



Contents lists available at ScienceDirect

International Journal of Research in Marketing

journal homepage: www.elsevier.com/locate/ijresmar

Overwhelming targeting options: Selecting audience segments for online advertising

Iman Ahmadi^a, Nadia Abou Nabout^{b,*}, Bernd Skiera^{c,d}, Elham Maleki^e, Johannes Fladenhofer^f

^a Assistant Professor of Marketing, Warwick Business School, University of Warwick, UK

^b Chaired Professor of Interactive Marketing & Social Media, WU Vienna, Austria

^c Chaired Professor of Electronic Commerce, Faculty of Economics and Business Administration, Goethe University Frankfurt, Germany

^d Professorial Research Fellow at Deakin Business School, 221 Burwood Highway, Burwood, VIC 3125, Australia

^e Senior Decision Scientist at Facebook, 9 Azure House, London NW10 7FE, UK

^f Founder and CEO, Netwings, Raiffeisenplatz 1, 4863 Seewalchen, Austria

ARTICLE INFO

Article history:

Received 2 June 2022

Available online xxxx

Keywords:

Targeting

Audience Segments

Online Advertising

Third-Party Data

Facebook

Spotify

Apple

App Tracking Transparency Framework

ABSTRACT

Even as online advertising continues to grow, a central question remains: Who to target? Yet, advertisers know little about how to select from the hundreds of audience segments for targeting (and combinations thereof) for a profitable online advertising campaign. Utilizing insights from a field experiment on Facebook (Study 1), we develop a model that helps advertisers solve the cold-start problem of selecting audience segments for targeting. Our model enables advertisers to calculate the break-even performance of an audience segment to make a targeted ad campaign at least as profitable as an untargeted one. Advertisers can use this novel model to decide whether to test specific audience segments in their campaigns (e.g., in randomized controlled trials). We apply our model to data from the Spotify ad platform to study the profitability of different audience segments (Study 2). Approximately half of those audience segments require the click-through rate to double compared to an untargeted campaign, which is unrealistically high for most ad campaigns. Our model also shows that narrow segments require a lift that is likely not attainable, specifically when the data quality of these segments is poor. We confirm this theoretical finding in an empirical study (Study 3): A decrease in data quality due to Apple's introduction of the App Tracking Transparency (ATT) framework more negatively affects the click-through rate of narrow (versus broad) audience segments.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The ongoing growth of online advertising (PwC, 2023) reached a new milestone at the end of 2018 when global digital ad spending surpassed global television ad spending for the first time (Bayer, Srinivasan, Riedl, & Skiera, 2020). For firms, the primary appeal of online ads is their capacity to target users more strategically (e.g., based on user demographics and online behavior). Targeted ads are touted as being more effective than untargeted ads—and indeed, recent research provides evidence that users are more likely to show interest in and click on targeted ads compared to untargeted versions (Goldfarb & Tucker, 2011b; Tucker, 2014; Yan et al., 2009). Consequently, many researchers have taken a keen interest in targeting

* Corresponding author.

E-mail addresses: iman.ahmadi@wbs.ac.uk (I. Ahmadi), nadia.abounabou@wu.ac.at (N. Abou Nabout), skiera@wiwi.uni-frankfurt.de, bernd.skiera@deakin.edu.au (B. Skiera), elham.maleki91@gmail.com (E. Maleki), johannes.fladenhofer@gmx.at (J. Fladenhofer).

<https://doi.org/10.1016/j.ijresmar.2023.08.004>

0167-8116/© 2023 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

(Choi & Mela, 2019; Goldfarb & Tucker, 2011a, 2011b; G. A. Johnson, Shriver, & Du, 2020; Lambrecht & Tucker, 2013; Rafeian & Yoganarasimhan, 2021).

Given the advantages of targeting, companies such as Google, Facebook/Instagram, Twitter, and Spotify have built digital advertising ecosystems that provide advertisers with self-service platforms to purchase and test hundreds of audience segments. Also known as “audience lists,” “user lists,” or “data segments,” these user groups share specific attributes such as demographics, income, region, interests, or behaviors. For instance, advertisers can currently target their Google ads using (1) affinity segments (based on people’s interests and habits), (2) in-market segments (based on recent purchase intent), (3) similar segments (based on interests similar to those of the advertiser’s website visitors or existing customers), (4) detailed demographic segments (based on long-term life facts), and (5) life-event segments (people who are amid important life milestones; see, for example, Google (2023)). Facebook’s core audiences similarly cover a wide range of location-, demographic-, interest-, and behavior-based segments for targeting (Meta, 2023), resulting in more than 600 segments.

Amidst this plethora of options in today’s campaign management systems, the literature is largely silent about selecting the most promising audience segments. For example, Marc S. Pritchard, the chief brand officer at Procter & Gamble, claimed: “We targeted too much, and we went too narrow. . . and now we’re looking at: What is the best way to get the most reach but also the right precision?” (Terlep & Seetharaman, 2016). Our interviews with five industry experts—whose job entails setting up (targeted) ad campaigns—echo this sentiment. These interviews also reveal the different ways these professionals approach the targeting decision (see Online Appendix A): One expert prefers behavioral segments, another one tends to select those that provide a medium reach, and yet another always targets broadly and lets Facebook find the “perfect” audience within those broadly defined boundaries.

Unfortunately, it is hard to predict how profitable an audience segment will be *before* testing it in a real-life campaign, as also highlighted by our industry experts (see Column (4) of Table A1). Intuitively, targeting should increase an advertiser’s profit by allowing the advertiser to stop wasting money on users who are not interested in the advertiser’s products (Iyer, Soberman, & Villas-Boas, 2005; J. P. Johnson, 2013; Skiera et al., 2022). However, it is difficult to ascertain the impact of targeting on an advertiser’s profit because targeting affects profit in three ways. First, targeting often comes with extra data costs for the advertiser (or an increased price of the ad impression), which negatively affects the advertiser’s profit. Second, targeting reduces the number of reachable users, which may decrease the total number of conversions and, in turn, the advertiser’s profit. Third, and despite the above, targeting should improve the performance of advertising campaigns via an increase in at least one of the following metrics: The probability of a click, the probability of a conversion, and the (long-term) margin per conversion (Beales, 2021; Yan et al., 2009). Those opposing effects make it difficult for advertisers to predict the profitability of audience segments.

This paper aims to help advertisers decide *whom* to target. More specifically, we develop a model that enables advertisers to systematically compare different targeting strategies with a no-targeting strategy as a benchmark, before running the corresponding ad campaigns. We choose no-targeting as a benchmark because, ideally, a targeted ad campaign is not only profitable, but also more profitable than an untargeted one.

To illustrate the problem of selecting audience segments in a real-life campaign, we investigate the profitability of different agency-selected targeting strategies in Study 1, a field experiment on Facebook. Building on the resulting insights, we develop a model that informs advertisers about the *break-even performance*¹ (i.e., the minimum lift in click-through rate (CTR), conversion rate (CR), and (long-term) margin per conversion) that makes targeting as profitable as our benchmark, no-targeting.

An important advantage of our approach is that, instead of testing countless different targeting strategies (and their combinations) in randomized controlled trials (RCTs), advertisers can use our model to simulate and systematically compare different targeting strategies *before* running them. While RCTs are costly and time-consuming, they are ideal for evaluating the causal impact of ads and the profitability of a limited number of different targeting strategies. Yet, running RCTs on hundreds of segments (and combinations thereof) is difficult, especially for smaller and medium-sized advertisers, given limited budgets. Thus, the information provided by our model narrows the set of options that advertisers need to test in an RCT, allowing them to solve the cold-start problem.

To show how advertisers can implement our model, we use a unique dataset from Spotify Ad Studio for targeted audio campaigns in Study 2. We find that the required increase in performance is surprisingly large for most audience segments: Around half of the available audience segments on Spotify would require an increase in CTR far larger than 100%. Yet, the literature on user targeting finds that targeting barely increases CTR by about 100% (Aziz & Telang, 2016; Farahat & Bailey, 2012; Rafeian & Yoganarasimhan, 2021). Thus, targeting those segments will likely be less profitable than no-targeting. Very narrow segments that reach no more than 5% of the population require an even bigger increase in CTR—more than approximately 150%. Thus, very narrow segments are often less profitable than no-targeting.

Our Spotify study also echoes one concern many scholars and practitioners have raised: Third-party data used to build audience segments are often inaccurate. For example, Neumann, Tucker, and Whitfield (2019) showed that the best (worst) third-party data provider presents ads to the right target market about 70% (40%) of the time. Thus, we extend our model to handle potentially inaccurate data and predict what happens when data quality decreases. In Study 3, we then investigate the decrease in data quality following Apple’s App Tracking Transparency (ATT) framework introduction, which affected tar-

¹ We differentiate between performance (herein, referred to drivers of profit, i.e., CTR, CR, and long-term margin per conversion) and profit.

getting practices on Facebook (Hercher, 2021; Rahmey, 2021). Using a difference-in-difference (DiD) approach, we find that this decrease in data accuracy harmed narrowly targeted ad campaigns more than broadly targeted ones, which aligns with predictions from our model.

2. Literature on targeting & contribution

The extant literature has focused on investigating the *effectiveness* of different types of targeting (Choi et al., 2020), namely: (1) *keyword* targeting in sponsored search advertising (e.g., Amaldoss, Jerath, & Sayedi, 2016; Du, Su, Zhang, & Zheng, 2017; Jerath, Ma, & Park, 2014; Skiera & Abou Nabout, 2013); (2) *location- or time-based* targeting (e.g., Hui, Inman, Huang, & Suher, 2013; Molitor, Spann, Ghose, & Reichhart, 2020; Phang, Luo, & Fang, 2019); (3) *contextual* targeting (i.e., based on the surrounding content; e.g., Ada, Abou Nabout, & Feit, 2022; Haan, Wiesel, & Pauwels, 2016; Lu & Zhao, 2014; Shehu, Abou Nabout, & Clement, 2021; Yaveroglu & Donthu, 2008), and (4) *behavioral* and *demographic* targeting (i.e., based on past behavior and age/gender, respectively; e.g., Ebbes & Netzer, 2022; Farahat & Bailey, 2012; Rafeian & Yoganarasimhan, 2021).

This study mainly contributes to the latter research stream on behavioral targeting, which relies on tracking technologies (e.g., third-party cookies or digital fingerprinting), allowing advertisers to target audiences based on previous browsing across multiple websites. For instance, a user surfing car-related pages may generally be interested in car content or even be in the market for a new car and thus particularly receptive to car ads.

Table 1 summarizes current research on behavioral targeting, which shows that behavioral information may increase ad effectiveness (Aziz & Telang, 2016; Farahat & Bailey, 2012; Goldfarb & Tucker, 2011b; Rafeian & Yoganarasimhan, 2021; Yan et al., 2009). As one of the earliest studies on behavioral targeting, Yan et al. (2009) reported that behaviorally targeting sponsored search ads aligned with a 670% increase in CTR. Yet, this study mixed the treatment and the selection effect of targeting. Therefore, their reported increase in CTR is too high because the targeted population is more likely to convert from advertising, even without seeing ads. On this point, Farahat and Bailey (2012) showed that naively estimating the targeting effect without accounting for selection bias leads to overestimating the lift from targeting brand-related searches by about 1,000%.

Of course, estimating the causal effect of behaviorally targeted advertising is challenging. Nonetheless, later studies have uncovered the positive effects of behavioral targeting on purchase intent and sales, with the former increasing by about 65% (Goldfarb & Tucker, 2011b) and the latter by 3.6% (G. A. Johnson, Lewis, & Reiley, 2017). Aziz and Telang (2016) also found that individuals with high baseline purchase probabilities responded positively to ads, increasing their purchase probability by up to 2.7 percentage points. Finally, Rafeian and Yoganarasimhan (2021) showed that their proposed targeting strategy—using behavioral and contextual information—improved CTR by 66.8% over the current system; these gains mainly stemmed from behavioral rather than contextual information. All in all, these studies show that targeting barely increases CTR by about 100%. While studying the causal effects of behaviorally targeted advertising is important, we still lack answers regarding how to select audience segments from an overwhelming range of options *before* data of such targeted ad campaigns comes in (i.e., before being able to perform any causal analysis based on RCTs can be done).

In addition, these studies point to the benefits of behavioral targeting, but they hardly consider its downsides: (1) the associated targeting cost and (2) the (potentially large) decrease in the campaign's reach. Thus, we know little about the profitability of different targeting strategies. An exception is a study by Neumann et al. (2019), which discovered that third-party audiences from various data brokers exhibit varying accuracy levels; given the extra costs of targeting and its relative inaccuracy, Neumann et al. (2019) find that third-party audiences are often economically unattractive, except for higher-priced media placements.

We thus contribute to the existing literature in several ways. First, we assist advertisers in narrowing down their options to strategically select segments worth testing in RCTs. Second, we shed light on the *degree* of effectiveness required for a targeted ad campaign to be as profitable as an untargeted one, by developing a novel model that considers the trade-off between reach, cost, and performance (i.e., effectiveness). Third, we demonstrate that targeting extremely narrow segments is highly unprofitable due to the large increase in performance needed to compensate for the loss in reach. Fourth, behavioral targeting faces challenges due to restrictions by major players like Google and Apple (Neumann, 2020; Schuh, 2019), making behavioral targeting almost impossible (Apple, 2022). Our model is applicable to any form of audience targeting that reduces the original audience size, irrespective of the underlying data (i.e., first- and third-party data). Fifth, we assess the profitability of targeting narrow and broad audience segments under conditions of poor data quality and empirically validate our model insights using data on Apple's ATT framework.

3. Study 1: Profitability of different, agency-selected targeting strategies on Facebook

To illustrate the problem of selecting audience segments in a real-life campaign, we investigate the profitability of different agency-selected targeting strategies in a field experiment on Facebook. Specifically, we asked an agency to design broad-, narrow-, and no-targeting strategies, allowing us to compare their profitability. It turns out that neither of the two targeted campaigns works better than the untargeted one, thus setting the stage to explore why and how advertisers can make better targeting decisions.

Table 1
Related Literature on Behavioral Targeting.

Study	Aim	Type of targeting	Data	Method	Targeting ... effectiveness (lift; in ...)	profitability	Guidance on how to select audience segments
Yan et al. (2009)	Effect of behavioral targeting on search engine advertising	Behavioral targeting	Click-through log from commercial search engine	User segmentation algorithms (e.g., clustering)	Yes (670%; CTR)*	No	No
Goldfarb and Tucker (2011b)	Effect of privacy regulations on ad effectiveness	Data-driven targeting	Survey + campaign data	DiD	Yes (65%; purchase intent)	No	No
Farahat and Bailey (2012)	Decompose the impact of targeting into selection and treatment effects	Behavioral targeting	Campaign data	DiD	Yes (4.8%; search-through rate)	No	No
Aziz and Telang (2016)	Effect of using more cookie information on accuracy of purchase predictions	Retargeting	Impression/bid-level data	Regression analysis + counterfactual simulations	Yes (<35%; purchase intent)	No	No
G. A. Johnson et al. (2017)	Effect of online display advertising on retailer's sales	Behavioral targeting	Campaign data	Field experiment	Yes (%3.6; sales)	Yes	No
Rafieian and Yoganarasimhan (2021)	Machine learning framework for targeting	Behavioral & contextual targeting	In-app ad network data	Machine learning + analytical model + counterfactual simulations	Yes (66.8%; CTR)	No	No
Neumann et al. (2019)	Accuracy of third-party audiences from different data brokers	Third-party audience targeting	Survey + campaign data	Comparison of means	No	Yes	No
This study	Model to calculate break-even performance of targeting	Audience segment targeting	(Pre-) Campaign data from Facebook + Spotify Ad Studio	Field experiments/ DiD + analytical model + counterfactual simulations	No	Yes	Yes

Notes: CTR: Click-through rate; DiD: Difference-in-difference; we compare the profitability of targeting to the profitability of no-targeting.* Note that this estimate is subject to a potential selection bias.

3.1. Experimental design

We collaborated with a leading Austrian ad agency and one of its clients, a car dealer in the Vienna region. In our experiment, we used Facebook's *A/B split test* feature, which allows for the random assignment of users into non-overlapping conditions in the Vienna region (i.e., Facebook assigns a specific user from Vienna only to one targeting strategy, not multiple ones; Meta, 2022a). We aim to study the effect of different targeting strategies on impressions, cost, CTRs, and profit after acquisition cost (since car sales are hard to measure, we calculate profit after acquisition cost as: $\text{Profit}_{it} = \text{€}1,000 \times \text{CTR}_{it} \times \text{CR}_{it} - \text{Acquisition Cost}_{it}$, where conversions correspond to registrations for test drives at the car dealer's website, valued at €1,000).

Specifically, given the rather limited overall budget of €5,700, we asked the agency to create one condition without targeting and two conditions with behaviorally targeted ad campaigns: One using "broad-targeting" and another using "narrow-targeting." The agency set up the following strategies: The no-targeting strategy consisted of all adult users (i.e., at least 18 years old) in the Vienna region (as test drives could only happen there). The broad-targeting strategy applied behavioral targeting to adult users interested in cars, electric cars, or vehicles (in the Vienna region). The narrow-targeting strategy also behaviorally targeted users but focused on adult users in the Vienna region who were interested in the advertised car brand *and* in at least one of the following car types: (1) sporty cars, (2) off-road cars, (3) station cars, (4) middle-class cars, (5) compact class cars, and (6) small cars (matching the dealer's offering).

We consider these targeting strategies as two ends of the targeting spectrum. Our limited budget prevented us from testing more strategies, which is exactly the managerial problem we tackle: How can advertisers decide which segments are worth testing in an RCT?

The experiment started on January 24, 2022, and ended on March 26, 2022. To avoid having other confounds (than our targeting strategies) affect our results, we assigned an equal budget to all targeting strategies (i.e., €5,700/3 = €1,900). While advertisers may choose different campaign settings for different targeting strategies in practice (e.g., a higher budget for broad-targeting), we must ensure that all conditions are as comparable as possible.

In addition, we used the same ad creative and bidding strategy for all conditions. The ad creatives covered the main range of the brand's car models: They featured the brand's main body types (i.e., Huckaback, SUV, and saloon), fuel types (petrol, plug-in hybrid, and full electric), and car designs (e.g., Fiesta and Mustang Mach-E; see Figure B1 of Online Appendix B.1). We employed a "lowest cost" (later called "highest volume") bidding strategy, allowing us to maximize delivery and conversions from the given budget (Meta, 2022b). At the time of our experiment, other bid strategies were available on Facebook, namely "bid cap" and "goal-based bidding." But using the same bid cap or cost goal across all conditions would disadvantage narrow-targeting strategies, which are usually more expensive.

3.2. Model specification

To formally analyze the differences in our outcome variables (i.e., impressions/reach, CPMs, CTRs, and profit after acquisition cost) induced by different targeting strategies, we estimated models of the following type (for summary statistics, see Table B1 of Online Appendix B.1):

$$Y_{it} = \beta_0 + \beta_1 \times \text{broad-targeting}_i + \beta_2 \times \text{narrow-targeting}_i + \tau_t + \varepsilon_{it}, \quad (1)$$

where Y_{it} is the dependent variable for strategy $i \in \{\text{no-targeting, broad-targeting, narrow-targeting}\}$ on day $t \in \{1, 2, \dots, 62\}$. Broad-targeting_{*i*} and narrow-targeting_{*i*} are indicator variables for the type of targeting. We included day fixed effects to control for changes induced by time (τ_t).²

3.3. Discussion of results

Table 2 shows the results for all our outcome variables: As expected, broad- and narrow-targeting led to significantly lower reach, lower impressions, and higher CPMs than no-targeting. Specifically, no-targeting achieved the highest (lowest) reach and impressions (CPM), followed by broad-targeting and then narrow-targeting. Yet, neither CTR nor profit was significantly higher under broad- and narrow-targeting than no-targeting. Specifically, narrow-targeting showed a larger increase in CTR and profit than broad-targeting (over no-targeting), though this increase was insignificant.

These findings raise the following questions: How large would the performance increase for our targeted strategies need to be for them to be more profitable than no-targeting. Was it even worth testing these behavioral segments in an experiment and spending money on such a test? Instead, would it be possible to *predict* that these two targeting strategies will likely not be more profitable than the untargeted campaign? Can we propose segments for testing that would more likely lead to a profit increase? Our model seeks to answer these questions. (To provide a quick preview: Our model reveals a break-even performance of 1.21 (1.33) for broad- (narrow-) targeting, which is higher than the realized performance increase of 1.02 (1.29) for broad- (narrow-) targeting, explaining why these selected strategies turned out to be less profitable than no-targeting (see Online Appendix B.2 for details)).

² We also test alternative specifications of this model: Specifically, we replaced our day fixed effects with (1) a deterministic daily time trend, (2) day-of-the-week fixed effects, and (3) a weekend (versus weekday) dummy. The results are robust across those alternative specifications.

Table 2

Field Experiment on Facebook: Effect of Narrow- and Broad-Targeting on Reach, Impressions, CPM, CTR, and Profit.

	Reach	Impressions	CPM	CTR	Profit
Broad-targeting	-660.71*** (0.000)	-1,092.02*** (0.000)	0.87*** (0.000)	0.01 (0.806)	-16.13 (0.885)
Narrow-targeting	-1,136.94*** (0.000)	-1,588.15*** (0.000)	1.62*** (0.000)	0.08 (0.178)	112.90 (0.287)
Day fixed effects	yes	yes	yes	yes	yes
Constant	2,299.88*** (0.000)	2,779.72*** (0.000)	4.57*** (0.000)	1.10*** (0.003)	-42.09 (0.562)
N	186	186	186	186	186
R ²	0.95	0.95	0.69	0.51	0.45

Notes: p -values in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; CPM: Cost per mille; CTR: Click-through rate; robust standard errors were used for estimations; the baseline is no-targeting; the number of observations is the three experimental conditions times the number of days (i.e., $3 \times 62 = 186$).

4. Model to determine break-even performance

To answer the question of which targeting strategies are most promising and worth testing in an RCT, we calculate the minimum performance increase for a targeted ad campaign to be at least as profitable as an untargeted one (i.e., we calculate the break-even performance). As our benchmark, we use a no-targeting strategy because, ideally, a targeted ad campaign is more profitable than an untargeted one.

We use the following information to develop the model: (1) the reach of a specific audience segment plus the size of the original, untargeted population; (2) the corresponding targeting/data cost³; (3) the expected CPM under targeting and no-targeting (either RTB or fixed price); and (4) the performance (CTR and CR) under no-targeting. Self-service platforms typically provide advertisers with most of this information (i.e., reach of a segment, size of the untargeted population, data cost, and expected CPMs). In addition, advertisers know other inputs from past campaigns (e.g., performance under no-targeting) or may be able to produce estimates of them based on their expert knowledge.

4.1. Modeling the effect of targeting on advertiser's profit

To derive the minimum performance increase for a targeted ad campaign to be at least as profitable as an untargeted one (i.e., the break-even performance), we first need to understand how targeting affects an advertiser's profit. Consequently, we set up the profit functions under both targeting and no-targeting. The advertiser's profit (π) is a function of the number of users who purchase the advertiser's product (Q), the (long-term) margin per conversion (m , which can represent customer lifetime value (CLV)), and the cost per conversion. Herein, we use subscript 0 when referring to no-targeting and subscript i when we refer to targeting audience segment i , where i belongs to the set of available audience segments I , offered on the platform for targeting (i.e., $i \in I$).

Note: Even if advertisers do not know the exact profit margins on their conversions (or find conversions difficult to measure, like in our Facebook field experiment), they can often value other outcomes, such as registrations, leads, or other actions on the website (e.g., booking a test drive, requesting a sales quote, sharing information with other users). Thus, they can still use our model to reduce the available targeting options even if they do not expect to generate instant profits.

4.2. Profit under no-targeting and targeting

We start by formalizing the advertiser's profit under no-targeting as follows:

$$\pi_0 = N_0 \text{CTR}_0 \text{CR}_0 \left(m_0 - \frac{\text{CPM}_0}{1,000 \text{CTR}_0 \text{CR}_0} \right), \quad (2)$$

where N_0 is the number of ad impressions, CTR_0 is click-through rate, CR_0 is the conversion rate, m_0 is the absolute (long-term) margin per conversion, and CPM_0 is the advertising cost in CPM (i.e., the cost an advertiser pays for 1,000 non-targeted ad impressions; see Online Appendix C.1 for more details).

Like the profit under no-targeting, we calculate the advertiser's profit when targeting audience segment $i \in I$ (π_i) as follows:

$$\pi_i = N_i \text{CTR}_i \text{CR}_i \left(m_i - \frac{\text{CPM}_i + P_i}{1,000 \text{CTR}_i \text{CR}_i} \right), \quad (3)$$

³ Targeting using audience segments typically requires an advertiser to pay the CPM for using the audience segment in addition to the media cost CPM for the individual impression (for more information, see for instance: <https://support.google.com/displayvideo/answer/6212219?hl=en>).

where CPM_i is the advertising cost in CPM for targeting audience segment i , $P_i \geq 0$ is the data cost in CPM (i.e., the extra cost that an advertiser may need to pay to target the respective audience segment i). We summarize all these variables in Table 3.

Table 3
Description of Variables in our Model.

Variable under no-targeting	Variable under targeting	Description
0		No-targeting subscript
i		Audience segment subscript in case of targeting
Input to Model		
π_0	π_i	Profit after acquisition cost
acq_0	./.	Acquisition cost
Q_0	./.	Number of conversions
N_0	N_i	Number of ad impressions
CTR_0	CTR_i	Click-through rate
CR_0	CR_i	Conversion rate
m_0	m_i	Margin per conversion
CPM_0	CPM_i	Cost per mille
./.	P_i, DC_i	Specifications of data cost through targeting
./.	$\hat{\lambda}$	Estimated parameter reflecting how costly targeting is
./.	α_i	Change in CTR through targeting
./.	β_i	Change in CR through targeting
./.	γ_i	Change in margin through targeting
./.	δ_i	Change in reach through targeting
./.	$\alpha_i \beta_i \gamma_i$	Performance multiplier
Output of Decision Model		
$(\alpha_i \beta_i \gamma_i)^*$		Break-even performance (i.e., minimum performance multiplier)

Notes: $CTR_i = \alpha_i CTR_0$; $CR_i = \beta_i CR_0$; $m_i = \gamma_i m_0$; $N_i = \delta_i N_0$.

Intuitively, the advertiser expects a performance lift (i.e., in CTR, CR, or margin per conversion) when targeting users who are more likely to be interested in the advertiser's products (Aziz & Telang, 2016; Farahat & Bailey, 2012; Goldfarb & Tucker, 2011b; Rafeian & Yoganarasimhan, 2021; Yan et al., 2009). Yet, targeting may decrease the number of ad impressions by reducing the number of available users because it excludes those who do not meet the targeting criteria. Moreover, targeting may come with additional costs for the advertiser because of increased performance from targeted users.

In line with these arguments, we define α_i , β_i , γ_i , and δ_i as multipliers, representing the changes in CTR, CR, margin, and reach, respectively, through targeting audience segment i . As a result, $CTR_i = \alpha_i CTR_0$, $CR_i = \beta_i CR_0$, $m_i = \gamma_i m_0$, and $N_i = \delta_i N_0$, where we expect an increase in CTR, CR, and margin (i.e., $\alpha_i \beta_i \gamma_i > 1$) and a decrease in reach (i.e., $0 < \delta_i < 1$) when targeting audience segment i . For example, a value of $\alpha_i = 1.1$ means that CTR_i (i.e., the click-through rate when targeting audience segment i) is 1.1 times CTR_0 and thereby 10% higher than CTR_0 (i.e., the click-through rate under no-targeting). Similarly, a value of $\delta_i = 0.4$ means that N_i (i.e., the number of available ad impressions when targeting audience segment i) is 0.4 times N_0 , which is 60% lower than N_0 (i.e., the number of available ad impressions under no-targeting). Finally, changes in α_i , β_i , and γ_i have similar effects on the advertiser's profit (see Equation (4); see Online Appendix C.2 for details), enabling us to introduce and use the *performance multiplier* of audience segment i , $\alpha_i \beta_i \gamma_i$, for our next steps.

Moreover, we assume: (1) the narrower the audience segment i , the more expensive it is to target (i.e., the higher the data cost), (2) the increase in data cost becomes smaller for lower reach with a decreasing marginal effect, and (3) data cost may vary on different platforms (e.g., the cost of targeting audience segment i may be more expensive on Facebook than on Spotify). These assumptions align with our observations from the field experiment on Facebook and those from our Spotify study (shown below). Therefore, we define data cost as a function of reach: $P_i = 1/\delta_i^{\hat{\lambda}}$ (where, in line with the first two assumptions, $dP_i/d\delta_i < 0$ and $d^2P_i/d\delta_i^2 > 0$). More specifically, $\hat{\lambda} > 0$ is an estimated parameter, indicating an audience segment is more expensive with a larger $\hat{\lambda}$. Advertisers can estimate $\hat{\lambda}$ using data from previous ad campaigns (see Online Appendix C.3 for estimations of $\hat{\lambda}$ in our field experiment on Facebook and the Spotify study).

Taken together, inserting $CTR_i = \alpha_i CTR_0$, $CR_i = \beta_i CR_0$, $m_i = \gamma_i m_0$, and data cost $P_i = 1/\delta_i^{\hat{\lambda}}$ in Equation (3) produces the following changes to the profit function of the advertiser targeting audience segment i :

$$\pi_i = \delta_i N_0 CTR_0 CR_0 \left(\alpha_i \beta_i \gamma_i m_0 - \frac{CPM_i + 1/\delta_i^{\hat{\lambda}}}{1,000 CTR_0 CR_0} \right). \quad (4)$$

4.3. Derivation of break-even performance

We now derive the break-even performance of audience segment i (i.e., $(\alpha_i \beta_i \gamma_i)^*$), such that targeting audience segment i is as profitable as no-targeting. We start by deriving the conditions under which the advertiser would realize equal profits under targeting audience segment i (π_i) and no-targeting (assuming the untargeted profit to be positive, i.e., $\pi_0 > 0$).⁴ We do so by setting Equation (4) equal to Equation (2):

$$\pi_i = \pi_0 \rightarrow \delta_i N_0 \text{CTR}_0 \text{CR}_0 \left(\alpha_i \beta_i \gamma_i m_0 - \frac{\text{CPM}_i + 1/\widehat{\delta}_i^\lambda}{1,000 \text{CTR}_0 \text{CR}_0} \right) = N_0 \text{CTR}_0 \text{CR}_0 \left(m_0 - \frac{\text{CPM}_0}{1,000 \text{CTR}_0 \text{CR}_0} \right). \quad (5)$$

We then solve for the performance multiplier (i.e., $\alpha_i \beta_i \gamma_i$) of audience segment i with a given reach (δ_i):

$$(\alpha_i \beta_i \gamma_i)^* = \frac{1}{\delta_i} - \frac{\text{CPM}_0 - \delta_i (\text{CPM}_i + 1/\widehat{\delta}_i^\lambda)}{1,000 \delta_i \text{CTR}_0 \text{CR}_0 m_0}. \quad (6)$$

Equation (6) describes the break-even performance (i.e., $(\alpha_i \beta_i \gamma_i)^*$) of audience segment i . Stated differently, it reveals the break-even performance related to targeting audience segment i that makes targeting as profitable as no-targeting (with the reach δ_i). In Equation (6), any performance multiplier lower than $(\alpha_i \beta_i \gamma_i)^*$ (i.e., $\alpha_i \beta_i \gamma_i < (\alpha_i \beta_i \gamma_i)^*$) makes targeting less profitable than no-targeting, and vice versa. Also note that this is the *causal effect* of targeting, i.e., advertisers need to evaluate whether they can achieve this break-even performance after accounting for potential selection bias (as outlined in Section 2).

4.4. Effect of reach on break-even performance

Making an audience segment narrower (i.e., decreasing its reach) requires a higher break-even performance. The crucial question is: How much higher? We can use our model to find the answer: The reach of an audience, which is indicative of the narrowness of an audience segment, has a negative and non-linear effect on the break-even performance because $\partial(\alpha_i \beta_i \gamma_i)^*/\partial\delta_i < 0$ and $\partial^2(\alpha_i \beta_i \gamma_i)^*/\partial\delta_i^2 > 0$ (see Online Appendix C.5). In Fig. 1, we use Equation (6) and illustrate the effect of the reach of an audience segment i on its break-even performance for a hypothetical advertiser (for different data cost functions).

Fig. 1 shows that the break-even performance of audience segment i increases with a decrease in its reach. This finding is intuitive because it states that the platform should promise a higher performance for smaller audience segments to achieve equal profits under targeting audience segment i , compared to no-targeting.

More importantly, Fig. 1 reveals a non-linear relationship between the break-even performance of an audience segment i and its reach, where the required magnitude of change in the break-even performance (for a specific change in reach, say $\Delta\delta_i$) becomes larger for lower reach with an increasing marginal effect. This finding is essential since ad platforms allow advertisers to narrow their targeting with respect to various demographic traits, interests, etc. (see Figure D1, Online Appendix D).

In other words, an audience segment's break-even performance must increase dramatically with a reduction in its reach due to targeting narrower audiences. This large effect of reach on break-even performance indicates that very narrow audiences are unlikely to be profitable.

4.5. Break-even performance when data cost is known

In practice, advertisers often have a good understanding of the data cost associated with targeted ad campaigns because of their previous campaigns or because the campaign management system directly reports costs. For example, on Spotify, an advertiser with an advertising budget of \$25,000 must pay an additional \$1.00 to target users with iOS devices (see Table 5). Spotify Ad Studio provides the advertiser with this information during the setup process, i.e., before running the campaign.

Therefore, we can further simplify our model by replacing $P_i (= 1/\widehat{\delta}_i^\lambda)$ in Equation (4) with the fixed data cost for targeting audience segment i , $\text{DC}_i \geq 0$:

$$\pi_i = \delta_i N_0 \text{CTR}_0 \text{CR}_0 \left(\alpha_i \beta_i \gamma_i m_0 - \frac{\text{CPM}_i + \text{DC}_i}{1,000 \text{CTR}_0 \text{CR}_0} \right). \quad (7)$$

We derive the break-even performance of audience segment i (i.e., $(\alpha_i \beta_i \gamma_i)^*$) by setting Equation (7) equal to Equation (2), and solving for the performance multiplier for audience segment i :

$$(\alpha_i \beta_i \gamma_i)^* = \frac{1}{\delta_i} - \frac{\text{CPM}_0 - \delta_i (\text{CPM}_i + \text{DC}_i)}{1,000 \delta_i \text{CTR}_0 \text{CR}_0 m_0}. \quad (8)$$

⁴ If $\pi_0 \leq 0$, it would be sufficient for advertisers to calculate the break-even performance that leads to (at least) zero profits (see Online Appendix C.4).

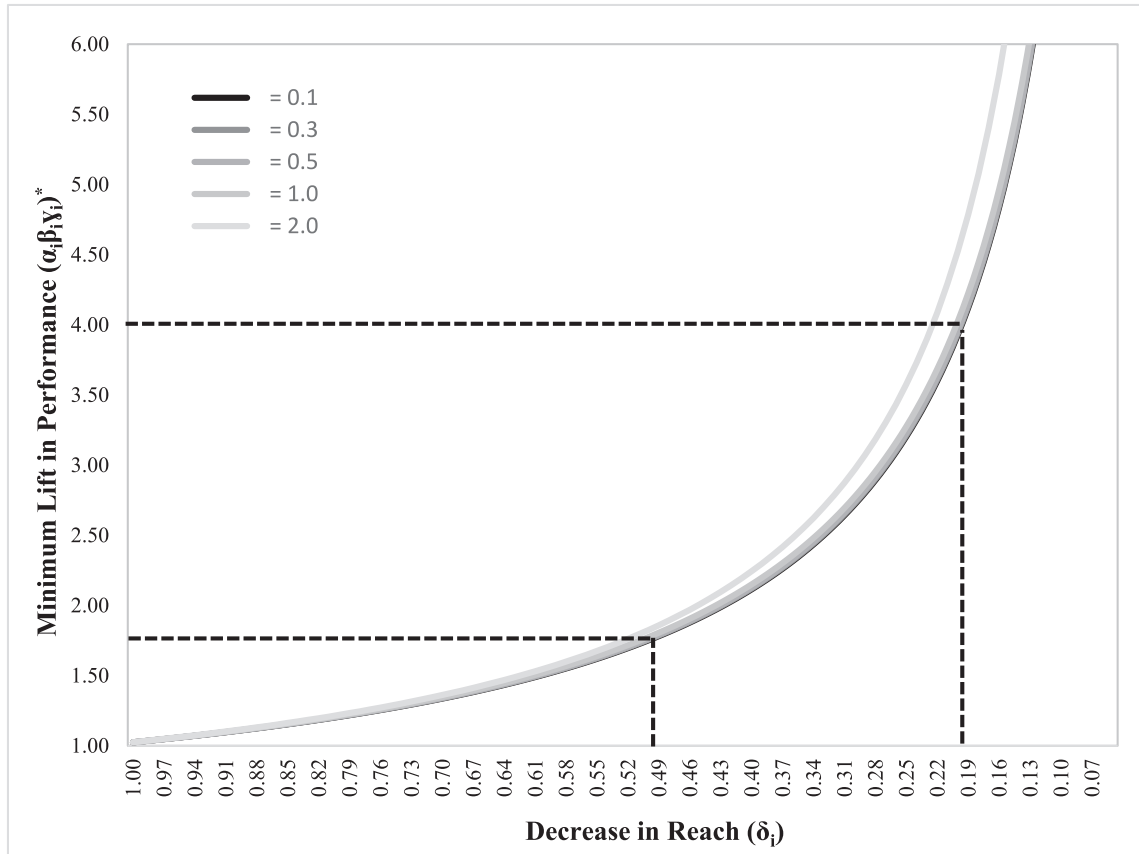


Fig. 1. Model Insights: Relationship between Reach and Break-Even Performance for Audience Segments with Varying Data Cost ($0.1 \leq \hat{\lambda} \leq 2.0$). Notes: We use the following numbers (which rely on our observations from Spotify, see Section 5): Number of impressions (N_0) = 1,800,000; click-through rate (CTR_0) = 1.00%; conversion rate (CR_0) = 2.00%; margin per conversion (m_0) = \$200.00; CPM under no-targeting (CPM_0) = \$11.00; CPM under targeting an audience segment i (CPM_i) = \$11.00; α_i , β_i , and γ_i are the multipliers of CTR_0 , CR_0 , and m_0 , respectively; $(\alpha_i \beta_i \gamma_i)^*$ represents the break-even performance; subscript 0 refers to no-targeting, and subscript i refers to targeting audience segment $i \in I$; $\delta_i = 1$ represents no-targeting. Reading example: For $\hat{\lambda} = 0.3$ (see the dashed lines, which represent data cost on Spotify), (1) if the reach of an audience segment i decreases from 1.00 to 0.50, its break-even performance must increase to 1.76 to leave the advertiser with the same profit for targeting and no-targeting; (2) if the reach of an audience segment i decreases from 1.00 to 0.20 instead of 0.50 (i.e., a further decrease in reach by 30 percentage points compared to (1)), its break-even performance must increase to 3.94 (i.e., a further increase by 218 percentage points compared to (1)) to leave the advertiser with the same profit for targeting and no-targeting.

Put differently, Equation (8) is the simplified version of Equation (6), with a given total cost of the audience segment ($CPM_i + DC_i$), provided by the platform. Similar to our findings in Section 4.4, Online Appendix C.6 illustrates that with a given cost of $CPM_i + DC_i$, the break-even performance of audience segment i (see Equation (8)) increases non-linearly with a decrease in reach. Thus, depending on whether the advertiser knows $CPM_i + DC_i$ or not, they can either use Equation (8) to derive the break-even performance or Equation (6).

5. Study 2: Break-even performance of different audience segments on Spotify

To illustrate how advertisers can implement our model, we apply it to a unique dataset from Spotify Ad Studio, which allows advertisers to target non-paying users (i.e., users who stream music for free with ads) on Spotify. Advertisers can use the platform to create their ads in an audio/video format and build a targeted ad campaign. Figure D1 of Online Appendix D shows a screenshot of Spotify Ad Studio at the time of our study.

Advertisers begin by setting the schedule and budget for their campaign; next, they determine their target audience. The maximum budget for a campaign was limited to \$25,000 at the time of data collection; Spotify distributes this money evenly throughout the campaign period, which is limited to twelve months. The campaign time and its budget also affect the likelihood that the advertiser will spend the budget (indicated by *budget delivery likelihood*), as well as the estimated number of unique users who will be served (indicated by *estimated reach*, which we use to determine the reach of an audience segment). Once an advertiser defines the campaign's budget and schedule, alongside the target audience's nation, the platform provides the advertiser with a cost per ad served (indicated by *\$ per impression*).

With the information provided by Spotify, advertisers can evaluate which targeting strategies are at least as profitable as a no-targeting strategy. The question then is: For an audience segment $i \in I$, is the required break-even performance—which makes its targeting as profitable as the no-targeting benchmark—attainable?

To calculate the break-even performance for various segments on Spotify, we retrieved the number of unique users available for targeting (i.e., *estimated reach*; see Figure D1 of Online Appendix D) by setting the campaign budget and time to their maximum values (i.e., \$25,000 and twelve months, respectively). Doing so resulted in 1.8 million UK users (i.e., $N_0 = 1,800,000$). Spotify also provides the CPM for an untargeted campaign (i.e., CPM_0) and the total cost for a targeted campaign i (i.e., $CPM_i + DC_i$, where $DC_i \geq 0$). Consequently, we used these total costs (i.e., $CPM_i + DC_i$) in Equation (8) to form expectations about the minimum lift in performance that should result from targeting audience segment i .

To this end, we run two types of analyses: In our first analysis, we simulate different advertisers with varying baseline CTRs, CRs, and margins in their untargeted campaigns. Our second analysis aims to generate deeper insights for an exemplary (hypothetical) advertiser from our simulation study.

5.1. Simulation study design

Our simulation study examines various scenarios that different advertisers may face in their campaigns. For instance, CTRs, CRs, and margins might vary widely for different products and advertisers. Thus, we vary baseline CTRs, CRs, and margins under no-targeting using four different factor levels for each (low, mid-low, mid-high, and high; see Table 4 for the summary of the simulation study design). We pick the corresponding ranges based on insights from our field experiment (average CTR_0 and CR_0 of 0.50% and 1.50%, respectively) and informal discussions with industry experts. In total, our simulation study covers $4^3 = 64$ scenarios. We then randomly draw ten values from the respective uniform distributions of CTR_0 , CR_0 , and m_0 for each scenario, resulting in a total of 640 ($= 64 \times 10$) different ad campaigns (i.e., $r \in \{1, 2, \dots, 640\}$).

Next, we collect information about all 71 audience segments from Spotify. These segments cover age, gender, platform/device, interests, real-time contexts, and genres (see Table 5 and Fig. 3 for details of audience segments). As a final step, we use Equation (8) to calculate the minimum lift in performance for each audience segment $i \in \{1, 2, \dots, 71\}$ in each campaign r (i.e., $(\alpha_i^r \beta_i^r \gamma_i^r)^*$), resulting in 45,440 ($= 64 \times 10 \times 71$) observations. Each observation corresponds to a minimum lift in our simulation study.

Table 4
Simulation Study on Spotify: Study Design.

Performance level	Click-through rate (CTR_0)	Low \in	[0.50%, 0.75%]
		Mid-low \in	[0.75%, 1.00%]
		Mid-high \in	[1.00%, 1.25%]
		High \in	[1.25%, 1.50%]
	Conversion rate (CR_0)	Low \in	[1.50%, 1.75%]
		Mid-low \in	[1.75%, 2.00%]
		Mid-high \in	[2.00%, 2.25%]
		High \in	[2.25%, 2.50%]
	Margin per conversion in \$ (m_0)	Low \in	[150, 250]
		Mid-low \in	[250, 500]
		Mid-high \in	[500, 750]
		High \in	[750, 1000]
Total number of scenarios		$4^3 =$	64
Number of replications per scenario			10
Total number of campaigns		$64 \times 10 =$	640
Number of audience segments per campaign			71
Number of observations		$640 \times 71 =$	45,440

5.2. Discussion of overall results

In the following, we first calculate the changes in $(\alpha_i \beta_i \gamma_i)^*$ for different types of audience segments, which represents the combined performance lift of CTR_0 , CR_0 , and m_0 . However, this combined performance lift might be difficult to interpret and compare to something advertisers know and understand. For example, a minimum lift in performance of 800% when targeting audience segment i may mean a minimum lift in CTR_0 of 300% (i.e., $\alpha_i = 4$), a minimum lift in m_0 of 100% (i.e., $\gamma_i = 2$), and no change in conversion rate (i.e., $\beta_i = 1$). But any other combination could also lead to a lift in performance of 800%. For illustration purposes, we assume $\alpha_i > 1$, $\beta_i > 1$ since the literature on targeting shows that CTR_0 and CR_0 might indeed increase (Aziz & Telang, 2016; Farahat & Bailey, 2012; Goldfarb & Tucker, 2011b; Rafieian & Yoganarasimhan, 2021; Yan et al., 2009), and $\gamma_i > 1$ since the targeted audience segment may purchase more frequently or in higher quantities. To trans-

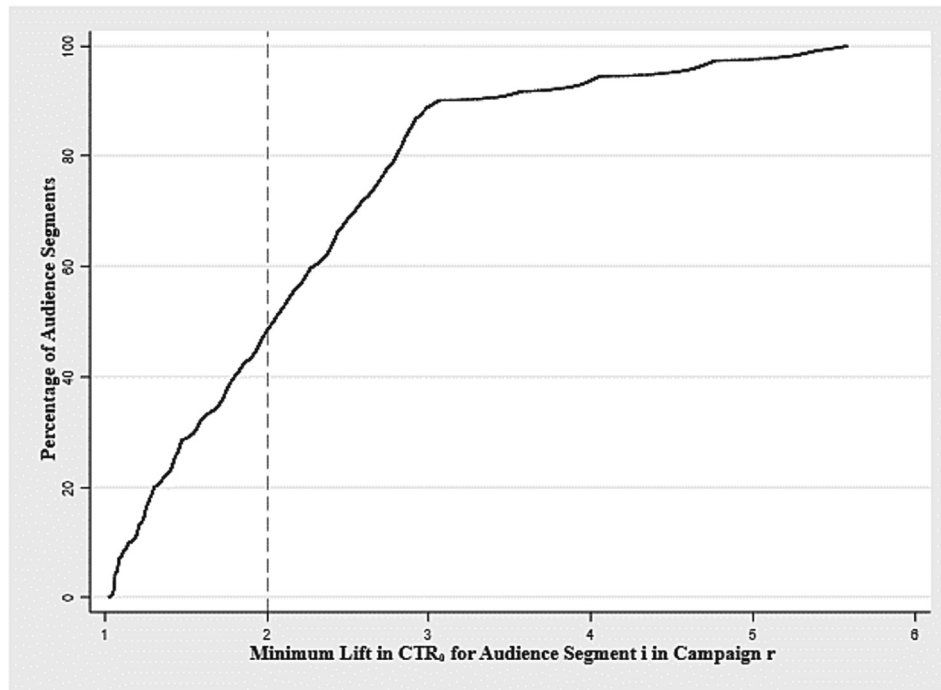


Fig. 2. Simulation Study Results: Minimum Lift in CTR₀ Across Different Audience Segments. Notes: Results are based on a total of 45,440 (= 71 × 640) audience segments (i.e., audience segment $i \in \{1, 2, \dots, 71\}$ in campaign $r \in \{1, 2, \dots, 640\}$); we assume $\alpha_i = \beta_i = \gamma_i$. Reading example: The dashed vertical line represents that around 50% of audience segments require a minimum increase in CTR₀ of more than 100% to be at least as profitable as no-targeting.

late the performance lift into something advertisers can easily relate to, we will assume that $\alpha_i = \beta_i = \gamma_i$ and mainly interpret the minimum lift in CTR₀ (i.e., α_i^*). (Note: We could easily adjust this assumption.⁵)

Our simulation study reveals that more than 50% of audience segments (i.e., 23,407 out of 45,440 audience segments on the right-hand side of the dashed vertical line in Figure D2 in Online Appendix D) require a minimum increase in performance larger than 700% to be at least as profitable as no-targeting. Assuming $\alpha_i = \beta_i = \gamma_i$ in Equation (8), we find that more than half of the audience segments require an increase in CTR₀, CR₀, and m_0 larger than 100% to be at least as profitable as no-targeting (see Fig. 2). These segments should deliver an average increase of 118% in CTR₀, with substantial variation across all 640 campaigns (SD = 0.95). Such an increase in CTR₀ is rather high: For example, the literature on user targeting reports that potential increases in CTR₀ rarely even reach 100% (Aziz & Telang, 2016; Farahat & Bailey, 2012; Rafeian & Yoganarasimhan, 2021). Thus, this finding raises doubts about the profitability of targeting (compared to no-targeting), at least on Spotify Ad Studio.

5.3. Discussion of results for a hypothetical advertiser

To explore which audience segments could potentially yield higher profits than no-targeting, we now zoom in on the results for one exemplary advertiser. This advertiser has a CTR₀ of 1.00%, a CR₀ of 2.00%, and a long-term margin per conversion (i.e., m_0) of \$200. Such an advertiser would resemble an online retailer like Zalando.

We start with the two most widely used demographic attributes—namely, age range and gender—and then continue with targeting a user's device type. Later, we include additional targeting options to explore narrower segments.

⁵ Assuming $\alpha_i = \beta_i = \gamma_i$ does not alter our later conclusion of doubtful profitability of targeting (compared to no-targeting). For instance, if we assume targeting leads to a lower (higher) increase in CR₀ and m_0 than CTR₀, say $\alpha_i = 1.5\beta_i = 1.5\gamma_i$ ($1.5\alpha_i = \beta_i = \gamma_i$), then around 70% (30%) of audience segments require a minimum increase in CTR₀ of more than 100% to be at least as profitable as no-targeting. However, with the assumption of $1.5\alpha_i = \beta_i = \gamma_i$, more than 60% of audience segments would need a minimum increase in CR₀ and m_0 of over 100% to match the profitability of no-targeting.

Table 5Hypothetical Advertiser on Spotify: Minimum Lift in CTR_0 of Different Audience Segments Based on Age, Gender, and Platform/Device.

Audience i	Reach in 000s (% of all users)	Cost of no-targeting (i.e., CPM_0)	Total cost of targeting (i.e., $CMP_i + DC_i$)	Min. lift in CTR_0 (α_i^*)
<i>Age Range</i>				
13–24	1,400 (77.78%)	\$ 11.00	\$ 11.00	1.06
25–34	1,000 (55.56%)	\$ 11.00	\$ 11.00	1.16
35–44	800 (44.44%)	\$ 11.00	\$ 11.00	1.24
45–65+	1,000 (55.56%)	\$ 11.00	\$ 11.00	1.16
<i>Gender</i>				
Male	1,500 (83.33%)	\$ 11.00	\$ 11.00	1.05
Female	1,500 (83.33%)	\$ 11.00	\$ 11.00	1.05
<i>Platform/Device</i>				
iOS	1,400 (77.78%)	\$ 11.00	\$ 12.00	1.07
Android	1,300 (72.22%)	\$ 11.00	\$ 12.00	1.09
Desktop	1,500 (83.33%)	\$ 11.00	\$ 12.00	1.05
Average	1,267 (70.00%)	\$ 11.00	\$ 11.33	1.10

Notes: Results are based on 1.8 million available users in the UK (i.e., $N_0 = 1.8$ million); we use the following numbers to derive the results: Click-through rate (CTR_0) = 1.00%; conversion rate (CR_0) = 2.00%; margin per conversion (m_0) = \$200.00; for α_i^* , we assume $\alpha_i = \beta_i = \gamma_i$; subscript 0 refers to no-targeting, and subscript i refers to targeting audience segment $i \in I$; the reported breakdown of available users under an audience segment i may be higher than the overall number of available users; e.g., the sum of users in each age range (i.e., $\sum_a \text{Age}_a$) is higher than the overall number of available users (i.e., 1.8 million). This discrepancy highlights a lack of accuracy in how users are profiled (see Section 5.5 for how to account for inaccuracies in audience segments).

5.3.1. Age range targeting

Spotify allows advertisers to target users aged 13 years to 65+ years without additional cost (i.e., $CPM_0 = CPM_a + DC_a = \$11.00$, where $a \in \{13-24, 25-34, 35-44, 45-65+\}$). Table 5 shows that users in the 13-to-24 category represent 77.78% of available users on Spotify. We use Equation (8) to calculate the minimum lift in performance (i.e., $(\alpha_i \beta_i \gamma_i)^*$) and then derive the minimum lift in CTR_0 (i.e., α_i^*) when targeting a specific age range. Table 5 shows that the minimum lift in CTR_0 increases non-linearly from 1.06 to 1.24 for different age ranges.

5.3.2. Gender targeting

Spotify Ad Studio also allows advertisers to target users based on their gender at no additional cost (i.e., $CPM_0 = CPM_g + DC_g = \$11.00$, where $g \in \{\text{male, female}\}$) and reports an equal number of available males and females (i.e., 1.5 million of each). Using the same parameters as those in Section 5.3.1, Equation (8) leads to a minimum lift in CTR_0 of 1.05 for males and females. In other words, when targeting a specific gender, the advertiser needs to see an increase of at least 5% in CTR_0 to achieve the same profit as under no-targeting.

5.3.3. Platform/device targeting

Spotify also allows targeting users based on their listening platform/device (i.e., iOS (iPhone/iPad), Android, or Desktop), albeit for an additional cost (i.e., $CPM_p + DC_p = \$12.00$, where $p \in \{\text{iOS (iPhone/iPad), Android, Desktop}\}$). Table 5 shows that targeting iOS, Android, and Desktop users requires a lift of at least 1.07, 1.09, and 1.05 in CTR_0 , respectively.

5.3.4. Additional targeting options

Finally, advertisers can target audiences at a more granular level based on their general interests (e.g., cooking, fitness, travel), current playlists (e.g., chill, study, workout), and recently listened to genres (see Table D1 of Online Appendix D). Fig. 3 illustrates the distribution of the minimum lift in CTR_0 for those audience segments (see Figure D3 in Online Appendix D for detailed results). The insights from Fig. 3 match the findings from our model: The minimum required lift in CTR_0 increases dramatically (see the bars and their respective values on the left) as the respective reach of the audience segment decreases (see the dashed lines and their respective values on the right).

Our advertiser must achieve a targeted CTR_0 lift of at least 2.25 (with considerable variation across audience segments), indicating an average increase of 125% in CTR_0 when targeting one of those audience segments. This increase is notably high compared to findings from previous studies, where the potential increase in CTR_0 barely reaches 100% (Aziz & Telang, 2016; Farahat & Bailey, 2012; Rafeian & Yoganarasimhan, 2021). Approximately half of the audience segments on Spotify require a higher increase in CTR_0 for the advertiser, suggesting they might be less profitable than no-targeting (see Table 5 and the bar charts exceeding the grey area in Fig. 3). Additionally, achieving a 100% CTR_0 increase is only feasible if audience reach exceeds 10% of the untargeted population. Narrow segments ($\leq 5\%$ reach) would require an increase in CTR_0 of over 146%, making them unprofitable in most settings.

These insights leave our advertiser with around half of the audience segments as potentially profitable for targeting. Specifically, our advertiser may want to target users based on age range, gender, platform/device, interests (i.e., interested in podcasts, tech, in-car listening, culture & society, parenting, comedy, gaming, and health & lifestyle), and genres (i.e., pop, hip hop, rock, indie rock, alternative, and electronica).

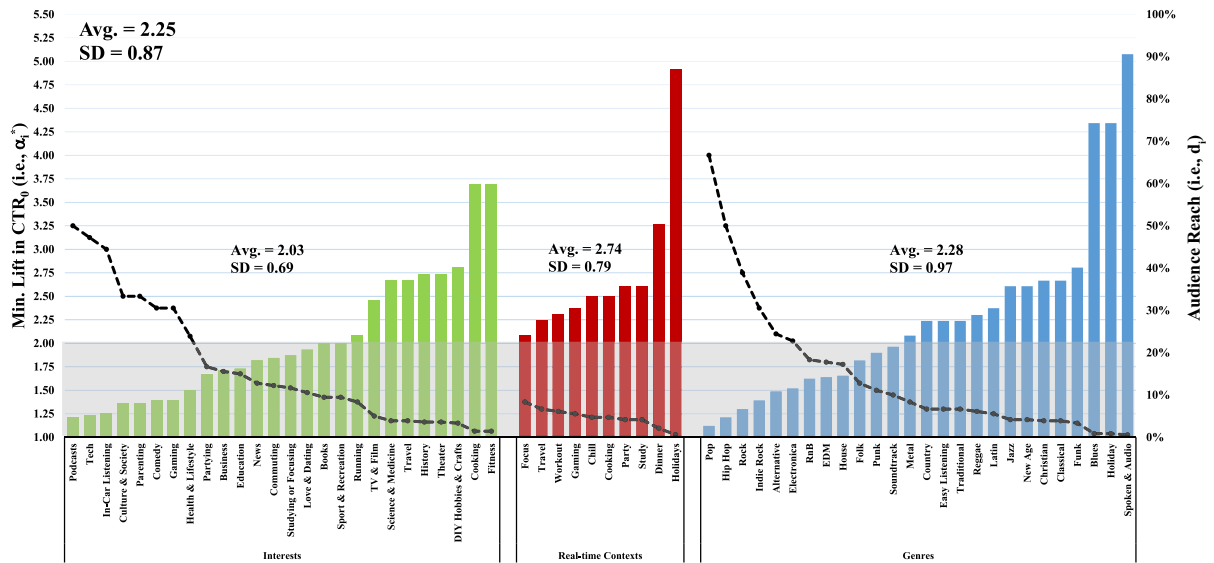


Fig. 3. Hypothetical Advertiser on Spotify (with $\alpha_i = \beta_i = \gamma_i$): Distribution of Minimum Lift in CTR_0 (i.e., α_i^*) and Reach Across Different Audience Segments. Notes: Values of the bar charts are on the left-hand side; values of the dashed lines are on the right-hand side. Avg.: Average; SD: Standard deviation; Min.: Minimum. We use the following values to derive the results: Number of impressions (N_0) = 1,800,000; click-through rate (CTR_0) = 1.00%; conversion rate (CR_0) = 2.00%; margin per conversion (m_0) = \$200.00; CPM under no-targeting (CPM_0) = \$11.00; CPM under targeting an audience segment i ($CPM_i + DC_i$) varies from \$11.00 to \$15.00; subscript 0 refers to no-targeting, and subscript i refers to targeting audience segment $i \in I$.

5.4. Top-up targeting

One way to avoid problems associated with narrow audience segments is to use top-up segments: Advertisers can combine (i.e., top-up) multiple (similar) narrow audience segments so that they are targeting all users who are part of *at least one* of the segments (top-up combines segments via an “OR”).

To better illustrate how topping-up can help lower the minimum required lift in CTR_0 , we combined segments from Fig. 3 (see Table D2 of Online Appendix D). For instance, Spotify Ad Studio allows advertisers to combine groups of users who are interested in “Fitness” with those interested in “Fitness/Health & Lifestyle/Running,” thereby generating a broader audience (in this case, at no extra cost).⁶ This combination of narrow audiences increases the reach from 1.44% to 30.00%, such that the minimum increase in CTR decreases from 270% (i.e., $\alpha_{Fitness}^* = 3.70$) to a more realistic level of 40% (i.e., $\alpha_{Fitness/Health \& Lifestyle/Running}^* = 1.40$). A similar interpretation applies to other combinations of segments (see examples in Table D2 of Online Appendix D). Consequently, topping-up multiple narrow audience segments may help advertisers achieve a higher campaign reach and decrease the break-even performance necessary to make targeting as profitable as no-targeting.

However, topping-up multiple narrow audience segments may aggravate an already existing problem: poor data quality. For example, Table 5 shows that the sum of available users in each (mutually exclusive) age range is higher than the overall number of available users (i.e., 1.8 million), which is impossible. This discrepancy indicates a lack of accuracy in audience segments when Spotify Ad Studio profiles its users, which aligns with the results from Neumann et al. (2019). When combining multiple, inaccurate audience segments, the true reach of the audience may be far lower than expected or reported by the campaign management system. The following section describes how a modified version of our model can account for such inaccuracies.

5.5. Break-even performance when reported segment size is inaccurate

Targeting narrow segments becomes less profitable when the segments are of poor quality and do not contain the intended users (e.g., when a female segment includes males; also see Neumann et al., 2019). Our derived break-even performance then presents the lower bound for the minimum lift in performance (or CTR_0 , CR_0 , and m_0). In other words: The presence of such inaccuracy in the reach of the corresponding audience segments means that the true decrease in reach when targeting audience segment i may be even higher than reported by the platform. In turn, we would require an even higher lift in performance than that derived from our original model.

To account for potential inaccuracies resulting from poor data quality, we extend our model by introducing a parameter that allows for scaling down the reported reach using a modified version of Equation (8):

⁶ The joint reach of two segments is typically smaller than the sum of users in each segment individually. Therefore, advertisers will use the joint reach provided by the campaign management system and plug it into our model to calculate the break-even performance for topped-up segments.

$$(\alpha_i \beta_i \gamma_i)^* = \frac{1}{(1 - \theta_i) \delta_i - \frac{\text{CPM}_0 - (1 - \theta_i) \delta_i (\text{CPM}_i + \text{DC}_i)}{1,000(1 - \theta_i) \delta_i \text{CTR}_0 \text{CR}_0 \text{m}_0}} \quad (9)$$

where $\theta_i > 0$ reflects the level of inaccuracy in the reach of audience segment i in %. For example, when targeting only females, a $\theta_{\text{female}} = 20\%$ indicates that 20% of the users in the provided female audience segment are male. Advertisers may obtain such estimates from running, for example, surveys on the targeted population.

Equation (9) thus gives advertisers a view of the break-even performance they should expect when audience segment quality is poor. For instance, utilizing Equation (9) and setting the true decrease in the reach of our audience segments in Fig. 3 to an “optimistic” average of 20% (i.e., $\theta_i = 20\%$; see Neumann et al., 2019), we need a performance lift of 2.42 instead of 2.25.

Importantly, the level of inaccuracy in the reach of an audience segment becomes even more important for narrow audience segments compared to broader ones: For example, for the same level of inaccuracy in the reach of an audience segment (again, say $\theta_i = 20\%$), the minimum lift in CTR_0 of narrow audience segments (e.g., those segments with a reach of $\leq 5\%$), on average, increases from 214% to 238% (i.e., an increase of 24 percentage points in CTR_0). However, broader audience segments (e.g., those segments with a reach of greater than 5%) only require an average increase of 10 percentage points in CTR_0 . In the next section, we will test this theoretical finding in an empirical study that involves Apple’s introduction of the ATT framework.

6. Study 3: Apple’s App Tracking Transparency (ATT) framework and decreased data quality

The introduction of Apple’s ATT framework on April 26, 2021—which allows users to opt-out of third-party tracking—decreased data quality and massively affected targeting practices on Facebook (Kraft, Skiera, & Koschella, 2023; Target Internet, 2021). Specifically, ATT limits the use of third-party data under iOS 14.5 since it requires all apps to obtain the user’s consent for being tracked on third-party apps (Kesler, 2022). Using data from 86 advertising campaigns on Facebook’s platforms (i.e., Facebook, Instagram, Messenger, Facebook Audience Network), we study how the introduction of ATT affected CPMs and CTRs of broad versus narrow targeting strategies.

6.1. Data

Our data on these 86 campaigns comes from six advertisers representing diverse industries (e.g., education, food e-commerce, telecommunication, and investments). The data span from April 12, 2021 to May 9, 2021, which covers two weeks before and two weeks after the introduction of ATT. Not all campaigns ran for four weeks, making the data unbalanced. Yet, we made sure that all campaigns were active before and after the introduction of ATT (on avg., 9.90 (8.79) days before (after) the introduction of ATT per campaign). In addition, our dataset has information on ad prices (in € and measured as CPM; Avg. = 12.62, SD = 6.10) and CTRs (in %; Avg. = 0.71, SD = 3.50). It is daily data for each campaign and has information on the device (i.e., iOS (= 1) versus Android/Desktop (= 0); Avg. = 0.48 and SD = 0.50) and the “narrowness” of the targeting strategy (i.e., narrow (= 1) versus broad (= 0), as defined by the agency; Avg. = 0.47, SD = 0.50). We provide further summary statistics on the number of impressions, CPM, and CTR in Online Appendix E, Table E1.

6.2. Model specification

Our dataset thus allows us to differentiate (1) whether the campaigns ran *before* or *after* the introduction of ATT (i.e., pre- and post-introduction of ATT), (2) whether the introduction of ATT affected the campaigns (i.e., campaigns on iOS versus Android/Desktop devices, allowing us to create *treatment* and *control* groups), and (3) whether the targeting strategy is *narrow* or *broad* (allowing us to estimate heterogeneous treatment effects).

We run a DiD analysis to investigate the effect of ATT on CPM and CTR under the parallel trend assumption. Figure E1 (Online Appendix E) indicates that the parallel trend assumption is reasonable. Specifically, we estimate the following regression:

$$Y_{jadt} = \beta_0 + \beta_1 \times \text{narrow}_j + \beta_2 \times \text{iOS}_{jd} + \beta_3 \times \text{post}_t + \beta_4 \times \text{iOS}_{jd} \times \text{post}_t + \beta_5 \times \text{iOS}_{jd} \times \text{narrow}_j + \beta_6 \times \text{post}_t \times \text{narrow}_j + \beta_7 \times \text{iOS}_{jd} \times \text{post}_t \times \text{narrow}_j + \delta_a + \epsilon_{jadt}, \quad (10)$$

where the dependent variable, Y_{jadt} , is either the CPM or CTR of campaign $j = 1, 2, \dots, 86$, by advertiser $a = 1, 2, \dots, 6$, targeting device d (iOS versus Android/Desktop), on day t (from April 12, 2021 to May 9, 2021); iOS_{jd} is a dummy variable that takes a value of 1 when impressions (of campaign j) occurred on iOS devices and 0 on Android/Desktop devices; post_t is an indicator variable for the period in which the ATT treatment occurs (where $\text{post}_t = 1$ if date t is after April 26, 2021 and $\text{post}_t = 0$ otherwise); narrow_j is equal to 1 when the campaign’s targeting strategy was defined as narrow and 0 when defined as broad; δ_a are advertiser fixed effects.

6.3. Discussion of results

Table 6 shows the results for both outcome variables (first three columns for CPM, last four columns for CTR): Models (1) and (4) show the results without our narrowness moderator; Models (2) and (5) depict the results that include narrowness of the campaign as a moderator. Because of the skewed nature of CTR, we also estimate the effect of ATT on CTR using a pooled Poisson regression (Wooldridge, 2022) in Model (6). Finally, since the main effect of ATT may only arise a few days after its introduction (Kraft et al., 2023), we run Models (3) and (7) (using similar settings as Models (2) and (6), respectively) on a sub-sample of our data, using only campaigns that have at least one week of post-observations.

Columns (1) and (4) of Table 6 indicate that campaigns targeted at iOS devices have higher CPMs and lower CTRs (though these coefficients are not significant). We also observe heterogeneous treatment effects on narrowly versus broadly targeted campaigns: We see significantly higher CPMs of targeting narrow audience segments on iOS devices after the introduction of ATT (see Column (2) of Table 6). In addition, we find that campaigns that targeted narrow audience segments on iOS devices have significantly lower CTRs after the introduction of ATT (compared to those that targeted broad audiences; see Columns (5) and (6) of Table 6). Notably, the effects on CPM and CTR become stronger and more significant when we restrict our sample to campaigns that ran at least one week after the introduction of ATT, thereby increasing the likelihood that more users installed the iOS update (see Models (3) and (7) of Table 6).

After ATT's introduction, the higher CPMs and lower CTRs of narrow (versus broad) targeting strategies suggest that advertisers benefit less from targeting narrower audiences with growing restrictions on third-party data. This finding will most likely also apply to the 'US Banning Surveillance Advertising Act of 2022 bill' and other regulations restricting third-party data access and, in turn, decreasing data quality. In addition, this finding supports the predictions from our model: With higher levels of inaccuracy in audience segments (here, because of increased restrictions on utilizing third-party data), advertisers must carefully assess the performance of narrow audience segments. Using our extended model, which explicitly captures data inaccuracies, advertisers can make more informed decisions about which segments to select for their targeted advertising campaigns under such circumstances.

Table 6
Difference-in-difference Results: Effect of ATT on CPM and CTR of Narrow versus Broad Audience Segments.

Model	CPM in €			CTR in %			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Narrow	–	2.35* (0.032)	3.19* (0.033)	–	–0.28 (0.406)	–0.55 (0.166)	–0.67 (0.129)
iOS	1.57*** (0.000)	1.65*** (0.000)	1.16* (0.073)	0.10 (0.734)	–0.35 (0.330)	–0.42 (0.240)	–0.73* (0.088)
Post	0.35 (0.366)	1.01* (0.041)	1.63** (0.003)	–0.39 (0.114)	–0.61 (0.139)	–0.68* (0.040)	–0.89* (0.024)
iOS × post	0.07 (0.863)	–0.67 (0.142)	–0.45 (0.335)	–0.02 (0.945)	0.48 (0.202)	0.61 (0.101)	0.98* (0.029)
iOS × narrow	–	–0.15 (0.767)	0.38 (0.598)	–	0.99 (0.107)	1.15* (0.044)	1.66* (0.012)
post × narrow	–	–1.82* (0.032)	–2.05* (0.028)	–	0.44 (0.345)	0.48 (0.326)	0.37 (0.419)
iOS × post × narrow	–	1.50* (0.068)	1.90* (0.025)	–	–1.08* (0.087)	–1.30* (0.044)	–1.45* (0.027)
Advertiser fixed effects	yes	yes	yes	yes	yes	no	no
Constant	11.68*** (0.000)	10.66*** (0.000)	9.89*** (0.000)	0.86*** (0.000)	1.00** (0.004)	0.03 (0.928)	0.26 (0.528)
N	2,225	2,225	1,610	2,225	2,225	2,225	1,610
R ² / Log Likelihood	0.30	0.32	0.34	0.01	0.01	–3,492.19	–2,581.91

Notes: *p*-values in parentheses; **p* < 0.10, ***p* < 0.05, ****p* < 0.01, *****p* < 0.001; ATT: Apple's App Tracking.

Transparency, CPM: Cost per mille; CTR: Click-through rate; clustering standard errors at the campaign level were used for estimations (similar conclusions occur when using robust standard errors).

7. Conclusion

By highlighting the questionable profitability of many audience segments, this paper aims to help advertisers decide whom to target. While the targeting decision lies at the core of marketing (Narayanan & Manchanda, 2009), managers are still struggling to recognize and pick the most promising audience segments for their targeted ad campaigns. Our discussions with industry professionals confirm this impression because campaign managers tackle the targeting decision in different and often opposing ways (see Column (3) of Table A1).

We propose a model that enables advertisers to select the most promising audience segments for their campaigns and allows them to overcome the cold-start problem. We suggest using our model to calculate the break-even performance

for many segments and then order these segments by break-even performance from smallest to largest. The literature on targeting suggests that an increase in CTR larger than 100% is likely not attainable. Thus, we suggest the advertiser to focus on testing segments that require a smaller lift. While our model empowers advertisers to more strategically settle on the segments worth testing in an RCT, our findings also reveal that untargeted campaigns may yield higher profits. This finding is important for ad agencies who often face pressure from clients to target narrow audience segments, as we learned from our interviews with industry professionals. Our paper challenges this practice, as it overlooks the important role that data cost and reach play in profitability. Therefore, our paper provides ad agencies with important arguments for why they must meticulously evaluate the profitability of narrow audience segments before committing to them.

Finally, our model reveals that targeting narrow audiences is often problematic, as smaller reach leads to dramatically higher break-even performance. Poor data quality amplifies the problem since the true reach of those segments might be overestimated and require an even higher increase in performance. Thus, growing restrictions on third-party data make narrow audiences even less attractive since their CTR decreases more than that of broad audience segments.

Our proposed model has limitations, as it relies on inputs that may not be readily available to all advertisers, especially small/new advertisers with limited experience. Some inputs, such as performance estimates under no targeting, may require a relationship with data providers or running costly surveys on the targeted population. Missing information on these inputs can make it challenging to apply our model. However, even with basic estimates, our model gives advertisers a rough idea of the required performance lift for different audience segments. Furthermore, while the model helps limit current targeting options, it requires constant updating for future campaigns. Nevertheless, advertisers can use insights gained from the current campaign to refine inputs for future targeting decisions. Finally, we decided to compare targeted and untargeted campaigns based on profits, requiring implementation of some form of conversion tracking. We thereby show how advertisers can apply our model when conversions, as in sales, are hard to measure (like in Study 1). However, advertisers may seek other goals, such as brand building, or sales. Consequently, evaluating targeting strategies solely based on profits may not be meaningful for all advertisers, particularly given the stage of the product life cycle and the product category.

Data availability

We provide the dataset necessary to replicate Study 2. Yet, we cannot share the dataset for the other two studies (Study 1 & 3) due to NDAs that we have signed for them.

Declaration of Competing Interest

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 833714).

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.ijresmar.2023.08.004>.

References

- Ada, S., Abou Nabout, N., & Feit, E. M. (2022). Context information can increase revenue in online display advertising auctions: Evidence from a policy change. *Journal of Marketing Research*, 59(5), 1040–1058. <https://doi.org/10.1177/00222437211070219>.
- Amaldoss, W., Jerath, K., & Sayedi, A. (2016). Keyword management costs and “broad match” in sponsored search advertising. *Marketing Science*, 35(2), 259–274. <https://doi.org/10.1287/mksc.2015.0919>.
- Apple (2022). User Privacy and Data Use. Retrieved from <https://developer.apple.com/app-store/user-privacy-and-data-use/>.
- Aziz, A., & Telang, R. (2016). What is a Digital Cookie Worth? SSRN Working Paper Series. Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2757325.
- Bayer, E., Srinivasan, S., Riedl, E. J., & Skiera, B. (2020). The impact of online display advertising and paid search advertising relative to offline advertising on firm performance and firm value. *International Journal of Research in Marketing*, 37(4), 789–804. <https://doi.org/10.1016/j.ijresmar.2020.02.002>.
- Beales, H. (2021). The Value of Behavioral Targeting. Retrieved from http://delivery.acm.org/10.1145/2960000/2956586/a4-hummel.pdf?ip=137.205.11.221&id=2956586&acc=OA&key=4D4702B0C3E38B35%2E4D4702B0C3E38B35%2E4D4702B0C3E38B35%2E5945DC2EABF3343C&__acm__=1557226973_29f1fdbf6ce5cb39a17ede7a256e138b.
- Choi, H., & Mela, C. F. (2019). *Display Advertising Pricing in Exchange Markets*.
- Choi, H., Mela, C. F., Balseiro, S. R., & Leary, A. (2020). Online display advertising markets: A literature review and future directions. *Information Systems Research*, 31(2), 556–575. <https://doi.org/10.1287/isre.2019.0902>.
- Du, X., Su, M., Zhang, X., & Zheng, X. (2017). Bidding for multiple keywords in sponsored search advertising: keyword categories and match types. *Information Systems Research*, 28(4), 711–722. <https://doi.org/10.1287/isre.2017.0724>.
- Ebbes, P., & Netzer, O. (2022). Using social network activity data to identify and target job seekers. *Management Science*, 68(4), 3026–3046. <https://doi.org/10.1287/mnsc.2021.3995>.
- Farahat, A., & Bailey, M. C. (2012). How Effective is Targeted Advertising? In *Proceedings of the 21st International Conference on World Wide Web* (pp. 111–120).
- Goldfarb, A., & Tucker, C. (2011a). Online display advertising: targeting and obtrusiveness. *Marketing Science*, 30(3), 389–404. <https://doi.org/10.1287/mksc.1100.0583>.
- Goldfarb, A., & Tucker, C. (2011b). Privacy regulation and online advertising. *Management Science*, 57(1), 57–71. <https://doi.org/10.1287/mnsc.1100.1246>.
- Google (2023). About Audience Reporting. Retrieved from <https://support.google.com/google-ads/answer/10566573>.
- de Haan, E., Wiesel, T., & Pauwels, K. (2016). The effectiveness of different forms of online advertising for purchase conversion in a multiple-channel attribution framework. *International Journal of Research in Marketing*, 33(3), 491–507. <https://doi.org/10.1016/j.ijresmar.2015.12.001>.

- Hercher, J. (2021). How Facebook is Overhauling its Attribution Standards to Deal with Apple's ATT. Retrieved from <https://www.adexchanger.com/mobile/how-facebook-is-overhauling-its-attribution-standards-to-deal-with-apples-att/>.
- Hui, S. K., Inman, J. J., Huang, Y., & Suher, J. (2013). The effect of in-store travel distance on unplanned spending: applications to mobile promotion strategies. *Journal of Marketing*, 77(2), 1–16. <https://doi.org/10.1509/jmr.11.0436>.
- Iyer, G., Soberman, D., & Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3), 461–476. <https://doi.org/10.1287/mksc.1050.0117>.
- Jerath, K., Ma, L., & Park, Y.-H. (2014). Consumer click behavior at a search engine: The role of keyword popularity. *Journal of Marketing Research*, 51(4), 480–486. <https://doi.org/10.1509/jmr.13.0099>.
- Johnson, G. A., Lewis, R. A., & Reiley, D. H. (2017). When less is more: Data and power in advertising experiments. *Marketing Science*, 36(1), 43–53. <https://doi.org/10.1287/mksc.2016.0998>.
- Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer privacy choice in online advertising: who opts out and at what cost to industry? *Marketing Science*, 39(1), 33–51. <https://doi.org/10.1287/mksc.2019.1198>.
- Johnson, J. P. (2013). Targeted advertising and advertising avoidance. *The RAND Journal of Economics*, 44(1), 128–144. <https://doi.org/10.1111/1756-2171.12014>.
- Kesler, R. (2022). The Impact of Apple's App Tracking Transparency on App Monetization. SSRN Working Paper Series. Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4090786.
- Kraft, L., Skiera, B., & Koschella, T. (2023). *Economic impact of opt-in versus opt-out requirements for personal data usage: The case of apple's app tracking transparency (ATT)*. Working Paper.
- Lambrecht, A., & Tucker, C. (2013). When does retargeting work? information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561–576. <https://doi.org/10.1509/jmr.11.0503>.
- Lu, X., & Zhao, X. (2014). Differential effects of keyword selection in search engine advertising on direct and indirect sales. *Journal of Management Information Systems*, 30(4), 299–326. <https://doi.org/10.2753/MIS0742-1222300411>.
- Meta (2022a). About A/B Testing. Retrieved from <https://www.facebook.com/business/help/1738164643098669?id=445653312788501>.
- Meta (2022b). About Highest Volume. Retrieved from <https://www.facebook.com/business/help/721453268045071?id=2196356200683573>.
- Meta (2023). Core Audience Targeting. Retrieved from <https://www.facebook.com/business/news/Core-Audiences>.
- Molitor, D., Spann, M., Ghose, A., & Reichhart, P. (2020). Effectiveness of location-based advertising and the impact of interface design. *Journal of Management Information Systems*, 37(2), 431–456. <https://doi.org/10.1080/07421222.2020.1759922>.
- Narayanan, S., & Manchanda, P. (2009). Heterogeneous learning and the targeting of marketing communication for new products. *Marketing Science*, 28(3), 424–441. <https://doi.org/10.1287/mksc.1080.0410>.
- Neumann, N. (2020). After Google Finally Really did it, a Necessary System Cleanup. Retrieved from <https://www.adexchanger.com/data-driven-thinking/after-google-finally-really-did-it-a-necessary-system-cleanup/>.
- Neumann, N., Tucker, C., & Whitfield, T. (2019). Frontiers: How effective is third-party consumer profiling? *Evidence from Field Studies*. *Marketing Science*, 38(6), 918–926. <https://doi.org/10.1287/mksc.2019.1188>.
- Phang, C. W., Luo, X., & Fang, Z. (2019). Mobile time-based targeting: Matching product-value appeal to time of day. *Journal of Management Information Systems*, 36(2), 513–545. <https://doi.org/10.1080/07421222.2019.1598696>.
- PwC (2023). Internet Advertising Revenue Report: Full-year 2022 Results. Retrieved from <https://www.iab.com/insights/internet-advertising-revenue-report-full-year-2022/>.
- Rafieian, O., & Yoganarasimhan, H. (2021). Targeting and privacy in mobile advertising. *Marketing Science*, 40(2), 193–218. <https://doi.org/10.1287/mksc.2020.1235>.
- Rahmey, M. (2021). How Apple's ATT Framework and the iOS 14 Changes are Impacting Advertisers & Paid Social Platforms. Retrieved from <https://disruptivedigital.agency/how-apples-att-framework-and-the-ios-14-changes-are-impacting-advertisers-paid-social-platforms/>.
- Schuh, J. (2019). Building a More Private Web. Retrieved from <https://www.blog.google/products/chrome/building-a-more-private-web/>.
- Shehu, E., Abou Nabout, N., & Clement, M. (2021). The risk of programmatic advertising: effects of website quality on advertising effectiveness. *International Journal of Research in Marketing*, 38(3), 663–677. <https://doi.org/10.1016/j.ijresmar.2020.10.004>.
- Skiera, B., & Abou Nabout, N. (2013). Practice prize paper—PROSAD: A bidding decision support system for profit optimizing search engine advertising. *Marketing Science*, 32(2), 213–220. <https://doi.org/10.1287/mksc.1120.0735>.
- Skiera, B., Miller, K., Jin, Y., Kraft, L., Laub, R., & Schmitt, J. (2022). *The impact of the General Data Protection Regulation (GDPR) on the online advertising market*. Frankfurt.
- Target Internet (2021). How Apple iOS14 Changes Have Affected Facebook Third-Party Tracking. Retrieved from <https://www.targetinternet.com/resources/how-apple-ios14-changes-have-affected-facebook-third-party-tracking/>.
- Terlep, S., & Seetharaman, D. (2016). P&G to Scale Back Targeted Facebook Ads. Retrieved from <https://www.wsj.com/articles/p-g-to-scale-back-targeted-facebook-ads-1470760949>.
- Tucker, C. (2014). Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research*, 51(5), 546–562. <https://doi.org/10.1509/jmr.10.0355>.
- Wooldridge, J. M. (2022). Simple Approaches to Nonlinear Difference-in-Differences with Panel Data. SSRN Working Paper Series. Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4183726.
- Yan, J., Liu, N., Wang, G., Zhang, W., Jiang, Y., & Chen, Z. (2009). How Much Can Behavioral Targeting Help Online Advertising? In: *Proceedings of the 18th International Conference on World Wide Web* (pp. 261–270). <https://doi.org/10.1145/1526709.1526745>.
- Yaveroglu, I., & Donthu, N. (2008). Advertising repetition and placement issues in on-line environments. *Journal of Advertising*, 37(2), 31–44. <https://doi.org/10.2753/JOA0091-3367370203>.