Multidimensional Targeting and Consumer Response

Stylianos Despotakis and Jungju Yu*

November 18, 2022

Abstract

Advancements in targeting technology have allowed firms to engage in more precise targeting based on several aspects of consumers' preferences. Exposed to more targeted ads, consumers are becoming increasingly aware of being targeted and respond accordingly. This paper provides a theoretical analysis of multidimensional targeting under which consumers can draw inferences about multiple components of their utility from the advertised product. We show that the firm can be worse off under multidimensional targeting than under single-dimensional targeting, in which the firm targets consumers based only on a single component of their utility. This is because, with multidimensional targeting, targeted consumers may face greater uncertainty about on which specific dimension(s) they can expect to enjoy the advertised product. Therefore, they may be less willing to exert a costly effort of clicking the ad and making a purchase decision. When this result holds, the firm may want to adopt a single-dimensional targeting strategy. However, we show that the firm cannot credibly commit to such a strategy once given access to multiple dimensions of customer data. Interestingly, a higher unit cost of advertising can mitigate the firm's commitment problem for utilizing customer data and, thus, increase the firm's profit. Moreover, the firm can sometimes lower the price to recover some of, but not entirely offset, the drawbacks of multidimensional targeting. We discuss the implications of our results regarding the current practice of targeted advertising and data privacy protection policies.

Keywords: Multidimensional targeting; Data granularity; Endogenous consumer response; Click-through rate; Commitment problem; Targeted online advertising.

^{*}Stylianos Despotakis: sdespota@cityu.edu.hk (City University of Hong Kong); Jungju Yu: jungju.yu@kaist.ac.kr (Korea Advanced Institute of Science and Technology). The authors thank Liang Guo, Woocheol Shin, and Jiwoong Shin for helpful feedback and the seminar participants at the 2021 INFORMS Marketing Science Conference, City University of Hong Kong, Korea Advanced Institute of Science and Technology, Korea University, Seoul National University, and the Marketing Theory Group in Hong Kong. The authors would also like to thank Wangnong Liao for providing excellent research assistance.

1 Introduction

Firms today have access to increasingly richer customer data from various sources. By applying advanced analytic tools to the granular data, firms can understand several distinct aspects of underlying customer preferences, which is then used for targeted advertising and personalized recommendations. For example, on Facebook and Instagram, restaurants can decide whom to target based on information about consumers' favorite cuisine, location, and atmosphere. Netflix provides personalized recommendations for movies based on the viewers' preferred genre (romantic, comedy, horror), plot (sequence, suspense), cast, and nationality.¹ Targeted advertising based on more granular customer data allows for significant improvements in targeting precision and thus reach consumers who are more likely to benefit from the advertised product or service. Therefore, it may seem evident that firms can only benefit from more granular customer data and a more comprehensive understanding of consumers' preferences.

In fact, firms have invested a massive amount of resources into collecting more data and adopting advanced analytics to their marketing practice. However, in reality, many firms have failed to reap the benefits of their investments in online advertising (e.g., Lambrecht and Tucker 2015). One of the oft-cited reasons for this failure is that most organizations do not know how exactly to use their vast amount and various types of customer data for better outcomes (Redman, 2017). For instance, should a restaurant send its ads selectively to those who enjoy the restaurant for both cuisine and location? Or, should it relax the targeting criterion and send an ad as long as a consumer will like either one of cuisine and location? As the granularity of customer data increases, figuring out how to best utilize the data has turned out to be a daunting task for firms.

The consumer-side story is not so hopeful, either. Despite the substantial advancements in targeting technology, both industry reports and academic research have found that online targeted ads may not be as effective in generating consumer responses as commonly believed. In 2018, across all industries, the average click-through rate of online display ads was only 0.46% and sponsored search ads 3.17%.² A growing body of recent research on digital advertising using field experiments has found that online targeted advertising may not be effective at all (e.g., Blake et al. 2015).

¹ "How Netflix's Recommendation Engine Works" (article): https://medium.com/@springboard_ind/how-netflixs-recommendation-engine-works-bd1ee381bf81

² "Understanding Click-Through Rate for PPC" (report): https://www.wordstream.com/click-through-rate

Given these challenges of optimizing and monetizing on customer data of intricate and multidimensional nature, this paper seeks to analyze a firm's optimal multidimensional targeting strategy. Moreover, we analyze the marginal value of having more granular customer data and answer an important question, "Will the firms necessarily be more profitable when given access to more dimensions of customer data for targeted advertising?" In particular, we examine the effect of multidimensional targeting on consumers' response to ads, i.e., their click decision, and how this effect may depend on the granularity of customer data. Our analysis of consumers' response to being targeted offers a rationale for why consumers might be reluctant to respond to ads even though the advertisers target the ads based on several components of customer preferences.

To answer these questions, we analyze the firm's optimal targeting under the environment of multidimensional consumer data. A firm is given access to data on two distinct components of consumer's utility, such as consumers' preferences for location and cuisine, and decides to whom to send an ad. We identify the firm's equilibrium strategy for multidimensional targeting and compare the equilibrium profits with the profit under a benchmark case of a single-dimensional environment where the firm only knows consumers' preferences for a single component.

Specifically, a monopolist sells a new product to consumers who draw utility from two distinct components. Consumers are uncertain about their component-wise utilities, but the firm might know them due to the informational advantage of detailed data on customers and product information, which consumers do not know ex-ante. So, the firm perceives consumers in two dimensions in the multidimensional targeting environment and one dimension in the single-dimensional environment. Based on each consumer's perceived type, the firm makes a targeting decision, i.e., to which individuals to send an ad at a positive marginal unit cost. Each consumer's consumption utility is determined by a weighted average (heterogeneous across consumers) of her componentwise utilities on the two dimensions. The weighted average represents the relative importance that a consumer places between the two components. The relative weight captures the consumer's idiosyncratic needs in a particular purchase occasion, which is unlikely to be predictable from the firm's perspective and, hence, private information to the consumer. Each consumer knows that the firm can use customer data for targeting, but she cannot observe how exactly the firm will use it. Thus, when targeted, a consumer updates her prior beliefs about her component-wise utilities according to her expectations about the firm's targeting strategy. Based on the posterior beliefs, the consumer decides whether to exert a costly effort of clicking the ad and learn about her exact utility from the product. She will then buy the product if her realized utility is no less than the price.

Identifying an equilibrium of multidimensional targeting can be analytically and computationally challenging because the strategy can be any arbitrary subset of the multidimensional consumer space. The firm chooses the set of consumers to whom to send ads and maximize its expected profits, given the consumers' anticipation for the firm's decision and the unit cost of advertising. The targeted consumers first decide whether to click the ad, considering their expectations about the firm's targeting strategy and the effort cost of clicking. Then, those who click the ad decide whether to buy the product, given their realized utility and the product's price. In equilibrium, the firm's optimal choice of targeting must coincide with consumers' expectations.

We show that the firm's equilibrium profit under multidimensional targeting can be lower than that under single-dimensional targeting. The intuition is as follows: In equilibrium, the firm targets a relatively larger set of consumers whose component-wise utility is high on at least one of the two dimensions. In contrast, under single-dimensional targeting, the firm sends add to a smaller set of consumers whose utility is sufficiently high on the particular single dimension. Therefore, when the firm's targeting is based on both dimensions of consumer information, a targeted consumer may be more uncertain about her expected utility from the product because her updated beliefs diffuse on a larger region of the firm's targeting set. In particular, she may be uncertain about which specific dimension (e.g., location vs cuisine of a restaurant) she might enjoy the advertised product more. The targeted consumers' uncertainty translates into their lower expected utility, which discourages consumers from exerting the costly effort of clicking. This reduces the clickthrough rate and the conversion rate under the multidimensional targeting environment. However, under single-dimensional targeting (e.g., location), a targeted consumer understands that she can expect to enjoy her matching value in that respective dimension, which leaves less uncertainty for the targeted consumer. Thus, consumers under the single-dimensional targeting environment can make more informed clicking decisions, thus resulting in higher click-through rate and conversion rate. Consequently, the firm's ad effectiveness and profit can be higher when the firm has access to less granular customer data.

When this counter-intuitive result holds, one might ask, "Why can't the firm under the multidi-

mensional targeting environment commit to sending ads to a small group of consumers whose utility from the product is high enough on both dimensions?" In this environment, the firm knows various aspects of consumer utility and, therefore, finds it tempting to target a larger group of consumers who could like the product for different reasons. On the other hand, under the single-dimensional targeting environment, the firm can only rely on a particular aspect of consumer utility (e.g, location), and must choose the targeting set with the other aspect (e.g., atmosphere) unknown. Then, it is in the firm's best interest to keep its targeting set small, consisting of consumers who are likely to enjoy the product enough. In other words, the firm's lack of access to granular customer data can act as a commitment device to send targeted ads to a select few, thus enhancing consumers' response to the ads.

The firm's commitment problem regarding its advertising decision, which makes the multidimensional targeting environment less appealing, is alleviated when it is expensive to send more ads. A higher cost of advertising will provide the firm with more commitment power and make multidimensional targeting more profitable than single-dimensional targeting. The extent of friction on the consumer side also plays a vital role in our main results. Suppose the effort cost of consumers clicking the ad is minimal. In that case, the benefit of the single-dimensional targeting in its ability to encourage the consumers' decision to click the ad by targeting a more select few diminishes. Therefore, the consumers' click cost should not be too low.

We first draw these insights on the value of multidimensional targeted advertising under exogenous pricing. Later, we extend our main model with the firm's endogenous pricing decision and show that our main results are robust. In this extension the price is announced in the advertisement, therefore the firm can commit to a price before consumers visit the firm's website. In this case, we find that sometimes the firm charges a lower price under multidimensional targeting than under single-dimensional targeting to increase the conversion rate. This way, pricing can partly recover, but not fully offset, the drawbacks of multidimensional targeting.

In summary, this paper offers a formal analysis of the optimal targeted advertising and identifies the marginal value of the granularity of consumer data. It also provides new insights into how the additional granularity of customer data affect consumers' inferences from the targeted ads and their willingness to respond to them. Moreover, this paper provides a rationale for why consumers today may demonstrate hesitance to respond to the ads even though the ads are based on increasingly



Figure 1: Tools that help consumers become more aware of being targeted

granular data on consumers. The current practices in online targeted advertising are consistent with the implications of our findings. On social media platforms like Facebook and Instagram, which together play dominant roles in online display advertising, ads carry a small label "sponsored" right below the name of the advertiser, as shown in Figure 1. Such a seemingly subtle and unnoticeable label has been shown to induce the viewers' understanding that the ad has been targeted to them (Summers et al., 2016). Our finding suggests that the firm can be worse off with multidimensional targeting with this label alone. However, viewers can see additional information by clicking the link, "Why am I seeing this ad?" This additional function on display ads improves the transparency of the firm's targeting strategy, which reduces consumers' uncertainty, albeit imperfectly. So, the advertisers may mitigate the potential pitfalls of, and better monetize on, more granular customer data through targeted advertising.

Next, we discuss the related literature. The rest of the paper is organized as follows: Section 2 introduces the model, and Section 3 analyzes the single-dimensional and multidimensional targeting environments. Section 4 compares the two environments and presents the paper's main results. Section 5 extends the main analysis to the firm's endogenous pricing problem and an alternative environment of consumer information and establishes robustness of the main results. Section 6 concludes. Proofs of all the results in the main text are in Appendix A.³

³Appendix A also includes additional extensions and numerical examples.

Related Literature

The early theoretical research has analyzed firms' optimal use of customer information in various settings, such as targeted pricing (e.g., Chen et al. 2001; Iver et al. 2005), pricing of advertisements (Bergemann and Bonatti, 2011), and consumers' nuisance costs from receiving ads (Johnson, 2013). This paper is related to the body of research investigating the effect of more accurate targeting on the equilibrium outcomes. Chen et al. (2001) analyzed a model of targeted pricing and showed that less precise targeting could improve the firms' profits. This is because firms are more likely to mistake price-sensitive consumers (switchers) for price-insensitive consumers (loyal customers), which leads the firms to charge higher prices, thus softening price competition between the firms. Moreover, Rafieian and Yoganarasimhan (2021) found that the ad network's revenue can be nonmonotonic in targeting accuracy because, as posited by Levin and Milgrom (2010), more accurate targeting further differentiates advertisers, thus reducing the bidding competition between them. In this stream of research, the nature of the firm's customer data is captured by a single parameter of targeting precision. However, in reality, customer data is composed of several different dimensions (e.g., Ghose et al. 2012; Trusov et al. 2016; Fong et al. 2019; Long et al. 2022). To our knowledge, this paper offers the first theoretical analysis of multidimensional targeting and investigates the marginal value of additional dimensions of customer data.⁴

More recently, the literature has accumulated a substantial amount of empirical evidence that using customer information for marketing activities can be effective in various situations, whether it is mobile targeting (Andrews et al., 2016), recommendation systems (Ghose et al., 2012), or display advertising (Manchanda et al., 2006; Braun and Moe, 2013; Hoban and Bucklin, 2015). However, some field experiments (Blake et al., 2015; Lewis and Rao, 2015) designed to estimate the causal effect of targeting have not been able to identify the significant effect of targeting compared to non-targeting. In Blake et al. (2015), sponsored search ads are shown to have heterogeneous

⁴Several other studies have found theoretically that firms unconditionally benefit from more precise consumer information. Iyer et al. (2005) analyzed two extreme cases of perfect targeting and no targeting for both advertising and pricing in a duopoly market and found that perfect targeting of advertising is always more profitable than when ads are not targeted at all (irrespective of whether prices are targeted or not). Bergemann and Bonatti (2011) studied a market of advertising and showed that more precise information on consumers always improves social welfare. However, it may decrease the price of ads as a more precise targeting rotates (instead of shifting) the demand curve for advertising. Finally, Johnson (2013) considered a model of targeted ads with consumers' nuisance cost from viewing ads and found that a higher precision in targeting always increases firms' profits, though its impact on consumers is ambiguous.

effects on consumers, and in particular, most consumers who click the ads are those who would have clicked a firm's organic link anyways. This yields a striking result that search ads are not effective at all on average, thus casting doubts on the conventional wisdom that online targeted ads are effective. This paper examines consumer inferences and decision processes. It offers an alternative explanation for why consumers' might be reluctant to respond to targeted ads based on a rich set of customer data, especially when it is composed of several distinct data types.

Previous research has shown that a major reason for consumers' hesitance to engage with targeted ads online is that customers may not trust firms with customer data (Goldfarb and Tucker, 2011; Kim et al., 2019). Morey et al. (2015) noted that the flood of customer data presents enormous opportunities for firms to abuse customer information. They also find that consumers are mostly unaware of which types of information are used. In other words, firms' accurate advertising decisions are primarily unobservable to consumers. This paper explores the firms' incentives for reneging consumers' expectations and using customer data in opportunistic ways.⁵ It further investigates how these incentives affect consumers' trust in and engagement with targeted ads.

Unlike the previous literature in targeted advertising which treats consumers as largely unaware agents, a few recent studies have examined the role of consumers being more conscious about being targeted. Summers et al. (2016) showed experimental evidence that consumers react differently to targeted ads than they do to the otherwise-equivalent non-targeted ads, thus indicating that consumers can distinguish the situations when they are targeted and when they are not. Shin and Yu (2021) provided a theoretical analysis of a setting in which consumers may make favorable inferences from the mere fact that they are targeted and thereby become more interested in a product category. This paper follows this line of research. In contrast to previous work, we consider multiple components of consumer utility (e.g., food taste and atmosphere of a restaurant). Hence, consumers update their beliefs about their expected utility in all dimensions and decide whether to click the ad. We show that targeted advertising on multiple dimensions of consumer preferences may leave consumers more uncertain about their match utility with the product. Thus, consumers may be hesitant to click the ad, which impairs the effectiveness of multidimensional targeting.

⁵Firms' opportunistic use of customer data more commonly refers to behaviors that negatively influence consumers' privacy, such as selling customer data to a third-party data broker. However, this paper focuses on a different type of opportunistic behavior that arises from a firm's lack of commitment to using the data for targeting purposes.

One can view the mechanism identified in this paper as related to the mechanisms in other research that show that adding noise to the information between the firm and the consumers can benefit the firm. Mayzlin and Shin (2011) studied a monopolist that can be one of three vertically differentiated types (high, medium and low) and identified an equilibrium in which the high-type firm pools with the low type firm by restraining from advertising any of its strengths. This advertising strategy makes consumers' information about the firm's type noisier, encouraging consumers to incur a search cost and learn about the firm's true type. This makes the consumers' discovery of the high-type firm more likely. Rhodes (2015) found that a monopolist's retailing (i.e., a decision to carry multiple products as opposed to selling a single product) leads to more heterogeneous valuations among the visiting consumers, which provides the firm offers a large selection of products encourages more consumers to incur the cost of visiting the store. In this paper, we investigate the effect of the granularity of consumer information in the context of targeted advertising.

2 Model

We analyze a model of targeted advertising in which a firm may have access to customer data of different levels of granularity and utilize it in costly targeted advertising. Thereby, we investigate the marginal value of granularity in customer data to the firm and its effect on consumers' information and decisions. Consider the following scenario: A user sees a sponsored ad for a restaurant on Instagram. The user understands that Instagram has an abundance of information on her given her extensive use and activities on the social platform, which means that this ad could be targeted based on a comprehensive understanding of what she likes in general. However, in this case, she did not know, among several aspects that she cares about in a restaurant (e.g., cuisine and location), which components she can expect to enjoy dining there more. Given this uncertainty, she was reluctant to pay closer attention to, and click, the ad and look for more information about the advertised restaurant. This scenario, which consumers commonly experience in various settings, such as ads for streaming contents, and mobile apps, demonstrates a phenomenon for which we want to provide a rationale. Though the targeted ads may be based on granular customer data, consumers may remain hesitant to respond to the ads. The main ingredients of the model are as follows: a firm and a unit mass of consumers interact in the market. The utility of a consumer for the firm's product consists of two distinct components, and it is of the following form:

$$u = \omega \cdot \alpha + (1 - \omega) \cdot \beta, \tag{1}$$

where α and β denote the consumer's utility from the two components. The component-wise utility measures the extent to which the firm's product fits the consumer's preferences for the attribute. In the restaurant example, α can be how much the consumer likes the restaurant's location, and β is the same for the food and cuisine. For each consumer, α and β are uniformly and independently drawn from an interval [0, 1]. The firm has information about the product that consumers do not have before clicking on an ad. Moreover, the firm can apply advanced analytics to the detailed data aggregated across the customer base and predict how much the consumers will like the product. Therefore, the firm has an ex-ante informational advantage over each consumer about the consumer's component-wise utility α and β .⁶ For instance, when a consumer sees a display ad or search ad of a new restaurant, the consumer does not know how close or far the restaurant is located from her. However, the restaurant's ad is based on relatively accurate information about the consumer's location based on GPS, nearby Wi-Fi hotspots, cell towers, and IP addresses. So, we capture this informational advantage of the firm by having the firm know the consumer's component-wise utilities α and (sometimes) β , whereas consumers do not know ex-ante.⁷

Between the two components of consumer utility, consumers place their weight $\omega \in [0, 1]$. We consider ω as the consumer's private information, determined by the consumer's idiosyncratic needs in a particular purchase occasion, which is unlikely to be predictable from the firm's perspective. For example, which component is more important between location and cuisine can depend on whether the consumer is out for an ordinary gathering with close friends or a romantic date dinner with their spouse.⁸ Each consumer's heterogeneous weight ω is drawn from a uniform distribution

 $^{^{6}}$ Long et al. (2022) also formulates two distinct components in consumer utility, the product quality and the consumer's preferences. They study a problem different from ours: a platform's placement of organic links. The sellers have private information about their products' quality, whereas the platform knows about consumers' preferences that the sellers do not know.

⁷This is a simple way to capture the informational advantage of the firm. One can consider an intermediate level of the firm's information advantage where the firm holds a more precise distribution for α and β than the consumer. However, as long as the extent of the firm's information advantage is sufficiently large, the analysis and results of this paper are robust.

⁸In practice, the firm might have some information about the ω of a particular consumer. However, the results in the current analysis are robust as long as the extent of information asymmetry regarding ω is significant enough.

on [0, 1].

Based on the information the firm has about each consumer, the firm decides whether to send an advertisement to each consumer by paying a unit cost $c_{ad} > 0$. In a simpler benchmark model, the firm only knows α and not β and, therefore, engages in *single-dimensional* (*sd*) targeting strategy. In general, the firm's single-dimensional targeting strategy is defined by a subset of consumers of matching value α for which the firm sends an ad, i.e., $\Sigma^{sd} \subseteq [0, 1]$. In the main model, the firm has access to more granular customer data and knows both α and β . Accordingly, the firm implements a *multidimensional* (*md*) targeting strategy. The firm's strategy is defined by the set of consumers, identified by their matching values α and β , for whom the firm sends an ad, i.e., $\Sigma^{md} \subseteq [0, 1] \times [0, 1]$.⁹

Advertising raises consumers' awareness of the firm's product.¹⁰ Consumers are not able to directly observe the firm's advertising decision Σ , and instead have an expectation $\tilde{\Sigma}$. Based on this expectation, the consumers decide whether to click the ad by exerting an effort cost $c_{\text{click}} > 0$. Consumers are aware of the targeting environment. So, under the *sd*-environment, consumers understand that targeting is on one dimension, say α without loss of generality.¹¹ Accordingly, upon being targeted, they update beliefs about their utility on that dimension, while their beliefs about the other dimension β remain the same as the prior. Under the *md*-environment, targeted consumers will update their beliefs on both dimensions of their utility according to their expectation about the firm's targeting strategy $\tilde{\Sigma}$.¹² The consumers' posterior distributions are denoted by

⁹Having access to more granular customer data leads to the firm's more comprehensive understanding of distinct components of each consumer's utility. As discussed in Section 1, this approach is different from the approach in most prior research in targeted advertising where a greater data granularity results in a more precise knowledge about a consumer's single component of utility.

¹⁰If a consumer does not see an ad, then the consumer would not be aware of the firm's product, so she will no longer participate in the game. More precisely, every consumer is initially aware of the firm's existence and has a common prior about her expected utility, such as α and β . However, without the ad, the consumer remains unaware of the firm's specific product.

¹¹In Section 5.2, we consider an alternative setting in which consumers under the single-dimensional environment are uncertain about which one dimension the firm knows. We show that the results from the main analysis are robust.

¹²Our assumption on the consumers' information that they are generally aware of which dimensions of customer information the firm has access to reflects reality. While the nature of customer data available to advertisers has become a lot more complex, it is often up to the consumers to decide which information will be accessible to the businesses. For instance, when signing up for a new Google account, each user can choose whether to share their contacts, locations, Web and App activities. Moreover, consumers are largely aware that their activities on different platforms or e-commerce, such as Facebook and Amazon, are used for targeted advertising and personalized recommendations. For instance, Amazon knows which products consumers searched for and purchased. So, Amazon shows more ads about products related to what the users have previously purchased, which consumers know. Facebook knows more about consumers' lifestyles and social connections. Therefore, the users are exposed to more ads for products, events, and pages reasonably consistent with the users' lifestyles and cultural affiliations, which again consumers understand.

 $f^{sd}(\cdot)$ and $f^{md}(\cdot)$, whose supports are both $[0,1] \times [0,1]$. In equilibrium, this expectation must coincide with the firm's optimal decision for targeted advertising, $\Sigma = \widetilde{\Sigma} = \Sigma^*$.

Given the consumers' expectation for the firm's advertising decision, each consumer's click decision depends on ω , her relative importance between the two components of her utility. Anticipating their purchase decision and realized utility, a subset of the consumers of type ω who decide to click among those targeted is denoted by $\Omega^{sd}(\widetilde{\Sigma}), \Omega^{md}(\widetilde{\Sigma}) \subseteq [0, 1]$. After clicking the ad, the consumer is taken to the product webpage and obtains more information about the product from which she learns her realized utility $u = \omega \cdot \alpha + (1-\omega) \cdot \beta$. She then purchases if $u \ge p$, where p is the product's price. Otherwise, she leaves the market, securing the outside option normalized to 0 utility. The firm gets the expected revenue from the consumers to whom it sends ads (Σ) and subsequently click ($\Omega(\widetilde{\Sigma})$) and buy the product. Subtracting the total ad cost from the expected revenue yields the firm's expected profit.

The sequence of the *sd*-targeting game (and the *md*-targeting game) is as follows. In Stage 0, the firm learns about each consumer's matching value α (and β), and each consumer realizes her relative importance between the two components, ω . In Stage 1, the firm decides to whom to send an ad at a unit cost $c_{ad} > 0$. A targeted consumer updates her unknown matching value α (and β) and decides whether to click the ad for more information at a cost $c_{click} > 0$. Then, in Stage 2, the consumer who clicks the ad and visits the firm's website finds out about her α and β and makes a purchasing decision.

In each of the targeting environments $dim \in \{sd, md\}$, an equilibrium denoted by a tuple $(\Sigma^{*dim}, \Omega^{*dim}, f^{dim})$ is defined as follows: (i) Σ^{*dim} is the firm's optimal targeting strategy that maximizes the firm's expected profit given the consumers' click decision characterized by Ω^{*dim} ; (ii) Ω^{*dim} is the subset of consumers of their private weight $\omega \in [0, 1]$ between the two components of consumers' utility, each of whom makes the click decision and purchase decision optimally given her expectation for the firm's targeting strategy Σ^{*dim} ; and (iii) upon being targeted, a consumer's beliefs about the components of her utility are updated according to Bayes rule $f^{dim}(\cdot)$, to be defined more precisely in the next section, and are consistent with the firm's equilibrium strategy.

We note that our main analysis focuses on a model with exogenous pricing. This allows us to maintain the tractability of the analysis and focus on consumers' inferences through targeted advertising. In Section 5.1, we extend the model with the firm's optimal pricing decision and show the robustness of the paper's main results.

3 Analysis

3.1 Single-Dimensional Targeting Environment

In the *sd*-targeting model, the firm knows only one dimension of the consumer's utility, α , the information it uses for advertising. For instance, a restaurant makes an advertising decision based on the customer's preferred location only, and consumers are aware of such a practice. All else equal, the firm expects a higher profit from a consumer with a higher value of α because, conditional on clicking the ad, she is more likely to purchase the product than another consumer with a lower value of α . Therefore, the firm's optimal targeting must follow a cutoff strategy in which the firm targets a consumer if and only if $\alpha \geq \overline{\alpha}$. Without loss of generality, we consider the firm's strategy $\Sigma = [\overline{\alpha}, 1]$ where $\overline{\alpha} \in [0, 1]$. If consumers expect $\widetilde{\Sigma} = [\widetilde{\alpha}, 1]$, then upon being targeted a consumer's posterior belief about α is a uniform distribution on the interval $[\widetilde{\alpha}, 1]$. On the other hand, her beliefs about β remain the same as the prior distribution, which is a uniform distribution on [0, 1], i.e.,

$$f^{sd}(\alpha,\beta;\widetilde{\Sigma}) = \frac{1}{|\widetilde{\Sigma}|} = \frac{1}{1-\widetilde{\alpha}} \text{ if } \alpha \in [\widetilde{\alpha},1]; \text{ otherwise } f^{sd}(\alpha,\beta;\widetilde{\Sigma}) = 0 \text{ if } \alpha \notin [\widetilde{\alpha},1], \tag{2}$$

where $|\cdot|$ denotes the size of the set.

With the updated beliefs, the consumer decides whether to click the ad by incurring a positive cost c_{click} and find out her exact utility from the product. The consumer will click the ad if and only if $\mathbb{E}[\max\{u-p,0\}; \widetilde{\Sigma}] \ge c_{\text{click}}$. The expectation is taken over the unknown α and β according to the posterior distribution $f^{sd}(\cdot, \cdot)$. Given that all targeted consumers have common posterior beliefs, a consumer's click decision hinges on her heterogeneous private weight between the two components of her utility ω :

$$\Omega^{sd}(\widetilde{\Sigma}) = \{ \omega \in [0,1] : \mathbb{E}[\max\{u-p,0\}; \widetilde{\Sigma}] \ge c_{\text{click}} \}.$$
(3)

The size of this set $|\Omega^{sd}(\widetilde{\Sigma})|$ corresponds to the click-through rate of targeted ads.

The consumers' optimal click decision can be characterized as follows:

Lemma 1 (Optimal click decision for sd-targeting). Let $\widetilde{\Sigma}$ be the consumers' expectation for the firm's advertising strategy under the sd-environment. There exist $\underline{\omega}$ and $\overline{\omega}$ with $0 \leq \underline{\omega} \leq \overline{\omega} \leq 1$ such that the set of consumers of type ω who click is $\Omega(\widetilde{\Sigma}) = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$. Also, if $\widetilde{\Sigma} = [\widetilde{\alpha}, 1]$, then $1 - \overline{\omega}(\widetilde{\alpha}) \geq \underline{\omega}(\widetilde{\alpha})$ so that a greater mass of consumers close to $\omega = 1$ click the ad than those close to $\omega = 0$. Lastly, $\frac{\partial \omega}{\partial \widetilde{\alpha}} \geq 0$ and $\frac{\partial \overline{\omega}}{\partial \widetilde{\alpha}} \leq 0$, such that the sizes of both intervals increase for a higher $\widetilde{\alpha}$.

The lemma states that a consumer with either sufficiently small or large values of ω is more likely to click the ad. This type of consumer will make a purchase, i.e., the realized utility will be higher than the price, if the consumer's utility is large enough in at least one dimension, either α (if ω is large) or β (if ω is small). On the other hand, if a consumer's type ω is in an intermediate range, then the consumer will buy the product only if her utility is high on both dimensions, which is less likely to occur. Therefore, a consumer with an extreme value of ω expects a greater utility from clicking the ad than a consumer whose ω is intermediate. Consequently, the set of ω types who click the ad takes the form of $\Omega(\widetilde{\Sigma}) = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$, where $\underline{\omega}$ and $\overline{\omega}$ are the smaller and larger roots of the indifference condition $\mathbb{E}[\max\{u-p,0\};\widetilde{\Sigma}] = c_{\text{click}}$.¹³ This is a general result that does not depend on the consumer's expectation for the firm's strategy and thus equally applies to the mdtargeting environment. That said, given the firm's cutoff-strategy on the α -dimension $\widetilde{\Sigma} = [\widetilde{\alpha}, 1]$, a targeted consumer's posterior beliefs about her unknown matching value on this dimension is more positive than those about the β -dimension. Therefore, in comparison to a consumer whose ω -type is smaller, a consumer of a higher ω -type has a greater expected utility from clicking the ad, which implies that there are more consumers who click the ad whose ω is larger, than consumers who click with smaller ω , i.e., $1 - \overline{\omega}(\widetilde{\alpha}) \geq \underline{\omega}(\widetilde{\alpha})$.

The firm considers the consumers' optimal click decisions $\Omega^{sd}([\tilde{\alpha}, 1]) = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$ and makes its advertising decision for each consumer to maximize the firm's total expected profit. At the individual consumer level, the firm finds it optimal to send an ad to a consumer whom the firm recognizes by her α type if and only if the expected revenue is no less than the unit cost of advertising, i.e.,

$$\pi^{sd}(\alpha; p, \Omega^{sd}([\widetilde{\alpha}, 1])) = p \cdot \mathbb{E}[\mathbf{1}\{u(\alpha; \beta, \omega) - p \ge 0 \& \omega \in \Omega^{sd}([\widetilde{\alpha}, 1])\}] - c_{\mathrm{ad}} \ge 0.$$
(4)

¹³The thresholds $\underline{\omega}$ and $\overline{\omega}$ are defined more precisely in Appendix A where we show that $\mathbb{E}[\max\{u-p,0\}; \widetilde{\Sigma}]$ is convex in ω , which guarantees that there can exist at most two distinct roots of the indifference condition.

The expectation is over the consumers' private type ω and the prior distribution over β . The expectation term computes the probability (from the firm's perspective) that the customer will make a purchase, which requires that the consumer will click the ad (i.e., $\omega \in \Omega^{sd}([\tilde{\alpha}, 1])$) and subsequently purchase the product (i.e., $u - p \ge 0$).

It is important to note that the consumer's click decision depends on her expectation for the firm's strategy, characterized by $\tilde{\alpha}$. In contrast, the consumer's purchase decision depends on the actual value of α , which is realized only after having incurred the cost of clicking the ad. This suggests that if consumers expect a highly selective targeting rule (i.e., a very high $\tilde{\alpha}$), then the firm may have an incentive to exploit such positive expectations of consumers. Provided a high $\tilde{\alpha}$ (i.e., close to 1), the targeted consumers posterior $f(\cdot, \cdot)$ upon being targeted will be very high, resulting in a large set Ω of consumers clicking the ad. Then, the firm can deviate by sending ads to consumers whose $\alpha \approx \tilde{\alpha} - \epsilon$ (for a small $\epsilon > 0$) who do not quite meet the targeting criterion expected by consumers. In turn, the same set of consumers (Ω) will unknowingly click the ad, some of whom will make a purchase.¹⁴ This is the commitment problem the firm faces in targeted advertising when consumers do not observe exactly how the firm utilizes information about consumers for targeted advertising. Due to the firm's lack of commitment, the firm may not be able to credibly communicate to consumers that it will select a high enough cutoff $\tilde{\alpha}$. This may discourage consumers from exerting efforts to click the ad, thus hurting the firm's profits. This type of commitment problem is present in both the sd- and md-targeting environments, but in Section 4 we show that it is more severe under the *md*-targeting environment.

Then, how should the firm set the threshold for its advertising? In equilibrium, the firm's optimal advertising strategy must be consistent with the consumers' expectation for the firm's strategy. Therefore, an equilibrium advertising strategy characterized by $\alpha = \overline{\alpha}^* \in [0, 1]$ must solve the indifference condition $\pi(\alpha; p, \Omega^{sd}([\overline{\alpha}^*, 1])) = 0$.

If the consumers' effort cost of clicking the ad c_{click} is too high, then no consumers will click the ad regardless of their beliefs about the firm's targeting strategy, i.e., $\Omega^{*sd} = \emptyset$, which, in turn, implies that $\Sigma^{*sd} = \emptyset$. Also, if the unit cost of advertising c_{ad} is prohibitively high, then the firm will not send any ads irrespective of the consumers' beliefs and click decisions, i.e., $\Sigma^{*sd} = \emptyset$. Both

¹⁴To some extent, this temptation can be disciplined by two effects in the model. One, sending each ad is costly, so, two, bringing in consumers who are less likely to buy the product by relaxing its targeting criterion can hurt the firm. These tradeoffs will balance out in equilibrium.

cases lead to *trivial equilibria* in which the firm engages in no advertising, i.e., $\Sigma^{*sd} = \emptyset$. This type of trivial equilibria always exist as long as the consumers' off-equilibrium beliefs when seeing an ad are sufficiently pessimistic. However, we focus on the existence of *non-trivial equilibria* in which the firm targets a non-empty set of consumers, i.e., $\Sigma^{*sd} \neq \emptyset$. The following proposition identifies the necessary and sufficient conditions under which a non-trivial equilibrium exists and shows that, when it exists, it is unique.

Proposition 1 (Equilibrium under sd-targeting). If $c_{click} \leq 1 - p$ and $c_{ad} \leq \overline{c_{ad}}(c_{click}, p)$, then there exists a unique non-trivial sd-targeting equilibrium in which the firm targets a non-empty set of consumers $\Sigma^{*sd} = [\overline{\alpha}^*, 1]$ (where $0 \leq \overline{\alpha}^* \leq 1$) and a positive fraction of the targeted consumers click the ad, $\Omega^{*sd} = [0, \underline{\omega}^{*sd}] \cup [\overline{\omega}^{*sd}, 1]$ (where $1 - \overline{\omega}^{*sd} \geq \underline{\omega}^{*sd}$). The cutoffs $\overline{c_{ad}}(c_{click}, p), \overline{\alpha}^*, \underline{\omega}^{*sd}$ and $\overline{\omega}^{*sd}$ are defined in the proof in Appendix A.

For the non-trivial equilibrium to exist for the *sd*-environment, the cost of advertising c_{ad} should not be too large because otherwise, the firm will find it too costly to engage in any advertising and be left only with the trivial equilibrium. Also, the consumers' effort cost of clicking the ad should not be too high because otherwise, no targeted consumers will find it optimal to click the ad.¹⁵

3.2 Multidimensional Targeting Environment

In the case of the multidimensional targeting environment, the firm has access to more granular customer data and, thus, can implement a targeting strategy Σ^{md} using both components of consumers' utility, namely α and β . The firm can choose how to use each dimension of customer information as it sees fit, so the firm's strategy is represented by any arbitrary subset of the unit square, i.e., $\Sigma^{md} \subseteq [0,1] \times [0,1]$. Unlike in the case of *sd*-targeting, a consumer targeted under the *md*-targeting environment makes inferences about both components of her utility so that her posterior distribution is a uniform distribution over the set $\tilde{\Sigma}$, the consumer's expectation about

¹⁵Corollary 1 in Appendix A presents a comparative static analysis on $\overline{\alpha}^*$ with respect to the key model primitives. As it is intuitive, $\frac{\partial \overline{\alpha}^*}{\partial c_{ad}} \geq 0$ and $\frac{\partial \overline{\alpha}^*}{\partial c_{click}} \geq 0$. In the former, the firm sends add more selectively to fewer consumers due to the direct unit cost of advertising increase. In the latter, the firm does the same in order to indirectly enhance targeted consumers' click decisions as their cost of click increases. The cutoff $\overline{\alpha}^*$ is also shown to be non-monotonic in p. This result is not specific to the *sd*-environment, and thus one can perform similar analysis for γ_1^* and γ_2^* which characterize the equilibrium targeting strategy in the *md*-targeting case. We omit this result from the main text in order to draw the readers' attention to the comparison between the two targeting environments.

the firm's multidimensional targeting strategy:

$$f^{md}(\alpha,\beta;\widetilde{\Sigma}) = \frac{1}{|\widetilde{\Sigma}|} \text{ if } (\alpha,\beta) \in \widetilde{\Sigma}, \text{ and otherwise } f^{md}(\alpha,\beta;\widetilde{\Sigma}) = 0.$$
 (5)

Given consumers' inferences from being targeted, consumers decide whether to click the ad and subsequently whether to make a purchase. Following similar steps of the analysis of the *sd*-targeting environment, we first establish the firm's optimal targeting strategy and then illustrate the consumers' optimal click decision, followed by the equilibrium characterization.

What the firm's optimal strategy should look like under the *md*-targeting environment is far less clear compared to *sd*-targeting because any subset of the unit square in the (α, β) -space is a feasible set for the firm's strategy.¹⁶ Should the firm target consumers with at least one high component? Or, consumers with high $\alpha + \beta$? To understand the structure of the firm's optimal strategy, it is useful to define a class of strategies as follows:

$$\mathcal{S} := \left\{ \Sigma^{md}(\gamma_1, \gamma_2) = [0, 1]^2 \cap \left(\{ (\alpha, \beta) \mid \beta \ge \gamma_1 \cdot \alpha + p(1 - \gamma_1) \text{ and } \alpha \le p \} \right) \\ \cup \left\{ (\alpha, \beta) \mid \beta \ge \gamma_2 \cdot \alpha + p(1 - \gamma_2) \text{ and } \alpha > p \} \right\} : \gamma_1, \gamma_2 \le 0 \right\}.$$
(6)

The next lemma states that the firm's optimal targeting strategy should belong to this class of strategies.

Lemma 2 (Optimal *md*-targeting strategy). In the non-trivial case where the firm targets a nonempty set of consumers (i.e., $\Sigma^{md} \neq \emptyset$), the firm's optimal targeting strategy under the *md*-targeting environment must satisfy $\Sigma^{md} \in S$. Therefore, the optimal targeting set Σ^{md} is characterized by the slopes of two straight lines, $\gamma_1, \gamma_2 \leq 0$.

As illustrated in Figure 2(a), the firm should never send an ad to consumers in the lower-left square where consumers' utility in both dimensions are lower than p (i.e., $\alpha, \beta < p$) because even if they click the ad, their realized utility $\omega \cdot \alpha + (1-\omega) \cdot \beta$ will be lower than p for all $\omega \in [0, 1]$. So, the consumers will never make a purchase, and the firm's ad cost will always be wasted. On the other

¹⁶Under the *sd*-targeting environment, consumers' attractiveness from the firm's perspective can be ordered, leading to a simple cutoff strategy for the firm's advertising. However, in this *md*-environment, there is no clear ranking, or monotonicity, in different consumers' value to the firm. For instance, a consumer whose matching values are $(\alpha, \beta) = (0.7, 0.5)$ is clearly more appealing to the firm than (0.6, 0.5), but not necessarily better than (0.5, 0.7).

hand, if the firm sends any ad, they must send it to all the consumers in the upper-right square where the consumers' utility is higher than p in both dimensions, i.e., $\alpha, \beta \geq p$. These consumers will always buy conditional on clicking the ad, which makes them the most profitable consumers from the firm's perspective. For other consumers, utility is higher than p in only one dimension and lower than p in the other dimension, so conditional on clicking the ad, a consumer's purchase decision will depend on her ω type. Therefore, the firm will decide to send ads to some consumers in the two rectangles in the upper-left and lower-right corners. This decision will determine the boundary of the set Σ^{sd} , which has to pass through (p, p).

The lemma above states that the boundary of Σ^{sd} consists of two straight lines of non-positive slopes. All else equal, a consumer with a higher α and β is a more profitable consumer to the firm because her purchase probability conditional on clicking the ad is greater, which explains why γ_1 and γ_2 cannot be positive. Moreover, the boundary has to be a straight line to the left of the point (p, p) (i.e., $\alpha). It can be shown that, given the same <math>\omega$ type, all consumers on the line have the same expected utility for the firm. So, the firm's expected profit is the same across consumers along the line, which implies that the firm does not discriminate with its advertising decision. That is, there exists a line such that above the line, the firm finds it profitable to send ads, and below the line, the firm chooses not to advertise. By the same logic, the boundary to the right of the point (p, p) must also be a straight line.¹⁷

Given the symmetry between the setup of the two dimensions, it is natural to focus on symmetric strategies in which the advertising strategy set Σ^* is symmetric about the 45-degree line, which implies $\gamma_1^* \cdot \gamma_2^* = 1$. In turn, $\Omega^* = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$ is symmetric about $\omega = 1/2$, i.e., $\underline{\omega} = 1 - \overline{\omega}$. As for the consumers' optimal click decision, we invoke the first part of Lemma 1, which is a general result that also applies to the *md*-targeting environment.¹⁸

Lemma 3 (Optimal click decision for *md*-targeting). Let $\tilde{\Sigma}$ be the consumers' expectation for the firm's targeting strategy under the *md*-environment. There exist $\underline{\omega}$ and $\overline{\omega}$ with $0 \leq \underline{\omega} \leq \overline{\omega} \leq 1$ such

¹⁷Note that γ_1^* and γ_2^* need not be the same. Consumers' purchase decision rule changes discontinuously at (p, p). Recall that on the boundary are the types of consumers (α, β) for whom the firm is indifferent to its advertising decision. From the firm's standpoint, the net profit from a consumer of type (α, β) is affected by their purchase decision conditional on clicking the ad. For a consumer on the left-hand side of (p, p), we have $\beta > \alpha$. So, her purchasing rule becomes $\omega \leq \frac{\beta-p}{\beta-\alpha}$. For a consumer on to the right of (p, p), her purchasing rule is, and the consumer will purchase if and only if $\omega \geq \frac{\alpha-p}{\alpha-\beta}$.

 $^{^{18}\}mathrm{Appendix}\ \mathrm{A}$ discusses asymmetric strategies and the resulting equilibria.



Figure 2: Optimal targeting strategy: sd- vs md-targeting

that the set of consumers of type ω is $\Omega(\widetilde{\Sigma}) = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$. Also, if $\widetilde{\Sigma} \in S$ and symmetric, i.e., $\gamma_1 \cdot \gamma_2 = 1$, then $1 - \overline{\omega} = \underline{\omega}$, i.e., the set $\Omega(\widetilde{\Sigma})$ is symmetric about $\omega = 1/2$.

In equilibrium, given the consumer's anticipation for the firm's advertising strategy, characterized by γ_1^*, γ_2^* , and the corresponding click decision Ω^{*md} , the firm should find it optimal to set $\gamma_1 = \gamma_1^*$ and $\gamma_2 = \gamma_2^*$. The indifference conditions along the two linear parts of the boundary of Σ are $\pi^{md}(\alpha, \beta; \Omega(\gamma_1^*, \gamma_2^*)) = p \cdot \mathbb{E}[\mathbf{1}\{u(\gamma_1, \gamma_2; \omega) - p \ge 0 \& \omega \in \Omega^{md}(\gamma_1^*, \gamma_2^*)\}] - c_{ad} = 0$. Next, we characterize the unique non-trivial symmetric equilibrium and identify necessary and sufficient conditions under which the equilibrium exists under the *md*-targeting environment.

Proposition 2 (Equilibrium under *md*-targeting). If $c_{ad} \leq p$ and $c_{click} \leq c_{click}^{max}(c_{ad}, p)$, where $c_{click}^{max}(c_{ad}, p)$ is defined in Equation (7), then there exists a unique non-trivial symmetric equilibrium in which the firm's targeting set $\Sigma^{*md} \in S$ is characterized by $\gamma_1^*, \gamma_2^* \leq 0$, where $\gamma_1^* \cdot \gamma_2^* = 1$. The set of consumers of type ω who click the ad is $\Omega^{*md} \in [0, \underline{\omega}^{*md}] \cup [1 - \underline{\omega}^{*md}, 1]$. More specifically, $\gamma_1^* = -\frac{c_{ad}}{p-c_{ad}}$, and the threshold $\underline{\omega}^{*md}$ is defined in the proof in Appendix A.

Figure 2 illustrates the firm's equilibrium targeting strategy under the *md*-environment next to the equilibrium strategy under the *sd*-environment. Under the former environment, the firm targets consumers with high matching value with at least one of the two components. Therefore, when targeted, consumers update their beliefs about α and β . More specifically, they understand



Figure 3: Parameter region where a non-trivial symmetric equilibrium exists, for p = 0.6.

that their α or β (or both) could be high. However, they remain uncertain about which is the case. Moreover, compared to the *sd*-environment, the firm targets a substantially larger set of consumers in the (α, β) -space. Therefore, the targeted consumers under the *md*-environment may face greater uncertainties about how much they will enjoy the product, thus reducing their expected utility from clicking the ad.

Figure 3 shows the parameter region where a non-trivial symmetric equilibrium exists. If the unit cost of advertising is too large (i.e., $c_{ad} > p$), then it can never be profitable to send an ad, and hence only the trivial equilibria exist. If the consumers' effort cost of clicking the ad is too high, then no consumer will click the ad ($\Omega^{*md} = \emptyset$), so no non-trivial symmetric equilibria exist.

It is interesting to note that for a given $c_{\rm ad} < p$, the scope of the non-trivial equilibrium, captured by $c_{\rm click}^{\rm max}(c_{\rm ad})$, is non-monotonic. This is because an increase in $c_{\rm ad}$ generates two opposing effects on the existence of the non-trivial equilibrium. The direct effect of an increase in $c_{\rm ad}$ is that the firm will send its ad more selectively. This effect reduces the targeting set Σ^{*md} , which makes the non-trivial equilibrium less likely to exist. Its indirect effect is that, given the firm's more selective targeting strategy, the targeted consumers' beliefs are updated more positively. This effect results in a larger click-through rate $|\Omega^{*md}|$, which means that the non-trivial equilibrium is more likely to hold. For the positive indirect effect to prevail, $c_{\rm ad}$ has to be large enough because,







(b) Click-through-rates $|\Omega^{*od}| = 1 - \omega^{*od}$ (dashed line) and $|\Omega^{*md}| = 2(1 - \overline{\omega}^{*md})$ (solid line).

Figure 4: How Ω^* changes as a function of the price, for $c_{\text{click}} = 0.1$ and $c_{\text{ad}} = 0.05$.

otherwise, for a small $c_{\rm ad}$ the firm cannot help but target a large customer base. Therefore, the scope of the non-trivial equilibrium $c_{\rm click}^{\rm max}(c_{\rm ad})$ decreases for smaller values of $c_{\rm ad}$ and then increases for its larger values.

A related interesting observation is that sometimes an increase in c_{ad} can increase the firm's equilibrium profit. This is because as c_{ad} increases, the firm can be more disciplined at targeting a smaller set of consumers Σ^{*md} , which can result in a higher click-through rate $|\Omega^{*md}|$ (e.g., see Example 1 in Appendix A.4).

Figure 4 illustrates the effect of price on the consumer's click decision. As depicted in Figure 4(a), consumers with high enough ω whose value for the targeting dimension is high will click the ad under the *sd*-targeting environment. On the other hand, under the *md*-targeting environment, consumers with either low or high values of ω click the ad. As demonstrated in Figure 4(b), at a lower range of prices, the click-through rate is higher under the *md*-environment than under the *sd*-environment, i.e., $|\Omega^{*md}| > |\Omega^{*sd}|$. However, for sufficiently large prices, the opposite is true, i.e., $|\Omega^{*md}| < |\Omega^{*sd}|$. When the price is high, consumers will click if they are more certain that they are likely to enjoy the product enough and make a purchase. As discussed before, consumers can draw more information about their expected utility under *sd*-targeting than under *md*-targeting. Consequently, the *sd*-targeting environment can induce a higher click-through rate even though the firm's targeting is based on less granular customer data.

4 The Value of Additional Dimensions of Customer Data

Intuitively, one may believe that the firm should be better off when it has access to more dimensions of customer data because the firm can engage in targeted advertising based on a more comprehensive understanding of the consumers' utility along several dimensions. However, as we show in this section, having more granular data can backfire.¹⁹

The equilibrium analysis in Section 3 illustrates a critical distinction that arises from the firm's access to an additional dimension of customer data. Proposition 1 shows that under the singledimensional environment, the firm adopts a cutoff strategy in one dimension of consumer utility. Consequently, a targeted consumer positively updates her expectations about her utility in that dimension. In contrast, Proposition 2 shows that under the multidimensional environment, the firm wants to target consumers who have a high enough matching value in at least one of the two components, and each targeted consumer will accordingly update her beliefs about both dimensions α and β . However, she remains uncertain whether her utility would be sufficiently high for the α or the β (or both) dimension(s). This consumer uncertainty can lower her overall expected utility upon visiting the product web page, discouraging the consumer's decision to click the ad and make a purchase. This leads to the following key result, which characterizes a sufficient condition under which the firm is more profitable under the *sd*-targeting than the *md*-targeting environment.

Proposition 3 (sd- vs md-targeting environment). If $c_{ad} < c_{ad}^{max}(p)$ and $c_{click}^{min}(c_{ad}, p) < c_{click} \leq c_{click}^{max}(c_{ad}, p)$, then $\Pi^{*sd} > \Pi^{*md}$, i.e., the firm's expected profit under the md-targeting environment is lower than that under the sd-targeting environment, where Π^{*sd} is the firm's profit in the unique non-trivial equilibrium under sd-targeting, and Π^{*md} the same in the unique symmetric non-trivial equilibrium under md-targeting. All the thresholds, $c_{ad}^{max}(p)$, $c_{click}^{min}(c_{ad}, p)$, and $c_{click}^{max}(c_{ad}, p)$ are defined in the proof in Appendix A.

Figure 5 compares the equilibrium targeting set sizes $|\Sigma^*|$ and consumers' click-through rates $|\Omega^*|$ between *sd*- and *md*-targeting, and it provides an intuition behind Proposition 3. As shown in Figure 5(a), the firm targets a larger set of consumers under *md*-targeting (i.e., $|\Sigma^{*md}| > |\Sigma^{*sd}|$).²⁰

¹⁹Based on this finding, one may posit that an additional dimension of customer data always backfires. However, our analysis of a zero-dimensional targeting environment (presented in Appendix A.3) where the firm has no information about consumers whatsoever shows that sd-targeting is always (weakly) better than the no-information case. Therefore, we show that the marginal value of additional dimensions of customer data can be non-monotonic.

²⁰Note that sometimes we can have the opposite, where $|\Sigma^{*md}| \leq |\Sigma^{*sd}|$. This can be the case if c_{click} is sufficiently

Therefore, the targeted consumers under the *md*-environment face greater uncertainty about their expected utility $\mathbb{E}[\max\{u - p, 0\}; \omega]$ and thus have to make less informed clicking decisions. Accordingly, consumers may be less willing to incur the cost of clicking, thus resulting in a lower click-through rate under the *md*-environment than under the *sd*-environment (i.e., $|\Omega^{*md}| < |\Omega^{*sd}|$) as shown in Figure 5(b). Moreover, the conversion rates, or the fraction of consumers who buy the product conditional on clicking the ad, can be higher under the *sd*-environment. This is because the firm under this environment targets a small set of consumers who are more likely to enjoy the product and therefore eventually convert upon clicking the ad.²¹ Hence, the lower response rate of the targeted ads under the *md*-environment can make the firm worse off despite the firm's access to more granular customer data, as depicted in Figure 5(c).

The counter-intuitive result of Proposition 3 depends on the key model parameters c_{ad} , c_{click} , and the price p. The result corresponds to the darker-shaded region in Figure 6. If the unit cost of advertising c_{ad} is too high, the firm will be disciplined to send ads selectively. This mitigates the firm's targeting decision problem, which is more severe under the md-environment. Therefore, the md-targeting becomes more appealing. Moreover, c_{click} must be not too low or too high, i.e., $c_{click} \in (c_{click}^{\min}(c_{ad}, p), c_{click}^{\max}(c_{ad}, p)]$. If it is too low $(c_{click} < c_{click}^{\min}(c_{ad}, p))$, then clicking becomes too cheap for consumers and the uncertainty of the md-equilibrium does not hurt the firm as much because consumers click despite of it. If it is too high $(c_{click} > c_{click}^{\max}(c_{ad}, p))$, then clicking becomes too expensive for consumers and the non-trivial md-equilibrium no longer exists.

From the consumers' perspective, the effect of the exogenous price is similar to the cost of click because a higher price lowers the benefit to the consumer upon clicking the ad, just as a higher click cost does. In that sense, the counter-intuitive result holds for prices sufficiently large. However, if the price becomes too high, then again, fewer and fewer consumers will click the ad, thus making the two targeting environments increasingly less distinguishable. For a complete intuition, one needs to understand the effect of price on the firm's advertising strategy and profits, which we discuss in Section 5.1 on endogenous pricing.

Given that targeting a smaller group of consumers can improve consumers' responses and hence

small. However, Proposition 3 focuses on the more interesting case when $|\Sigma^{*md}| > |\Sigma^{*sd}|$ holds and consequently we have $\Pi^{*sd} > \Pi^{*md}$. Corollary 2 in Appendix A formalizes this.

 $^{^{21}}$ For this reason, our main result is robust to an alternative ad payment scheme, namely pay-per-click. More details are provided in Appendix A.2.



Figure 5: Comparing the equilibria under *sd*-targeting (dashed line) vs *md*-targeting (solid line): the targeting set sizes $|\Sigma^*|$, click-through rates $|\Omega^*|$, and the profits, for $c_{ad} = 0.05$ and p = 0.6.

the firm's profits, one might ask an interesting question, "Can the firm under the multidimensional environment keep its targeting set small, for instance by mimicking a single-dimensional cutoffstrategy, and thereby secure a higher response rate?" The following proposition states that the answer is negative. That is, once given access to multiple dimensions of customer data, the firm cannot credibly commit to a single-dimensional strategy.

Proposition 4. A single-dimensional targeting strategy $\Sigma = [\overline{\alpha}, 1] \times [0, 1]$ cannot be part of an equilibrium under the multidimensional targeting environment.

Under the multidimensional environment, the firm knows various aspects of consumer utility and, therefore, finds it tempting to target a large group of consumers who would like the product on at least one of the distinct dimensions. For instance, the firm knows who would like a particular local restaurant on one dimension, such as location, and who else might like it on another dimension,



Figure 6: Parameter region (darker-shaded region) where a non-trivial equilibrium exists for both sd- and md-environments and the firm's profit under sd-targeting is higher than that under md-targeting, for p = 0.6.

such as the atmosphere. On average, both groups would equally enjoy the restaurant. Therefore, the firm cannot credibly commit to focusing on only one group and keeping its targeting set small. On the other hand, under single-dimensional targeting, the firm can only rely on a particular aspect of consumer utility, say location, and therefore must choose the targeting set with the other aspect (atmosphere) unknown. Then, it is in the firm's best interest to keep its targeting set small by selecting a selective targeting criterion. In other words, the firm's limited access to fewer dimensions of customer data acts as a commitment device to send targeted ads to a small group of consumers who are more likely to enjoy and thus buy the product. This increases the targeted consumers' response rate to the ad.²²

It is worth noting that the paper's main results rely on two critical elements of the model: (1) consumers are aware that ads are targeted, and (2) consumers are unable to observe the firm's exact decision for how to use the available customer data. These elements portray today's consumers who are increasingly becoming aware that platforms and advertisers have access to their information and use it for targeted advertising. As noted by Morey et al. (2015), it is still difficult for these

²²Even in an alternative model of *sd*-targeting where consumers do not know which one dimension the firm knows about them, the firm's limited knowledge about consumers on a single dimension still acts as a commitment device to keep its targeting set Σ^{*sd} small. Therefore, the click-through rate can be higher under the *sd*-environment and, thus, our main results are robust to this alternative model. Section 5.2 discusses more details about this.

aware consumers to know how the firm precisely uses different types of customer data for targeting purposes. However, as discussed about Figure 1 in Section 1, platforms have recently adopted technologies that help consumers become more aware of targeted ads. At the same time, there is often a link that provides users with context information about why they are targeted. This makes firms' targeting strategy more observable to consumers, suggesting that such devices that increase consumers' awareness and reduce their uncertainty about the firm's targeting decision can improve their response to targeted ads (Rafieian and Yoganarasimhan, 2021; Ada et al., 2021). Therefore, the firm may be able to overcome the potential drawbacks of coping with several dimensions of customer data that we have identified thus far.

Our findings further provide regulatory implications about the current policies on data privacy protection, such as the European Union's General Data Protection Regulation (GDPR), which prohibits firms from collecting and transferring personal data or tracking users' online behaviors without eliciting users' consent.²³ Such a privacy-protection policy will likely result in a reduced data granularity. Nevertheless, we show that this aspect of GDPR may increase response to targeted ads among consumers who value the product enough, and thereby benefit the firms. Also, another facet of GDPR, namely making the firms' use of consumer data more transparent, will likely lessen consumers' uncertainty about the firms' targeting practice and enhance consumers' response to the ads. These implications are mainly consistent with recent empirical research on the impact of GDPR on consumers and businesses (Aridor et al., 2020; Godinho de Matos and Adjerid, 2021).

5 Extensions

Though we analyze a stylized model, the mechanism at play for our main result, Proposition 3, is robust to various alternative models. In general, when the firm has more information about consumers, the consumers can face a greater uncertainty about the firm's strategy, i.e., how the firm will utilize that information. In turn, the firm can face a greater temptation to exploit the information asymmetry, which limits the firm's ability to credibly commit to its targeting strategy.²⁴

²³Apple's recent adoption of a new policy on consumer data privacy has a similar spirit. For more details, see https://www.forbes.com/sites/forbestechcouncil/2021/10/12/three-ways-apples-privacy-changes-will-impact-your-business.

²⁴These are critical effects universal to several different settings, irrespective of the exact information of the firm (e.g., whether the firm knows or does not know about ω) and the consumers (e.g., whether the consumers know on which one dimension they are targeted), or whether the price is exogenous or endogenous.

In this section, we present two extensions—the firm's endogenous pricing decision and an alternative information setting for consumers—and establish the robustness of our results. Appendix A.2 discusses an additional extension for an alternative ad payment scheme.

5.1 The Firm's Pricing Decision

Thus far, we have abstracted away from the firm's pricing decision. In this subsection, we extend our main model and consider endogenous pricing.

In this extension model, there is an additional stage of the firm's pricing decision before learning about individual consumers' matching values. The firm announces the price, which consumers observe.²⁵ So, the firm sets the price p and then plays the subgame, which coincides with the main model with exogenous pricing where the firm makes an advertising decision Σ . Therefore, an equilibrium of this game, denoted by $(p^*, \Sigma^*, \Omega^*, f^*)$, is naturally defined by the equilibrium advertising strategy, the consumers' click decisions (which is an equilibrium of the main model), the price which maximizes the firm's profit in the subgame, and consumers' beliefs about their unknown matching values which are determined by their expectation for the firm's advertising strategy.²⁶ Moreover, we assume that there is some marginal production cost $mc \geq 0$ for each unit of product sold. So, for each transaction, the consumer pays price p and the firm receives profit $p - mc - c_{ad}$. The rest of the game is the same as in the main model.

Consumers' optimal click decisions follow the same rule as the main model. Given the set Ω of consumers who click, the expected profit of the firm when it targets consumers in the set Σ , and it sets the price to p is

$$\Pi(p,\Sigma;\Omega) = (p - mc) \cdot \int_{\Sigma} \left(\int_{\Omega} \mathbf{1} \{ \omega \alpha + (1 - \omega)\beta \ge p \} d\omega \right) d\alpha \, d\beta - c_{\mathrm{ad}} \cdot |\Sigma|.$$

The pair $(p^*, \Sigma^*_{\rm price})$ is an equilibrium pair if $(p^*, \Sigma^*_{\rm price}) \in \operatorname{argmax}_{p, \Sigma} \Pi(p, \Sigma; \Omega(p, \Sigma^*_{\rm price})).$

Figure 7 is a graphical illustration which shows that our main result in Proposition 3 is robust to the firm's endogenous pricing. Note that under the sd-environment (dashed line), the firm

 $^{^{25}}$ Alternatively, one can consider a model where the firm does not announce the price. Then, the firm's lack of commitment to its price in addition to its advertising decision makes the analysis more involved. However, we can verify numerically that the results remain robust.

²⁶The firm will choose a price that maximizes its expected profit according to the distribution of consumers without foreknowledge about their matching values. Therefore, the price does not affect consumers' beliefs about their expected utility.



Figure 7: The effect of price on the equilibrium profits for both sd- (dashed line) and md-environments (solid line), for $c_{\text{click}} = 0.08$, $c_{\text{ad}} = 0.001$, and mc = 0.6.

charges price $p^{*sd} \approx 0.688$ and obtains a greater profit than the *md*-environment (solid line) where the firm charges $p^{*md} \approx 0.653$ (see more details in Example 3 in Appendix A.4). Consistent with the case of the main analysis with exogenous pricing, the firm sends add to fewer consumers, i.e., $|\Sigma_{\text{price}}^{*sd}| < |\Sigma_{\text{price}}^{*md}|$. Moreover, even though the optimal click-through-rate under the *md*-environment is higher than that under the *sd*-environment $(|\Omega_{\text{price}}^{*md}| > |\Omega_{\text{price}}^{*sd}|)$, this is not enough to compensate for the lower conversion rate. Thus, the main result and the intuition behind it in Proposition 3 is robust to the firm's endogenous pricing decision.

Proposition 5 (Endogenous pricing). If the firm's pricing decision is endogenous, for a sufficiently small c_{ad} , there exists an intermediate range of c_{click} in which $\Pi_{price}^{*sd} > \Pi_{price}^{*md}$, i.e., the firm's expected profit under the sd-targeting environment is higher than the profit under the md-targeting environment.

Figure 7 shows that the firm charges a higher (optimal) price under the *sd*- than under the *md*-targeting environment. In the *sd*-environment, the firm sends add to a select group of consumers whose matching value is high on a single dimension. Therefore, the firm faces a group of visitors whose expected value from the product is somewhat homogeneous and high, so it charges a reasonably high price. However, under the *md*-targeting environment, the firm sends adds to more consumers whose value from the product is more heterogeneous along distinct dimensions and, on average, somewhat low. Therefore, the firm sets a relatively lower price, which helps increase the consumers' click-through rates, thus partly recovering, but not fully offsetting, the drawbacks of multidimensional targeting.

5.2 Consumer Uncertainty about the Firm's Access to Customer Data under the Single-Dimensional Environment

Our analysis of the single-dimensional targeting case assumed that consumers know on which of the two dimensions they are targeted. Suppose an alternative scenario where consumers do not know which one of the two dimensions the firm will be targeting. Then, consumers would be more uncertain about their expected utility, which may lead to a lower equilibrium profit than under the sd-targeting. This section analyzes this alternative scenario and shows that our main result is still robust. This is because, just as in the sd-targeting environment analyzed in Section 3.1, the firm's knowledge about only one dimension of consumer utility act as a commitment device that allows the firm to keep a relatively small targeting set, which leads to a greater click-through rate and conversion rate.

In the model we analyze here, the firm has access to only one dimension of consumer utility, either α or β , which is randomly determined with probability 1/2. Without loss of generality, we assume that the realized dimension is α . However, consumers do not know the realized dimension, so they believe that each dimension is equally likely. We refer to this as *u*-sd-targeting.

The firm's optimal advertising strategy here again follows a threshold rule. Let x^* be the threshold in equilibrium such that the firm will send to a consumer if and only if $\alpha \ge x^*$. We focus on an equilibrium in which consumers have symmetric beliefs about the firm's advertising strategy along the two dimensions. Given the firm's equilibrium strategy and the consumers' uncertainty about the targeted dimension, the targeted consumers' posterior belief is as follows:

$$f^{u\text{-}sd}(\alpha,\beta;x^*) = \begin{cases} \frac{1}{2} \cdot \frac{1}{1-x^*}, & \text{if } \alpha \in [x^*,1] \text{ and } \beta \in [0,x^*] \text{ or } \alpha \in [0,x^*] \text{ and } \beta \in [x^*,1], \\ \frac{1}{1-x^*}, & \text{if } \alpha \in [x^*,1] \text{ and } \beta \in [x^*,1]. \end{cases}$$

Repeating the same analysis as in Section 3, we characterize the equilibrium and show the same qualitative result.

Proposition 6. In the unique non-trivial equilibrium of the u-sd-targeting model,

 The firm adopts a cutoff strategy Σ^{*u-sd} = [ᾱ^{*u-sd}, 1] and consumers with extreme values of private weight ω click, i.e., Ω^{*u-sd} = [0, ω̂^{*u-sd}] ∪ [1 − ω̂^{*u-sd}, 1]. The equilibrium profit can be higher than that under md-targeting, i.e., Π^{*u-sd} > Π^{*md}, if c_{ad} is sufficiently small and c_{click} is in an intermediate range.

The *u-sd*-targeting can be more profitable than *md*-targeting because the firm under the *u-sd*targeting environment has access to only one dimension of consumer data, which can serve as an effective commitment device regarding the firm's advertising decision, just as in the *sd*-targeting case. Consequently, as we see in Example 4(a) in Appendix A.4, the size of the targeting set is smaller under *u-sd*-targeting than under *md*-targeting ($|\Sigma^{*u-sd}| < |\Sigma^{*md}|$). This leads to a higher click-through rate ($|\Omega^{*u-sd}| > |\Omega^{*md}|$) and a higher conversion rate under *u-sd*-targeting. As a result, the firm can be better off under the *u-sd*-targeting environment, i.e., $\Pi^{*u-sd} > \Pi^{*md}$.

6 Conclusion

This paper examines how firms should cope with the increasing granularity of customer data in the context of targeted advertising. We provide a framework to formally investigate firms' optimal use of multidimensional targeting and its implications. In particular, we analyze how consumers who are aware that ads are targeted will respond to targeted ads depending on the number of dimensions in customer data. Surprisingly, we show that the firm can be better off when it has less granular customer data. This is because a firm with access to more granular customer data can adopt a more complicated targeting strategy along distinct dimensions, making targeted consumers more uncertain about how much they can expect to enjoy the product along each dimension. Moreover, the firm's comprehensive understanding of consumers tempts the firm to target a larger pool of consumers who may enjoy the product for various reasons. This means that the firm cannot commit to a selective targeting criterion. For both reasons, targeted consumers face more significant uncertainties about their realized utility from the product, which discourages them from incurring the effort cost of clicking the ad.

Our findings explain why consumers' response rates (e.g., click-through rates) may still be disappointingly low despite the astonishing improvements in targeting technology. Today, targeted advertising involves more granular customer data and convoluted analytic tools, and consumers may feel unsure about what to expect from the advertised product. Therefore, consumers can be reluctant to respond to the ads, and the firm's profits can suffer. Moreover, firms may overestimate the value of big data analysis if consumers' endogenous response to the ads is not considered.

Our paper suggests that firms may have incentives to self-discipline in acquiring and utilizing an excessive amount and dimensions of customer data. In relation to that, it would be interesting to analyze the effect of multidimensional targeting on consumers' ad avoidance (e.g., Anderson and Gans 2011; Johnson 2013) and adoption of ad blockers (e.g., Chen and Liu 2021; Despotakis et al. 2021; Gritckevich et al. 2021) by explicitly modeling consumers' nuisance cost of seeing irrelevant ads. Lastly, further studies on consumers' perception and AI recommendations, will be valuable.

References

- Ada, S., Nabout, N. A., and Feit, E. M. (2021). Express: Context information can increase revenue in online display advertising auctions: Evidence from a policy change. *Journal of Marketing Research*, page 00222437211070219.
- Anderson, S. P. and Gans, J. S. (2011). Platform siphoning: Ad-avoidance and media content. American Economic Journal: Microeconomics, 3(4):1–34.
- Andrews, M., Luo, X., Fang, Z., and Ghose, A. (2016). Mobile ad effectiveness: Hyper-contextual targeting with crowdedness. *Marketing Science*, 35(2):218–233.
- Aridor, G., Che, Y.-K., and Salz, T. (2020). The effect of privacy regulation on the data industry: Empirical evidence from gdpr. Working Paper 26900, National Bureau of Economic Research.
- Bergemann, D. and Bonatti, A. (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3):417–443.
- Blake, T., Nosko, C., and Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1):155–174.
- Braun, M. and Moe, W. W. (2013). Online display advertising: Modeling the effects of multiple creatives and individual impression histories. *Marketing Science*, 32(5):753–767.
- Chen, Y. and Liu, Q. (2021). Signaling through advertising when ad can be blocked. *Forthcoming,* Marketing Science.
- Chen, Y., Narasimhan, C., and Zhang, Z. J. (2001). Individual marketing with imperfect targetability. *Marketing Science*, 20(1):23–41.
- Despotakis, S., Ravi, R., and Srinivasan, K. (2021). The beneficial effects of ad blockers. Management Science, 67(4):2096–2125.
- Fong, N., Zhang, Y., Luo, X., and Wang, X. (2019). Targeted promotions on an e-book platform: Crowding out, heterogeneity, and opportunity costs. *Journal of Marketing Research*, 56(2):310–323.
- Ghose, A., Ipeirotis, P. G., and Li, B. (2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science*, 31(3):493–520.
- Godinho de Matos, M. and Adjerid, I. (2021). Consumer consent and firm targeting after gdpr: The case of a large telecom provider. *Management Science*.
- Goldfarb, A. and Tucker, C. (2011). Privacy regulation and online advertising. *Management science*, 57(1):57–71.
- Gritckevich, A., Katona, Z., and Sarvary, M. (2021). Ad blocking. Management Science.
- Hoban, P. R. and Bucklin, R. E. (2015). Effects of internet display advertising in the purchase

funnel: Model-based insights from a randomized field experiment. *Journal of Marketing Research*, 52(3):375–393.

- Iyer, G., Soberman, D., and Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3):461–476.
- Johnson, J. P. (2013). Targeted advertising and advertising avoidance. The RAND Journal of Economics, 44(1):128–144.
- Kim, T., Barasz, K., and John, L. K. (2019). Why am i seeing this ad? the effect of ad transparency on ad effectiveness. *Journal of Consumer Research*, 45(5):906–932.
- Lambrecht, A. and Tucker, C. E. (2015). Can big data protect a firm from competition? Available at SSRN 2705530.
- Levin, J. and Milgrom, P. (2010). Online advertising: Heterogeneity and conflation in market design. American Economic Review, 100(2):603–07.
- Lewis, R. A. and Rao, J. M. (2015). The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics*, 130(4):1941–1973.
- Long, F., Jerath, K., and Sarvary, M. (2022). Designing an online retail marketplace: Leveraging information from sponsored advertising. *Marketing Science*, 41(1):115–138.
- Manchanda, P., Dubé, J.-P., Goh, K. Y., and Chintagunta, P. K. (2006). The effect of banner advertising on internet purchasing. *Journal of Marketing Research*, 43(1):98–108.
- Mayzlin, D. and Shin, J. (2011). Uninformative advertising as an invitation to search. Marketing Science, 30(4):666–685.
- Morey, T., Forbath, T., and Schoop, A. (2015). Customer data: Designing for transparency and trust. Harvard Business Review, 93(5):96–105.
- Rafieian, O. and Yoganarasimhan, H. (2021). Targeting and privacy in mobile advertising. Marketing Science, 40(2):193–218.
- Redman, T. (2017). Does your company know what to do with all its data. Harvard Business Review, Junio.(https://hbr.org/2017/06/does-your-company-know-what-to-do-with-all-its-data).
- Rhodes, A. (2015). Multiproduct retailing. The Review of Economic Studies, 82(1):360–390.
- Shin, J. and Yu, J. (2021). Targeted advertising and consumer inference. *Marketing Science*, 40(5):900–922.
- Summers, C. A., Smith, R. W., and Reczek, R. W. (2016). An audience of one: Behaviorally targeted ads as implied social labels. *Journal of Consumer Research*, 43(1):156–178.
- Trusov, M., Ma, L., and Jamal, Z. (2016). Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting. *Marketing Science*, 35(3):405–426.

A Appendix

A.1 Analysis and Proofs

Proof of Lemma 1

Part (i). We first show that the consumer's expected utility $\mathbb{E}[\max\{0, u - p\}; \omega, \Sigma]$ is convex in ω , and therefore the set Ω is of a form $\Omega = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$. Then, we identify a necessary and sufficient condition that $\underline{\omega} = 0$. Note that $\mathbb{E}[\max\{0, u - p\}; \omega, \Sigma] = \frac{1}{|\Sigma|} \cdot \int_{\Sigma} \max\{\omega \cdot \alpha + (1 - \omega) \cdot \beta - p, 0\} d\alpha d\beta$, where $|\Sigma| = \int_{\Sigma} 1 d\alpha d\beta$. Define $f(\alpha, \beta; \omega) := \max\{\omega \cdot \alpha + (1 - \omega) \cdot \beta - p, 0\}$. Notice that f is convex with respect to ω (as a maximum of two linear functions). In other words, for $\omega_1, \omega_2, t \in [0, 1]$, it holds that $f(\alpha, \beta; t \cdot \omega_1 + (1 - t) \cdot \omega_2) \leq t \cdot f(\alpha, \beta; \omega_1) + (1 - t) \cdot f(\alpha, \beta; \omega_2)$. Therefore $\int_{\Sigma} f(\alpha, \beta; t \cdot \omega_1 + (1 - t) \cdot \omega_2) d\alpha d\beta \leq \int_{\Sigma} (t \cdot f(\alpha, \beta; \omega_1) + (1 - t) \cdot f(\alpha, \beta; \omega_2)) d\alpha d\beta = t \cdot \int_{\Sigma} f(\alpha, \beta; \omega_1) d\alpha d\beta + (1 - t) \cdot \int_{\Sigma} f(\alpha, \beta; \omega_2) d\alpha d\beta$. This means that $\mathbb{E}[\max\{0, u - p\}; \omega]$ is convex w.r.t. ω . It follows that $\Omega(\Sigma) = \{\omega \in [0, 1] : \mathbb{E}[\max\{0, u - p\}; \omega] \geq c_{\text{click}}\}$ should be of the form $[0, \underline{\omega}] \cup [\overline{\omega}, 1]$, where $\underline{\omega}$ and $\overline{\omega}$ are the smaller and larger root of the equation. We ensure that the two cutoffs are well-defined as follows: If two distinct roots exist and the smaller root $\underline{\omega}$ is ≤ 0 , then we redefine $\underline{\omega} := 0$ and the size of the interval $[0, \underline{\omega}]$ becomes zero. Likewise, if $\overline{\omega} \geq 1$, then we redefine $\overline{\omega} := 1$. If $\min_{\omega} \mathbb{E}[\max\{0, u - p\}; \omega] \geq c_{\text{click}}$, then everyone clicks. In this case, we define $\underline{\omega} = \overline{\omega} := \operatorname{argmin}_{\omega} \mathbb{E}[\max\{0, u - p\}; \omega], \text{ and } \Omega = [0, \underline{\omega}] \cup [\overline{\omega}, 1] = [0, 1]$. Note that $\underline{\omega} = 0$ if and only if $\mathbb{E}[\max\{u - p, 0\}; \omega = 0, \overline{\alpha}] = \int_0^1 \max\{\beta - p, 0\} d\beta = \frac{(1 - p)^2}{2} \leq c_{\text{click}}$.

Part (ii). For a given $\overline{\alpha}$, a consumer's expected payoff for clicking on an ad is $\mathbb{E}[\max\{u-p,0\}; \omega, \overline{\alpha}] - c_{\text{click}}$. Differentiating this with respect to $\overline{\alpha}$, $\partial \mathbb{E}[\max\{u-p,0\}; \omega, \overline{\alpha}] / \partial \overline{\alpha} = \frac{1}{(1-\overline{\alpha})^2} \int_{\overline{\alpha}}^1 \int_0^1 [f(\alpha, \beta; \omega) - f(\overline{\alpha}, \beta; \omega)] d\beta d\alpha \ge 0$, because $f(\alpha, \beta; \omega)$ is weakly increasing in α , which makes the integrand ≤ 0 .

The cutoffs $\underline{\omega}$ and $\overline{\omega}$ solve $\mathbb{E}[\max\{u-p,0\}; \omega, \overline{\alpha}] - c_{\text{click}} = 0$. Differentiating the equation with respect to $\overline{\alpha}$ gives $\partial \mathbb{E}[\max\{u-p,0\}; \underline{\omega}, \overline{\alpha}] / \partial \overline{\alpha} + \partial \mathbb{E}[\max\{u-p,0\}; \underline{\omega}, \overline{\alpha}] / \partial \underline{\omega} \cdot \partial \underline{\omega} / \partial \overline{\alpha} = 0$. This shows that $\partial \underline{\omega} / \partial \overline{\alpha} \ge 0$ because $\partial \mathbb{E}[\max\{u-p,0\}; \underline{\omega}, \overline{\alpha}] / \partial \overline{\alpha} \ge 0$ and $\partial \mathbb{E}[\max\{u-p,0\}; \underline{\omega}, \overline{\alpha}] / \partial \underline{\omega} \le 0$. Likewise, it is straightforward to show that $\partial \overline{\omega} / \partial \overline{\alpha} \le 0$ because $\partial \mathbb{E}[\max\{u-p,0\}; \overline{\omega}, \overline{\alpha}] / \partial \overline{\omega} \ge 0$.

It remains to show that $1 - \overline{\omega} \geq \underline{\omega}$. Note that if $\overline{\alpha} = 0$, then $\mathbb{E}[\max\{u - p, 0\}; \underline{\omega}, \overline{\alpha}]$ is symmetric about $\omega = 1/2$. For any $\overline{\alpha} > 0$, it shifts to the right, which implies that $\overline{\omega} > 1/2$. Therefore, we have $\overline{\omega} > 1 - \overline{\omega}$, which implies $\mathbb{E}[\max\{u - p, 0\}; \overline{\omega}, \overline{\alpha}] \geq \mathbb{E}[\max\{u - p, 0\}; 1 - \overline{\omega}, \overline{\alpha}]$ because the former places a higher weight to α , whose distribution is superior to that of β . Therefore, $\mathbb{E}[\max\{u-p,0\};\underline{\omega},\overline{\alpha}] \geq \mathbb{E}[\max\{u-p,0\};1-\overline{\omega},\overline{\alpha}]$ must hold because the consumer's expected values at $\omega = \overline{\omega}$ and $\underline{\omega}$ coincide with c_{click} . Given that $\underline{\omega} \leq \overline{\omega}$, we must have $\partial \mathbb{E}[\max\{u-p,0\};\omega,\overline{\alpha}]/\partial \omega \leq 0$ for $\omega \leq \underline{\omega}$. This proves that $1-\overline{\omega} \geq \underline{\omega}$.

The expected utility is computed as follows (to be used to characterize the equilibrium later):

- 1. If $\overline{\alpha} < 2 \frac{1}{p}$, then $\mathbb{E}[\max\{u p, 0\}; \omega, \overline{\alpha}] = \frac{1}{1 \overline{\alpha}} \cdot \frac{(1 p)^3 (1 p \omega(1 \overline{\alpha}))^3}{6(1 \omega)\omega}$ if $\omega < \frac{1 p}{1 \overline{\alpha}}; \frac{1}{1 \overline{\alpha}} \cdot \frac{(1 p)^3}{6(1 \omega)\omega}$ if $\frac{1 p}{1 \overline{\alpha}} \le \omega \le p$; and $\frac{1}{1 \overline{\alpha}} \cdot \frac{1 + 3p^2 + \omega + \omega^2 3p(1 + \omega)}{6\omega}$ if $\omega > p$.
- 2. If $2 \frac{1}{p} \leq \overline{\alpha} \leq p$, then $\mathbb{E}[\max\{u p, 0\}; \omega, \overline{\alpha}] = \frac{1}{1 \overline{\alpha}} \cdot \frac{(1 p)^3 (1 p \omega(1 \overline{\alpha}))^3}{6(1 \omega)\omega}$ if $\omega < p$; $\frac{1}{1 \overline{\alpha}} \cdot \left(\frac{3(1 p)(\omega p)}{6\omega} + \frac{(1 \omega)^3 (1 p \omega(1 \overline{\alpha}))^3}{6(1 \omega)\omega}\right)$ if $p \leq \omega \leq \frac{1 p}{1 \overline{\alpha}}$; and $\frac{1}{1 \overline{\alpha}} \cdot \frac{1 + 3p^2 + \omega + \omega^2 3p(1 + \omega)}{6\omega}$ if $\omega > \frac{1 p}{1 \overline{\alpha}}$.
- 3. If $\overline{\alpha} > p$, then $\mathbb{E}[\max\{u-p,0\}; \omega, \overline{\alpha}] = \frac{1}{1-\overline{\alpha}} \cdot \frac{(1-p)^3 (1-p-\omega(1-\overline{\alpha}))^3}{6(1-\omega)\omega}$ if $\omega < p$; $\frac{1}{1-\overline{\alpha}} \cdot \left(\frac{3(1-p)(\omega-p)}{6\omega} + \frac{(1-\omega)^3 (1-p-\omega(1-\overline{\alpha}))^3}{6(1-\omega)\omega}\right)$ if $p \le \omega \le \frac{p}{\overline{\alpha}}$; and $\frac{1}{1-\overline{\alpha}} \cdot \frac{(1-\overline{\alpha})(1-2p+\omega\overline{\alpha})}{2}$ if $\omega > \frac{p}{\overline{\alpha}}$.

Proof of Proposition 1

We prove this proposition in two parts. First, given the optimal click rule $\Omega(\tilde{\alpha})$, we identify the firm's optimal choice of the cutoff $\bar{\alpha}$. Second, we identify the equilibrium cutoff $\bar{\alpha}^*$ as a stable cutoff level which is consistent with consumer's anticipation.

Part (i). For a given $\Omega(\widetilde{\alpha}) = [0, \underline{\omega}(\widetilde{\alpha})] \cup [\overline{\omega}(\widetilde{\alpha}), 1]$, the firm's expected profit after sending an ad to a consumer of type α is $\pi(\alpha; \Omega(\widetilde{\alpha})) = p \cdot (G(\alpha; \underline{\omega}(\widetilde{\alpha})) + F(\alpha; \overline{\omega}(\widetilde{\alpha}))) - c_{ad}$, where $G(\alpha; \underline{\omega}(\widetilde{\alpha})) := \int_{0}^{\underline{\omega}(\widetilde{\alpha})} \int_{0}^{1} \mathbf{1} \{\omega\alpha + (1-\omega)\beta - p \ge 0\} d\beta d\omega$ and $F(\alpha; \overline{\omega}(\widetilde{\alpha})) := \int_{\overline{\omega}(\widetilde{\alpha})}^{1} \int_{0}^{1} \mathbf{1} \{\omega\alpha + (1-\omega)\beta - p \ge 0\} d\beta d\omega$. Note that G is increasing in $\underline{\omega}$ and F decreasing in $\overline{\omega}$, properties which will be used later. Moreover, both F and G are continuous and weakly increasing in α on [0, 1]. Consequently, so is $\pi(\alpha; \Omega(\widetilde{\alpha}))$, and $\pi(0; \Omega(\widetilde{\alpha})) \le \pi(\alpha; \Omega(\widetilde{\alpha})) \le \pi(1; \Omega(\widetilde{\alpha}))$. We consider three cases: (1) If $\pi(0; \Omega(\widetilde{\alpha})) \ge 0$, then $\pi(\alpha; \Omega(\widetilde{\alpha})) \ge 0$ for any $\alpha \in [0, 1]$. Therefore, it is optimal for the firm to send an ad to every consumer, i.e. $\overline{\alpha} = 0$; (2) If $\pi(1; \Omega(\widetilde{\alpha})) \le 0$ then $\pi(\alpha; \Omega(\widetilde{\alpha})) \le 0$ for any $\alpha \in [0, 1]$. So, the firm will not send an ad to any consumer, i.e. $\overline{\alpha} = 1$; (3) Otherwise, if $\pi(0; \Omega(\widetilde{\alpha})) < 0 < \pi(1; \Omega(\widetilde{\alpha}))$, then there exists a unique $\overline{\alpha} \in (0, 1)$ for which $\pi(\overline{\alpha}; \Omega) = 0$.

We can characterize $F(\cdot; \cdot)$ and $G(\cdot; \cdot)$ as follows: $F(\alpha; \overline{\omega}) = 0$ if $\alpha < 1 - \frac{1-p}{\overline{\omega}}; 1 - \overline{\omega} - p + \overline{\omega}\alpha - (p - \alpha)\log\left(\frac{(1-\alpha)(1-\overline{\omega})}{p-\alpha}\right)$ if $1 - \frac{1-p}{\overline{\omega}} \le \alpha < p; (1-\overline{\omega})(1-p)$ if $\alpha = p; 1 - \overline{\omega} - p + \overline{\omega}\alpha + (\alpha - p)\log\left(\frac{\alpha(1-\overline{\omega})}{\alpha-p}\right)$

 $\begin{array}{l} \text{if } p < \alpha \leq \frac{p}{\overline{\omega}} \text{; and } 1 - \overline{\omega} \text{ if } \alpha > \frac{p}{\overline{\omega}} \text{. The function } G(\alpha; \underline{\omega}) = 1 - p + (p - \alpha) \log\left(\frac{p - \alpha}{1 - \alpha}\right) \text{ if } \alpha < 1 - \frac{1 - p}{\underline{\omega}} \text{; } \\ \underline{\omega}(1 - \alpha) - (\alpha - p) \log\left(1 - \underline{\omega}\right) \text{ if } 1 - \frac{1 - p}{\underline{\omega}} \leq \alpha \leq \frac{p}{\underline{\omega}} \text{; and } \underline{\omega} - p - (\alpha - p) \log\left(\frac{\alpha - p}{\alpha}\right) \text{ if } \alpha > \frac{p}{\underline{\omega}}. \end{array}$

Recall that F is decreasing in $\overline{\omega}$ and G increasing in $\underline{\omega}$. The function $G(\alpha)$ for the case of $\underline{\omega} = 0$ can be obtained by taking the limit of the function as $\underline{\omega} \to 0$, such that $G(\alpha) \equiv 0$ everywhere. Recall that in equilibrium $\pi(\overline{\alpha}^*;\Omega) = 0$. By the chain rule, $0 = \frac{d\pi(\overline{\alpha}^*;\Omega)}{d\overline{\omega}} = \frac{\partial\pi}{\partial\overline{\omega}} + \frac{\partial\pi}{\partial\alpha}|_{\alpha=\overline{\alpha}^*} \cdot \frac{\partial\overline{\alpha}^*}{\partial\overline{\omega}}$. We know that $\frac{\partial\pi}{\partial\overline{\omega}} \leq 0$ because the targeted consumer's click-through rate goes down, and $\frac{\partial\pi}{\partial\alpha} \geq 0$ because a consumer with a higher α has a higher expected utility, all else equal. Therefore, $\frac{\partial\overline{\alpha}^*}{\partial\overline{\omega}} \geq 0$. Likewise, $0 = \frac{d\pi(\overline{\alpha};\Omega)}{d\underline{\omega}} = \frac{\partial\pi}{\partial\underline{\omega}} + \frac{\partial\pi}{\partial\alpha}|_{\alpha=\overline{\alpha}^*} \cdot \frac{\partial\overline{\alpha}^*}{\partial\underline{\omega}}$, and it is straightforward to show that $\frac{\partial\overline{\alpha}}{\partial\underline{\omega}} \leq 0$. *Part (ii).* The firm's expected profit when the firm's cutoff strategy characterized by α is consistent with the consumers' beliefs, $\pi(\alpha;\overline{\omega}(\alpha),\underline{\omega}(\alpha))$, changes in α as follows: $\frac{d\pi(\alpha;\overline{\omega},\underline{\omega})}{d\alpha} = \frac{\partial\pi}{\partial\alpha} + \frac{\partial\pi}{\partial\overline{\omega}} + \frac{\partial\pi}{\partial\overline{\omega}} + \frac{\partial\pi}{\partial\overline{\omega}} \leq 0$, $\frac{\partial\pi}{\partial\overline{\omega}} \leq 0$, $\frac{\partial\pi}{\partial\overline{\omega}} \leq 0$, $\frac{\partial\pi}{\partial\overline{\omega}} \geq 0$. In equilibrium, $\pi(\overline{\alpha}^*;\overline{\omega}(\overline{\alpha}^*),\underline{\omega}(\overline{\alpha}^*)) = 0$. If $\pi(1;\Omega(1)) < 0$, then the firm does not send any ad in equilibrium, i.e., $\Sigma^{*sd} = \emptyset$. Otherwise if $\pi(1;\Omega(1)) \geq 0$, then there exists a unique equilibrium $\overline{\alpha}^* \geq 0$. This requires c_{ad} to be sufficiently small. Specifically, c_{ad} should be at most $\overline{c_{ad}}(c_{click}, p) := p \cdot (G(1;\underline{\omega}(1)) + F(1;\overline{\omega}(1)))$. By considering the different cases for c_{click} , we obtain that $\overline{c_{ad}}(c_{click}, p) = p(1-p)(1-\log(1-p))$ if $c_{click} \leq \frac{(1-p)^2}{2}$; $p(1-p)\left(1+\log\left(\frac{1-p}{2c_{click}}\right)\right)$ if $\frac{(1-p)^2}{2} < c_{click} \leq \frac{1-p}{2}$; and $2p(1-p-c_{click})$ if $\frac{1-p}{2} < c_{click} \leq 1-p$.

Corollary 1. The cutoff for the firm's equilibrium targeting strategy $\overline{\alpha}^*$ satisfies $\frac{\partial \overline{\alpha}^*}{\partial c_{ad}} \geq 0$ and $\frac{\partial \overline{\alpha}^*}{\partial c_{click}} \geq 0$. Moreover, $\overline{\alpha}^*$ can be non-monotonic in p; For p sufficiently small, we may have $\frac{\partial \overline{\alpha}^*}{\partial p} \leq 0$, whereas for a sufficiently large $p, \frac{\partial \overline{\alpha}^*}{\partial p} \geq 0$ holds.

Proof of Corollary 1

First, differentiating the equation $\pi^{sd}(\overline{\alpha}^*; p, \Omega^{sd}([\overline{\alpha}^*, 1])) = 0$ defined in Equation (4) with respect to c_{click} , we have $0 = \frac{d\pi^{*sd}}{dc_{\text{click}}} = \frac{\partial\pi^{*sd}}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial c_{\text{click}}} + \frac{\partial\pi^{*sd}}{\partial\underline{\omega}} \cdot \frac{\partial\omega}{\partial c_{\text{click}}} + (\frac{\partial\pi^{*sd}}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial\overline{\alpha}} + \frac{\partial\pi^{*sd}}{\partial\underline{\omega}} \cdot \frac{\partial\omega}{\partial\overline{\alpha}} + \frac{\partial\pi^{*sd}}{\partial\overline{\omega}}) \cdot \frac{\partial\overline{\alpha}}{\partial c_{\text{click}}}.$ Note that $\frac{\partial\pi}{\partial\underline{\omega}} \ge 0$, $\frac{\partial\pi}{\partial\overline{\omega}} \le 0$, $\frac{\partial\overline{\omega}}{\partial c_{\text{click}}} \ge 0$, $\frac{\partial\omega}{\partial\overline{c}_{\text{click}}} \le 0$, $\frac{\partial\overline{\omega}}{\partial\overline{\alpha}} \ge 0$, $\frac{\partial\overline{\omega}}{\partial\overline{\alpha}} \le 0$, and $\frac{\partial\pi^{*sd}}{\partial\overline{\alpha}} \ge 0$. So, it must be $\frac{\partial\overline{\alpha}^*}{\partial c_{\text{click}}} \ge 0$.

Second, differentiating the equation with respect to $c_{\rm ad}$, we have $0 = \frac{d\pi^{*sd}}{dc_{\rm ad}} = \frac{\partial\pi^{*sd}}{\partial c_{\rm ad}} + \left(\frac{\partial\pi^{*sd}}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial\overline{\omega}} + \frac{\partial\pi^{*sd}}{\partial\overline{\omega}} + \frac{\partial\pi^{*sd}}{\partial\overline{c}_{\rm ad}}\right) \cdot \frac{\partial\overline{\alpha}}{\partial c_{\rm ad}}$, where $\frac{\partial\pi}{\partial c_{\rm ad}} = -1 \leq 0$. Together with the properties identified previously, this shows that $\frac{\partial\overline{\alpha}^{*}}{\partial c_{\rm ad}} \geq 0$ must hold.

Lastly, differentiating the equation with respect to p, we have $0 = \frac{d\pi(\overline{\alpha}^*;\Omega)}{dp} = \frac{\partial\pi}{\partial p} + (\frac{\partial\pi}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial p} + \frac{\partial\pi}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial p} + \frac{\partial\pi}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial\overline{\alpha}} + \frac{\partial\pi^{*sd}}{\partial\overline{\alpha}}) \cdot \frac{\partial\overline{\alpha}^*}{\partial p}$. Note that $\frac{\partial\overline{\omega}}{\partial p} \ge 0$, and $\frac{\partial\omega}{\partial p} \le 0$. The expression in the first parentheses is $\frac{\partial\pi}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial p} + \frac{\partial\pi}{\partial\overline{\omega}} \cdot \frac{\partial\omega}{\partial p} \le 0$, and the expression in the second parentheses is $\frac{\partial\pi}{\partial\overline{\omega}} \cdot \frac{\partial\overline{\omega}}{\partial\overline{\alpha}} + \frac{\partial\pi^{*sd}}{\partial\overline{\alpha}} \ge 0$. However, the sign of the first term $\frac{\partial\pi}{\partial p}$ can be positive or negative. Note that if p = 0, $\frac{\partial\pi}{\partial p} = \int_{\Omega(\overline{\alpha}^*)} \int_0^1 \mathbf{1}\{\omega \cdot \alpha + (1-\omega) \cdot \beta - p \ge 0\} d\beta d\omega \ge 0$. And, if p = 1, then $\frac{\partial}{\partial p}(\int_{\Omega(\overline{\alpha}^*)} \int_0^1 \mathbf{1}\{\omega \cdot \alpha + (1-\omega) \cdot \beta - p \ge 0\} d\beta d\omega) \le 0$. This implies that $\frac{\partial\pi}{\partial p} \ge 0$ for small enough p and $\frac{\partial\pi}{\partial p} \ge 0$ for p large enough. Note that if $\frac{\partial\pi}{\partial p} \le 0$ can hold.

Proof of Lemma 2

If $p < c_{ad}$, then $\pi(\alpha, \beta; \Omega^*) < 0$ for all consumers and any Ω^* . Therefore, $\Sigma^* = \emptyset$ and $\Omega^* = \emptyset$. We need $p \ge c_{ad}$ for $\Sigma^* \neq \emptyset$, which we assume for the rest of the proof. We prove the proposition by proving the following two claims.

Claim 1. If the optimal $\Sigma^* \neq \emptyset$, then the upper closed square is a subset of Σ^* , i.e., $\{(\alpha, \beta) : \alpha, \beta \ge p\} \subset \Sigma^*$. The lower open square has empty intersection with Σ^* , i.e., $\{(\alpha, \beta) : \alpha, \beta < p\} \cap \Sigma^* = \emptyset$. Lastly, $(p, p) \in \Sigma^*$ and, in particular, in the boundary of the set, $(p, p) \in \partial \Sigma^*$.

Proof. It needs to be shown that for any neighborhood of (p, p), there exists at least one consumer in the set Σ^* and another consumer not in the set. First, for any $\alpha < p$ and $\beta < p$, the integrand in Equation (4) vanishes with probability 1. Therefore, $\pi(\alpha, \beta; \Omega^*) = 0 - c_{ad} < 0$, and hence $(\alpha, \beta) \notin$ Σ^* . Second, for any $\alpha, \beta \ge p$, the integrand is 1 for all $\omega \in [0, 1]$, i.e., the probability of all consumers who click the ad is 1. The expected profit from such a consumer is $\pi(\alpha, \beta; \Omega) = p \cdot |\Omega^*| - c_{ad}$. This is the maximal attainable profit from an individual consumer, so if this is negative, $\Sigma^* = \emptyset$. By the assumption that $\Sigma^* \neq \emptyset$, this implies that $\pi(\alpha, \beta; \Omega) = p \cdot |\Omega^*| - c_{ad} \ge 0$, and hence $(\alpha, \beta) \in \Sigma^*$.

In summary, for any small neighborhood of a consumer of type $(\alpha, \beta) = (p, p)$, there exists an arbitrarily small $\epsilon > 0$ $(p - \epsilon, p - \epsilon) \notin \Sigma^*$ and $(p + \epsilon, p + \epsilon) \in \Sigma^*$. This proves that $(p, p) \in \partial \Sigma^*$.²⁷

Claim 2. If the optimal $\Sigma^* \neq \emptyset$, then the boundary of Σ^* is composed of two lines. Then, it must be that $\Sigma^{*md} \in S$, as defined in Equation 6.

²⁷Recall that, as introduced in Claim 1, $\partial \Sigma^*$ denotes the boundary of set Σ^* .

Proof. We begin by proving that if $(\alpha, \beta) \in \partial \Sigma^*$, then $\pi(\alpha, \beta; \Omega) = 0$. Suppose not. First, if $\pi(\alpha, \beta; \Omega) > 0$, then by continuity there exists a small neighborhood of the consumer in which the firm makes a positive profit from sending an ad. This contradicts with the presumption that $(\alpha, \beta) \in \partial \Sigma^*$. Likewise, if $\pi(\alpha, \beta; \Omega) < 0$, then there is a small neighborhood of the consumer in which the firm should not send an ad, which again is contradictory.

Suppose $\alpha \geq p \geq \beta$ and $(\alpha, \beta) \in \partial \Sigma^*$. Conditional on clicking the ad, the consumer will buy the product if and only if $u - p \geq 0$, or $\omega \geq \frac{p-\beta}{\alpha-\beta}$. Therefore, $\pi(\alpha, \beta; \Omega^*) = p \cdot \int_{\Omega^* \cap [\frac{p-\beta}{\alpha-\beta}, 1]} 1 \, d\omega - c_{\rm ad} = 0$ must hold. The expression is weakly decreasing in $\frac{p-\beta}{\alpha-\beta}$. Also, the consumer will buy if indifferent between buying and not buying. So, there exists the unique constant k such that $\frac{p-\beta}{\alpha-\beta} = k$ that satisfies the indifference condition. More precisely, k is pinned down by the assumed $\Omega = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$. This equation can be re-written as $Y = -\frac{k}{1-k}X + \frac{p}{1-k}$, which is a line that passes through (α, β) and (p, p). So, if $(\alpha, \beta) \in \partial \Sigma^*$, then all points (α', β') on this line with $\alpha' > \beta'$ is in $\partial \Sigma^*$.

By symmetry, if $\alpha \leq p \leq \beta$ and $(\alpha, \beta) \in \partial \Sigma^*$, then all points (α'', β'') on the line connecting (α, β) and (p, p) with $\alpha' \leq p \leq \beta'$ is in $\partial \Sigma^*$. Therefore, the lower bound of the set Σ^* consists of two line segments: $Y = \gamma_1 \cdot X + (1 - \gamma_1)p$ for $X \in [0, p]$ and $Y = \gamma_2 \cdot X + (1 - \gamma_2)p$ for $X \in (p, 1]$. \Box

This completes the proof of the proposition.

Proof of Lemma 3

The proof of Lemma 1, part (i), does not assume anything about Σ , the consumers' expectation about the firm's targeting strategy. Therefore, the proof extends to the case of multidimensional targeting strategies.

Proof of Proposition 2

The indifference conditions for the firm's advertising decision along the boundary of the set Σ^{*md} can be re-written in integrals as follows: $p \cdot \int_{\Omega^*(\gamma_1^*, \gamma_2^*) \cap [0, -\frac{\gamma_1}{1-\gamma_1}]} 1 \, d\omega = c_{\rm ad}, \, p \cdot \int_{\Omega^*(\gamma_1^*, \gamma_2^*) \cap [-\frac{\gamma_2}{1-\gamma_2}, 1]} 1 \, d\omega = c_{\rm ad}$. In equilibrium, the solutions of these equations must be $\gamma_1 = \gamma_1^*$ and $\gamma_2 = \gamma_2^*$. This equilibrium condition then becomes $|([0, \underline{\omega}] \cup [\overline{\omega}, 1]) \cap [-\frac{\gamma_2}{1-\gamma_2}, 1]|_{\gamma_2 = \gamma_2^*} = |([0, \underline{\omega}] \cup [\overline{\omega}, 1]) \cap [0, -\frac{\gamma_1}{1-\gamma_1}]|_{\gamma_1 = \gamma_1^*} = c_{\rm ad}/p$. Note that $\underline{\omega}$ and $\overline{\omega}$ depend on γ_1^* and γ_2^* .

Claim 3 characterizes the firm's optimal advertising choice γ_1 and γ_2 given an arbitrary click

decision denoted by $\Omega = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$. Then, Claim 4 analyzes the consumers' optimal click decision given that they expect the firm to use a symmetric advertising strategy with $\gamma_1^* \cdot \gamma_2^* = 1$. Finally, Claim 5 shows the existence and uniqueness of the symmetric equilibrium.

Claim 3 (Characterization of Σ^*). Given $\Omega = [0, \underline{\omega}] \cup [\overline{\omega}, 1]$, if $c_{ad} > p(1 - \overline{\omega} + \underline{\omega})$, it is $\Sigma^* = \emptyset$. Otherwise, Σ^* is as characterized in Claim 2 (and Lemma 2), where γ_1^* and γ_2^* are defined as follows:

$$\gamma_{1}^{*} = \begin{cases} -\frac{c_{ad}}{p-c_{ad}}, & \text{if } c_{ad} \leq p \cdot \underline{\omega}, \\ -\frac{p(\overline{\omega}-\underline{\omega})+c_{ad}}{p(1-\overline{\omega}+\underline{\omega})-c_{ad}}, & \text{if } p \cdot \underline{\omega} < c_{ad} < p(1-\overline{\omega}+\underline{\omega}), \\ -\infty, & \text{if } c_{ad} \geq p(1-\overline{\omega}+\underline{\omega}), \end{cases}, \quad \gamma_{2}^{*} = \begin{cases} -\frac{p-c_{ad}}{c_{ad}}, & \text{if } c_{ad} \leq p \cdot (1-\overline{\omega}), \\ -\frac{p(1-\overline{\omega}+\underline{\omega})-c_{ad}}{p(\overline{\omega}-\underline{\omega})+c_{ad}}, & \text{if } p \cdot (1-\overline{\omega}) < c_{ad} < p(1-\overline{\omega}+\underline{\omega}), \\ 0, & \text{if } c_{ad} \geq p(1-\overline{\omega}+\underline{\omega}). \end{cases}$$

We have $\frac{\partial \gamma_1^*}{\partial \underline{\omega}} \ge 0$, $\frac{\partial \gamma_1^*}{\partial \overline{\omega}} \le 0$, $\frac{\partial \gamma_2^*}{\partial \underline{\omega}} \le 0$ and $\frac{\partial \gamma_2^*}{\partial \overline{\omega}} \ge 0$.

Proof. Recall that $\pi(\alpha, \beta; \Omega) = p \cdot \int_{\omega \in \Omega} \mathbf{1}\{\omega \alpha + (1-\omega)\beta \ge p\} d\omega - c_{\mathrm{ad}}$. Note that if $c_{\mathrm{ad}} > p(1-\overline{\omega}+\underline{\omega})$, then $\pi(\alpha, \beta; \Omega) < 0$ for all (α, β) . So, the firm does not want to target any consumer, i.e., $\Sigma^* = \emptyset$.

Next, we consider the case $c_{ad} \leq p(1 - \overline{\omega} + \underline{\omega})$. We first derive the optimal γ_1^* . Let β_0 be the minimum value for which $\pi(0, \beta_0; \Omega) \geq 0$. Notice that we allow β_0 to be > 1. Also, for $c_{ad} > 0$, it is always $\beta_0 > p$. Moreover, note that $\gamma_1 = 1 - \beta_0/p$.

We have that $\pi(0, \beta_0; \Omega) = p \cdot \int_{\omega \in \Omega} \mathbf{1}\{(1-\omega)\beta_0 \ge p\} d\omega - c_{\mathrm{ad}} = p \cdot \int_{\omega \in \Omega} \mathbf{1}\{\omega \le 1 - p/\beta_0\} d\omega - c_{\mathrm{ad}} = p \cdot \int_{\omega \in \Omega \cap [0, 1-p/\beta_0]} 1 d\omega - c_{\mathrm{ad}} = 0$. There are two cases: (1) If $0 \le 1 - p/\beta_0 \le \overline{\omega}$, then $p \cdot \int_{\omega \in \Omega \cap [0, 1-p/\beta_0]} 1 d\omega = p \cdot \min\{1 - p/\beta_0, \underline{\omega}\}$; (2) If $\overline{\omega} < 1 - p/\beta_0 \le 1$, then $p \cdot \int_{\omega \in \Omega \cap [0, 1-p/\beta_0]} 1 d\omega = p \cdot (\underline{\omega} + 1 - p/\beta_0 - \overline{\omega})$.

Suppose that $c_{\rm ad} \leq p \cdot \underline{\omega}$. For $\beta_0 = \frac{p^2}{p - c_{\rm ad}}$, we are in the first case and we get that $\pi(0, \beta_0; \Omega) = p \cdot \min\{1 - p/\beta_0, \underline{\omega}\} - c_{\rm ad} = 0$. For $\beta_0 < \frac{p^2}{p - c_{\rm ad}}$, we are still in the first case and $\pi(0, \beta_0; \Omega) < 0$. Therefore, it is $\beta_0 = \frac{p^2}{p - c_{\rm ad}}$. From this, we obtain the value of γ_1 to be $\gamma_1 = 1 - \beta_0/p = -\frac{c_{\rm ad}}{p - c_{\rm ad}}$.

Now suppose that $p \cdot \underline{\omega} < c_{\rm ad} < p(1 - \overline{\omega} + \underline{\omega})$. For $\beta_0 = \frac{p^2}{p(1 - \overline{\omega} + \underline{\omega}) - c_{\rm ad}}$, we are in the second case and $\pi(0, \beta_0; \Omega) = p \cdot (\underline{\omega} + 1 - p/\beta_0 - \overline{\omega}) - c_{\rm ad} = 0$. For $\beta_0 \in \left(\frac{p}{1 - \overline{\omega}}, \frac{p^2}{p(1 - \overline{\omega} + \underline{\omega}) - c_{\rm ad}}\right)$, we are still in the second case and $\pi(0, \beta_0; \Omega) < 0$. Therefore, it is $\beta_0 = \frac{p^2}{p(1 - \overline{\omega} + \underline{\omega}) - c_{\rm ad}}$. From this, we obtain the value of γ_1 to be $\gamma_1 = 1 - \beta_0/p = -\frac{p(\overline{\omega} - \underline{\omega}) + c_{\rm ad}}{p(1 - \overline{\omega} + \underline{\omega}) - c_{\rm ad}}$.

Similarly, by considering the minimum value α_0 for which $\pi(\alpha_0, 0; \Omega) \ge 0$, we can derive γ_2 . It is straightforward to show the following properties (which will be used later to prove uniqueness of a symmetric equilibrium): $\frac{\partial \gamma_1^*}{\partial \omega} \ge 0$, $\frac{\partial \gamma_1^*}{\partial \overline{\omega}} \le 0$, $\frac{\partial \gamma_2^*}{\partial \omega} \le 0$ and $\frac{\partial \gamma_2^*}{\partial \overline{\omega}} \ge 0$. Claim 4 (Characterization of $\underline{\omega}$ and $\overline{\omega}$). Suppose that consumers' anticipation for the firm's advertising strategy is symmetric, i.e., $\gamma_1^* \cdot \gamma_2^* = 1$. Then, Ω^* is symmetric about $\omega = 1/2$, i.e., $\overline{\omega} = 1 - \underline{\omega}$.

Proof. For given γ_1^* , γ_2^* , a consumer's expected payoff for clicking on an ad is $\mathbb{E}[\max\{u-p,0\}; \omega, \gamma_1^*, \gamma_2^*] - c_{\text{click}} = \frac{1}{|\Sigma|} \cdot \left(\int_{\Sigma} \max\{\omega\alpha + (1-\omega)\beta - p, 0\} \, d\alpha \, d\beta\right) - c_{\text{click}} = \frac{1}{|\Sigma|} \cdot \left(\int_{\Sigma_1} \max\{\omega\alpha + (1-\omega)\beta - p, 0\} \, d\alpha \, d\beta + \int_{\Sigma_2} \max\{\omega\alpha + (1-\omega)\beta - p, 0\} \, d\alpha \, d\beta\right) - c_{\text{click}} = \frac{1}{|\Sigma|} \cdot \left(H_1(\omega) + H_2(\omega)\right) - c_{\text{click}},$

where $\Sigma_1 = \{(\alpha, \beta) \in [0, 1]^2 \mid \beta \geq \gamma_1^* \cdot \alpha + p(1 - \gamma_1^*) \text{ and } \alpha \leq p\}$ and $\Sigma_2 = \{(\alpha, \beta) \in [0, 1]^2 \mid \beta \geq \gamma_2^* \cdot \alpha + p(1 - \gamma_2^*) \text{ and } \alpha > p\}$. Note that H_1 and H_2 (and hence $H_1 + H_2$) are continuous and convex. Given the firm's symmetric strategy $\gamma_1^* \cdot \gamma_2^* = 1$, $\mathbb{E}[\max\{u - p, 0\}; \omega, \gamma_1^*, \gamma_2^*]$ must be symmetric about $\omega = 1/2$.

We explicitly compute $H_1(\omega)$ and $H_2(\omega)$ as follows: If $p \cdot (1 - \gamma_1^*) \leq 1$, then $H_1(\omega) = \frac{p(p^2(-\gamma_1^2(1-\omega)-2\gamma_1^*\omega+3)-3p(2-\omega)-3\omega+3)}{6}$ if $\omega < -\frac{\gamma_1^*}{1-\gamma_1^*}$; $\frac{(1-p)^3(1-\omega)^3-(1-p-\omega)^3}{6(1-\omega)\omega}$ if $-\frac{\gamma_1^*}{1-\gamma_1^*} \leq \omega \leq 1-p$; and $\frac{(1-p)^3(1-\omega)^2}{6\omega}$ if $\omega > 1-p$. If $p \cdot (1-\gamma_1^*) > 1$, then $H_1(\omega) = \frac{(1-p)^3(-2\gamma_1^*(1-\omega)-\omega)}{6\gamma_1^2}$ if $\omega \leq -\frac{\gamma_1^*}{1-\gamma_1^*}$; and $\frac{(1-p)^3(1-\omega)^2}{6\omega}$ if $\omega > -\frac{\gamma_1^*}{1-\gamma_1^*}$. If $-p \cdot (1-\gamma_2^*) \geq \gamma_2^*$, then $H_2(\omega) = \frac{(1-p)^3(\omega^2-3\omega+3)}{6(1-\omega)}$ if $\omega < p$; $\frac{p^3(1-\omega)^2+3p^2(2-\omega)\omega-3p(3-\omega)\omega+3\omega}{6\omega}$ if $p \leq \omega \leq -\frac{\gamma_2^*}{1-\gamma_2^*}$; and $\frac{3\gamma_2^2(1-p)(-p(2-\omega)+1)-2\gamma_2^*p^3(1-\omega)-p^3\omega}{6\gamma_2^2}$ if $\omega > -\frac{\gamma_2^*}{1-\gamma_2^*}$. If $-p \cdot (1-\gamma_2^*) < \gamma_2^*$, then $H_2(\omega) = \frac{(1-p)^3(\omega^2-3\omega+3)}{6(1-\omega)}$ if $\omega < -\frac{\gamma_2^*}{1-\gamma_2^*}$. If $-p \cdot (1-\gamma_2^*) < \gamma_2^*$, then $H_2(\omega) = \frac{(1-p)^3(\omega^2-3\omega+3)}{6(1-\omega)}$ if $\omega \leq -\frac{\gamma_2^*}{1-\gamma_2^*}$; and $\frac{(1-p)^3(-\gamma_2^2(1-\omega)-2\gamma_2^*\omega+3)}{6}$ if $\omega > -\frac{\gamma_2^*}{1-\gamma_2^*}$.

It is easy to verify (by plugging in $1 - \omega$ in place of ω) that if $\gamma_1^* \cdot \gamma_2^* = 1$, then $H_1(\omega) + H_2(\omega)$ is symmetric about $\omega = 1/2$, where the function attains its minimum. So, if $\mathbb{E}[\max\{u-p,0\}; \omega = 1/2, \gamma_1^*, 1/\gamma_1^*] > c_{\text{click}}$, then we can define $\underline{\omega} = \overline{\omega} = 1/2$, and hence $\Omega^* = [0, 1]$. If $\mathbb{E}[\max\{u-p,0\}; \omega = 0, \gamma_1^*, 1/\gamma_1^*] < c_{\text{click}}$, then $\Omega^* = \emptyset$ (i.e., $\underline{\omega} < 0$ and $\overline{\omega} > 1$) which leads to trivial equilibria. Lastly, if $c_{\text{click}} \leq \mathbb{E}[\max\{u-p,0\}; \omega = 0, \gamma_1^*, 1/\gamma_1^*]$, then $\Omega^* \neq \emptyset$.

Claim 5. There exists a unique symmetric equilibrium if and only if $c_{click} \leq p$ and $c_{ad} \leq c_{click}^{max}(c_{ad}, p)$, where $c_{click}^{max}(c_{ad}, p)$ is defined in the proof of Lemma 5.

Proof. For $-\frac{\gamma_2}{1-\gamma_2} \leq \omega \leq -\frac{\gamma_1}{1-\gamma_1}$, we can see from the proof of Claim 4 that $H_1(\omega) + H_2(\omega)$ is independent of ω . This means in a symmetric equilibrium, it will either be $\overline{\omega} \geq -\frac{\gamma_1}{1-\gamma_1}$ or $\overline{\omega} = \underline{\omega}$.

We will show that in a symmetric equilibrium it has to be $\gamma_1^* = -\frac{c_{\rm ad}}{p-c_{\rm ad}}$. Suppose that there is a symmetric equilibrium with $\gamma_1^* = -\frac{p(\overline{\omega}-\underline{\omega})+c_{\rm ad}}{p(1-\overline{\omega}+\underline{\omega})-c_{\rm ad}}$. If $\overline{\omega} \ge -\frac{\gamma_1^*}{1-\gamma_1^*}$, then this means that $\overline{\omega} \ge \frac{c_{\rm ad}}{p} + \overline{\omega} - \underline{\omega}$, or equivalently $p \cdot \underline{\omega} \ge c_{\rm ad}$, and from the definition of γ_1^* in Claim 3 this implies that $\gamma_1^* = -\frac{c_{\rm ad}}{p-c_{\rm ad}}$. If $\overline{\omega} = \underline{\omega}$, then again we conclude that $\gamma_1^* = -\frac{c_{\rm ad}}{p-c_{\rm ad}}$. In other words, when a symmetric equilibrium exists it is the unique one derived from $\gamma_1^* = -\frac{c_{\rm ad}}{p-c_{\rm ad}}$.

For existence, we need $c_{ad} \leq p$ (because otherwise the firm sends no ads) and $c_{click} \leq c_{click}^{max}(c_{ad}, p)$ (because otherwise no consumer clicks; see also Equation (7) in the proof of Lemma 5 for the expression for $c_{click}^{max}(c_{ad}, p)$).

This completes the proof of the proposition.

Proof of Proposition 3

The result follows from the following lemmas.

Lemma 4. The profit under sd-targeting is decreasing in the interval $((1-p)^2/2, (1-p)/2]$ with respect to c_{click} , and it is 0 for $c_{click} = (1-p)/2$.

Proof. From the proof of Lemma 1 we know that in the interval $((1-p)^2/2, (1-p)/2]$ it is $\underline{\omega} = 0$. As c_{click} increases in this interval, fewer consumers want to click, i.e. $\overline{\omega}$ increases. Because of this, $\overline{\alpha}$ also increases (see proof of Proposition 1). Fewer targeted consumers and a smaller click-through rate causes a decrease in the profit. At the point where $c_{\text{click}} = (1-p)/2$, the equilibrium is $\overline{\alpha} = p^2/(p-c_{\text{ad}})$ and $\overline{\omega} = 1 - c_{\text{ad}}/p, \underline{\omega} = 0$. This holds because $\mathbb{E}[\max\{u-p,0\}; \omega = 1 - c_{\text{ad}}/p, \overline{\alpha} = p^2/(p-c_{\text{ad}})] - c_{\text{click}} = \frac{1-p}{2} - c_{\text{click}} = 0$. Notice also that for any $\alpha \ge \overline{\alpha}$, it is $F(\alpha; \overline{\omega}) = 1 - \overline{\omega} = c_{\text{ad}}/p$. Therefore, $\pi(\alpha; \underline{\omega}, \overline{\omega}) = p \cdot (G(\alpha; \underline{\omega}) + F(\alpha; \overline{\omega})) - c_{\text{ad}} = p \cdot (0 + c_{\text{ad}}/p) - c_{\text{ad}} = 0$. As a result, the profit at $c_{\text{click}} = (1-p)/2$ becomes 0.

Lemma 5. Let $L(c_{ad}, p) := \lim_{c_{click} \to c_{click}^{max}(c_{ad}, p)^{-}} \Pi^{*md}(c_{ad}, c_{click}, p)$ be md-targeting profit's limit as c_{click} tends to $c_{click}^{max}(c_{ad}, p)$ from below. Then, $L(c_{ad}, p)$ approaches 0 as c_{ad} tends to 0.

Proof. We define the threshold c_{click}^{\max} as $c_{\text{click}}^{\max}(c_{\text{ad}}, p) := (H_1(c_{\text{ad}}/p) + H_2(c_{\text{ad}}/p))/|\Sigma|$ (where H_1 and H_2 are defined as in the proof of Proposition 2). For $c_{\text{click}} = c_{\text{click}}^{\max}$, the equilibrium $\underline{\omega}$ becomes equal to c_{ad}/p (by definition). If we look at the definition of γ_1^* , we see that this is the smallest value $\underline{\omega}$ can take in a non-trivial symmetric equilibrium. As a result, for $c_{\text{click}} > c_{\text{click}}^{\max}$, there is no

non-trivial symmetric equilibrium. If we now consider all the possible cases for $c_{\rm ad}$, we find out that

$$c_{\text{click}}^{\max}(c_{\text{ad}}, p) = \begin{cases} \frac{(1-p)(4c_{\text{ad}}-p)}{6c_{\text{ad}}} & \text{if } c_{\text{ad}} > p(1-p) \text{ and } c_{\text{ad}} > \frac{p}{2}, \\ \frac{(1-p)p}{6(p-c_{\text{ad}})} & \text{if } c_{\text{ad}} > p(1-p) \text{ and } c_{\text{ad}} \le \frac{p}{2}, \\ \frac{c_{\text{ad}}^2 - 3c_{\text{ad}}(1-p)p + 3(1-p)^2 p^2}{6p(p(1-p)(1+p)-c_{\text{ad}})} & \text{if } c_{\text{ad}} > p(1-p) \text{ and } c_{\text{ad}} \le \frac{p}{2}, \\ \frac{3c_{\text{ad}}(1-2p)(p-c_{\text{ad}}) + p^5}{6c_{\text{ad}}(p(1-p)(1+p)-c_{\text{ad}})} & \text{if } c_{\text{ad}} \le p(1-p) \text{ and } c_{\text{ad}} \le p^2, \\ \frac{3c_{\text{ad}}^2(1-2p)(p-c_{\text{ad}}) + p^5}{6c_{\text{ad}}(p(1-p)(1+p)-c_{\text{ad}})} & \text{if } c_{\text{ad}} \le p(1-p) \text{ and } c_{\text{ad}} > p^2 \text{ and } c_{\text{ad}} \le \frac{p}{2}, \\ \frac{3c_{\text{ad}}^2(1-2p)-2c_{\text{ad}}p(2p^3-6p+3) + 3p^2(p^3-2p+1)}{6(p-c_{\text{ad}})(p(1-p)(1+p)-c_{\text{ad}})} & \text{if } c_{\text{ad}} \le p(1-p) \text{ and } c_{\text{ad}} > p^2 \text{ and } c_{\text{ad}} > \frac{p}{2}. \end{cases}$$

Because we assume that c_{ad} is sufficiently small (and therefore $c_{ad} < \min\{p^2, p(1-p)\}$), it is $c_{click}^{\max}(c_{ad}, p) = \frac{c_{ad}^2 - 3c_{ad}(1-p)p + 3(1-p)^2p^2}{6p(p(1-p)(1+p)-c_{ad})}$. At the point $c_{click} = c_{click}^{\max}$ the profit is positive for $c_{ad} > 0$. However, as c_{ad} tends to 0, $\underline{\omega}$ tends to 0 and therefore the profit approaches 0.

Lemma 6. For $c_{ad} < \min\{p^2, p(1-p)\}$, it holds that $(1-p)^2/2 < c_{click}^{max}(c_{ad}, p) < (1-p)/2$. Proof. The inequality becomes $\frac{(1-p)^2}{2} < \frac{c_{ad}^2 - 3c_{ad}(1-p)p + 3(1-p)^2p^2}{6p(p(1-p)(1+p)-c_{ad})} < \frac{1-p}{2}$. The first inequality is equivalent to $c_{ad}^2 - 3c_{ad}(1-p)p^2 + 3(1-p)^2p^4 > 0$, which is true for any c_{ad} . The second inequality is equivalent to $c_{ad}^2 < 3(1-p)^2p^3$. This is true because $c_{ad}^2 < \min\{p^4, p^2(1-p)^2\} \le 3(1-p)^2p^3$, for any $p \in [0, 1]$.

Combining the lemmas above, we can conclude that there exist δ , $\delta'(\delta) > 0$ such that for $c_{\rm ad} < \delta$, $c_{\rm click} = c_{\rm click}^{\rm max}(c_{\rm ad}, p) - \delta'$, the equilibria under both *sd*- and *md*-targeting are non-trivial and the profit under *sd*-targeting is larger.

We can now define $c_{\text{click}}^{\min}(c_{\text{ad}}, p)$ as the root of the equation $\Pi^{*sd}(p, c_{\text{click}}, c_{\text{ad}}) = \Pi^{*md}(p, c_{\text{click}}, c_{\text{ad}})$ w.r.t. to c_{click} , i.e., the minimum c_{click} such that $\Pi^{*sd} \geq \Pi^{*md}$. Finally, we can define $c_{\text{ad}}^{\max}(p)$ as the root of the equation $c_{\text{click}}^{\min}(c_{\text{ad}}, p) = c_{\text{click}}^{\max}(c_{\text{ad}}, p)$ w.r.t. to c_{ad} , i.e., the maximum c_{ad} such that $c_{\text{click}}^{\min} \leq c_{\text{click}}^{\max}$. This completes the proof of the proposition.

Asymmetric Equilibria. In some cases, under *md*-targeting, in addition to the symmetric equilibrium there are also asymmetric equilibria. Due to the symmetric structure of the game, the asymmetric equilibria come in pairs. Using the analysis in the proof of Proposition 2, we can characterize and find the asymmetric equilibria of the game when they exist. Note that Claim 3 and 4 in that proof do not require a symmetric strategy or a symmetric set for consumers' click decision, respectively. Figure 8 shows one example with two asymmetric equilibria and a symmetric one.



Figure 8: All the different equilibrium *md*-targeting strategies for p = 0.61, $c_{\text{click}} = 0.1$, and $c_{\text{ad}} = 0.05$.

The set for the firm's targeted advertising (shaded in gray) is smaller for an asymmetric equilibrium than for a symmetric one, which implies that the targeted consumers' uncertainty is mitigated in asymmetric equilibria. Therefore, the firm's equilibrium profit can be greater under md-targeting than under sd-targeting. In cases with multiple equilibria like this, an equilibrium-selection rule is needed. We focus our analysis on the symmetric equilibrium of the game, as the more natural one (from Proposition 2 we know that there can be only one symmetric equilibrium). However, our results are robust under different equilibrium-selection rules. Example 2 in Appendix A.4 shows a case where asymmetric equilibria do not exist, and hence the symmetric equilibrium is the only non-trivial equilibrium of the game. As sd-targeting has higher profit than md-targeting in this example, Proposition 3 is robust under different equilibrium-selection rules, even though for a more limited parameter region.

Corollary 2. If c_{ad} is sufficiently small and c_{click} is in an intermediate range, then the firm sends more ads under the md-environment compared to the sd-environment (i.e., $|\Sigma^{*md}| > |\Sigma^{*sd}|$), while the conversion rate (number of purchases per ad) is higher under the sd-environment.

Proof of Corollary 2

First, we will show that for a sufficiently small c_{ad} and for a c_{click} sufficiently close to (and at most) c_{click}^{max} , it holds that $|\Sigma^{*md}| > |\Sigma^{*sd}|$. Note that as c_{ad} approaches 0, $|\Sigma^{*md}|$ approaches $1 - p^2$, i.e. the firm targets every consumer in the set $[0, 1]^2 \setminus [0, p]^2$. For sd-targeting, as c_{ad} approaches 0

and when c_{click} is close to c_{click}^{\max} , we have that (1) $\underline{\omega}^{*sd} = 0$ (see Lemma 1 and Lemma 6), and (2) $\overline{\alpha}^*$ approaches a point where $\overline{\omega}^{*sd} \cdot \overline{\alpha}^* + (1 - \overline{\omega}^{*sd}) \cdot 1 = p$. This is because at this point a consumer with $\alpha < \overline{\alpha}^*$ will never buy the product even in the best-case scenario where $\beta = 1$, therefore the firm doesn't target those consumers even if c_{ad} is very small. Thus, $|\Sigma^{*sd}|$ approaches $1 - \overline{\alpha}^* = \frac{1-p}{\overline{\omega}^{*sd}}$. Now, for $c_{\text{ad}} \to 0$, we have that $c_{\text{click}}^{\max} \to \frac{1-p}{2(1+p)}$ (see Lemma 5). For this value of c_{click} and using the expressions for $\mathbb{E}[\max\{u-p,0\}; \omega, \overline{\alpha}]$ from the proof of Lemma 1, we can also find that $\overline{\omega}^{*sd} = \frac{3p^2 + \sqrt{3(1-p)^2(3+2p-p^2)} + 2p-1}{2(1+p)}$. Therefore, to show that $|\Sigma^{*sd}| < |\Sigma^{*md}|$, it is enough to show that $\frac{2(1-p^2)}{3p^2 + \sqrt{3(1-p)^2(3+2p-p^2)} + 2p-1} < 1-p^2$, or equivalently $3p^2 + \sqrt{3(1-p)^2(3+2p-p^2)} + 2p-3 > 0$, which is true for any $p \in (0, 1)$.

Finally, let CR^{*md} and CR^{*sd} denote the conversion rates of md- and sd-targeting respectively. For the profits, it holds that $\Pi^{*md} = p \cdot |\Sigma^{*md}| \cdot CR^{*md} - c_{ad} \cdot |\Sigma^{*md}| = |\Sigma^{*md}| \cdot (p \cdot CR^{*md} - c_{ad})$ and similarly $\Pi^{*sd} = |\Sigma^{*sd}| \cdot (p \cdot CR^{*sd} - c_{ad})$. Under the conditions of Proposition 3, it holds that $\Pi^{*md} < \Pi^{*sd}$ and $|\Sigma^{*md}| > |\Sigma^{*sd}|$, therefore it must be $CR^{*sd} > CR^{*md}$.

Proof of Proposition 4

From Lemma 2, the only candidate among the possible *sd*-targeting strategies to be an equilibrium for the *md*-environment is $\Sigma^{*md} = [p, 1] \times [0, 1]$, which corresponds to the setting $\gamma_1 = \gamma_2 = -\infty$. This is because in an equilibrium of the *md*-environment, the boundary of Σ^{*md} must pass through the point (p, p). Therefore, the consumer of type $(\alpha, \beta) = (p, 0)$ is on the boundary, which means that given the equilibrium set of consumers who click the ad Ω^{*md} , the firm's expected profit from this consumer must be non-negative. However, we have that $\pi(p, 0; \Omega^{*md}) = p \cdot \int_{\omega \in \Omega^{*md}} \mathbf{1}\{\omega \geq 1\} d\omega - c_{\mathrm{ad}} = -c_{\mathrm{ad}} < 0$. Therefore, the consumer (p, 0) cannot be in the set Σ^{*md} .

Proof of Proposition 5

Let $p^{*md}(c_{ad}, c_{click}, mc)$ be the optimal price under *md*-targeting, given the values of c_{ad} , c_{click} , and *mc*. Consider a sufficiently small c_{ad} . From Proposition 3, we know that if the inequality

$$c_{\text{click}}^{\min}(c_{\text{ad}}, mc, p^{*md}(c_{\text{ad}}, c_{\text{click}}, mc)) \le c_{\text{click}} \le c_{\text{click}}^{\max}(c_{\text{ad}}, mc, p^{*md}(c_{\text{ad}}, c_{\text{click}}, mc))$$
(8)

is satisfied for c_{click} , then the profit under *sd*-targeting with price p^{*md} is higher than the profit under *md*-targeting with price p^{*md} , i.e. higher than the optimal *md*-profit. Therefore, the optimal *sd*-profit will also be higher than the optimal *md*-profit. The only thing that remains is to show that the inequality (8) about c_{click} can actually be satisfied, i.e. that there is a non-empty region of the parameter space where inequality (8) holds (notice that c_{click} is involved in all three parts of the inequality). This can be shown with a numerical example.

First, we consider a very small value of c_{ad} , let's say 0.001 and a sufficiently large mc (but not too large to make the profit zero), let's say 0.6. We will show that for $c_{click} = 0.08$ inequality (8) holds. From Figure 7 we can see that the optimal price p^{*md} is ≈ 0.653 . From the expression for c_{click}^{max} in Lemma 5 and for $p = p^{*md}$ we can find that $c_{click}^{max} = \frac{c_{ad}^2 - 3c_{ad}(1-p)p + 3(1-p)^2p^2}{6p(p(1-p)(1+p)-c_{ad})} \approx 0.105$. On the other end, for c_{click}^{min} , we can find that the lowest value of c_{click} for which $\Pi^{sd}(p^{*md}) = \Pi^{md}(p^{*md})$ is ≈ 0.0791 . Since it holds that 0.0791 < 0.08 < 0.105, inequality (8) is satisfied, and therefore the result holds. In Example 3 in Appendix A.4, we can also see the optimal price p^{*sd} under sd-targeting and the optimal profits for both environments, which are not needed for this proof.

Proof of Proposition 6

The analysis of the *u-sd*-targeting model is very similar to that of *sd*-targeting in the proofs of Lemma 1 and Proposition 1. The main difference is that the expected utility of a consumer who clicks $\mathbb{E}^{u-sd}[\max\{u-p,0\};\omega,\overline{x}^{u-sd}]$ is different than $\mathbb{E}^{sd}[\max\{u-p,0\};\omega,\overline{\alpha}^{sd}]$ as defined in Lemma 1. More specifically, if we define the function $H(\omega;\overline{\alpha}^{sd}) := (1-\overline{\alpha}^{sd}) \cdot \mathbb{E}^{sd}[\max\{u-p,0\};\omega,\overline{\alpha}^{sd}]$, then it is $\mathbb{E}^{u-sd}[\max\{u-p,0\};\omega,\overline{x}^{u-sd}] = \frac{1}{2} \cdot \frac{H(\omega;\overline{x}^{u-sd})}{1-\overline{x}^{u-sd}} + \frac{1}{2} \cdot \frac{H(1-\omega;\overline{x}^{u-sd})}{1-\overline{x}^{u-sd}} =: \frac{H^{u-sd}(\omega;\overline{x}^{u-sd})}{1-\overline{x}^{u-sd}}$. In contrast to $H(\omega)$, the function $H^{u-sd}(\omega)$ is symmetric about the line $\omega = 1/2$, which results in the symmetric nature of the equilibrium click set $\Omega^{*u-sd} = [0, \widehat{\omega}^{*u-sd}] \cup [1 - \widehat{\omega}^{*u-sd}, 1]$.

For the comparison of the profit under *u-sd*-targeting and *md*-targeting, the proof is similar to the proof of Proposition 3. First, we need to identify the point where the *u-sd*-profit becomes 0 as c_{click} increases, and then compare this point to $c_{\text{click}}^{\max}(c_{\text{ad}}, p)$. Let ω^* be the root of the equation $\omega - (1-p) \cdot \ln(1-\omega) = \frac{c_{\text{ad}}}{p}$ in [0, 1/2].²⁸ As c_{click} increases, the profit under *u-sd*-targeting becomes 0 for $c_{\text{click}} = c_{\text{click}}^{u-sd}$:= $\lim_{\overline{x}^{u-sd} \to 1} \frac{H^{u-sd}(\omega^*; \overline{x}^{u-sd})}{1-\overline{x}^{u-sd}} = \frac{\omega^{*2} - 3\omega^* + 2p\omega^* + p^2 - 4p + 3}{4(1-\omega^*)}$. Due

²⁸This is the equation $\pi(1;\omega) = p \cdot (G(1;\omega) + F(1;1-\omega)) - c_{ad} = 0$, where F and G are defined as in the proof of Proposition 1.

to Lemma 6, it is enough to show that for sufficiently small c_{ad} , it holds that $c_{click}^{u-sd,\max} \ge \frac{1-p}{2}$. This is true because as c_{ad} approaches 0, ω^* approaches 0 and therefore $c_{click}^{u-sd,\max}$ approaches $\frac{p^2-4p+3}{4}$, which is greater than $\frac{1-p}{2}$ for any p < 1.

A.2 Pay-Per-Click Advertising

In our main model, the firm pays a fixed amount of c_{ad} for every ad it sends to a consumer, i.e. we use a pay-per-impression (PPM) payment scheme. In this section, we consider an alternative advertising payment scheme of pay-per-click (PPC) where the firm pays for an ad a cost $c_{ppc} > 0$ only if a consumer clicks on it. Through the analysis here, we want to establish robustness of our main result that the firm can sometimes be better off under *sd*-targeting than under *md*-targeting. The analysis is the same as the PPM scheme except that now the firm incurs ad cost for each ad clicked. So, the firm's expected profit function given that it chooses a targeting set Σ and faces consumers whose click decision is represented by the set Ω (which depends on consumers' anticipation for the firm's targeting set) is $\Pi(\Sigma; \Omega) = p \cdot \int_{\Sigma} \left(\int_{\Omega} \mathbf{1} \{ \omega \alpha + (1 - \omega)\beta \ge p \} d\omega \right) d\alpha d\beta - c_{ppc} \cdot |\Omega| \cdot |\Sigma|$. The set Σ^* is an equilibrium targeting set if $\Sigma^* \in \operatorname{argmax}_{\Sigma} \Pi(\Sigma; \Omega(\Sigma^*))$.

One may posit that the result may not be robust to the PPC payment scheme if most of the benefit of *sd*-targeting is realized through its ability to induce higher click-through-rate. However, we find that the result is still robust here. The reason is that the main benefit of *sd*-targeting is not only in its ability to induce a higher click-through-rate, but also more generally yielding a higher overall conversion rate. By maintaining the set of consumers who are targeted small enough, the firm is able to convert a higher fraction of the targeted consumers into making a purchase. Example 5 is an analogue of Example 2 for the PPC scheme, both in Appendix A.4.

A.3 Zero-Dimensional Targeting

Here we analyze a case in which the firm has no access to customer information, i.e., the zerodimensional (zd) targeting environment. In this model, the firm does not know anything about α or β . As a result, in equilibrium the firm will either target every consumer or no consumer, i.e., $\Sigma^{*zd} = [0, 1]^2$, or \emptyset .

Proposition 7 (zero-dimensional vs single-dimensional targeting). The firm's expected profit under

sd-targeting is at least equal to the profit under zd-targeting.

Proof. Notice that if in equilibrium $\Sigma^{*zd} = \emptyset$, then the result is trivially true. Therefore, we focus on the case where $\Sigma^{*zd} = [0,1]^2$. We make two simple observations. First, for a fixed targeting set Σ , a larger set Ω of consumers who click is weakly better for the firm. This is true because when Σ is fixed but more consumers click, then potentially more consumers buy the product, without any extra cost for the firm. Second, the set Ω^{*sd} of consumers who click under *sd*-targeting is at least as large as the set Ω^{*zd} of consumers who click under *zd*-targeting, i.e. $\Omega^{*sd} \supseteq \Omega^{*zd}$. This follows from the fact that $\Sigma^{*sd} \subseteq \Sigma^{*zd} = [0, 1]^2$.

Now suppose that there are equilibria such that $\Pi^{*zd} > \Pi^{*sd}$. Since the two profits are different, it is $\Sigma^{*sd} \neq [0,1]^2$. However, this means that under *sd*-targeting the firm has a profitable deviation; it can target every consumer, while Ω^{*sd} does not change and it is larger than Ω^{*zd} , therefore it can strictly increase its profit, which is a contradiction. It follows that $\Pi^{*sd} \geq \Pi^{*zd}$ always holds. \Box

A.4 Numerical Examples

In this section, we present some numerical examples that illustrate the various results in the paper.

Example 1 (The effect of c_{ad} on profit, Section 3.2). Let p = 0.78, $c_{click} = 0.025$, and $c_{ad} = 0.09$. Under *md*-targeting, we have that $\Omega^{*md} \approx [0, 0.349] \cup [0.651, 1]$ and $\gamma_1^* \approx -0.130$. The firm's expected profit is $\Pi^{*md} \approx 0.0556$. If we increase the advertising cost to $c'_{ad} = 0.1$, we get that $\Omega^{*md'} \approx [0, 0.377] \cup [0.623, 1]$, and $\gamma_1^{*'} \approx -0.147$. The firm's expected profit is $\Pi^{*md'} \approx 0.0571$.

Notice that when c_{ad} is higher, $\gamma_1^{*'} < \gamma_1^*$, i.e. the targeting set becomes smaller, but $|\Omega^{*md'}| > |\Omega^{*md}|$. This is what leads to $\Pi^{*md'}$ being larger.

Example 2 (A comparison between the *sd*- and *md*-targeting environments, Section 4, Figure 2, Proposition 3). Let p = 0.6, $c_{\text{click}} = 0.1$, and $c_{\text{ad}} = 0.05$. The *sd*-targeting equilibrium is characterized by $\Sigma^{*sd} \approx [0.568, 1]$ and $\Omega^{*sd} \approx [0.612, 1]$. The firm's expected profit is $\Pi^{*sd} \approx 0.0626$. On the other hand, the equilibrium of *md*-targeting is $\Sigma^{*md} \approx \{(\alpha, \beta) \mid \beta \geq -0.0909 \cdot \alpha + 0.655 \text{ or } \alpha \geq -0.0909 \cdot \beta + 0.655\}$ and $\Omega^{*md} \approx [0, 0.175] \cup [0.825, 1]$. The firm's equilibrium profit is $\Pi^{*md} \approx 0.0507$. We can see that $\Pi^{*sd} > \Pi^{*md}$. Notice that $|\Sigma^{*sd}| \approx 0.432 < 0.607 \approx |\Sigma^{*md}|$ and $|\Omega^{*sd}| \approx 0.388 > 0.350 \approx |\Omega^{*md}|$.

For comparison, we can also consider the case where the firm can commit to the targeting strategy and announce it before the consumers' click decision. Under this full-commitment (fc)case, it is $|\Sigma^{*sd,fc}| \approx 0.330 < |\Sigma^{*sd}|$ and $|\Sigma^{*md,fc}| \approx 0.427 < |\Sigma^{*md}|$. For the profits, it is $\Pi^{*sd,fc} \approx$ $0.0746 > \Pi^{*sd}$ and $\Pi^{*md,fc} \approx 0.169 > \Pi^{*md}$. Thus, the lack of commitment hurts the firm more under *md*-targeting than under *sd*-targeting.

Example 3 (Endogenous pricing, Section 5.1). Let $c_{\text{click}} = 0.08$, $c_{\text{ad}} = 0.001$, and mc = 0.6. Under sd-targeting, it is $p^{*sd} \approx 0.688$, $\Sigma^{*sd} \approx [0.647, 1]$, and $\Omega^{*sd} \approx [0.775, 1]$. So, the firm's advertising amount is $|\Sigma^{*sd}| \approx 0.353$ and the click-through rate is $|\Omega^{*sd}| \approx 0.225$. The firm's expected profit is $\Pi^{*sd} \approx 0.00534$. Under md-targeting, had the firm set the same price as the equilibrium price under sd-targeting, p^{*sd} , then its profit is $\Pi^{md}(p = p^{*sd}) = 0.00308 < \Pi^{*sd}$, i.e. the firm will have a lower profit compared to sd-targeting. Under md-targeting, the optimal price is $p^{*md} \approx 0.653$, and then $|\Sigma^{*md}| \approx 0.566$, and $|\Omega^{*md}| \approx 0.289$. Firm's profit is $\Pi^{*md} \approx 0.00452$ ($< \Pi^{*sd}$).

Example 4 (The *u-sd*-environment, Section 5.2). The following examples demonstrate that $\Pi^{*u-sd} > \Pi^{*md}$, whereas both $\Pi^{*u-sd} > \Pi^{*sd}$ and $\Pi^{*u-sd} < \Pi^{*sd}$ can hold. Let p = 0.61 and $c_{\text{click}} = 0.1$.

- (a) For $c_{\rm ad} = 0.12$, under *sd*-targeting, we have $|\Sigma^{*sd}| \approx 0.380$, $|\Omega^{*sd}| \approx 0.405$, and $\Pi^{*u-sd} \approx 0.0367$. Under *u-sd*-targeting, we have $|\Sigma^{*u-sd}| \approx 0.393$, $|\Omega^{*u-sd}| \approx 0.539$, and $\Pi^{*u-sd} \approx 0.0394$. Under *md*-targeting, we have $|\Sigma^{*md}| \approx 0.537$, $|\Omega^{*md}| \approx 0.421$, and $\Pi^{*md} \approx 0.0265$. Notice that $\Pi^{*u-sd} > \Pi^{*sd} > \Pi^{*md}$.
- (b) For $c_{ad} = 0.05$, under *sd*-targeting, we have $|\Sigma^{*sd}| \approx 0.414$, $|\Omega^{*sd}| \approx 0.361$, and $\Pi^{*u-sd} \approx 0.0569$. Under *u-sd*-targeting, we have $|\Sigma^{*u-sd}| \approx 0.442$, $|\Omega^{*u-sd}| \approx 0.387$, and $\Pi^{*u-sd} \approx 0.0438$. Under *md*-targeting, we have $|\Sigma^{*md}| \approx 0.595$, $|\Omega^{*md}| \approx 0.300$, and $\Pi^{*md} \approx 0.0391$. Notice that $\Pi^{*sd} > \Pi^{*u-sd} > \Pi^{*md}$.

Example 5 (Pay-per-click advertising, Appendix A.2). Let p = 0.6, $c_{\text{click}} = 0.1$, and $c_{\text{ppc}} = 0.1$. In the *sd*-targeting equilibrium, $|\Sigma^{*sd}| \approx 0.445$ and $|\Omega^{*sd}| \approx 0.371$. The firm's expected profit is $\Pi^{*sd} \approx 0.0650$. In the *md*-targeting equilibrium, $|\Sigma^{*md}| \approx 0.619$ and $|\Omega^{*md}| \approx 0.326$. The resulting expected profit of the firm is $\Pi^{*md} \approx 0.0559$. We can see that $\Pi^{*sd} > \Pi^{*md}$. Notice that $|\Sigma^{*sd}| < |\Sigma^{*md}|$ and $|\Omega^{*sd}| > |\Omega^{*md}|$.

B Online Appendix: A Simpler Model - Discrete Distributions

In this section, we consider a simpler version of the main model where α , β , and ω for each consumer are distributed uniformly in the set $\{0, 1/2, 1\}$. The following proposition verifies the robustness of the result in Proposition 3 in this new setting with discrete values.

Proposition 8 (Discrete distributions). If $\left(\frac{5}{8} or <math>\left(\frac{3}{4} , then$

- Under the sd-targeting environment, the firm targets the consumers in the set $\Sigma^{*sd} = \{(\alpha, \beta) : \alpha = 1\}$ and the consumers who click are those in the set $\Omega^{*sd} = \{\omega : \omega \in \{0, 1/2, 1\}\}.$
- Under the md-targeting environment, the firm targets the consumers in the set $\Sigma^{*md} = \{(\alpha, \beta) : \alpha = 1 \text{ or } \beta = 1\}$ and the consumers who click are those in the set $\Omega^{*sd} = \{\omega : \omega \in \{0, 1\}\}.$
- It holds that $\Pi^{*sd} > \Pi^{*md}$.

As we can see in Proposition 8, for intermediate values of the parameters, in equilibrium the firm targets fewer consumers under the *sd*-environment $(|\Sigma^{*sd}| < |\Sigma^{*md}|)$ and the click-through rate is higher $(|\Omega^{*sd}| > |\Omega^{*md}|)$, which leads to a higher profit for the *sd*-environment. This is consistent with the intuition behind Proposition 3 of the main model.

Asymmetric Dimensions. Next, we consider an alternative version of the discrete-values model where the distributions of the dimensions α and β can be different from each other. More specifically, we assume that β takes the values 0, 1/2, and 1 with probabilities q, 1 - 2q, and q, respectively, where q is a parameter in [0, 1/2]. If q = 1/3, then the distribution of β is uniform (same as α). If q = 0, then β is always equal to 1/2, therefore there is no uncertainty about β for consumers. If q = 1/2, β can be either 0 or 1 with probability 1/2 each.

We want to explore how the parameter q (and as a result the variance of the distribution of β) can affect the result of Proposition 8. As shown in the following proposition, for any value of $q \in [0, 1/2]$, there are conditions under which we get a result similar to Proposition 8.

Proposition 9 (Asymmetric dimensions). If $\left(\max\left\{\frac{8q^2+3q-3}{8q^2+4q-4}, \frac{1}{2}\right\} or <math>\left(\frac{3}{4} , then$

the equilibria in the two targeting environments are as specified in Proposition 8 and it holds that $\Pi^{*sd} > \Pi^{*md}$. As q increases in the interval (0, 1/3], the parameter region under which this result holds becomes larger.

In general, we observe that for $q \leq 1/3$, the lower the value of q is, the more restrictive the conditions of the result are (more specifically for c_{click}). This is expected, because as q decreases below 1/3, the prior of the β distribution becomes more precise, therefore the uncertainty of consumers about their utility is reduced. Reducing consumer uncertainty can help the firm alleviate its commitment problem and make md-targeting better. This is consistent with the insights described in Section 4 about the main model. Notice also that in the extreme case where q = 0, md-targeting becomes equivalent to sd-targeting, and therefore the two profits become equal.