $-0.06$  and  $-0.05$ , suggesting the inclusion of this control variable has little impact on the structural parameters.

serves as both i) advertiser  $a$ 's prior expectation for its unknown CTR at any new site, and ii) the baseline for judging how successful its advertising has been. For example, if the true CTR at site  $s$  turns out to be greater than what the advertiser previously expected, then  $c_{as}/\gamma_a$  will be greater than 1, reflecting better than anticipated returns to advertising at site  $s$ . To the extent initial beliefs about click through rates are too high (low) at a given site, one would expect advertisers' initial probabilities of advertising to be higher (lower).

We assume prior beliefs about the true CTR,  $c_{as}$ , denoted  $\tilde{c}_{as}$ , follow a beta distribution, owing to the beta's flexibility in representing distributions bounded between 0 and 1. Specifically,

$$\tilde{c}_{as} \mid \gamma_a \sim \text{Beta}(1, (1 - \gamma_a) / \gamma_a), \quad (3)$$

which implies  $\mathbb{E}[\tilde{c}_{as} \mid \gamma_a] = \gamma_a$ , and places the bulk of probability density between 0 and  $4\gamma$ . Hence,  $\tilde{c}_{as}$ , like most CTRs, is a priori skewed towards 0. The parameter  $\gamma_a$  reflects advertiser  $a$ 's prior belief about its CTR at a previously unused site.


Conditional on the true CTR,  $c_{as}$ , the likelihood for the cumulative number of clicks as of week  $w$ ,  $n_{asw}^C$ , follows a binomial distribution with likelihood

$$n_{asw}^C \mid n_{asw}^I, c_{as} \sim \text{Binomial}(n_{asw}^I, c_{as}), \quad (4)$$

where  $n_{asw}^I$  is the cumulative number of advertising impressions served. The number of clicks received,  $n_{asw}^C$ , depends on the true CTR at site  $s$ ,  $c_{as}$ , hence impressions and clicks provide information about the true, but unknown CTR. Accordingly, we specify a Bayesian updating process on advertiser beliefs about the CTR, leading to a beta-binomial posterior CTR belief distribution for the number of clicks. The mean of this updated posterior distribution has an expected value of

$$\mathbb{E}[\tilde{c}_{as} \mid n_{asw}^I, n_{asw}^C, \gamma_a] = \gamma_a \frac{1 + n_{asw}^C}{1 + \gamma_a n_{asw}^I}. \quad (5)$$

Accordingly, the expected ratio of CTR to  $\gamma_a$  is  $\mathbb{E}[\tilde{c}_{as}/\gamma_a \mid n_{asw}^I, n_{asw}^C, \gamma_a] = (1 + n_{asw}^C)/(1 + \gamma_a n_{asw}^I)$ . As previously noted, Equation (5) implies that the advertiser initially expects a CTR of  $\mathbb{E}[\tilde{c}_{as} \mid \gamma_a] = \gamma_a$  prior to advertising (when  $n_{asw}^C = 0$  and  $n_{asw}^I = 0$ ). As  $n_{asw}^C \rightarrow \infty$  and  $n_{asw}^I \rightarrow \infty$ , the ratio  $n_{asw}^C/n_{asw}^I \rightarrow c_{as}$ , and thus  $\mathbb{E}[\tilde{c}_{as} \mid \gamma_a] \rightarrow c_{as}$ . Hence, initial beliefs about the CTR begin as  $\gamma_a$ , but eventually converge with enough advertising to the true CTR,  $c_{as}$ . If initial beliefs are correct, the expected value of the ratio  $\tilde{c}_{as}/\gamma_a$  remains at one and never changes (in expectation). If initial beliefs are low (high), this ratio increases above (decreases below) one. For example, if the true CTR is less than  $\gamma_a$ , then the number of clicks received,  $n_{asw}^C$ , will grow at a slower rate than the a priori expected number of clicks received,  $\gamma_a n_{asw}^I$ . Hence  $\gamma_a$  in Equation (5) will be multiplied by a ratio that converges to a value between 0 and 1. If initial beliefs are too high, the



**Figure 5:** Simulated Posterior CTR Beliefs. The evolution of posterior beliefs for  $\mathbb{E}[\tilde{c}_{as}/\gamma_a | n_{asw}^I, n_{asw}^C, \gamma_a]$  is simulated under a true click through rate of 0.0075. The advertiser prior is either pessimistic (blue plot,  $\gamma_a = 0.005$ ), accurate (green plot,  $\gamma_a = 0.0075$ ) or optimistic (red plot,  $\gamma_a = 0.01$ ).

advertiser spends too much on advertising in early periods, with the added expense not recovered by sales. Conversely, if initial beliefs are too low, the advertiser advertises too little, or not at all. Figure 5 shows a hypothetical learning process for the ratio of CTR to  $\gamma_a$  for optimistic, accurate, and pessimistic prior beliefs, and illustrates how these beliefs change over time.

### 3.2.3 Learning About Match

The expression for true match values in Equation (2) is conditioned on the true CTR,  $c_{as}$ . Using the standard beta-binomial distribution, as applied to the prior for  $\tilde{c}_{as}$  from Equation (3) and the likelihood for  $n_{asw}^I$  from Equation (4), and then taking the expectation of  $\delta_{asw}$  with respect to  $\tilde{c}_{as}$  in Equation (2), the advertiser's updated expected match with site  $s$ ,  $\tilde{\delta}_{asw}$ , is given by

$$\begin{aligned} \mathbb{E}[\tilde{\delta}_{asw} | n_{asw}^I, n_{asw}^C, \epsilon_{asw}, \theta] &= \check{\mu}_{asw} \cdot \exp(\epsilon_{asw}) \\ &= \frac{1 + n_{asw}^C}{1 + \gamma_a n_{asw}^I} \exp(\xi_a + \eta_s + \phi_{\tau[a,s,w]} + \psi_{m[w]}) \cdot \exp(\epsilon_{asw}), \end{aligned} \quad (6)$$

where  $\theta$  represents the model parameters to be estimated. Defining  $\check{\mu}_{asw} \equiv \mathbb{E}[\tilde{\delta}_{asw} | n_{asw}^I, n_{asw}^C] / \exp(\epsilon_{asw})$  will prove useful when deriving the likelihood function. This expression for expected match has a simple interpretation: Prior to advertising, the advertiser would have expected  $\gamma_a n_{asw}^I$  clicks after the network served  $n_{asw}^I$  impressions. Instead, it received  $n_{asw}^C$  clicks. The ratio of these two quantities determines whether the advertiser now expects higher or lower match, relative to when it placed the first ad. Combining Equation (6) with Equation (1), where true match is replaced by expected match, yields an expression for expected advertiser payoffs conditional on past impressions ( $n_{asw}^I$ ) and clicks ( $n_{asw}^C$ ).

### 3.3 Estimation

The likelihood formulation uses the method described in Lee and Allenby (2014). The central idea behind this approach is to derive a set of inequality constraints on  $\varepsilon$  that rationalize the observed set of choices. For example, if an advertiser buys a 7 day ad run at a given price, the advertiser must expect a higher payoff compared to a shorter or longer subscription at an alternative price. Many advertiser choices are observed over many periods, yielding a large number of inequality constraints from which one can derive a likelihood.

Advertiser site choices are observed at the weekly level. Hence, we consider the idiosyncratic demand shifters,  $\varepsilon_{asw}$ , in Equation (6), and their implications for advertiser payoffs given by Equation (1). Replacing the true match with expected match, the expected payoff equation becomes

$$\begin{aligned} V_{asw}^{\theta}(x, \varepsilon_{asw}) &\equiv \mathbb{E} [\pi_{asw}(x) | \theta, \varepsilon_{asw}, n_{asw}^I, n_{asw}^C] \\ &= \check{\mu}_{asw} \exp(\varepsilon_{asw}) \zeta_a^{-1} \log(1 + t_s x \zeta_a) - p_{sw}(x), \end{aligned} \quad (7)$$

where the set of parameters to estimate is  $\theta = \{\gamma, \xi, \eta, \phi, \psi, \zeta\}$ . Of special interest in this set is  $\gamma$ , the vector of advertisers' prior beliefs about the efficacy of their advertising, which is informative about whether their naivety induces them to advertise too much or too little.

Conditional on a set of parameters  $\theta$ , the likelihood of  $x_{asw}$  is given by the probability density of the unobserved  $\varepsilon_{asw}$  after integrating it over the region of  $\varepsilon$  that can rationalize the observed choice of  $x_{asw}$ . This region is defined by upper and lower bounds, which are themselves determined by a pair of inequality constraints. Let  $\uparrow x$  denote an alternative subscription that is longer than the subscription of length  $x$  that was purchased, and let  $\uparrow p$  denote its (higher) price. Similarly, let  $\downarrow x$  denote an alternative, shorter subscription, and  $\downarrow p$  its (lower) price. Moreover, assume for the moment that both shorter and longer alternatives to  $x_{asw}$  were offered by site  $s$ .

The lower bound for the region of  $\varepsilon$  that can rationalize  $x_{asw}$  is obtained from the observation that the advertiser did not buy the shorter subscription,  $\downarrow x$ . Hence,  $V_{asw}^{\theta}(x, \varepsilon_{asw}) > V_{asw}^{\theta}(\downarrow x, \varepsilon_{asw})$ . Substituting Equation (7) into this inequality and isolating  $\varepsilon_{asw}$  leads to the following lower bound for  $\varepsilon_{asw}$  (Lee and Allenby 2014):

$$\varepsilon_{asw} > \ell b^{\theta}(x_{asw}, p_{sw}), \quad \ell b^{\theta}(x, p) \equiv \log(p - \downarrow p) - \log \left[ \check{\mu}_{asw} \zeta_a^{-1} \log \left( \frac{t_s x \zeta_a + 1}{t_s \downarrow x \zeta_a + 1} \right) \right] \quad (8)$$

Similarly, because the advertiser did not buy the longer subscription,  $V_{asw}^{\theta}(x, \varepsilon_{asw}) > V_{asw}^{\theta}(\uparrow x, \varepsilon_{asw})$ , and thus

$$\varepsilon_{asw} < ub^{\theta}(x_{asw}, p_{sw}), \quad ub^{\theta}(x, p) \equiv \log(\uparrow p - p) - \log \left[ \check{\mu}_{asw} \zeta_a^{-1} \log \left( \frac{t_s \uparrow x \zeta_a + 1}{t_s x \zeta_a + 1} \right) \right]. \quad (9)$$



Finally, if  $x_{asw} = 0$ , meaning the advertiser did not buy a subscription, then there is no lower bound, and thus  $\epsilon_{asw} > -\infty$ . Similarly, if a longer alternative to  $x_{asw}$  was not offered, then there is no known upper bound, and thus  $\epsilon_{asw} < \infty$ . Defining  $\ell b$  or  $ub$  as negative or positive infinity when  $x_{asw}$  lies at one of these boundaries, the resulting likelihood for each observation is obtained by integrating over the joint density of  $\epsilon$ , denoted  $f(\epsilon_{asw})$ , over the regions indicated by Equations (8) and (9):

$$L(x, p | \theta) = \int_{\ell b^\theta(x, p)}^{ub^\theta(x, p)} f(\epsilon) d\epsilon. \quad (10)$$

The intuition behind this likelihood is that the observed choice probability is maximized if the interval of  $\epsilon$  that can rationalize the observed choice is wide, and minimized if the interval that can rationalize the observed choice is narrow. As noted by Lee and Allenby (2014), this discrete likelihood approach admits the possibility that intermediate options between the observed choice and the closest available alternatives might have been preferred, had they been offered. It therefore does not assume that the observed choice  $x_{asw}$  was optimal, but rather that it was simply better than the closest alternatives.<sup>12</sup>

To complete the likelihood, we assume  $f(\epsilon)$  is a normal pdf with mean 0 and variance  $\sigma^2$ , and that the  $\epsilon$ s are independent. Prior distributions for the fixed effects  $\eta$ ,  $\xi$ ,  $\phi$ , and  $\psi$  are independent Student- $t$  with 4 degrees of freedom, mean 0, and unit variance. The prior distribution for  $\sigma$  is exponential, and the penalized complexity approach of Simpson et al. (2017) is used to choose the exponential rate parameter. This entails choosing a rate such that  $\Pr[\sigma > U] = \alpha$ . We set  $U = 1$  and  $\alpha = .1$ , so that  $\Pr[\sigma > 1] = .1$ . The resulting prior distribution is  $\text{Exponential}(-\log(\alpha)/U)$ , leading to  $\sigma \sim \text{Exponential}(\log 10)$ . The prior distributions for  $\zeta_a$  and  $\gamma_a$  are defined hierarchically to allow pooling across advertisers, and both parameters' prior distributions are derived by transforming exponential variates. The conditional prior for  $\zeta_a$  is exponential, shifted by .01 to improve numerical stability during estimation, with mean  $.01 + \bar{\zeta}$ .  $\bar{\zeta}$  is exponential with prior mean  $10/\log 10$ , so that  $\Pr[\bar{\zeta} > 10] = .1$ . Accordingly, the marginal prior mean for  $\zeta_a$  is  $.01 + 10/\log 10$ . The conditional prior for  $\gamma_a$  resembles a unimodal beta distribution with most of the mass near zero, but is derived from  $\gamma_a | \bar{\gamma} = g_a / (\bar{\gamma}^{-1} + g_a)$ , with  $g_a \sim \text{Exponential}(1)$ . The prior distribution for  $\bar{\gamma}$  is exponential with rate  $-\log(.1)/.002$ , so that  $\Pr[\bar{\gamma} > .002] = .1$ . The resulting marginal prior mean for  $\gamma_a$  is .00087, with  $\Pr[\gamma_a > .002] \approx .12$  (for reference, the average CTR is .00045). We sample from the model's posterior distribution using Hamiltonian Monte Carlo (HMC), as implemented in the cmdstanr package for R (Stan

<sup>12</sup>Conditioned upon first observing an advertiser purchase on the network, all subsequent non-purchase decisions ( $x = 0$ ) are included in the advertiser's choice set. Prior to the advertiser's first purchase on any site, non-choices do not enter the likelihood. Not including the zero choices before the first purchase implies that either the advertiser was unaware of the sites prior to their first purchase, or they purchased a subscription in the week they joined the network. The timing of the advertisers' first purchases on the network are assumed to be exogenous (Balseiro and Candogan 2017, Balseiro and Gur 2019, Bimpikis et al. 2020, Wu 2015).

Development Team 2017).

### 3.4 Subscription Prices

Subscription prices,  $p_{sw}(x)$ , vary by week, publisher, and subscription length.<sup>13</sup> Although these prices are generally stable over time within publisher-subscription length, the series can sometimes exhibit perturbations around their means. While their effect on estimation is limited, these perturbations induce counterfactual analyses in some instances that yield results with unusual valuations (for example, in a week where the longer ad subscription is priced more than the shorter one). To alleviate this problem, we use OLS to regress the log of subscription prices on week fixed effects (to account for seasonality or other time-varying trends), site fixed effects (to account for differences in audience attractiveness or site popularity), and the log of subscription length,

$$\mathbb{E}[\log(p_{sw}(x))] = \mu_w + \mu_s + \nu \log(x). \quad (11)$$

The  $R^2$  of the log pricing regression is 0.94, evidencing high fit. When estimating the likelihood from Equations (8) and (9), price is set to be

$$p_{sw}(x) = \exp\left(\frac{\hat{\xi}^2}{2} + \hat{\mu}_w + \hat{\mu}_s + \hat{\nu} \log(x)\right) \quad (12)$$

where  $\hat{\xi}^2$  is the residual variance from the pricing regression for  $\log(p_{sw}(x))$ .<sup>14</sup>

## 4 Results

### 4.1 Model Fit

In addition to estimating the advertiser learning model described in Section 3, we estimate three simpler models. These omit learning (setting the ratio  $c_{as}/\gamma_a = 1$  for all choices) and/or the time-varying fixed effects (setting  $\phi_\tau$  and  $\psi_m$  to zero), and help to ascertain the predictive value of modeling advertiser learning. The fit of each model is presented in Table 2. Results suggest that modeling advertiser learning leads to a substantial improvement in fit. The improvement in fit from modeling advertiser learning (a decrease in  $ELPD_{LOO}$  of

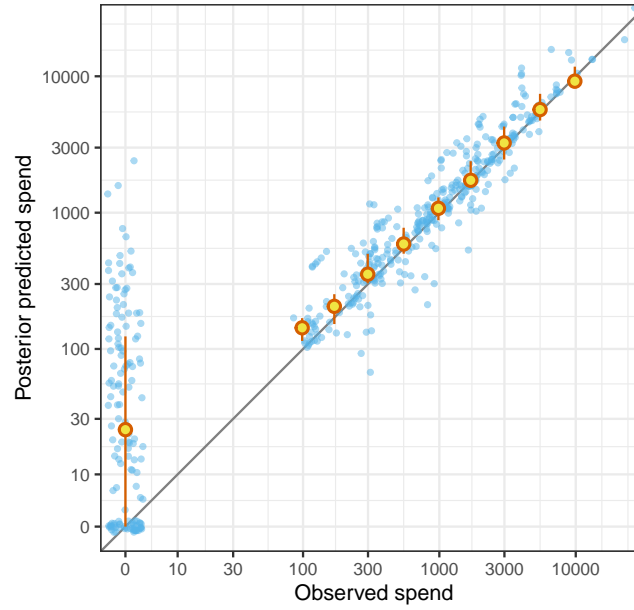
---

<sup>13</sup>Subscription lengths  $x$  are observed only when purchased in a given week. We impute unobserved subscription lengths in a given week for a given site using a nearest neighbor approach. If a subscription of length  $x$  was offered within 4 weeks before or after week  $w$ , we assume it was also offered in week  $w$ .

<sup>14</sup>Because prices are inferred from observed sales, one cannot rule out the possibility of negotiated instead of posted prices. To ascertain the potential for negotiated prices, we checked whether prices paid for the same inventory from the same publisher differ across advertisers in the same week. The coefficient of variation (bias corrected) for price for the same ad inventory sold to different advertisers in the same week is 12.4%, meaning the price variation is small. In addition, in 2/3 of all instances for which we observe multiple advertisers at a site, they pay exactly the same price. In other words, some price negotiation may exist, but its degree is slight, so prices are set to be constant across advertisers within a given week at a given publisher.

	Model 1	Model 2	Model 3
Time-Varying			
Fixed Effects	No	Yes	Yes
Learning	No	No	Yes
$p_{LOO}$	95.4	114.7	148.1
$ELPD_{LOO}$	-3671.8 (138.8)	-3520.9 (135.6)	-3386.9 (132.3)
Incremental			
Improvement		-150.9 (18.5)	-134.0 (18.1)


**Table 2:** Model Fit and Comparison.  $ELPD_{LOO}$  is the Bayesian leave-out-one estimate of log pointwise predictive density, an estimate of model fit; and  $p_{LOO}$  is the effective number of parameters, an estimate of model complexity (Vehtari et al. 2016). Incremental improvements compare Model 2 to Model 1, and Model 3 to Model 2. Values in parentheses show standard errors of point estimates.



**Figure 6:** Posterior Predictive Fit for Advertising Spend. Expected total spend for all advertiser-site pairs shown with binned medians and inner 50th percentiles. Points corresponding with advertiser-site pairs with zero observed spend are jittered. The observed spend of zero does not represent the inaugural campaign (where spend is necessarily positive), but instead weeks in which no spend is observed in the data. The distribution of positive forecasted spend amounts when observed spend is zero reflects left truncated prediction errors.

134.0) is almost as large as the improvement in fit from including time-varying fixed effects (a decrease in  $ELPD_{LOO}$  of 150.9).

The posterior predictive fit for advertiser spending is depicted in Figure 6, where each point represents the total amount an advertiser spent with a publisher over 27 weeks (the procedure for simulating from the posterior description is described in Section 5.1.1). Overall, the model predicts spending well, including cases of no spending.



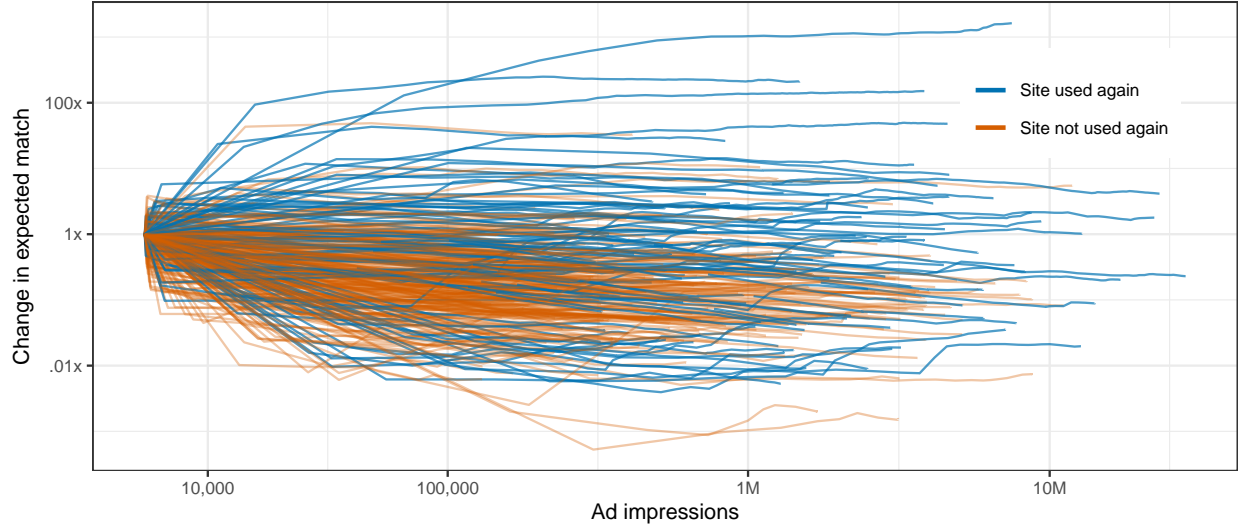
**Figure 7:** Realized and Prior/Reference Click Through Rates. Each row represents an advertiser and shows the posterior mean of  $\gamma_a$  and their observed CTRs for sites at which they placed ads. The median CTR is .045%, and the median prior/reference CTR is .232%; both are depicted as vertical dashed lines.

## 4.2 Parameter Estimates

Recall, the key parameter estimates pertain to the  $\gamma_a$ 's, which represent advertisers' prior beliefs about CTRs, and a central question is whether these beliefs are optimistic or pessimistic. Given that Section 2.2 shows that the number of sites used by advertisers tends to decrease over time and that advertisers tend to switch away from those initial choices, we hazard that advertisers tend to be over-optimistic about site match. Hence, the parameter  $\gamma_a$ , which reflects prior beliefs about match (in terms of CTR), should reflect this (see Section 3.2.2). Estimates indicate that for nearly all advertisers, the  $\gamma_a$ s are indeed higher than their empirical CTRs. Figure 7 plots the estimates for  $\gamma_a$  against the realized CTRs for all observed advertiser-site pairs. Note that the median prior belief for click through rates is .232%, whereas the median true CTR is about one fifth that amount, .045%. This suggests that advertisers are, on average, initially over-optimistic. There is considerable heterogeneity in initial beliefs, as well as in how accurate those beliefs are, as some are correct, some are optimistic, and some are pessimistic.

Figure 8 depicts the changes in advertiser CTR beliefs over time, scaled by the prior belief ( $\tilde{c}_{as}/\gamma_a$ ). Prior to advertising at a site, this ratio is equal to one. With more information, the ratio approaches  $c_{as}/\gamma_a$ . For example, if the true CTR is twice the initial belief, then this ratio will approach two. Values less than one imply advertisers were initially optimistic.

Because most of the belief paths in Figure 8 decrease from their starting point of 1, advertisers appear largely optimistic about the likelihood of clicks they will obtain on a site when they first advertise. Accordingly,



**Figure 8:** Advertiser Match Beliefs Over Time. Depicts outcomes for the estimation sample, with each line representing an advertiser-site combination. Change in expected match is the ratio of an advertiser’s CTR beliefs to their initial prior,  $\tilde{c}_{as}/\gamma_a$ , which updates based on impressions and clicks observed each day. The horizontal axis presents the cumulative number of impressions (rather than time) to allow comparison across advertiser-site pairs, which differ in impressions served per day and the timing of ad purchases.

advertisers often spend too much, but sometimes spend too little. Moreover, beliefs take several thousand impressions before they begin to stabilize, meaning an advertiser could misspend a considerable sum on advertising before learning its true site match. Whether optimistic or pessimistic, advertisers’ best advertising choices—were they to know their true match with each site up front—would frequently differ from their actual advertising choices. Advertisers thus face losses relative to being fully informed.

One might expect that publishers and the platform have little incentive to address over-optimistic advertisers. Were one to correct advertisers’ prior beliefs about their match with an initial set of sites, one might expect ad spend across those sites and the platform as a whole to be lower. However, this expectation ignores the possibility that an advertiser might switch to a better matched publisher, potentially increasing total advertising on the publisher sites and the platform, thereby increasing revenues. Section 5 quantifies these trade-offs and considers a potential solution that can benefit advertisers, sites, and the platform.

### 4.3 Robustness Checks

Several issues present with our specification, including the potential for forward-looking behavior by advertisers, supply-side considerations, and other concerns. Below, we detail these issues and offer some analyses to better rationalize our modeling choices. We highlight the findings of various robustness checks in this section and provide further detail in Online Appendices A–C. Model performance and insights are robust

across these checks.

### 4.3.1 Test and Learn Behavior

Test and learn behavior entails advertisers choosing to “test” sites with lower expected match, but higher match uncertainty. Resolving this uncertainty creates future value in the event the advertiser discovers a better than expected match. The demand model assumes advertisers are not forward-looking when making choices. One reason for this assumption is that advertisers are often unsophisticated. Recently, test and learn algorithms have been automated in ad exchange markets, showing great gains in ad performance (Tunuguntla and Hoban 2020, Schwartz et al. 2017); such gains would not be possible if advertisers were already using a test and learn strategy. However, to the degree test and learn behavior is present, a model that assumes otherwise would yield over-optimistic priors ( $\gamma_a$ ), because lower demand arising from lower match uncertainty would be rationalized as a lower-than-expected CTR.

To explore the question of whether advertisers are forward-looking, we consider advertiser switching across ad subscription lengths within a publisher. Forward-looking advertisers, all else equal, should tend to first use shorter subscriptions and, if the short trial works at the publisher, switch to longer subscriptions (a form of “dipping one’s toes in the water before taking a swim”). No evidence exists for this behavior. In 1,240 cases, advertisers switched from shorter to longer subscriptions, in 1,266 cases the next ad subscription was shorter, and in 3,007 the next ad subscription was the same length.

In addition, to more formally address this concern, the ad subscription model in Section 3.1 is extended to incorporate a UCB component for test and learn (Auer et al. 2002, Zhuo 2023). Results indicate that adding a test and learn component yields no change in model fit as indicated by the  $ELPD_{LOO}$ , and that the prior belief estimates are essentially identical. This suggests the model is robust to the assumption of myopic advertisers. Details of this analysis are reported in Online Appendix A.

### 4.3.2 Supply-side Considerations

The model and counterfactuals do not consider how aggregate advertiser demand might affect publisher prices or CTRs. This section briefly overviews these considerations next. Detailed analyses are in Online Appendix B.

**Pricing.** Equation (1) raises the potential for price endogeneity bias. Such a bias could manifest if site-week omitted factors i) correlate with the advertiser’s value for ads, and ii) are not captured by the observables in the estimation equation (these observables include week and site fixed effects to control for site popularity

and seasonality). If price endogeneity exists, one would expect prices to change in response to demand, affecting both inference (lowering the implied price effect) and counterfactual analyses. To explore the issue, we consider the degree to which site, week, and site-week components covary with log prices. The site-week component explains less than 1% of the total variation, meaning few pricing equation unobservables exist at the site-week level that could correlate with unobservable site-week demand shocks. The potential for endogeneity bias is thus limited. A concern with this analysis of variance approach is that the site fixed effects in a log pricing model can absorb site-specific average demand. Because average demand can change during the counterfactual analysis, the fixed effects might therefore not be counterfactually invariant. Exploring this issue, we find no significant link between average prices and the fixed effects.<sup>15</sup> In spite of these analyses, to the extent that unobserved variables do induce any remaining endogeneity, counterfactual welfare results would reflect a lower bound on the welfare gains, as imputed price sensitivity would increase after controlling price endogeneity, amplifying the welfare effects from changes in ad spend.

**Crowding.** Another concern pertains to crowding. When more advertisers appear concurrently on a site, it becomes possible that CTRs could decrease. If that were the case, it would be necessary to control for how the number of advertisers on a site affects clicks. However, in direct advertising contexts, such crowding effects can be small because ads appear over a long duration and are rotated. For example, the rate card in Figure 1 indicates that in-column side bars ads are rotated with each impression, with an estimated 50,000 impressions monthly. This means that, even if several ads are running simultaneously, there is still ample opportunity for repeat visitors to see an ad in a prominent position on the page. To test this conjecture more formally, site-week clicks are regressed on the number of advertisers, controlling for site fixed effects and average advertiser clicks. Results indicate a small and insignificant effect (and of the opposite sign than predicted by crowding) of the number of advertisers on clicks. This result suggests that crowding effects are not considerable in this context.

#### 4.3.3 Additional Considerations

We briefly consider several additional issues, including the sample selection criterion, budget constraints, heterogeneity in advertiser priors across publishers, and the assumed prior variance of CTR beliefs. A detailed analysis of these considerations is presented in Online Appendix C.

---

<sup>15</sup>Relatedly, we explore whether lag advertiser demand at a publisher site affects its prices. As there is essentially no price variation over time within site, no significant link is expected between lag demand and prices and we observe none.

**Sample Selection.** Results are based on the largest 20 publishers and 100 random advertisers. As the number of publishers increases, counterfactual increases in match values due to improving advertisers’ priors should be higher, because the opportunity for advertisers to sort into a better match is greater. Results indicate this to be the case when the number of publishers included in the sample increases by 25%. Hence, the reported overall gains in our paper are a conservative estimate.

**Budget Constraints.** To the extent that budget constraints exist, advertiser spending will be attenuated in any counterfactual analysis. Budget constraints are not observed, and there is no evidence advertisers are constrained. Moreover, it is unclear why advertisers would not raise funds to pay for positive investments. Because budget constraints are not observed, analyses pertaining to this issue are speculative. However, under the premise that advertiser budgets are constrained to some percentage increase over the maximum spend in the data, it is possible to assess how counterfactual insights would change in response to different budget constraints. As one might expect, the tighter the budget constraint, the less welfare gains are evidenced in our counterfactual analyses.

**Heterogeneity in Advertiser Priors Across Publishers and Variance in Learning Rates.** We conduct two additional robustness checks. First, we accommodate the potential for heterogeneity in advertisers’ prior beliefs across sites, and second we estimate rather than set learning rates. Though adding a large number of parameters, heterogeneity in prior beliefs has a negligible impact on model fit and an immaterial effect on prior belief estimates. Estimating prior variances (and thus the learning rate) degrades model fit and has no material effect on inferences about prior beliefs.

## 5 The Role of Information Provision on Advertiser and Publisher Welfare

The results in Section 4 indicate that many advertisers are over-optimistic about advertising efficacy and therefore tend to overspend. Holding all else fixed, redressing this overoptimism could cause advertisers to spend less, leaving advertisers better off, but publishers and the platform worse off. On the other hand, advertisers could migrate some of their ad spend to other sites that are a better match, increasing their welfare and possibly their total spend, thus increasing publisher and network revenue. Hence, whether endowing advertisers with better information increases or reduces overall welfare is an empirical question. This section explores how informing advertiser priors,  $\tilde{c}_{as}$ , affects advertisers’ valuations and spending levels, and explores the consequences for publisher and platform welfare. It first outlines approaches to affording better advertiser priors and then considers the welfare implications of these approaches.



## 5.1 Information Provision

We consider two mechanisms for counterfactual information provision: i) whether or not the advertiser is endowed with full information about the clicks they obtain from their chosen sites, and ii) whether or not advertisers have access to pooled data from other advertisers. We simulate advertising choices and compare advertising spend under each of these mechanisms to a common baseline. Appendix B describes the procedure used to simulate counterfactual data in the baseline and other scenarios.

### 5.1.1 Simulation Procedure and Baseline Scenario

The baseline and counterfactual scenarios differ in terms of i) the advertisers' initial information about their expected CTRs, and thus expected match utilities in week 1; and ii) the advertisers' choice sets, which are closely related to the information available for simulating clicks.

**Advertisers' initial information.** In the baseline scenario,  $B$ , advertisers have the same initial information as in the estimation sample. Thus, for sites at which the advertiser placed ads prior to week 1, their initial expected ratio of CTR to  $\gamma_a$  in the baseline scenario is  $\mathbb{E}[\tilde{c}_{as0}/\gamma_a | n_{as0}^I, n_{as0}^C, \gamma_a] = (1 + n_{as0}^C)/(1 + \gamma_a n_{as0}^I)$ , with  $n_{as0}^I > 0$ ; for all other sites,  $\mathbb{E}[\tilde{c}_{as0}/\gamma_a | \gamma_a] = 1$ . The counterfactual scenarios augment these initial information states with additional CTR information.

**Advertisers' choice sets.** In the baseline scenario, advertisers' choice sets are restricted to the sites with observed purchases in the estimation sample, as this allows us to simulate clicks using an approximation to  $c_{as}$  based on all advertising outcomes in the full data set. In the counterfactual simulations, we use these values of  $c_{as}$  to simulate clicks whenever they are available from the observed advertising data. In some counterfactual scenarios, we expand the advertisers' choice sets to include all sites in the estimation sample, and accordingly simulate clicks for these sites based on imputed values of  $c_{as}$ , which are obtained via a procedure described in Section 5.1.3.

### 5.1.2 Full Information Within Advertiser

The first counterfactual scenario considers endowing advertisers with full information about the sites for which we observe ad buys in the data. It is labeled the "full information" counterfactual and denoted  $C_F$ . In this counterfactual scenario, we set  $n_{as0}^I = 10^{12}$  and  $n_{as0}^C = \lfloor 10^{12} c_{as} \rfloor$ , so that  $\mathbb{E}[\tilde{c}_{as0}/\gamma_a | n_{as0}^I, n_{as0}^C, \gamma_a, C_F] \approx c_{as}/\gamma_a$ .<sup>16</sup> The value of  $c_{as}$  is the same value used in the baseline scenario, and advertisers' choice sets are

<sup>16</sup>We do not use the value of  $c_{as}/\gamma_a$  directly because i) it leads to numerical discrepancies between this ratio and simulated clicks and impressions related to floating point precision for large numbers, and ii) it completely eliminates learning. By setting  $n^I = 10^{12}$

restricted to observed advertiser-site pairs, as in the baseline. Hence the only difference between  $C_F$  and  $B$  lies in advertisers’ prior information. This counterfactual yields advertiser spending and revenue as if the advertiser had already known the response they would eventually receive from ads placed on the site (i.e., an oracle prior).

While this counterfactual can yield insights into the advertisers’ losses in the baseline scenario, relative to a hypothetical endowment with complete information about advertising responses at chosen sites, it does not consider how advertiser outcomes might change at sites where no spending occurred. As noted previously, this limitation is a consequence of only observing CTRs for sites that advertisers actually used in the data. Moreover, this approach to providing advertisers with better information is infeasible because full information about site CTRs is not revealed until after sites are chosen. By then, it is too late for advertisers to change course. Hence, we consider an alternative that involves sharing information across advertisers.

### 5.1.3 Pooling Information Across Advertisers

Given the ad network observes CTRs across all advertisers over time (as well as their ad creatives), it is well positioned to better inform advertisers about initial match by pooling information across advertisers. This counterfactual scenario updates advertisers’ prior information by considering the performance of similar advertisers and ads on the platform, and extrapolating this information to sites that advertisers did not purchase from in the data. Reflecting advertisers’ augmented information states, we also expand their choice sets to allow ad purchases at any site. When simulating clicks, we use the observed (i.e., full information) CTRs when available, and the imputed CTRs otherwise. Hence, for sites with observed purchases, there is a discrepancy between advertisers’ expectations, which are based on imputed CTRs, and realizations of clicks, which are based on observed CTRs.

For the purposes of this counterfactual simulation, we expect that ads with similar content generate similar CTRs when placed with a given publisher. More conceptually, we expect advertisers with similar ads to have similar levels of match with sites’ audiences, and thus generate similar match signals in terms of CTR. Inferences from this scenario depend on information that is readily available to platforms, hence this approach to pooling is feasible for platforms to implement.

**Estimating advertiser similarities.** Our approach entails i) deriving a measure of the similarity between any two advertisers, and ii) pooling outcomes among similar advertisers to predict ad performance. This raises the practical question of how to measure similarity between advertisers. As is common on many ad platforms, advertisers’ priors are strongly updated, but there is still sufficient opportunity for learning and adaptation.

platforms, there is no advertiser demographic information in the data. However, in addition to recording the number of impressions and clicks, the platform also retains the images used by advertisers in their display ad campaigns, and these images contain useful information about the advertisers.

To estimate similarity between advertisers, we first retrieve a set of concept tags from Google’s Cloud Vision API<sup>17</sup> describing the content of each display ad placed by advertisers at sites in the focal cluster of 165 liberal blogs. Based on those tags, we construct a term frequency/inverse document frequency (TF-IDF) matrix reflecting how prevalent and distinctive each tag is in characterizing advertisers. A square matrix of pairwise cosine similarities between advertisers is then computed from the TF-IDF measures. Finally, we use the similarities within each row of this matrix to impute click through rates for all advertiser-site pairs in the estimation sample, including those for which no ads were placed. Appendix C details these specific steps.

**Imputing CTRs.** Given a vector of similarities,  $r_a$ , between advertiser  $a$  and all other advertisers, as well as the set of advertisers,  $\mathcal{A}_s$ , for which we can approximate  $c_{as}$  for site  $s$ , we impute advertiser  $a$ ’s CTR at site  $s$  as

$$\bar{c}_{as} = \frac{\sum_{j \in \mathcal{A}_s \setminus a} r_{aj} c_{js}}{\sum_{j \in \mathcal{A}_s \setminus a} r_{aj}}. \quad (13)$$

In this counterfactual scenario, which we refer to as “pooling” and denote  $C_P$ , we set  $n_{as0}^I = 10^{12}$  and  $n_{as0}^C = \lfloor 10^{12} \bar{c}_{as} \rfloor$ , so that  $\mathbb{E}[\tilde{c}_{as0}/\gamma_a | n_{as0}^I, n_{as0}^C, \gamma_a, C_P] \approx \bar{c}_{as}/\gamma_a$ .

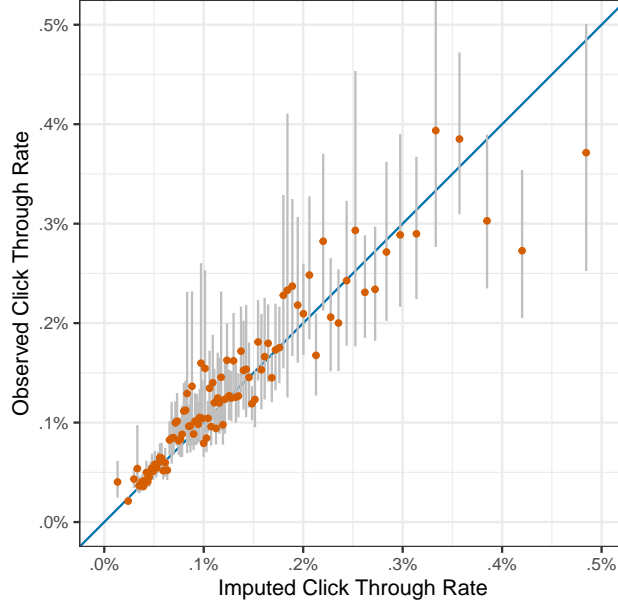
For the subset of advertiser-site pairs with observed ad buys, the imputed CTRs,  $\bar{c}_{as}$ , can be compared against the observed CTRs,  $c_{as}$ . Figure 9 shows this comparison, grouping imputed CTRs into bins of roughly equal size, and plotting the average observed CTRs corresponding with these advertiser-site pairs. The observations in the plots largely hew to a 45 degree line, meaning that the information from like advertisers can be used by the platform to predict initial click through rates.<sup>18</sup>

#### 5.1.4 Full Information Combined with Pooling

We also consider a third counterfactual that combines  $C_F$  and  $C_P$ . In this scenario, advertisers are endowed with full information about their CTRs at sites with observed purchases, as in  $C_F$ ; as well as imputed CTRs

<sup>17</sup><https://cloud.google.com/vision> (Archive.org)

<sup>18</sup>A linear regression of observed CTRs on imputed CTRs yields an intercept of .00014 (SE = .000037), a slope of .957 (.0244), and an  $R^2$  of .1436. The computation borrows observed CTRs from future outcomes in the data as an approximation to the knowledge the platform would have accumulated based on past experiences with campaigns across a large number of advertisers. A robustness check calculates expected CTRs using only impressions and clicks obtained in the six months prior to the estimation period, obtaining an intercept of .00018 (.000039), a slope of .945 (.0273), and an  $R^2$  of .1152. Thus, using past data only, imputed CTRs vary little in expectation, but are noisier, as one might expect. Given both approaches yield unbiased estimates of CTRs, CTRs are imputed based on the full data, as they are less noisy for the purposes of counterfactual evaluation.



**Figure 9:** Click Through Rate Imputation. The horizontal axis shows imputed CTRs for advertiser-publisher pairs observed in the estimation sample, grouped into bins of equal numbers of observations and truncated above the 99th percentile. The vertical axis shows the distribution of actual CTRs for each bin, with points indicating means and vertical lines indicating bootstrapped 95% CIs for the means. A 45 degree diagonal line is shown; observations closer to this line indicate more accurate imputations.

at all other sites, as in  $C_P$ . We refer to this as the “combined” scenario and denote it  $C_{F+P}$ . Comparing  $C_{F+P}$  to  $C_P$  provides information about the value of better match information (i.e., oracle versus approximate information about CTRs). Comparing  $C_{F+P}$  to  $C_F$  provides information about the value of expanding advertisers’ choice sets.

### 5.1.5 Combining Information Within and Across Advertisers

Integrating the respective approaches outlined in Sections 5.1.2 and 5.1.3 implies a fully crossed two-by-two counterfactual design, as indicated in Table 3. Cell  $B$  represents the base case, wherein the advertiser has no additional information about other sites, nor its imminent clicks. In the no sharing cells ( $B$  and  $C_F$ ), advertisers are constrained to choose only from the set of sites at which they advertised in the estimation data. Cell  $B$  involves naïve priors, uses observed CTRs to simulate clicks, and serves as the base case

		Own Information	
		No Information	Full Information
Shared Information	No Sharing	$B$	$C_F$
	Pooled Information	$C_P$	$C_{F+P}$

**Table 3:** Counterfactual Scenarios. The full information condition replaces the first periods’ advertiser CTR priors with the last periods’ oracle CTR posteriors. The shared information condition replaces the first period advertisers’ CTR priors with a similarity weighted average of other advertisers’ CTRs.

		Own Information			
		No Information		Full Information	
		Spend	Welfare	Spend	Welfare
Shared Information	No Sharing	\$0	\$0	−\$180 (−22.1%)	\$901 (3.5%)
	Pooled Information	\$671 (52.8%)	\$2,756 (15.5%)	\$546 (45.0%)	\$2,984 (16.6%)


**Table 4:** Advertiser Counterfactual Changes in Spend and Welfare. Each cell reports the median advertiser’s posterior mean counterfactual change in ad spend and welfare (relative to baseline), with the associated percentage increase reported in parentheses.

for measuring welfare gains from enhancing prior information. Cell  $C_F$  affords advertisers ex-ante full information on the clicks they obtained in the full data (i.e., their ex-post advertising outcomes), erasing their lack of prior information. Contrasting  $C_F$  with  $C_B$  yields theoretical insights into the cost advertisers bear owing to their lack of a priori knowledge about the sites on which they advertised.


Cell  $C_P$  pools advertiser information using the approach outlined in Section 5.1.3. Contrasting cell  $C_P$  to  $B$  yields the advertiser value of pooling information via the direct ad network platform (or any other information pooling entity), and expanding the advertisers’ choice set to those sites where the advertisers did not place ads in the data. Finally, cell  $C_{F+P}$  relative to  $B$  contrasts having information both from other advertisers (via pooling) and within advertiser (via full information). Other comparisons exist. For example, contrasting differences in  $C_P$  and  $B$  with differences in  $C_F$  and  $B$  yields a sense for which information sets generate greater improvements in advertiser outcomes.

## 5.2 Advertiser Welfare

Table 4 reports the results of the advertiser counterfactual outcomes in terms of relative gains compared to the baseline. Advertiser outcomes are computed using two metrics, change in ad spend, and change in welfare. Figure 10 reports, for each metric, the median and inner 50th percentile of advertisers’ posterior



**Figure 10:** Advertiser Counterfactual Changes in Spend and Welfare. Points and vertical lines indicate medians and inner 50th percentiles of advertisers’ posterior mean counterfactual changes. Panel (1) presents changes in ad spend relative to baseline across the three counterfactual conditions. Panel (2) reports changes in advertiser welfare.

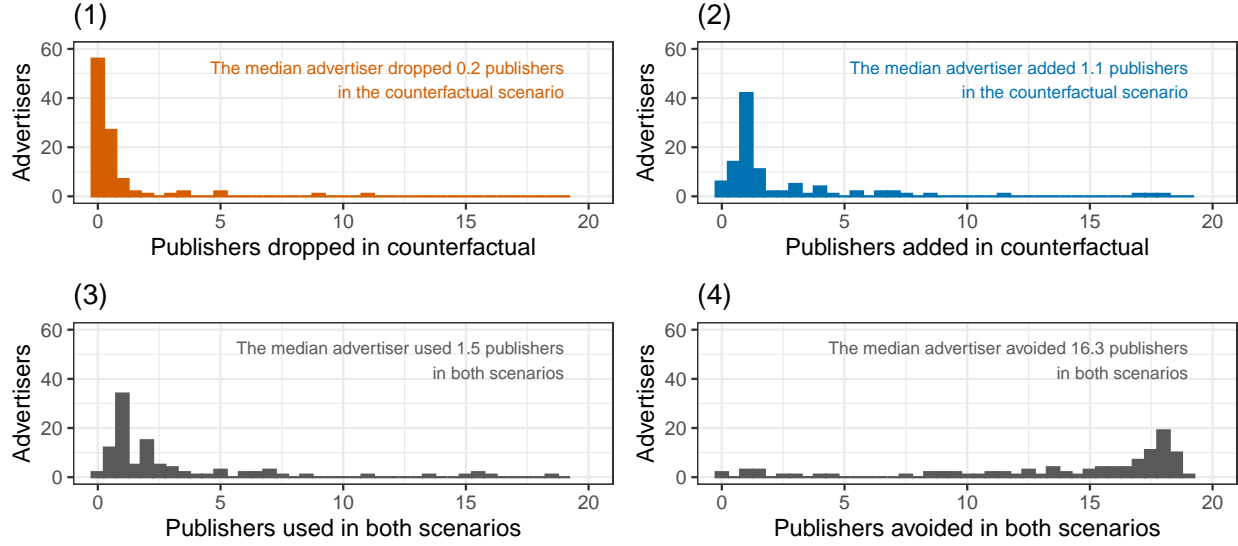


**Figure 11:** Change in Spend by Advertiser Under the Pooling Counterfactual. Inset shows the inner 60% of advertisers at higher resolution. Points indicate posterior means, horizontal bars are bootstrap 95% CIs for the means. Advertisers are sorted from lowest to highest mean change in spend. Vertical line indicates the median outcome.

mean counterfactual changes (relative to baseline) in each counterfactual cell in Table 3. Beginning with the full information comparison ( $C_F$  vs  $B$ ), results indicate that the median ad spend decreases by \$180 (22.1%). In  $C_F$ , the preponderance of advertiser priors tend to be updated negatively across the selected sites (i.e., sites with advertising in the data), hence advertisers spend less, or stop spending entirely at these sites. As advertisers' entire ad budgets can be shifted away from the lower match publishers in the set they ultimately selected, the median welfare increases by \$901 (3.5%).

Considering next the implications of pooled information ( $C_P$  vs  $B$ ), we observe a \$671 (52.8%) counterfactual increase in median spend. Unlike the  $C_F$  case, advertisers can shift spend to any publisher, not just those used in the data and baseline case. Hence pooling information enables advertisers to find better matches with new publishers, leading to an increase in ad spending in  $C_P$  compared to  $B$ . The median advertiser value increases \$2,756 (15.5%) under  $C_P$ . The final case,  $C_{F+P}$ , combines both within-advertiser oracle and pooled information. The combined case's results largely hew to the pooled information counterfactual, because the oracle priors apply only to a limited number of sites. Hence, the solutions across  $C_{F+P}$  and  $C_P$  are similar. The median advertising spend increases \$546 (45.0%) and the median advertiser welfare increases \$2,984 (16.6%). As the information requirements from the  $C_P$  case are closest to what could be implemented in practice, subsequent discussion focuses on the pooled information scenario,  $C_P$ .

Figure 11 reports the mean change in spend in the pooling counterfactual,  $C_P$ , for each of the 100



**Figure 12:** Change in Publisher Usage by Advertisers Under the Pooling Counterfactual. Each vertical bar represents the number of advertisers that evidenced a given level change in publisher usage from the baseline to the counterfactual condition. For example, the first bar in panel (1) indicates that nearly 60 advertisers dropped no publishers in the counterfactual. The third bar in panel (1) indicates that over 40 advertisers dropped one publisher.

advertisers. The figure suggests considerable heterogeneity in spending changes, presumably owing to considerable differences across advertisers in both information states and match values. Advertisers with higher levels of spending tend to show the greatest magnitude changes in counterfactual spend.

Figure 12 depicts changes in the number of publishers used by advertisers under the pooling counterfactual,  $C_P$ . If the high and incorrect priors dominate, then most advertisers will decrease the number of publishers used. If the assortive matching effect dominates, then advertisers will sort into using more sites. In this regard, the four quadrants of Figure 12 decompose the aggregate changes in the number of publishers used into partial effects, representing the number of publishers dropped or added in the counterfactual (panels 1 and 2), and the number used or avoided in both scenarios (panels 3 and 4). The lower right quadrant (panel 4) indicates that most sites not chosen by advertisers under  $B$  continue not to be chosen under  $C_P$ . For example, twenty of the 100 advertisers did not use 18 of the 20 publishers in both the baseline and counterfactual setting. Comparing the histogram in the upper left cell (panel 1) to the one in the upper right (panel 2) indicates that i) more new publishers were added by advertisers than were removed, as advertisers found good matches; and ii) pooling information induces considerable change in advertiser behavior. Overall, this figure suggests that, rather than concentrating on a single publisher or merely eliminating sites where prior beliefs were too high, advertisers seem to be sorting into better matches (that is, finding sites that generate higher advertising value).


		Own Information	
		No Information	Full Information
Shared Information	No Sharing	\$0	−\$4,684 (−21.0%)
	Pooled Information	\$9,618 (63.9%)	\$8,879 (46.3%)

**Table 5:** Publisher Counterfactual Revenue Outcomes

### 5.3 Publisher and Platform Welfare


Table 5 reports, and Figure 13 depicts how revenues change across the 20 publishers under the pooling counterfactual. The network collects a flat percentage of publisher revenues, thus these changes are proportional to the network’s gains and losses.

Two effects of information provision—that is, endowing advertisers with better priors—are possible. First, as advertisers are generally over-optimistic about the sites they chose in the estimation sample, demand at publisher sites could decrease and revenues could fall. Comparing  $C_F$  to  $B$  indicates that 18 of the 20 publishers lose revenue and that the median revenue loss is \$4,684 (21.0%), because the optimistic advertiser prior effect dominates under  $C_F$ . Advertisers realize that they were too optimistic, but because they are not provided information about better potential matches, they cut spending. In such a scenario, it is likely that the network and publishers would not seek to inform advertisers that their priors are too high, as the network and publisher sites would lose money. However, comparing  $C_P$  to  $B$  indicates that information pooling causes all publishers to gain revenue, with a median increase of \$9,618 (63.9%). As suggested by Figure 13, advertisers can better sort into a match across publisher sites, finding new options that generate greater value than the set of publishers selected in the baseline setting. The revenue from  $C_{F+P}$  is lower than  $C_P$  because the negative effect on spending from correcting advertisers’ optimistic priors slightly offsets the positive effect from better



**Figure 13:** Change in Publisher Revenue by Site Under the Pooling Counterfactual. Points indicate posterior means, horizontal bars are bootstrap 95% CIs for the means. Publishers are sorted from lowest to highest mean change in revenue. Vertical line indicates the median outcome.





**Figure 14:** Distribution of Advertisers Gained and Lost by Publishers Under the Pooling Counterfactual. Each vertical bar represents the number of publishers that evidenced a given level change in advertisers across the baseline and counterfactual condition. For example, the first bar in the left hand panel indicates that one publisher lost no advertisers. The second bar indicates that 3 publishers lost two advertisers.

matching. Under both  $C_P$  and  $C_{F+P}$ , the sorting effect dominates the prior effect, and the pooled information mechanism generates not only positive welfare outcomes for advertisers, but also for publishers.

Focusing again on the pooling counterfactual,  $C_P$ , we see that publishers not only gain revenue, but they also tend to serve more advertisers. Figure 14 reports the distribution (across publishers) of the number of advertisers lost (left panel) and gained (right panel) in  $C_P$  versus  $B$ . Comparing the two distributions, it is evident that more advertisers are gained than lost, as the gains from enhancing match (enabling advertisers to find higher value publishers) more than offset the losses from rectifying advertisers’ over-optimistic priors.

Collectively, the counterfactual results suggest that the ad network would share information to improve advertisers’ match information, because all publishers gain revenue under the pooling counterfactual and the network’s change in revenue is proportional to these gains. Moreover, the information needed to impute advertiser CTRs is readily available to the network so there would be few barriers to implementing such a system. Yet two considerations present. First, the network might be reticent to provide advertisers with imputed CTRs, as they might be construed as guarantees on ad performance. As an alternative, the network might choose to suggest a set of high-match sites. Second, although the largest publishers (i.e., the top 20) in the data set all see increases in revenue, the same might not be true for the smallest publishers. However, given the low baseline of revenues at these sites, there is a floor on how much these sites can be hurt by lower advertiser spend, but no ceiling on how much they can benefit from better sorting.

## 5.4 Counterfactual Summary

Based on the median publisher gain in revenue of \$9,618 and the median advertiser welfare gain of \$2,756, the combined welfare gains across the 100 advertisers and 20 publishers over the six-month window over our estimation data is on the order of \$468,000 under the pooling counterfactual. The median advertiser’s spending rises \$671. Extrapolating the spending increase from the random sample of 100 advertisers to the

population of 8,000 advertisers suggests the platform’s revenue could grow its total revenue on the order of \$5.4M (an analogous extrapolation across publishers is misleading, as our sample focuses on the 20 largest). To the extent a similar pooled approach could be used across multiple ad networks, or other ad channels such as retail media, gains would be larger.

As a caveat, it should be noted that our counterfactual analysis constitutes a partial equilibrium. Publisher sites, faced with increased demand, could raise subscription prices. Although, as Appendix B.1 indicates, little such evidence for price response is observed in the data. To the extent this does occur, advertiser gains could be somewhat diminished and publisher gains could be somewhat decreased. However, most gains in our analysis arise from better sorting, and there is little concentration in advertisers across sites (see Online Appendix B.2). As suggested by Figure 14, sites tend to face a change in advertiser composition owing to better sorting, rather than having all advertisers concentrate into a single site. To the extent sites are imperfect substitutes, this lack of concentration should offset the tendency of any dominant site to substantially raise prices. In sum, changes in advertising information could have downstream consequences for firms’ pricing strategies that offset some of their gains by increasing or decreasing price competition; these issues are beyond the scope of our research.

## **6 Conclusions**

This paper considers the implications of advertiser learning in direct display advertising markets. In contrast to exchange markets, direct advertising is bought in bulk (many thousands of impressions at a time) rather than sequentially, which limits advertisers’ opportunities to use test and learn strategies to efficiently identify publishers whose audiences are a good match for their ads (Tunuguntla and Hoban 2020). In direct markets, advertisers’ initial information states are economically consequential because, lacking the ability to test and learn incrementally, they are effectively guessing which publishers to use. Hence, there is a considerable potential to choose sites incorrectly and, as a corollary, to enhance advertiser (and possibly publisher) welfare by endowing advertisers with better information.

To ascertain whether advertisers have incorrect beliefs and assess the potential consequences thereof, we collect direct sales data from an ad exchange. The data describe ad purchases from the exchange’s inception, providing an ideal setting to explore the nature of advertiser learning as advertisers enter the network. Patterns in these data suggest evidence of advertiser learning. Advertisers initially run ads at relatively many sites and, after additional exploration, ultimately settle on a smaller number where they presumably find greater customer response. Using a direct utility model of advertiser choices of ad subscriptions across sites, in

which match is allowed to be learned from past advertising outcomes, we document that advertisers' initial priors are quite misinformed. Advertisers overestimate initial CTRs by a factor of nearly five, as median beliefs are 0.23%, compared to a median CTR of 0.045%. In short, advertisers are initially far too optimistic about the publishers they have chosen for advertising.

We consider various mechanisms to redress this problem and ascertain the impact of advertiser information on advertiser welfare and publisher revenues. We consider two such mechanisms. The first uses advertiser outcomes at the end of the data to inform their priors at the start of the data. While this suggests how much advertiser outcomes could potentially improve on the sites they actually explored, the approach cannot inform us about how well advertisers would fare on sites they did not explore. Moreover, the mechanism is infeasible, as it is impossible to know ad outcomes ex-ante. Hence, we consider a second mechanism wherein the platform can pool its information across advertiser on ad outcomes, drawing on similar advertisers who previously placed ads on the considered sites. We operationalize this information sharing mechanism using Google's Cloud Vision API to measure similarities between ads and advertisers, and then imputing CTRs for advertisers at new publisher sites. The approach is predictive of advertising performance.

Endowing advertisers with pooled information leads to a median advertiser increase of 52.8% (\$671) in ad spend, and a median advertiser welfare gain of 15.5% (\$2,756). On the publisher side, there exist two opposing forces: Rectifying optimistic priors should decrease ad spend and publisher revenue, while better sorting between advertisers and publishers should increase revenue. The latter effect dominates, generating a median publisher gain of 63.9% (\$9,618). All publishers gain in revenue. In other words, redressing advertiser uncertainty in direct ad networks leads to large welfare gains for publishers and advertisers alike. Extrapolating these gains across all 8,000 advertisers in our considered network suggests welfare gains up to \$5,400,000 over six months. The information pooling approach, if ported to other direct advertising networks, would yield even greater gains. Of note, direct ad purchasing also exists in non-digital channels as well, such as television, radio, and out of home. It stands to reason that advertisers also face a learning problem in those contexts, but the number of publisher alternatives in digital dwarfs these other channels, and likely exacerbates the learning problem and potential gains from redressing it.

Advertiser learning will become a more substantial research topic and practical opportunity in the coming years. Because advertisers are often ill-informed, poor advertiser decisions can induce advertisers to cease advertising, creating an incentive to shift decision rights from the advertiser to publishers and platforms. Examples of this shift already include platform-recommended ad delivery optimization, audience selection, and bidding, as well as generative AI to create ad content. In all these instances, platforms and publishers pool

information across advertisers to enhance advertisers' overall outcomes. As the ad ecosystem becomes more complex, the opportunities to pool information grow. In the face of cookie deprecation, these opportunities are especially relevant for first-party sites, including platforms and publishers such as this paper considers.

A related question of interest is why advertisers are ill-informed. In our context, their priors are far too optimistic. Possible reasons include advertisers' use of market research that lacks causal variation, and potentially overstated claims of ad effects by publishers and ad networks.<sup>19</sup>

Our analysis focuses on the advertiser demand side. An interesting potential extension is to consider how publisher ad pricing responds to changes in advertiser demand due to informing priors. To the extent that advertiser demand increases, publishers can raise ad prices and increase their share of total welfare gains at the cost of advertisers. That said, efficient sorting can reduce the concentration of advertisers on some sites and raise it on others, making the overall effect ambiguous and highly variable across publishers. Given that publishers face a large competitive set in the face of dynamic demand as a result of learning, the supply side problem would be a challenging, but useful extension in the analysis of direct display advertising markets. Another issue of interest is how to balance inventory and pricing across both direct and exchange markets (Balseiro et al. 2014), a challenge that is also exacerbated by advertiser learning. In sum, we hope this research sparks more interest on these and other topics in the large and economically consequential direct advertising market.

## References

- Ahmadi I, Abou Nabout N, Skiera B, Maleki E, Fladenhofer J (2023) Overwhelming targeting options: Selecting audience segments for online advertising. *International Journal of Research in Marketing* 41(1):24–40, ISSN 0167-8116.
- Alcobendas M, Zeithammer R (2021) Adjustment of bidding strategies after a switch to first-price rules. *SSRN Electronic Journal* ISSN 1556-5068.
- Allouah A, Besbes O (2017) Auctions in the online display advertising chain: A case for independent campaign management. *SSRN Electronic Journal* ISSN 1556-5068.
- Athey S, Ellison G (2011) Position auctions with consumer search. *The Quarterly Journal of Economics* 126(3):1213–1270, ISSN 1531-4650.
- Auer P, Cesa-Bianchi N, Fischer P (2002) Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47:235–256.

---

<sup>19</sup>The authors thank Nils Wernerfelt for sharing these observations.

- Bajari P, Benkard CL, Levin J (2007) Estimating dynamic models of imperfect competition. *Econometrica* 75(5):1331–1370, ISSN 1468-0262.
- Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61(4):864–884, ISSN 1526-5501.
- Balseiro SR, Candogan O (2017) Optimal contracts for intermediaries in online advertising. *Operations Research* 65(4):878–896, ISSN 1526-5463.
- Balseiro SR, Feldman J, Mirrokni V, Muthukrishnan S (2014) Yield optimization of display advertising with ad exchange. *Management Science* 60(12):2886–2907.
- Balseiro SR, Gur Y (2019) Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65(9):3952–3968, ISSN 1526-5501.
- Benkard CL (2000) Learning and forgetting: The dynamics of aircraft production. *American Economic Review* 90(4):1034–1054, ISSN 0002-8282.
- Berry ST (1992) Estimation of a model of entry in the airline industry. *Econometrica* 60(4):889, ISSN 0012-9682.
- Bhat CR (2008) The multiple discrete-continuous extreme value (mdcev) model: Role of utility function parameters, identification considerations, and model extensions. *Transportation Research Part B: Methodological* 42(3):274–303, ISSN 0191-2615.
- Bimpikis K, Elmaghraby WJ, Moon K, Zhang W (2020) Managing market thickness in online business-to-business markets. *Management Science* 66(12):5783–5822, ISSN 1526-5501.
- Bompaire M, Gilotte A, Heymann B (2021) Causal models for real time bidding with repeated user interactions. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 75–85.
- Cai H, Ren K, Zhang W, Malialis K, Wang J, Yu Y, Guo D (2017) Real-time bidding by reinforcement learning in display advertising. *Proceedings of the Tenth ACM International Conference on Web Search & Data Mining*, 661–670, WSDM 2017 (New York, NY, USA: Association for Computing Machinery), ISBN 9781450346757.
- Ching AT, Erdem T, Keane MP (2013) Learning models: An assessment of progress, challenges, and new developments. *Marketing Science* 32(6):913–938, ISSN 1526-548X.
- Choi H, Mela CF, Balseiro SR, Leary A (2020) Online display advertising markets: A literature review and future directions. *Information Systems Research* 31(2):556–575, ISSN 1526-5536.
- Choi WJ, Sayedi A (2019) Learning in online advertising. *Marketing Science* 38(4):584–608, ISSN 1526-548X.
- Dubé JP, Hitsch GJ, Manchanda P (2005) An empirical model of advertising dynamics. *Quantitative Marketing and Economics* 3(2):107–144, ISSN 1573-711X.
- Gope J, Jain SK (2017) A survey on solving cold start problem in recommender systems. *2017 International Conference on Computing, Communication and Automation (ICCCA)*, 133–138 (IEEE).

- Guan M, Jiang H (2018) Nonparametric stochastic contextual bandits. volume 32 (Association for the Advancement of Artificial Intelligence (AAAI)), ISSN 2159-5399.
- Hendel I, Spiegel Y (2014) Small steps for workers, a giant leap for productivity. *American Economic Journal: Applied Economics* 6(1):73–90, ISSN 1945-7790.
- Lam XN, Vu T, Le TD, Duong AD (2008) Addressing cold-start problem in recommendation systems. *Proceedings of the 2nd International Conference on Ubiquitous Information Management and Communication*, 208–211, ICUIMC08 (ACM).
- Lee S, Allenby GM (2014) Modeling indivisible demand. *Marketing Science* 33(3):364–381, ISSN 1526-548X.
- Levitt SD, List JA, Syverson C (2013) Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of Political Economy* 121(4):643–681, ISSN 1537-534X.
- Lika B, Kolomvatsos K, Hadjiefthymiades S (2014) Facing the cold start problem in recommender systems. *Expert Systems with Applications* 41(4):2065–2073, ISSN 0957-4174.
- Little JD, Lodish LM (1969) A media planning calculus. *Operations Research* 17(1):1–35, ISSN 1526-5463.
- Liu J, Toubia O (2018) A semantic approach for estimating consumer content preferences from online search queries. *Marketing Science* 37(6):930–952, ISSN 1526-548X.
- Perlich C, Dalessandro B, Hook R, Stitelman O, Raeder T, Provost F (2012) Bid optimizing and inventory scoring in targeted online advertising. *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 804–812, KDD 12 (ACM).
- Ren K, Zhang W, Chang K, Rong Y, Yu Y, Wang J (2018) Bidding machine: Learning to bid for directly optimizing profits in display advertising. *IEEE Transactions on Knowledge and Data Engineering* 30(4):645–659, ISSN 1041-4347.
- Schein AI, Popescul A, Ungar LH, Pennock DM (2002) Methods and metrics for cold-start recommendations. *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 253–260, SIGIR02 (New York, NY, USA: Association for Computing Machinery), ISBN 1581135610.
- Schwartz EM, Bradlow ET, Fader PS (2017) Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science* 36(4):500–522, ISSN 1526-548X.
- Scott SL (2010) A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry* 26(6):639–658, ISSN 1526-4025.
- Silge J, Robinson D (2016) tidytext: Text mining and analysis using tidy data principles in r. *The Journal of Open Source Software* 1(3):37, ISSN 2475-9066.
- Simpson D, Rue H, Riebler A, Martins TG, Sørbye SH (2017) Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science* 32(1), ISSN 0883-4237.
- Stan Development Team (2017) Rstan: the r interface to stan. r package version 2.16.2. <http://mc-stan.org>.

- Tadelis S, Hooton C, Manjeer U, Deisenroth D, Wernerfelt N, Dadson N, Greenbaum L (2023) Learning, sophistication, and the returns to advertising: Implications for differences in firm performance. Working Paper 31201, National Bureau of Economic Research.
- Theussl S, Hornik K (2023) *Rglpk: R/GNU Linear Programming Kit Interface*. <https://CRAN.R-project.org/package=Rglpk>, r package version 0.6-5.
- Toubia O, Netzer O (2017) Idea generation, creativity, and prototypicality. *Marketing Science* 36(1):1–20, ISSN 1526-548X.
- Tunuguntla S, Hoban PR (2020) A near-optimal bidding strategy for real-time display advertising auctions. *Journal of Marketing Research* 58(1):1–21, ISSN 1547-7193.
- Varian HR (2007) Position auctions. *International Journal of Industrial Organization* 25(6):1163–1178, ISSN 0167-7187.
- Vehtari A, Gelman A, Gabry J (2016) Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and Computing* 27(5):1413–1432, ISSN 1573-1375.
- Waisman C, Nair HS, Carrion C (2019) Online causal inference for advertising in real-time bidding auctions.
- Wu C (2015) Matching value and market design in online advertising networks: An empirical analysis. *Marketing Science* 34(6):906–921, ISSN 1526-548X.
- Xu B, Deng Y, Mela CF (2022) A scalable recommendation engine for new users and items. *SSRN Electronic Journal* ISSN 1556-5068.
- Yao Z, Kong D, Lu M, Bai X, Yang J, Xiong H (2023) Multi-view multi-task campaign embedding for cold-start conversion rate forecasting. *IEEE Transactions on Big Data* 9(1):280–293, ISSN 2372-2096.
- Zhou D, Li L, Gu Q (2020) Neural contextual bandits with ucb-based exploration. III HD, Singh A, eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 11492–11502 (PMLR), <https://proceedings.mlr.press/v119/zhou20a.html>.
- Zhuo R (2023) Exploit or explore? an empirical study of resource allocation in scientific labs. Technical report, Working Paper.

## Appendix A Identification

This section outlines how the parameters in the advertiser demand model are identified by different sources of variation in the data. First, following Lee and Allenby (2014), a two-parameter version of the advertiser payoff function in Equation (1), namely  $\delta \zeta_a^{-1} \log(1 + x_{asw} \zeta_a) - p_s(x_{asw})$ , is identified by repeated observations of ad subscription purchases lasting  $x$  days at a price of  $p(x)$ . Lee and Allenby (2014) show that  $\delta$  is primarily identified by the proportion of non-advertising choices, as a higher match parameter  $\delta$  increases the overall attractiveness of advertising. The satiation parameter  $\zeta_a$  is primarily identified by variation in observations with  $x > 0$ , in combination with non-linearity in pricing (longer subscriptions cost less per day than shorter subscriptions). A preponderance of shorter (longer) subscriptions in the data implies a larger (smaller) value of  $\zeta_a$ , as the higher the value of the satiation parameter  $\zeta_a$ , the sharper the decrease in marginal returns to more impressions.

Second, by observing choices from many advertisers at multiple sites over time,  $\delta$  can be factored multiplicatively into advertiser, site, and month fixed-effects, as given by Equation 2. Moreover, the fixed effects  $\phi_{\tau[asw]}$ , which are selected according to the number of weeks since an advertiser placed an ad at a site after the first purchase, are identified by differences in demand (common to all advertisers) along this advertiser-site time dimension.

Third, differences in demand at the advertiser-site level in the face of differences in the observed cumulative advertiser-site CTR allow separate identification of a dynamic match component, which we posit takes the form of a Bayesian learning model. In contrast with standard quality learning models in the literature (Ching et al. 2013), wherein the value of signals is latent, our empirical setting allows us to observe these signals in the data (in the form of impressions and clicks). One can, therefore, observe how advertisers' subscription purchase patterns change after obtaining a specific number of clicks and impressions, and this variation identifies  $\gamma_a$ .

Of note, Equations (1) and (2) include three parameters defined at the advertiser level. These are separately identified based on differences in advertiser demand along three dimensions of the data. The parameters for satiation ( $\zeta_a$ ) and initial CTR prior ( $\gamma_a$ ) are separately identified because  $\zeta_a$  applies in all periods, whereas  $\gamma_a$  (which cancels out of the payoff function when  $n^I = 0$ ) only applies to choices after using a site for the first time. A similar argument allows separate identification of  $\gamma_a$  and the advertiser fixed effect ( $\xi_a$ ). Finally, separate identification of  $\xi_a$  from  $\zeta_a$  is due to variation in choices  $x$  at different prices, as previously described



in the context of a two-parameter payoff function (i.e., replacing  $\delta$  with  $\xi_a$  in the earlier description).

## Appendix B Simulation Procedure

All counterfactual scenarios share a common simulation procedure based on 100 draws from the posterior distribution of the model parameters,  $\theta$ . For each draw of  $\theta$ ,  $\varepsilon$  is sampled from the posterior predictive distribution of  $\varepsilon|\theta$ , and the same  $\varepsilon$  vector is used for each baseline/counterfactual pair (each counterfactual is paired with a unique baseline simulation). Most simulated choices by advertiser  $a$  at site  $s$  in week  $w$  have a corresponding observed choice; for these simulated choices,  $\varepsilon_{asw}$  is sampled from a distribution truncated by the upper and lower bounds defined in Equations (8) and (9), conditional on  $\theta$ . If there is not a corresponding choice in the data (e.g., a subscription lasting  $x > 7$  days starting in week  $w$  at site  $s$  prevents a choice at site  $s$  in week  $w + 1$ )—then  $\varepsilon_{asw}$  is sampled from an unrestricted distribution. We average simulated choices over draws of  $\theta$  to obtain posterior predictive means for the outcomes of interest.

Using the set of available ad durations,  $x$ , and prices,  $p$ , from the estimation sample, purchases for advertiser  $a$  are simulated starting with week  $w = 1$  (or the first week the advertiser was observed in the network, if the advertiser joined during the estimation period). For each site  $s$ , advertiser  $a$  chooses an optimal number of days of advertising,  $x$ . If  $x = 0$ , advertiser  $a$  does not place an ad at site  $s$  in week  $w$ . Otherwise,  $x > 0$  and  $n^I = x \cdot t_s$  impressions are delivered, yielding  $n^C | n^I, c_{as} \sim \text{Binomial}(n^I, c_{as})$  clicks. If  $0 \leq x \leq 7$ , then the next advertising choice for site  $s$  occurs in week  $w + 1$ ; otherwise the next choice occurs in a later week, depending on the value of  $x$ . The procedure repeats for each site, and then advances to the next week. Choices in subsequent weeks are informed by Bayesian updating on  $\tilde{c}_{as}$  based on the cumulative number of impressions,  $n^I$ , and clicks,  $n^C$ , obtained from previous ad buys.

The outcomes of interest from these simulations are i) the amount each advertiser spends in total or at each site, which is obtained directly from the simulations; and ii) a measure of advertiser welfare. Regarding the latter: because advertisers choose sites based on net expected payoff (Equation (1)), expected net payoffs are transformed into a “true” net payoff by replacing the simulated value of  $\mathbb{E}[\tilde{c}_{asw} | n^I_{asw}, n^C_{asw}]$  with  $\mathbb{E}[\tilde{c}_{asw} | n^I = 10^{12}, n^C = \lfloor 10^{12} c_{as} \rfloor]$ . We use  $\mathbb{E}[\tilde{c}_{asw} | n^I = 10^{12}, n^C = \lfloor 10^{12} c_{as} \rfloor]$  rather than  $c_{asw}$  because the procedure for counterfactually manipulating advertisers’ prior beliefs entails altering their state space ( $n^I$  and  $n^C$ ) in a similar manner. Using  $c_{asw}$  instead would introduce simulation error due to differences in floating point precision for large numbers.<sup>20</sup>

<sup>20</sup>Owing to path dependencies in the simulation, it is possible for choices in the baseline and counterfactual with common values of  $a$ ,  $s$ , and  $w$ , to depend on different time-varying fixed effects,  $\phi_{\tau[a,s,w]}$ . In some cases, this leads to large differences in welfare that are not meaningfully related to the most important differences between the baseline and counterfactual scenarios. Hence, in these



**Figure 15:** Example Of An Ad Image In The Dataset.

All simulations incorporate the advertiser participation constraint. Economic profits must be positive (i.e.,  $\pi_{asw}(x) > 0$ ), else advertisers would not advertise. For example, should match become sufficiently low in the face of new information, an advertiser will cease to advertise.

## Appendix C Measuring Advertising Similarity

To tag ad images, we use Google’s Cloud Vision API, a pre-trained machine learning model developed by Google that ingests images as input and returns concept tags as output. Using this product, we query the set of 10 tags that best characterize each ad image in our dataset.<sup>21</sup> Figure 15 depicts an example ad from the dataset, and Google’s Cloud Vision API retrieves the following tags for this image: television program, television presenter, tie, news, font, display device, event, cable television, electric blue and public speaking.

Using the ad image tags returned by Google’s Cloud Vision API, we construct a TF-IDF matrix of advertisers by tags (Silge and Robinson 2016).<sup>22</sup> For the purpose of computing advertiser similarity distinctive characterizations of an image are most informative. Hence, TF-IDF is an attractive metric, as it upweights ad tags that are relatively unique, and down-weights commonly appearing tags.

Figure 16 reports the top 10 TF-IDF tags for the four advertisers appearing most frequently in the data. The most frequent advertiser (Advertiser A) is primarily characterized by tags related to food. Similarly, the third most frequent advertiser (Advertiser C) purveys products related to plants. In this sense, the TF-IDF describes advertisers based on the ads shown in their campaigns.


To measure the similarity between advertisers based on their ad images tags, we compute the cosine similarity of their TF-IDF vectors. The cosine similarity between two vectors  $a$  and  $b$  is  $f(a,b) = \frac{a \cdot b}{\|a\| \|b\|} = \frac{\sum_k a_k b_k}{\sqrt{\sum_k a_k^2} \sqrt{\sum_k b_k^2}}$  and ranges from 0 to 1: it is 0 when the ad vectors are orthogonal, and 1 when the ad vectors are identical.<sup>23</sup> Stacking the vectors of advertisers’ cosine similarities with other advertisers yields a square

cases, the baseline and counterfactual  $\phi$ s are replaced with their simple average.

<sup>21</sup>A robustness check uses 20 tags (when available from the API) and results are qualitatively the same.

<sup>22</sup>For a marketing application of TF-IDF, see Toubia and Netzer (2017).

<sup>23</sup>See Liu and Toubia (2018) for an application of cosine similarity in marketing.



**Figure 16:** Top 10 Tags For The Top 4 Advertisers. Advertiser A placed the most ads in the data set, Advertiser D the fourth most. Depicts the ten tags with the highest TF-IDF for each advertiser.

and symmetric matrix  $R$  with 1 in the main diagonal and entries  $r_{aj}$  containing the cosine similarity between advertiser  $a$  and advertiser  $j$ . The ad copy similarity matrix  $R$  is then used to generate a weighted prediction of CTRs for advertisers on sites where they have not previously advertised, per Equation (13), under the assumption that similar items have similar click through rates on a site.

# Online Appendix

## Online Appendix A Test and Learn

Recently, Zhuo (2023) developed an upper confidence bound (UCB) approach to address test and learn in dynamic models of choice, and the UCB approach has been shown to achieve nearly optimal regret relative to the use of exact continuation values (Auer et al. 2002, Guan and Jiang 2018, Zhou et al. 2020). This section discusses the implementation of the approach and the associated findings.

### A.1 Specification

Consider the choice of subscription length of  $x$  days made by advertiser  $a$  at site  $s$  in week  $w$ . Because choices of the optimal  $x$  at each site  $s$  are independent, the discussion that follows suppresses both the  $a$  and  $s$  subscripts. Denoting  $\bar{c} \equiv n^C/n^I$ , the period (myopic) objective function used in the paper (given in Equations (6) and (7)) can be written as

$$V(x; \varepsilon, n^I, \bar{c}, w) = \frac{1 + \bar{c}n^I}{1 + \gamma n^I} \exp(\xi + \eta + \phi_{\tau[w]} + \psi_{m[w]}) \cdot \exp(\varepsilon) \zeta^{-1} \log(1 + tx\zeta) - p_w(x).$$

The Bellman equation for the corresponding dynamic problem would be

$$W(\varepsilon, n^I, \bar{c}, w) = \sup_{x \in \mathcal{X}_w} V(x; \varepsilon, n^I, \bar{c}, w) + \beta \int W(\varepsilon', n'', \bar{c}', w+1) dF(\varepsilon', n'', \bar{c}' | x, \varepsilon, n^I, \bar{c}, w)$$

where  $\beta$  is the weekly discount rate, and  $\mathcal{X}_w$  is the set of available subscription lengths in week  $w$  (if a previously purchased subscription is still running in week  $w$ , then  $\mathcal{X}_w = \{0\}$ ). Because the  $\varepsilon$ s are i.i.d. across weeks and Bayes rule implies  $\mathbb{E}[\bar{c}' | n''] = \bar{c}$ , the Bellman equation can be simplified to

$$W(\varepsilon, n^I, \bar{c}, w) = \sup_{x \in \mathcal{X}_w} V(x; \varepsilon, n^I, \bar{c}, w) + \beta \int \mathbb{E}_{\varepsilon'} [W(\varepsilon', n'', \bar{c}, w+1)] dF(n'' | x, n^I),$$

which shows the link between forward-looking behavior and the choice of  $x$ : For larger values of  $x$  (i.e. more days of advertising),  $n^I$  increases (in expectation by  $tx$ , where  $t$  is expected daily traffic), increasing the available information about match. Hence, if there is a discrepancy between choices under myopic and forward-looking behavior, it is likely due to a trade-off between a higher value of  $x$  that produces a less favorable  $V(x)$  in the current period, but which maximizes the Bellman equation due to its positive influence on  $n^I$ .<sup>24</sup> The corresponding choice-specific value function for this Bellman equation is

$$v(x; \varepsilon, n^I, \bar{c}, w) = V(x; \varepsilon, n^I, \bar{c}, w) + \beta \int \mathbb{E}_{\varepsilon'} [W(\varepsilon', n'', \bar{c}, w+1)] dF(n'' | x, n^I)$$

<sup>24</sup>The evolution of  $w$  is also part of the Bellman recursion, as it determines the values of the time-varying fixed effects  $\phi_{\tau[w]}$  and  $\psi_{m[w]}$ . We do not consider these mean 0, time-varying control variables to be paramount for understanding dynamic behavior, and focus on the cumulative number of impressions.

Zhuo (2023) proposes using the UCB approach (Auer et al. 2002) to approximate the value function, because UCB-like indices achieve a nearly optimal regret guarantees compared to the optimal exact solution of the Bellman equation (Guan and Jiang 2018, Zhou et al. 2020). Applying Zhuo’s approach to approximating the UCB continuation value yields

$$\beta \int \mathbb{E}_{\epsilon'} [W(\epsilon', n^I, \bar{c}, w + 1)] dF(n^I | x, n^I) = 1(x > 0) \frac{\lambda}{\sqrt{j}}$$

where  $1(x > 0)$  indicates that a subscription is purchased, and thus new impressions will arrive,  $\lambda > 0$  is a parameter to be estimated, and  $j$  represents the number of signals received. In our case, ad impressions are signals, hence the choice-specific value function becomes<sup>25</sup>

$$v(x; \epsilon, n^I, \bar{c}, w) = V(x; \epsilon, n^I, \bar{c}, w) + 1(x > 0) \frac{\lambda}{\sqrt{1 + n^I}}.$$

The intuition behind this approach is that sites with no prior advertising at a site receive an increase in continuation value when  $n^I$  is small, encouraging forward-looking advertisers to explore new sites. The parameter  $\lambda$  indicates the relative importance of the continuation value; a large estimated value of  $\lambda$  suggests the data are consistent with forward-looking behavior. Variation in CTRs identify the match parameter, whereas the test and learn parameter is identified from variation in accumulated impressions, leading to a different likelihood of subscription purchase.

## A.2 Estimation

Conditional on the state variable for the cumulative number of impressions,  $n^I$ , the approximate UCB continuation value takes on one of two possible values. If  $x = 0$ , then the continuation value is 0, and if  $x > 0$ , then the continuation value is  $\lambda / \sqrt{1 + n^I}$ . The implications of this for estimation are the following.

First, when a value of  $x = 0$  is observed in the data, the upper bound on  $\epsilon$  given in Equation (9) changes to

$$ub^\theta(x, p) \equiv \log \left( \uparrow p - \frac{\lambda}{\sqrt{1 + n^I}} - p \right) - \log \left[ \check{\mu}_{asw} \zeta_a^{-1} \log \left( \frac{t_s \uparrow x \zeta_a + 1}{t_s x \zeta_a + 1} \right) \right],$$

and because  $p = 0$  and  $x = 0$  in this case, the above can be further simplified as

$$ub^\theta(x, p) \equiv \log \left( \uparrow p - \frac{\lambda}{\sqrt{1 + n^I}} \right) - \log \left[ \check{\mu}_{asw} \zeta_a^{-1} \log \left( \frac{t_s \uparrow x \zeta_a + 1}{1} \right) \right].$$

One can interpret  $-\frac{\lambda}{\sqrt{1 + n^I}}$  as an amount by which the price of the shortest subscription is effectively decreased (relative to the choice not to advertise) under forward-looking behavior. When  $n^I = 0$  (i.e. an advertiser has yet to try a site), this price decrease is equal to the monetary value of the time-discounted, net benefits

<sup>25</sup>Following Zhuo (2023),  $j = 1$  when no signals are observed (i.e.,  $n^I = 0$ ). In our context, this would be essentially no information because  $j = 1$  is the equivalent of receiving one impression (that is, a very small fraction of the typical 821K impressions in a one-week subscription). Setting  $j = 0$  is impracticable because it would imply an infinite continuation value.

from future advertising at this site due to having better information. By effectively lowering the price of the forgone ad subscription when observing  $x = 0$ , the UCB term lowers, compared to the myopic case, the upper bound on the range of  $\varepsilon$ s that can rationalize  $x = 0$ . As the advertiser grows more experienced advertising at the site,  $n^I$  will be higher, the value of future information will be diminished (because the value of reducing uncertainty is diminished), and the effective decrease in the price of advertising will eventually go to zero. Second, a similar logic applies when the observed  $x > 0$  is the smallest available, non-zero subscription length (i.e.,  $x = \inf \mathcal{X}_w \setminus \{0\}$ ). In this case (noting that  $\downarrow p = 0$  and  $\downarrow x = 0$ ), the lower bound on  $\varepsilon$  given in Equation (8) changes to

$$\ell b^\theta(x, p) \equiv \log \left( p - \frac{\lambda}{\sqrt{1 + n^I}} \right) - \log \left[ \check{\mu}_{asw} \zeta_a^{-1} \log \left( \frac{t_s x \zeta_a + 1}{1} \right) \right].$$

The interpretation of the UCB term as a price reduction applies here as well. Given  $x > 0$  was observed instead of  $x = 0$ , and in light of the information gains from advertising, the lower bound of the  $\varepsilon$  that rationalizes  $x > 0$  can be even lower than in the myopic case. Third, in all other cases, the UCB term differences out of the expressions for the upper and lower bounds, so these bounds are unchanged compared to the myopic case.

### A.3 Results

Adopting this approach, several findings emerge. First, the inclusion of test and learn behaviors into our model does not impact its overall predictive performance, with  $ELPD_{LOO}$  improving by .1, but with a standard error of this difference equal to .8. Posterior estimates of  $\gamma$  (prior beliefs on CTRs) are essentially identical to the null model of no forward-looking behavior, with the p-value from a two-sided Kolmogorov-Smirnov test equal to .994. A WLS regression ( $\hat{\gamma}_a^{UCB} \sim 1 + \hat{\gamma}_a$ ) of the 100 advertiser prior belief estimates in the UCB model,  $\hat{\gamma}_a^{UCB}$ , on the 100 prior beliefs in the non-UCB model,  $\hat{\gamma}_a$  (where weights are the inverse squared standard errors of the UCB prior belief estimates,  $(\hat{\sigma}^{UCB})^{-2}$ ), yields an intercept of  $-.0000000353$  (.000000151) and slope of .997 (.00415) with an  $R^2$  of .998. Hence, neither the intercept nor slope is significantly different from 0 or 1, respectively. Third, the  $\lambda$  parameter in the UCB approximation to the value function is estimated to be  $\hat{\lambda} = 2.12$  with a standard deviation of 2.06. Owing to the dollar metric for the advertiser utility function, the UCB approximation to the value function can be interpreted as a \$2.12 (\$2.06) incentive to use test and learn on new sites (against an average subscription price of \$1226). Thus, advertisers in our data place relatively little value on exploration (i.e., their willingness to pay for information is roughly 0.2% of the purchase price).

## Online Appendix B Supply-Side Considerations

### B.1 Pricing

It is possible that aggregate demand could affect publisher site prices in the counterfactual setting, even if there is essentially no price variation over time in the estimation data. To address this issue, we conduct a cross-sectional regression of the publisher fixed-effects from the price regression on aggregate advertiser demand levels (i.e., average weekly number of advertiser subscriptions at the publisher’s site). We use a cross-sectional WLS regression of the week fixed effects ( $\mu_w$  in Equation (11)) on advertiser demand (the weekly number of advertisers with ads running at each site), as there is essentially no publisher variation in price over time. Weights in the regression are inverse standard errors for estimated  $\mu_w$ ’s, and implicitly (via the cross-sectional data structure), the number of sites with any ads running each week. Table 1 reports the results, which indicate no significant relationship between the weekly number of advertisers on a site and the publisher fixed-effects.


Similarly, when adding the previous week’s cumulative CTR and number of advertisers to the price regression (i.e., equation (11) is changed to  $\mathbb{E}[\log(p_{sw}(x))] = \mu_w + \mu_s + \nu \log(x) + \omega_A A_{s,w-1} + \omega_C C_{s,w-1}$ , with  $A_{sw}$  the number of advertisers with ads running at site  $s$  in week  $w$ , and  $C_{sw}$  the cumulative CTR for all ads run at site  $C_{sw}$  as of week  $w$ ), we do not find significant coefficients for either the number of advertisers in the previous week ( $\omega_A = -.0079$  (.0078),  $p = .313$ ), nor the previous week’s cumulative site-level CTR ( $\omega_C = 3.91$  (4.24),  $p = .357$ ). This is likely because publisher prices vary little over time.

### B.2 Crowding

To address the effect of crowding, we conduct two analyses. First, we report the number of instances across sites and weeks with multiple subscriptions in both estimation and the counterfactual. As depicted in Figure 1, in most instances only one or two advertisers are present at any given point in time at a given publisher site.

DV: Site fixed effects, $\hat{\mu}_w$	Estimate	Std. error	t-statistic	p-value
Intercept	2.63	.00362	725.48	.00
Number of advertisers	.000822	.00154	.53	.59
$R^2$	.000127			
Observations	2,243			

**Table 1:** The Effect of The Number of Advertisers on a Site on Log Price Regression’s Site Fixed-Effects. The table reports the results of a weighted least squares regression of site fixed effects on the average weekly number of advertisers’ with active subscriptions at each site (where weights are the squared inverses of the fixed-effects’ standard errors). The site fixed-effects (i.e., the dependent measures) are obtained from the log price regression of prices on the site fixed-effects and other control variables as indicated in Equation (11).



**Figure 1:** The Number of Advertisers On Each Site. The horizontal axis lists the 20 publishers and the vertical axis is a count of the number of advertisers on the publisher site. Box plots show the distribution of the number of advertisers over the weeks in the data in the baseline and observed counterfactuals, along with the observed data. Boxes and horizontal lines show the 25th, 50th, and 75th percentiles. Lines extend to 1.5 times the inner-quartile range, and observations outside this range are shown as dots.

Hence, crowding is limited and not likely to have a large effect on ad effectiveness as defined by clicks.

Second, we regress advertiser clicks by site-week on the number of advertisers on the site, controlling for site fixed effects and advertiser-specific click propensity. The latter is defined as the sum of the average number of clicks over time received by the advertiser across sites each week. More specifically, the regression is  $q_{sw} \sim A_{sw} + F_s + \bar{q}_{sw}$ , where  $q_{sw} = \sum_a (n_{asw}^C - n_{asw-1}^C)$  is the number of incremental clicks across all advertisers with ads running at site  $s$  in week  $w$ ,  $A_{sw}$  is the number of advertisers with ads running at site  $s$  in week  $w$ , and  $F_s$  is a fixed effect for site  $s$ ;  $\bar{q}_{sw}$  is the average of  $\bar{q}_a$  for advertisers with ads running at site  $w$  in week  $w$ , and  $\bar{q}_a$  is advertiser  $a$ 's average weekly incremental clicks per site. As reported in Table 2, the effect of number of advertisers on a site and the clicks that the advertiser receives is not significant. Overall, little evidence of crowding exists in our context.

## Online Appendix C Additional Analyses

### C.1 Sample Selection

To ascertain the effect of increasing the number of publishers, we expand the number of sites by 25% from 20 to 25 and repeat the information pooling counterfactual analyses in Section 5. The median increase in publisher revenue rises from 63.9% to 84.3%. The median advertiser spends 113.8% more (compared to



DV: Weekly site-level clicks	Estimate	Std. error	t-statistic	p-value
Intercept	−311.47	173.07	−1.80	.07
Sum of average clicks among active advertisers	1.05	.06	16.81	.00
Number of active advertisers	15.09	31.02	.49	.63
Site fixed effects	yes			
$R^2$	.7461			
Observations	504			

**Table 2:** The Effect of Advertiser Crowding at a Publisher Site on Advertiser Clicks. The table reports a regression of advertiser clicks by site-week on the number of advertisers, controlling for site fixed effects and advertiser-specific average clicks across sites.

52.8% with a smaller set of sites).

## C.2 Budget Constraints

We consider a counterfactual simulation in which advertisers are constrained to a particular budget. For these simulations, rather than choosing subscription lengths for each site in each week independently, advertiser  $a$  solves an optimization problem constrained by the remaining budget in week  $w$ ,  $b_{aw}$ :  $\max_{\vec{x}} \sum_s \pi_{asw}(\vec{x}_s)$  with  $\vec{x}_s \in \mathcal{X}_{sw}$ , subject to  $\sum_s p_{sw}(\vec{x}_s) \leq b_{aw}$ . The GNU Linear Programming Kit, called via the Rglpk R package interface, performs the optimization quickly and efficiently (Theussl and Hornik 2023). When we constrain advertisers to spend no more than their observed budget in the data, the median advertiser’s spend drops by 9.3% and only 18% of advertisers increase their spend under the counterfactual (in these cases, both baseline and counterfactual spend fall below the assumed budget constraint); moreover, the median change in welfare is 0%. Because overall ad spend drops, the platform has no incentive to pool information. When budgets are allowed to increase to as much as 150% of the observed budget, 42% of advertisers increase their spend over their observed budget. And as the budget constraint rises to 300% of spend, about 55% of advertisers increase their spend and the median advertiser’s spend and welfare increase by 7.1% and 4.4% respectively. Hence total platform revenues increase when advertisers are allowed to spend as much as 300% of their observed budget and, at this point, the platform has a positive incentive to pool information. Notably, although the budget constraint is set to 300% of the observed budget, the average counterfactual spend among advertisers who increase spending is only 165% of the observed budget.

## C.3 Heterogeneity in Priors Within Advertiser and Across Publisher

It is possible for the advertiser to have different beliefs about match at the publisher level—that is, to have different prior beliefs across sites about the distribution of unknown factors affecting match. Relaxing this assumption requires the alternative assumption that advertisers are aware of and consider how match differs across all potential available sites, which would be an informationally demanding task. To explore whether

and how the assumption of common prior beliefs across sites affects inference, we estimate a model wherein advertisers have site-specific CTR priors:  $\gamma_{as} = g_{as} / (\bar{\gamma}^{-1} + g_{as})$ , with  $g_{as} \sim \text{Exponential}(1)$ . In the main model specification, there are 100 advertisers, hence there are 100 parameters representing the  $\gamma_{as}$ . In the augmented model, because there are 20 sites, we have 2000 parameters representing  $\gamma_{as}$ . The following results obtain: First, model fit is marginally improved, with  $ELPD_{LOO}$  decreasing by 20 points (the SE of the difference is 6.9) due to adding 1900 additional parameters (for comparison in terms of the model complexity and fit tradeoff, this gain can be contrasted to an improvement of 150 points when adding just 17 time-varying fixed effects). The two models recover similar prior beliefs, with a correlation in advertiser-level point estimates of .993 (i.e., a correlation between  $\hat{\gamma}_a$  in the main specification, and  $\hat{\gamma}_a = \frac{1}{20} \sum_s \hat{\gamma}_{as}$  in the augmented specification). A variance decomposition of the point estimates for the  $\gamma_{as}$  in the augmented model shows that 92% of the variance of the advertiser-site priors is explained by advertisers alone. In sum, this analysis largely supports our assumption that prior beliefs differ predominantly by advertiser and adding publisher heterogeneity has a largely inconsequential effect on model performance and inference.

#### C.4 Estimating Learning Rates (Prior Variances)

As an extension to the proposed learning model, one could incorporate a second parameter,  $\kappa$ , to determine the prior variance of  $\tilde{c}$ , and thus, the rate of advertiser learning. Prior beliefs in this model are specified as  $\tilde{c} | \kappa, \gamma \sim \text{Beta}(\kappa, \kappa(1 - \gamma) / \gamma)$ , leading to an updated expected CTR of  $\gamma \frac{\kappa + n^C}{\kappa + \gamma m^I}$ . The value of  $\kappa$ , however, drops out of the expectation prior to advertising (i.e., the prior expectation is still  $\gamma$  in this model), and  $\kappa$ 's effect on the expectation is likely to be dominated by  $n^C$  and  $n^I$  thereafter. Hence, the parameter  $\kappa$  is weakly identified in this setting, so it is normalized it to one in our main specification. As a robustness check, this section relaxes that assumption.

We estimate a model in which the prior variance parameter ( $\kappa$ ) is estimated rather than set to 1. Several findings emerge. First, the posterior mean is 2.72, and the inner 90% interval is [1.05, 6.78] (note: we constrain  $\kappa \geq 1$ ). Considering the prior mean for  $\kappa$  is 10, and the posterior differs from this prior mean, the prior variance parameter is indeed identified. However, the  $ELPD_{LOO}$  for this model is not better than the model in which  $\kappa$  is set to 1 (worsening by 3.4, with a standard error of 1.4). Second, the posterior means for the  $\gamma$ s across the two models are nearly identical. Using the same weighted regression approach described in the discussion of the  $\gamma$ 's under the UCB model, we regress the posterior mean  $\gamma$ 's from the model with  $\kappa > 1$  on those from the model with  $\kappa = 1$  and obtain an intercept of .00000145 (.000000294) and a slope of .985 (.00475), with an  $R^2$  of .998. In summary, adding  $\kappa$  to the model neither improves fit, nor affects inferences on  $\gamma$ .