Targeted Advertising as an Implicit Recommendation and Personal Data Opt-Out

Z. Eddie Ning, Jiwoong Shin, Jungju Yu*

October 5, 2022

Abstract

Advances in data collection and algorithms help advertisers to better target individual consumers by predicting each consumer's preferences. We first show that when consumers have uncertainties about their preferences, an ad targeted to a consumer carries an implicit message: the algorithm predicts that the product fits her preferences. This implicit recommendation influences the consumer's purchase decision but also introduces misaligned incentives. As the accuracy improves, consumer inference from targeted ads becomes stronger, but so does the advertiser's incentive to exploit it to affect the consumer's decision. Under exogenous price, when individual-level prediction becomes more accurate, the advertiser adopts a less targeted advertising strategy due to its enhanced incentive to exploit a stronger recommendation effect. Even if the firm's prediction is perfectly accurate, consumers still receive ads for "bad products" and make incorrect purchase decisions. Despite these negative consequences, the consumer surplus can remain positive because the firm can better identify consumers with a good fit for the product. In contrast, under endogenous price, a better prediction allows the advertiser to raise its price instead of exploiting its recommendation effect. Thus, it leads to more targeted advertising and lower consumer welfare, which may incentivize consumers to opt-out of data collection.

Keywords: AI, algorithms, targeted advertising, recommendation, persuasion, pricing, privacy choice, personal data opt-out

*Z. Eddie Ning is an Assistant Professor of Marketing at the University of British Columbia. Jiwoong Shin is a Professor of Marketing at the Yale University. Jungju Yu is an Assistant Professor of Marketing at the Korea Advanced Institute of Science and Technology. Emails: eddie.ning@sauder.ubc.ca, jiwoong.shin@yale.edu, jungju.yu@kaist.ac.kr.

1 Introduction

With the exploding amount of online data collected, predictive technologies like artificial intelligence (AI) empower advertisers to target consumers at an increasingly granular level. Advertisers today have information on individual consumers' online behaviors, such as browsing patterns or social media likes and posts, as well as offline behaviors, such as store visits and purchases collected through smart devices. Such individual-level data allows advertisers to predict which product is fit for which consumers and target them accordingly. Agrawal et al. (2018) frame the role of AI in business as a prediction machine that aids decision-making under uncertainty. Therefore, the advancement of AI enables more accurate predictions at cheaper costs. We can expect that, as technology improvement and data collection continue, algorithms used by advertisers or platforms will be able to predict the preferences of each consumer at a higher accuracy. This is the core promise for the bright future of the digital advertising industry. However, the reality is not necessarily playing out so well.

If consumers have uncertainty regarding their preferences for a product, the fact that they are targeted by a firm may convey some information. For example, a consumer may receive targeted ads about a new paid app on her phone. Without knowing all the features of the new app, she is uncertain about her utility from the app. However, she may be aware that the firm has information about the fit between her and the app, e.g., whether she downloaded other apps designed for a similar consumer segment or whether other users like her enjoyed the new app. In such a case, upon being targeted, the consumer may infer that the algorithm predicts that the new app could be is a good match for her, subsequently raising her willingness to pay for the product. Such consumer behavior has been observed in laboratory studies (Summers et al., 2016). In their studies, upon receiving the same ad, consumers who were told that the ad was targeted based on their browsing history showed higher purchase intentions than consumers who were told otherwise. Thus, the action of targeting can be persuasive beyond the advertising print itself.

It is not obvious, however, how a rational consumer should make such inferences. In reality, the consumer cannot observe the advertiser's prediction about her, nor can she observe the degree to which the ad she receives is targeted. The advertiser could have sent advertising to a broad demographic group of which she is a member, in which case being targeted reveals little information to the consumer. In contrast, the advertiser could have targeted a specific segment of consumers with behavioral patterns that best match the app's intended users, in which case being targeted is much more informative. Furthermore, suppose consumers make favorable inferences from being targeted and become more willing to buy. In that case, the advertiser may have incentives to target even consumers that are not predicted to have a good fit for the app and charge a higher price. Then, the consumer may also be concerned that she may face unnecessarily higher prices as the advertiser's algorithms become more accurate. This line of logic may prompt her to seek a way to protect their privacy and disable ad targeting, eliminating her need to make inferences from being targeted. In practice, recent regulations such as GDPR and CCPA allow consumers to choose whether to opt-in or out of such personal data collection, which advertisers need to target, empowering consumers to control their own data. As such privacy laws go into effect, consumers' decision to opt-out may discipline the advertiser's incentives, ultimately influencing the advertiser's targeting strategy and consumer welfare.

In this paper, we study the optimal targeting strategy of an advertiser with an imperfect prediction algorithm when consumers make inferences from being targeted. In particular, we focus on the implications of such targeting strategies on consumer welfare and privacy choices, which, in turn, affect the firm's targeting strategy. How do advertisers use this algorithmic prediction in their targeting and pricing strategy? As advertisers can make individual-level predictions more accurately, are consumers more likely to get ads for products that better match their underlying preferences? What happens in the hypothetical limit as AI becomes omniscient? Can powerful AI predictions eliminate imprecise targeting and post-purchase regret? When do consumers have incentives to opt-out of data collection to curb firms' ability to predict their preferences?

To address these questions, we study a model between one advertiser and consumers with two-sided private information. The advertiser first receives a noisy signal about each consumer's match with the product and decides whether to target that consumer. The consumer, upon being targeted and seeing the ad, observes the price and realizes a private impression about her match with the product which is also noisy. The consumer makes an inference on her match with the product and decides whether or not to buy. We allow both the advertiser's targeting strategy and the consumer's buying strategy to be mixed.

We first show that when a consumer has uncertainty about her preferences for an unknown product, an ad targeting that consumer carries an implicit message: the advertiser's algorithm predicts that the product matches her preferences. Thus, when an advertiser has individual-level data, a targeted ad can act as an implicit recommendation which influences a consumer's purchase decision. However, it also creates misaligned incentives; An advertiser may cheat by sending ads to a consumer even though the advertiser does not believe that the product is a good fit for that consumer. A higher prediction accuracy strengthens the recommendation role of a targeted ad but also worsens this incentives problem.

Under exogenous price, due to the misaligned incentives, a higher prediction accuracy often pushes the advertiser sometimes to send an ad to consumers with a bad fit, leading to a less targeted advertising strategy. This effect also becomes more prominent as the algorithm's prediction accuracy increases. As a result, consumers receive ads for unfit products and make incorrect purchase decisions even if the advertiser's AI can make perfectly accurate predictions. However, the consumer surplus can be positive despite the misaligned incentives because the firm can better identify consumers with a good fit for the product. These results illustrate the limit of prediction technology in reducing market friction. Mistargeting is not only an outcome of firms' technological capabilities but can also be an outcome of the advertiser's economic incentives.

Interestingly, under endogenous price, the results are reversed. The advertiser takes advantage of a higher prediction accuracy, which strengthens the recommendation effect of a targeted ad by raising the price instead of diluting its recommendation effect by sending ads to consumers not fit for the product. Consequently, consumer welfare decreases to zero as the prediction accuracy becomes perfectly accurate. Therefore, if consumers have a choice of whether to opt-out of data collection or ad targeting, they have incentives to do so if the advertiser's prediction accuracy is sufficiently high while the accuracy of the consumer's own impression from the ad is sufficiently low. It contrasts with the exogenous price case where consumers have no incentive to opt-out. This is because under exogenous pricing opting out of data collection may strip away a consumer's opportunity to buy an appropriate product, and consumers do not have to worry about being charged a higher price in case their predicted match for the product is good.

Next, we discuss the related literature. Section 3 presents the model. Section 4 studies the equilibrium targeting strategy and the effect of prediction accuracy under the exogenous price. Section 5 extends the model to the advertiser's optimal pricing decision. Section 6 concludes.

2 Literature Review

Our paper relates to several streams of research in online advertising targeting, recommendation system, and consumer privacy choice. First, the literature on online advertising has emphasized the importance of targetability, which refers to firms' ability to identify individual consumer preferences using various customer data, such as demographics, browsing behaviors or past purchases (e.g., Agarwal et al., 2009; Chen et al., 2009; Joshi et al., 2011; Shen and Villas-Boas, 2018). Such targeted advertising message based on customer characteristics improves the performance of communications and consumer response (Bleier and Eisenbeiss, 2015; Goldfarb and Tucker, 2011; Rafieian and Yoganarasimhan, 2021). While a stream of research focuses on the role of advertising content in persuasion (Anderson and Renault, 2006; Chakraborty and Harbaugh, 2010; Kamenica and Gentzkow, 2011; Mayzlin and Shin, 2011; Shin, 2005), several papers show an additional effect of targeting beyond the information contained in the advertising content. Iyer et al. (2005) study the impact of targeted advertising in competitive settings and find that firms can be better off with targeted advertising compared to no targeting because of differentiation. Anand and Shachar (2009) studies a model where competing firms target advertisements to particular consumers, and advertising is noisy. It shows that getting targeted can serve as a signal on product attributes. Thus, targeted ads can increase advertising effectiveness beyond the information contained directly in the advertising message. In comparison, our paper analyzes the effect of prediction accuracy on targeting strategy, pricing strategy, and consumers' privacy choices, which are not studied in Anand and Shachar (2009). Bergemann and Bonatti (2011) consider targeted advertising under imperfect targetability, but focus on the impact of targetability on the advertising market.

The recent paper by Shin and Yu (2021) considers a setup similar to ours where firms decide whether to target a consumer based on noisy information on the consumer's match to the product category. They provide the micro-mechanism of the consumer inference process from merely receiving a targeted ad, which influences the consumer's decision on whether to search for a competitor, creating spillover effects of one firm's advertising for competitors. In this setting, they focus on trade-offs between targeting accuracy and advertising intensity assuming the firm's targeting accuracy is exogenously given. In contrast, the current paper focuses on the credibility of targeted advertising incorporating the advertiser's strategic mistargeting incentives and its implications on the consumer's data privacy choice, which, in turn, affects targeting accuracy.[1] Thus, in our model, the firm's targeting ability is the endogenous outcome of the consumers' strategic decisions.

The paper also relates to the literature on platform design such as optimal recommender systems and content personalization. The marketing literature largely focuses on AI-based features and unstructured data extraction for improving the recommendation systems (Dzyabura and Hauser, 2019) or the effectiveness of data-driven content personalization (Yoganarasimhan, 2020). Recent studies turn to the platform's search design issues. Dukes and Liu (2015) investigates the platform's optimal search design when it can influence the consumers' search costs which affect their search behaviors for both between and within products. For the optimal search design, Jullien and Pavan (2019) highlights the broad trade-off between match quality and a firm's price, and the higher search precision has a non-monotonic effect on the firm's prices. Armstrong and Zhou (2022) further focuses on the platform's intermediary role that can control the information environment that consumers face. Then, it asks how much product information it should reveal and characterize the optimal information structure for the consumers and sellers. Also, Zhong (2022) investigate the platform's search design and its endogenous effects on the platform revenue jointly. Built on this strand of research on platform design and recommendation systems, we investigate the credibility of such recommendation systems and their implications on consumers' privacy choice decisions.

Finally, several papers have studied the consumer's privacy decision to opt-out of data collection (Acquisti and Varian 2005; De Corniere and De Nijs 2016; Montes et al. 2019). The literature largely focuses on the economic trade-off between the benefit and cost of disclosing customer information. The benefit is that consumers can receive a more relevant product recommendation or advertisement. However, at the same time, the firm can use this information for price discrimination. In these papers, consumers opt out of data collection to avoid such price discrimination. Ichihashi (2020) shows that a firm's commitment to not using consumer information in its pricing decisions alleviates price-discrimination concerns. This encourages consumers to disclose information, which, in turn, improves the firm's product recommendation quality. We contribute to this literature by showing that even under uniform pricing or no price discrimination, consumers may have incentives to opt out of data collection when firms can predict individual preferences with sufficient accuracy.

[1] Li and Xu (2022) also study the credibility of a firm's personalized pricing where the firm knows consumers' values, which are unknown to consumers before their costly inspection. However, they do not consider the effect of targeting accuracy, which in our model can be an endogenous outcome of consumers' strategic privacy choices.

3 Model

We consider the online interaction between a firm that advertises a product and a unit mass of consumers.[2] Consumers receive a utility of either 1 or 0, depending on whether it matches a consumer's need. Let p denote the price of the product. To focus on the equilibrium advertising strategy, we first assume the price to be exogenous. We later relax this assumption by endogenizing the price. A consumer may or may not receive a targeted ad from the advertiser ($A \in \{a, \varnothing\}$). Upon receiving the ad, a consumer decides whether or not to purchase the product. A consumer receives a utility $1 - p$ from purchasing if the product is a good match and a utility of $-p$ if the product is a bad match. A consumer receives a utility of 0 if she does not buy.

There are two types of consumers, $t \in \{G, B\}$. A μ fraction of consumers is good-types ($t = G$) and a $1 - \mu$ fraction is bad-types ($t = B$). The probability that the product is a good match for a consumer $\mu = \Pr(t = G)$ is common knowledge for both the firm and the consumer. A high μ can be interpreted as the product having mass market appeal, while a small μ represents a more niche product. The advertiser has data on the consumer i's online behaviors to predict t_i, the true match between the product and consumer i. The advertiser's algorithm produces a signal about t_i, denoted as $s \in \{s_G, s_B\}$. This private signal is accurate with probability $\alpha \geq \frac{1}{2}$. A higher α implies that the advertiser has more data on a consumer or a more advanced algorithm to predict t_i more accurately. Given the prediction algorithm, the advertiser decides whether to target a consumer. The cost of sending an ad is fixed and denoted as k, which is assumed to be smaller than p.

If a consumer is targeted, she receives the ad. She can click on the ad, and conduct a further search for product information. By doing so, a consumer forms her own impression of the product fit, which is a private signal, $m \in \{m_g, m_b\}$. The impression is accurate with probability $\beta \geq \frac{1}{2}$. A higher β is interpreted as there being more available information on the product so that the consumer's own impression of fit is more accurate. For simplicity, we assume that the advertiser's prediction and the consumer's impression are independent. A consumer then decides whether to buy the product. We assume that a consumer cannot buy the product if she does not receive an ad due to the lack of awareness.

We allow for the possibility of a mixed strategy. Let $\sigma(s|\mu, \alpha, \beta) \in [0, 1]$ denotes the probability

[2] Hereafter, we interchangeably use the term "firm" and "advertiser."

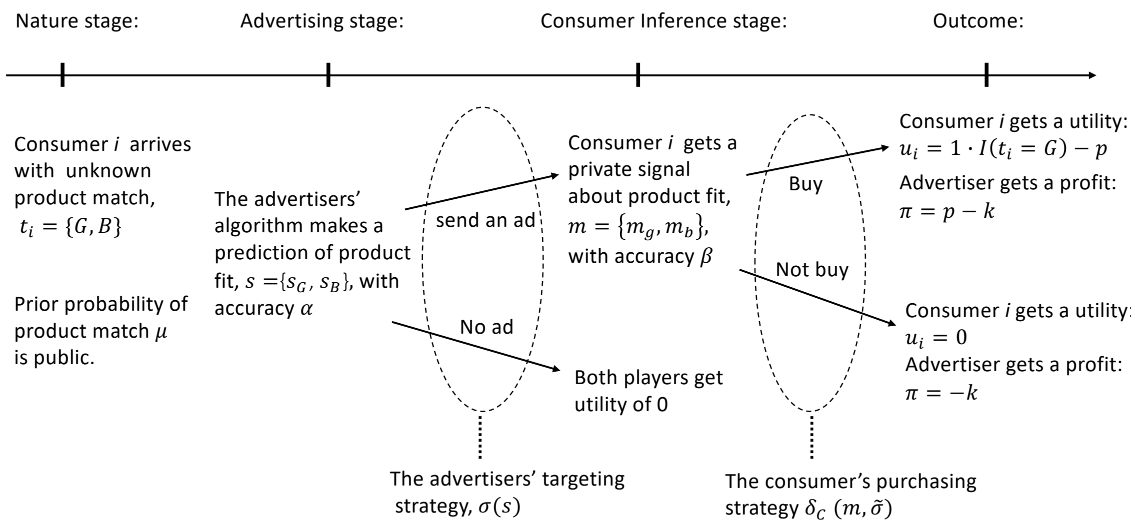

Figure 1: Timeline of the game sequence

that the advertiser targets the consumer, as a function of the algorithm's prediction of the consumer's product fit s. The consumer's strategy is to choose whether to purchase a product after observing the firm's targeted ad. When a consumer does not receive an ad ($A = \varnothing$), she remains unaware of the product. Thus, she does not participate in the game anymore. In contrast, when a consumer receives an ad ($A = a$), she becomes aware of the product and decides whether to purchase it or not. Importantly, consumers do not directly observe the firm's advertising strategy $\sigma(\cdot)$ and instead have rational beliefs about it. Therefore, the consumer's purchase decision depends on her belief about the firm's advertising strategy, denoted by $\widetilde{\sigma}$, not on the true σ.

Let $\delta_c(m, \widetilde{\sigma}|\mu, \alpha, \beta) \in [0, 1]$ denotes the probability that a consumer buys the product upon being targeted, as a function of her private impression of product fit m and her belief of the advertiser's targeting strategy $\widetilde{\sigma}$. Because μ, α, and β are parameters of the model, we will abbreviate $\sigma(s|\mu, \alpha, \beta)$ to $\sigma(s)$ and abbreviate $\delta_c(m, \widetilde{\sigma}|\mu, \alpha, \beta)$ to $\delta_c(m, \widetilde{\sigma})$ for the remaining of the paper. Figure 1 provides a graphical illustration of the game sequence.

Note that if the advertiser's targeting strategy is non-targeted such that $\sigma(s_G) = \sigma(s_B)$, it does not depend on its algorithmic prediction of the consumer's product fit. Only when $\sigma(s_G) \neq \sigma(s_B)$ does the advertiser's targeting strategy varies for a consumer type depending on its prediction of the consumer preference. We distinguish the two types of targeting strategies as follows:

Definition 1. *The advertiser's strategy, $\sigma(s)$, is* ***individually targeted*** *if $|\sigma(s_G) - \sigma(s_B)| \neq 0$. We say the advertisers' strategy is* ***more individually targeted*** *if $|\sigma(s_G) - \sigma(s_B)|$ is higher.*

We solve for the Perfect Bayesian Equilibrium of this game. On the off-equilibrium path node, which is the consumer's buying decision when the advertiser never sends the ad, we assume that consumers believe that the advertiser's targeting strategy does not depend on s.[3] We focus on equilibria with a positive ex-ante probability of transaction.

Consumer inference

Upon receiving an ad ($A = a$), a consumer updates her posterior belief about product fit given her own signal (m) and her belief about the advertiser's targeting strategy ($\widetilde{\sigma}$). We denote her posterior belief as $\widetilde{\mu}(\widetilde{\sigma}, m) = \Pr(t_i = G|\widetilde{\sigma}, m)$. The firm's advertising strategy is a function of its private signal about the consumer's product match, s, which a consumer does not observe. Hence, a consumer infers the firm's signal s based on the anticipated advertising strategy $\widetilde{\sigma}(s)$ for updating her belief about her product match t_i.

To derive a consumer's posterior belief, we first analyze how a consumer updates her belief about the product match based on her private impression m. With abuse of notation, we denote $\mu_c(m)$ for $m \in \{m_g, m_b\}$ as such updated belief, which we call the consumer's private prior after receiving her own private signal m:

$$\mu_c(m_g) = \frac{\beta\mu}{\beta\mu + (1-\beta)(1-\mu)}, \qquad \mu_c(m_b) = \frac{(1-\beta)\mu}{(1-\beta)\mu + \beta(1-\mu)}. \tag{1}$$

Next, the mere fact a consumer receives an ad ($A = a$) can influence the consumer's belief. When she receives an ad ($A = a$), she considers the following two possibilities: (1) the firm received a signal s_G with accuracy α (the probability of this event is $\mu_c\alpha + (1-\mu_c)(1-\alpha)$) and sends an ad following its strategy $\widetilde{\sigma}(s_G)$ or (2) the firm received a signal s_B with accuracy α (the probability

[3] We assume that the level of advertising is not observed by consumers. In any equilibrium where $0 < \sigma < 1$, both receiving and not receiving an advertisement are consistent with the firm's equilibrium advertising strategy. For example, if a consumer believes that the firm sent out an advertisement to 10% of the time for s_G, the fact that she did or did not receive an advertisement does not change her equilibrium belief. On the other hand, in an equilibrium in which $\sigma = 0$ for both $s = \{g, b\}$, consumers expect a zero probability of receiving an advertisement irrespective of the firm's signal s. Hence, Bayesian updating does not provide any guidance for pinning down consumers' off-equilibrium beliefs when consumers actually observe an advertisement in this case. Then, we impose the above off-equilibrium that the advertiser's targeting strategy does not depend on s.

of this event is $\mu_c(1-\alpha)+(1-\mu_c)\alpha)$, but the firm still sends an ad with probability $\widetilde{\sigma}(s_B)$. From the anticipated advertiser's strategy $\widetilde{\sigma}(s) = \big(\widetilde{\sigma}(s_G), \widetilde{\sigma}(s_B)\big)$, a consumer infer probabilities that the true product fit is $t_i = G$ or $t_i = B$, incorporating her own private prior $\mu_c = \mu_c(m)$, which is a function of her private signal $m \in \{m_g, m_b\}$.

Then, the consumer's posterior belief after observing an ad $(A = a)$ is:

$$\widetilde{\mu}(\widetilde{\sigma}, m) = \Pr(t_i = G|\widetilde{\sigma}, m) = \frac{\widetilde{\sigma}(s_G)\cdot\mu_c\cdot\alpha + \widetilde{\sigma}(s_B)\cdot\mu_c\cdot(1-\alpha)}{\widetilde{\sigma}(s_G)\big(\mu_c\alpha+(1-\mu_c)(1-\alpha)\big)+\widetilde{\sigma}(s_B)\big(\mu_c(1-\alpha)+(1-\mu_c)\alpha\big)}, \quad (2)$$

where $\mu_c = \mu_c(m)$ is from Equation (1) depending on consumer's private signal m.

As we can see from equation (2), the consumer's posterior belief about the product match $\widetilde{\mu}(\widetilde{\sigma}, m)$ depends on (i) her private signal through $\mu_c(m)$, (ii) anticipated firms' advertising strategy $\widetilde{\sigma}(s)$, and (iii) information accuracy of the firm and consumer α and β. The next lemma characterizes the consumer's belief updating process about the product match from a targeted ad.

Lemma 1. *For any $m \in \{m_g, m_b\}$, $\widetilde{\mu}(\widetilde{\sigma}, m) \geq \mu_c(m)$ if and only if $\widetilde{\sigma}(s_G) \geq \widetilde{\sigma}(s_B)$. The marginal change in the posterior beliefs is increasing in the firm's information accuracy α: $\frac{\partial\,[\widetilde{\mu}-\mu_c]}{\partial\,\alpha} \geq 0$.*

The lemma suggests that a consumer updates her beliefs about her product match more positively when the advertising strategy is individually targeted (i.e., $\widetilde{\sigma}(s_G) \geq \widetilde{\sigma}(s_B)$). Thus, the act of individual targeting can be *persuasive*. In contrast, suppose a consumer believes that the advertiser's strategy is not individually targeted, i.e., $|\widetilde{\sigma}(s_G) - \widetilde{\sigma}(s_B)| = 0$. Then, the consumer's posterior belief only depends on her own signal, $\widetilde{\mu}(\widetilde{\sigma}, m) = \mu_c(m)$; thus, the consumer's willingness-to-pay depends on her signal of product fit and will be either $\widetilde{\mu}(\widetilde{\sigma}, m_g) = \mu_c(m_g)$ or $\widetilde{\mu}(\widetilde{\sigma}, m_b) = \mu_c(m_b)$.

Persuadable consumer

If the advertiser's strategy is individually targeted, a targeted ad also performs the function of a recommendation to consumers: "you are only receiving this ad because we believe this product is a good fit for you." In equilibrium, consumers have to incorporate this information when making a purchase decision as shown in Equation (2).

Such an implicit recommendation may or may not influence consumers' buying decisions depending on the product price. For instance, suppose the price is significantly high such that

$p > \widetilde{\mu}(\sigma, m_g)$. Even under the case that $\widetilde{\sigma} = (\widetilde{\sigma}(s_G) = 1, \widetilde{\sigma}(s_B) = 0)$, which achieves the most favorable posterior belief, a consumer never purchases a product even after receiving an ad. Consequently, the advertiser never targets consumers. Similarly, if $p \leq \mu_c(m_b)$, then a consumer would purchase the product as long as she receives an ad regardless of whether advertising is targeted or not. We characterize the conditions where the targeted ad can persuade consumers' decisions. We define the upper limit of posterior belief that any individually targeted advertising can achieve $\widetilde{\mu}^U = \widetilde{\mu}(\widetilde{\sigma}, m \mid \sigma(s_G) = 1, \sigma(s_B) = 0)$.

Now, consider the case where the price is such that $\mu_c(m_b) < p \leq \widetilde{\mu}^U$. When a consumer receives a bad impression of product fit m_b, her belief about the product fit $\mu_c(m_b)$ is lower than the price. If she believes that the ad is not individually-targeted, $\widetilde{\sigma}(s_G) = \widetilde{\sigma}(s_B)$, then her final posterior belief of product fit is still $\widetilde{\mu}(\widetilde{\sigma}, m_b) = \mu_c(m_b)$. So, she does not buy upon receiving the ad. On the other hand, if she believes that the ad is individually targeted (i.e., $\widetilde{\sigma}(s_G) > \widetilde{\sigma}(s_B)$), then there always exists advertising strategy $\sigma = (\sigma(s_G), \sigma(s_B))$ whose $|\sigma(s_G) - \sigma(s_B)|$ is sufficiently large that consumer's posterior belief becomes $p \leq \widetilde{\mu}(\widetilde{\sigma}, m_b) \leq \mu^U$. Therefore, the consumer would consider buying upon receiving an individually targeted ad even if her private signal is m_b. Similarly, if $\mu_c(m_g) < p \leq \widetilde{\mu}^U$, then the advertiser's targeting strategy can persuade consumers to buy when a consumer receives a good signal of product fit, m_g.

Consumers are **persuadable** if, under some $m \in \{m_g, m_b\}$, consumers would not buy the product if they believe the ad is non-targeted mass-marketing but would buy the product if they believe the ad is sufficiently individually-targeted.

Lemma 2. *For any price $p \in \left(\mu_c(m_b),\ \widetilde{\mu}^U\right]$, there exists a threshold ξ^* such that for all advertising strategies which are more individually targeted $|\widetilde{\sigma}(s_G) - \widetilde{\sigma}(s_B)| \geq \xi^*$, consumers are persuadable.*

By persuadable consumers, we refer to the *price range* given the model primitives (such as μ, α, β) in which consumers can be persuaded. In our analysis of the exogenous price, we focus on the case where the price is such that consumers are persuadable, and thus, targeted advertising can influence consumers' behaviors. For the endogenous price case, we consider all the price ranges where consumers can be either persuadable or non-persuadable.

4 Analysis

4.1 Benchmark: A consumer without informative private signal

We first analyze a simple benchmark case without the consumer's private signal. The consumer's private signal of product fit is pure noise: $\beta = 1/2$. We show the firm still engages in individually targeted advertising, but it doesn't advertise to *only* the good-type consumers. Again, the price is exogenously given as $p \in \left(\mu,\ \widetilde{\mu}^U\right]$, so that consumers are persuadable.

Given $\beta = 1/2$, consumers cannot make an informed assessment of product fit from their private signals. Thus, we have $\mu_c(m) = \mu$ for all $m \in \{m_g, m_b\}$. A consumer can only update her posterior belief through her belief about the advertiser's targeting strategy. The consumer's posterior is different from her prior μ if and only if $\widetilde{\sigma}(s_G) \neq \widetilde{\sigma}(s_B)$. Also, note that a targeted consumer with $A = a$ has the same posterior belief regardless of their true types $t_i \in \{g, b\}$ since she does not directly observe the firm's signal $s \in \{s_G, s_B\}$. Hence, all the targeted consumers behave the same way and make a purchase with probability δ_c.

Suppose a consumer believes that the advertiser targets her only when s_G. Then, a consumer sees the ad as a recommendation, and her optimal response is always to buy when she receives the ad because her posterior belief $\widetilde{\mu}(\widetilde{\sigma}_A) = \widetilde{\mu}^U$, which is greater than price. However, if a consumer's response is to always purchase after receiving the ad, the advertiser prefers to take advantage of the implicit recommendation by targeting consumers even if its signal is s_B. By doing so, the advertiser effectively cheats on its message by sending a perverse recommendation. Thus, in equilibrium, we cannot have a pure-strategy separating equilibrium where $\sigma(s_G) = 1$ and $\sigma(s_B) = 0$.[4]

Also, there cannot be an equilibrium in which the adverting is not individually-targeted and mass-marketed to everyone, i.e., $\sigma(s_G) = \sigma(s_B)$. Because a non-targeted ad cannot serve as a recommendation, the consumer's posterior belief about product fit is the same as her prior $\widetilde{\mu}(\sigma) = \mu$, which is lower than the price p. Thus, there is no pure-strategy pooling equilibrium in which there is an ex-ante non-zero probability of transaction.[5]

We now consider a mixed strategy, where the advertiser targets the consumer with probability $\sigma(s) \in [0, 1]$, and a consumer purchases the product with the probability with $\delta_c(m, \widetilde{\sigma}) \in [0, 1]$.

[4] For the same reason, there is no equilibrium in which $\sigma(s_G) > 0$ and $\sigma(s_B) = 0$.

[5] It is straightforward to check that we cannot have $\sigma(s_G) = 0$ and $\sigma(s_B) = 1$ in an equilibrium.

Because $\beta = 1/2$, δ_c now only depends on the anticipated advertiser's targeting strategy, $\widetilde{\sigma}$.

For consumers to mix, they must be indifferent between purchasing and not purchasing after receiving an ad, which depends on their posterior beliefs. From equation (2), the posterior belief is

$$\widetilde{\mu}(\widetilde{\sigma}) = \frac{\widetilde{\sigma}(s_G) \cdot \mu\alpha + \widetilde{\sigma}(s_B) \cdot \mu(1-\alpha)}{\widetilde{\sigma}(s_G)\big(\mu\alpha + (1-\mu)(1-\alpha)\big) + \widetilde{\sigma}(s_B)\big(\mu(1-\alpha) + (1-\mu)\alpha\big)} \tag{3}$$

The consumer's expected utility from purchasing is $\mathbb{EU}(\text{purchasing}) = \widetilde{\mu}(\widetilde{\sigma})(1-p) + (1-\widetilde{\mu}(\widetilde{\sigma}))(-p)$, and no purchasing is $\mathbb{EU}(\text{no purchasing}) = 0$. In equilibrium, they must be the same so that the consumer is indifferent, which leads to $\widetilde{\mu}(\widetilde{\sigma}) = p$, which pins down the firm's mixed strategy.

Also, consumers mix in such a way that the firm is indifferent between sending an ad and no ad. The expected payoff from sending an ad $\mathbb{E}\Pi(A = a|s) = \delta_c(p-k) + (1-\delta_c)(-k)$ must be the same as not sending an ad $\mathbb{E}\Pi(A = \varnothing|s) = 0$. Therefore, we have that $\delta_c = k/p$.

Proposition 1. *Suppose consumers do not have an informative private signal about the product fit. Then, there is no equilibrium where the firm can obtain a positive expected profit. The set of equilibria* $\sigma^* = \big(\sigma^*(s_G), \sigma^*(s_B)\big)$*, and* δ_c^* *are characterized as follows. (i) ads are individually targeted such that* $|\sigma^*(s_G) - \sigma^*(s_B)| > 0$*, (ii) the advertising strategy* σ^* *satisfies* $\sigma^*(s_G) = \phi \cdot \sigma^*(s_B)$*, where* $\phi = \left(\frac{\alpha \cdot (1-\mu)p - (1-\alpha) \cdot \mu(1-p)}{\alpha \cdot \mu(1-p) - (1-\alpha) \cdot (1-\mu)p}\right) > 1$ *for persuadable consumers, and (iii)* $\delta_c^* = k/p$*.*

First, without a customer's informative private information, the firm's expected profit is zero. It is straightforward from the firm's indifference condition. Second, even under the mixed equilibria, the advertiser's strategy is individually targeted $|\sigma^*(s_G) - \sigma^*(s_B)| > 0$. In particular, $\sigma^*(s_G) = \phi \cdot \sigma^*(s_B)$, where $\phi = \phi(\mu, \alpha, p) > 1$, which implies that the firm sends an ad to a consumer that firm perceives as a good type s_G with a higher probability (or intensity) than to a bad type s_B.

Suppose the ads are not individually targeted in equilibrium, which is the case when $\phi = 1$. Given our focus on the case of persuadable consumers, consumers' posterior will be the same as the prior, and $\widetilde{\mu}(\widetilde{\sigma}) = \mu < p$ and thus, the ads can no longer be persuasive. That is, no consumer will be persuaded to make a purchase. Therefore, even though the firm is facing consumers without heterogeneous private information and thus, making the same purchasing decisions, it is disciplined to target its ad, treating consumers differently.

In this mixed strategy equilibrium, the posterior belief must be equal to price ($\widetilde{\mu}(\widetilde{\sigma}) = p$) from consumers' indifference condition that $\mathbb{E}\Pi(a) = \mathbb{E}\Pi(\varnothing)$. Then, as the advertiser's prediction accuracy improves, the role of recommendation from an individually targeted ad can be more prominent. Under a higher prediction accuracy, the same advertising strategy can be more persuasive, as it can move consumers' posterior beliefs further up. However, this introduces an incentive for the advertiser to cheat by targeting consumers with s_B until the posterior belief decreases back to p. The following proposition formally states the relationship between prediction accuracy and the advertiser's equilibrium targeting strategy.

Proposition 2. *For persuadable consumers, as prediction accuracy increases, the advertiser's strategy becomes less individually targeted. That is,* $|\phi| = |\frac{\sigma(s_G)}{\sigma(s_B)}| > 1$ *strictly decreases in* α*.*

When the advertiser's prediction of an individual's preferences becomes more accurate, we should expect the advertiser's targeting strategy to be more focused on the "right" consumers, all else being equal. However, all else is not equal. A targeted ad carries an implicit message of recommendation that affects the consumer's evaluation of the product. A better prediction accuracy strengthens the power of such implicit recommendations, creating the incentive for the advertiser to intentionally dilute the message to capture a bigger market. Thus, ironically the improvement in prediction can lead to less individual-targeted advertising.

We also consider the limit case as the firm's prediction accuracy α approaches 1. Even under this extreme case where AI is so powerful to predict preferences perfectly, the following corollary shows that the advertiser still targets "wrong" consumers for whom the product is a bad fit, and consumers sometimes must make the incorrect purchase resulting in post-purchase regret.

Corollary 1. *In the limit as* $\alpha \to 1$*, (i) the advertiser knowingly targets bad-type consumers with a positive probability:* $\sigma^*(s_B) > 0$*. (ii) Given that good-type consumers are targeted, they do not buy the product with a probability* $(p - k)/p$*, whereas when bad-type consumers are targeted, they buy with a probability* k/p*.*

The corollary highlights the fact that even under perfect prediction technology, the firm's strategic mistargeting will persist. Thus, prediction technology alone cannot eliminate the problem of matching between the consumer and the product.

4.2 A consumer with the informative private signal

The benchmark assumes consumers do not have any informative private signal. The situation changes when we relax this strong form of information asymmetry and allows consumers to have an informative private signal (i.e., $\beta > 1/2$). Nevertheless, as we show, the main findings that a better prediction accuracy worsens incentive misalignment and leads to a less individual-targeted advertising strategy are robust.

Unlike the benchmark case, the firm has a weakly greater incentive to send an ad to a consumer of type s_G than to another consumer of type s_B. Thus, the firm will not send an ad to any consumer of the latter type (s_B) unless the entire population of the former type (s_G) is already covered by the firm's advertising. This also implies that there cannot exist a totally mixed equilibrium where both (1) $0 < \sigma^*(s_B) < 1$ and (2) $0 < \sigma^*(s_G) < 1$ hold simultaneously.

Lemma 3. *Suppose consumers are persuadable. If $\beta > 1/2$, a totally mixed equilibrium with $0 < \sigma^*(s_B) \leq \sigma^*(s_G) < 1$ does not exist. Moreover, if $\sigma^*(s_B) > 0$, then it must be $\sigma^*_A(s_G) = 1$.*

The signals the firm and consumers individually observe are positively correlated through the consumer's true type $t \in \{G, B\}$. Therefore, all else equal, the firm has a greater incentive to send an ad to consumers with s_G than s_B. Likewise, consumers with a good impression m_g are more likely to have a good match than consumers with m_b. Hence, $\delta^*_c(m_g) \geq \delta^*_c(m_b)$. So, in equilibrium, the firm will exhaust the entire segment of consumers with s_G before it advertises to anyone in the other segment s_B. Hence, if $0 < \sigma^*(s_B) < 1$, it must be $\sigma^*(s_G) = 1$. Moreover, if $0 < \sigma^*(s_G) < 1$, it must be $0 = \sigma^*(s_B)$. So, $0 < \sigma^*(s_B) < 1$ and $0 < \sigma^*(s_G) < 1$ cannot hold at the same time.

Next, depending on the parameter values, there are two different types of pure-strategy equilibria with $\sigma^*(s_G) = 1$, which we analyze one by one: One, the pooling equilibria where the firm advertises to both types, and (2) the separating equilibria where the firm only advertises to the good type s_G.

Pure strategy equilibrium: separating and pooling

When consumers are persuadable ($\mu_c(m_b) < p \leq \widetilde{\mu}^U$), there can exist a pure-strategy *separating* equilibrium, where the firm only sends an ad only if a consumer is of a good type, $\sigma^s = (\sigma(s_G) = 1, \sigma(s_B) = 0)$. Here, the superscript s denotes ***separating*** equilibrium. Under the separating equilibrium, consumers believe that the firm only targets her because $s_i = s_G$. Thus, the ad serves as a clear recommendation. However, as we discussed in the benchmark case, if a consumer's optimal response is to always buy when she receives the ad irrespective of her own impression, the firm can cheat on its message by sending an ad to consumers with s_B. To prevent such a deviation, consumers' posterior belief after receiving a targeted ad is not too high to purchase: $\widetilde{\mu}(\sigma^s, m_b) = \frac{\mu(1-\beta)\alpha}{\mu(1-\beta)\alpha+(1-\mu)(1-\alpha)\beta} < p$. On the other hand, if a consumer has a good impression m_g, her posterior becomes high enough to purchase a product $p \leq \widetilde{\mu}(\sigma^s, m_g) = \widetilde{\mu}^U$. Therefore, there can exist a pure-strategy separating equilibrium where $\sigma^s = (\sigma^s(s_G) = 1, \sigma^s(s_B) = 0)$, and a consumer purchases a product only if her impression of product fit is good, $\delta_c^s(m_g) = 1$, $\delta_c^s(m_b) = 0$.

We also check that the firm indeed has incentives to target its ad to a good-type, but not to a bad-type consumer. The firm's incentive conditions for those with s_G and s_B are $\mathbb{E}\Pi(A = a|s_G) = p \cdot \Pr(m_g|s_G) - k > 0$ and $\mathbb{E}\Pi(A = a|s_B) = p \cdot \Pr(m_g|s_B) - k < 0$. These conditions correspond to $p \cdot \Pr(m_g|s_B) < k < p \cdot \Pr(m_g|s_G)$.[6] The consumer's extra information about the product fit disciplines the firm to be truthful.

Proposition 3. *Suppose consumers are persuadable:* $\mu_c(m_b) \leq p \leq \widetilde{\mu}(\sigma^s, m_g) = \widetilde{\mu}^U$. *If* $p \cdot \Pr(m_g|s_B) < k < p \cdot \Pr(m_g|s_G)$ *and* $p > \widetilde{\mu}(\sigma^s, m_b) = \frac{\mu(1-\beta)\alpha}{\mu(1-\beta)\alpha+(1-\mu)(1-\alpha)\beta}$, *there exists a separating equilibrium where* $\sigma^s = (\sigma^s(s_G) = 1, \sigma^s(s_B) = 0)$, *and* $\delta_c^s = (\delta_c^s(m_g) = 1, \delta_c^s(m_b) = 0)$.

Next, there can also exist another pure-strategy equilibrium, i.e. *pooling equilibrium*, where the firm's advertising strategy is not targeted, i.e., $\sigma^{\text{pool-}m_g} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$, and targeted consumers purchase only if their impression of the product fit is good, $\delta_c^{\text{pool-}m_g}(m_g) = 1$ and $\delta_c^{\text{pool-}m_g}(m_b) = 0$. Here, the superscript 'pool-m_g' denotes pooling equilibrium where only a consumer with a good impression m_g purchases (we call it ***pooling-g***). This implies that $\mu_c(m_b) < p < \widetilde{\mu}(\sigma^{\text{pool-}m_g}, g) = \mu_c(m_g)$ since the ads are not individually targeted: $\widetilde{\mu}(\sigma^{\text{pool-}m_g}, g) = \mu_c(m_g)$. Also, the firm should find it profitable to send an ad to even a bad-type consumer with s_B. That is, $\mathbb{E}\Pi(A = a|s_B) = p \cdot \Pr(m_g|s_B) - k > 0$, or equivalently, $k < p \cdot \Pr(m_g|s_B)$.

The condition is satisfied if the cost of advertising k is sufficiently low, or the advertiser's prediction α becomes sufficiently low: $\partial/\partial\,\alpha(\Pr(m_g|s_B)) = -\frac{\mu(1-\mu)(2\beta-1)}{(\mu(1-\alpha)+(1-\mu)\alpha)^2} < 0$ because $\beta > 1/2$. Therefore, there can be a pure strategy pooling equilibrium with $\sigma^{\text{pool-}m_g}$, $\delta_c(m_g) = 1$, and $\delta_c(m_b) = 0$ when the mass-marketing is cheap or the algorithm's prediction is too noisy.

[6]If $k > p \cdot \Pr(m_g|s_G)$, the firm's profit margin is so small that it is not profitable to even target the consumer when s_G.

On the other hand, if the advertising cost is too high ($k \geq p$), it is not worth advertising. Also, even when $k < p$ and it is worth targeting *only* consumer with s_G ($\widetilde{\mu}(\sigma^s, m_b) \leq p \leq \widetilde{\mu}^U$), if the advertising cost is just expensive enough that $p \cdot \Pr(m_g|s_G) \leq k$, it is not profitable to even target the consumer with s_G given consumers' purchasing decisions. Then, we get another type of pooling equilibrium in which the firm does not advertise at all $\sigma(s_G) = \sigma(s_B) = 0$ (we call it ***pooling-*$\varnothing$**).

Proposition 4. *Suppose consumers are persuadable: $\mu_c(m_b) < p \leq \widetilde{\mu}(\sigma^s, m_g) = \widetilde{\mu}^U$. There are two different types of pooling equilibria.*

1. *If $k < p \cdot \Pr(m_g|s_B)$ and $p < \mu_c(m_g)$, there exists a pooling equilibrium where $\sigma^{pool\text{-}m_g} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$ and $\delta_c^{pool\text{-}m_g} = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$. The parameter space in which this equilibrium exists shrinks as the advertiser's prediction accuracy α increases.*

2. *If $k \geq p$ or $p \cdot \Pr(m_g|s_G) \leq k < p$ and $p \geq \widetilde{\mu}(\sigma^s, m_b)$, there exists another pooling equilibrium where $\sigma^{pool-\varnothing} = (\sigma(s_G) = 0, \sigma(s_B) = 0)$.*

In the range of prices where the consumer is not persuadable (either $p > \widetilde{\mu}^U$ or $p < \mu_c(m_b)$), there also exist other pooling equilibria. If $p > \widetilde{\mu}^U$, the price is so expensive that no consumer can be persuaded to make a purchase, irrespective of the firm's targeting strategy. Therefore, we again have the pooling-$\varnothing$ equilibrium. If $p < \mu_c(m_b)$, a targeted consumer will make a purchase irrespective of her private impression m. Thus it is profitable for the firm to target all consumers, i.e., the firm's equilibrium strategy is $\sigma^{\text{pool-all}} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$. We refer to this as the ***pooling-all*** equilibrium.

Semi-separating equilibrium

When consumers are persuadable, there can also exist a "hybrid" type of PBE, where the advertiser mixes between sending an ad and no ad when s_B (i.e., $0 < \sigma^{\text{semi}}(s_B) < 1$) while the advertiser always sends an ad if s_G (i.e., $\sigma^{\text{semi}}(s_G) = 1$).[7] As a result, consumers imperfectly update their prior beliefs from the ad. Contrast this with a pooling equilibrium, where no updating is possible (the posterior

[7] In this paper, we do not analyze another possible semi-separating equilibrium in which $\sigma^*(s_B) = 0$ and $\sigma^*(s_G) \in (0, 1)$. In this equilibrium, the firm makes zero expected profit because the firm is indifferent between sending and not sending ads to consumers with s_G. Moreover, this equilibrium can only exist for a specific μ given model primitives, k, p, and the firm's strategy, i.e., $\widetilde{\mu}(\sigma^s, m_b) = p$ (if $\delta_c(m_b) \in (0, 1)$ and $\delta_c(m_g) = 1$) or $\widetilde{\mu}(\sigma^s, m_g) = p$ (if $\delta_c(m_b) = 0$ and $\delta_c(m_g) \in (0, 1)$). We focus on the other semi-separating equilibrium where the firm's expected profit can be positive, i.e., $\sigma(s_G) = 1$ and $\sigma(s_B) \in (0, 1)$.

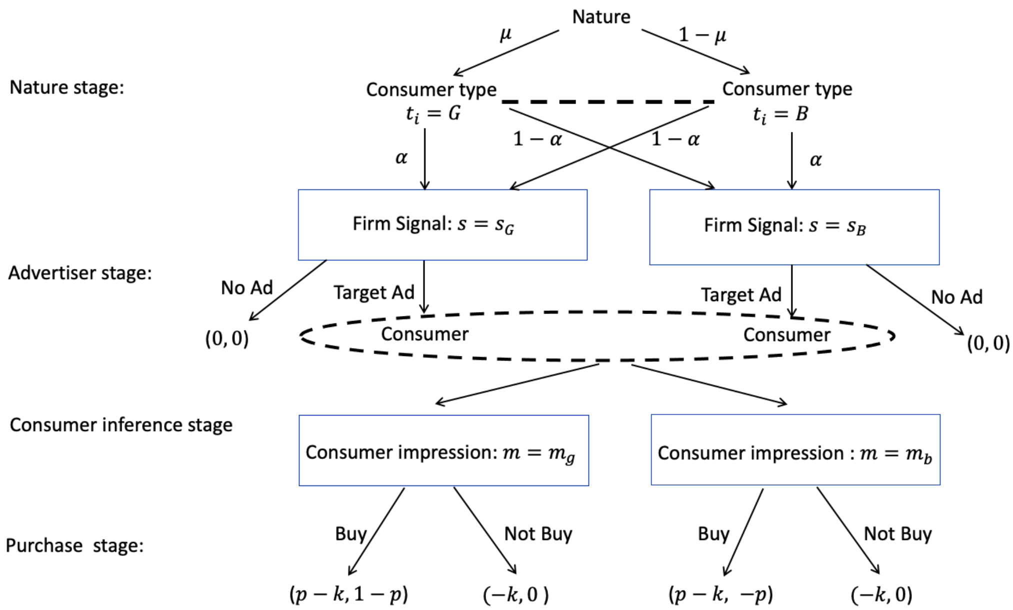

Figure 2: The game-tree for semi-separating equilibrium

beliefs along the equilibrium path are the same as the priors), and with a separating equilibrium, where updating is perfect (receiving an ad perfectly reveals the firm's signal s).

More specifically, when a consumer receives an ad, there are two possibilities: (1) the firm received a signal s_G with accuracy α (the probability of this event is $\mu\alpha + (1-\mu)(1-\alpha)$) or (2) the firm received a signal s_B with accuracy α (the probability of this event is $\mu(1-\alpha)+(1-\mu)\alpha$), but the firm "cheats" by sending an ad with probability $\sigma(s_B) > 0$. Therefore, the firm influences consumers' posterior beliefs through its strategy $\sigma(s_B)$. In equilibrium, given consumers' posterior beliefs $\widetilde{\mu}(\sigma^{\text{semi}}, m)$, they choose an optimal purchase behaviors about whether to purchase or not. Figure 2 shows the game structure of this semi-separating equilibrium.

In this equilibrium, facing a consumer with s_B, the firm mixes between sending an ad and not sending. The firm's expected payoff from sending an ad is $\mathbb{E}\Pi(A = a|s_B) = \Pr(m_g|s_B) \cdot \{p \cdot \delta_c(m_g) - k\} + \Pr(m_b|s_B) \cdot \{p \cdot \delta_c(m_b) - k\}$. Also, the expected payoff from not sending an ad is $\mathbb{E}\Pi(A = \varnothing|s_B) = 0$. For the existence of a semi-separating equilibrium, the firm should be indifferent, i.e., $\mathbb{E}\Pi(A = a|s_B) = \mathbb{E}\Pi(A = \varnothing|s_B) = 0$.[8] We first characterize the mixing behaviors of the consumer in a semi-separating equilibrium.

[8]Also, the firm must find it optimal to send an ad to consumers with s_G. It is confirmed that $\mathbb{E}\Pi(A = a|s_G) \geq \mathbb{E}\Pi(A = a|s_B)$ because $\Pr(m_g|s_G) > \Pr(m_g|s_B)$.

Lemma 4. *In a semi-separating equilibrium, the consumer with either $m = m_g$ or m_b mixes between buying and not buying after receiving an ad. If $\delta_c(m_b) > 0$, it must be $\delta_c(m_g) = 1$. Also, if $\delta_c(m_g) < 1$, it must be $\delta_c(m_b) = 0$.*

The lemma suggests that there are two different types of semi-separating. (1) Only consumers with a *bad* impression mix: $\delta_c(m_b) \in (0, 1)$ and $\delta_c(m_g) = 1$, which we call '*semi-b' case*,and (2) only consumers with a *good* impression mix: $\delta_c(m_g) \in (0, 1)$ and $\delta_c(m_b) = 0$, which we call '*semi-g' case*. But not both types mix their behaviors. More specifically, if a consumer with a bad impression *mixes* such that $\delta_c(m_b) > 0$, it must be that a consumer with a good impression *always* buys $\delta_c(m_g) = 1$. Also, if consumers with a good impression *mix* such that $\delta_c(m_g) < 1$, it must be that consumers with a bad impression *never* buy $\delta_c(m_b) = 0$. We characterize those two semi-separating equilibria in the following proposition.

Proposition 5. *Suppose consumers are persuadable: $\mu_c(m_b) \leq p \leq \widetilde{\mu}^U$. There exist two semi-separating equilibria where $\sigma^{semi}(s_G) = 1$ and $0 < \sigma^{semi}(s_B) < 1$:*

1. *If $\mu_c(m_b) < p < \widetilde{\mu}(\sigma^s, m_b)$ and $0 < k < p$, then $\sigma^{semi\text{-}b}(s_B) = \frac{p(1-\mu_c(m_b))(1-\alpha)-(1-p)\cdot\mu_c(m_b)\cdot\alpha}{(1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha}$, and only consumers with a bad impression m_b mix: $\delta_c^{semi\text{-}b}(m_g) = 1$, and $\delta_c^{semi\text{-}b}(m_b) = \frac{k}{p} - \left(1 - \frac{k}{p}\right) \frac{\mu\cdot(1-\alpha)\beta+(1-\mu)\cdot\alpha\cdot(1-\beta)}{\mu(1-\alpha)(1-\beta)+(1-\mu)\alpha\cdot\beta}$.*

2. *If $\mu_c(m_g) < p < \widetilde{\mu}^U$ and $k < p \cdot \Pr(m_g|s_B)$, then $\sigma^{semi\text{-}g}(s_B) = \frac{p(1-\mu_c(m_g))(1-\alpha)-(1-p)\cdot\mu_c(m_g)\cdot\alpha}{(1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha}$, and only consumers with a good impression m_g mix: $\delta_c^{semi\text{-}g}(m_g) = \frac{k}{p} \frac{\mu(1-\alpha)+(1-\mu)\alpha}{\mu(1-\alpha)\beta+(1-\mu)\alpha(1-\beta)}$ and $\delta_c^{semi\text{-}g}(m_b) = 0$. Note that $\widetilde{\mu}^U = \widetilde{\mu}(\sigma^s, m_g)$.*

Figure 3 demonstrates all equilibria of the game and shows how they change as the firm's targeting technology α improves.[9] First, a unique equilibrium exists in the entire parameter space except for a measure zero set.[10] Figure 3(a) depicts how the region of each equilibrium changes in response to an increase in α, indicated by bold arrows. The posterior belief $\widetilde{\mu}(\sigma^s, m_g)$ increases and converges to $\widetilde{\mu}^U$, eventually forcing $\widetilde{\mu}^U \to 1$. Moreover, the slope $1/\Pr(m_g|s_B)$ gets steeper. Therefore, as illustrated in Figure 3(b) where $\alpha \to 1$, the parameter region for semi-separating

[9] This figure depicts the case where $\widetilde{\mu}(\sigma^s, m_g) > \mu_c(m_g)$, i.e., the consumer's posterior belief is more positive when being targeted carries more information about the consumer's true match-type than the consumer's private signal. This is the case if the firm's signal is much more informative than the consumer's own signal, i.e., $\alpha \gg \beta$.

[10] They can exist multiple equilibria on the boundary between different equilibrium regions. We detail these boundaries in the proof of Lemma A-2 in the Appendix.

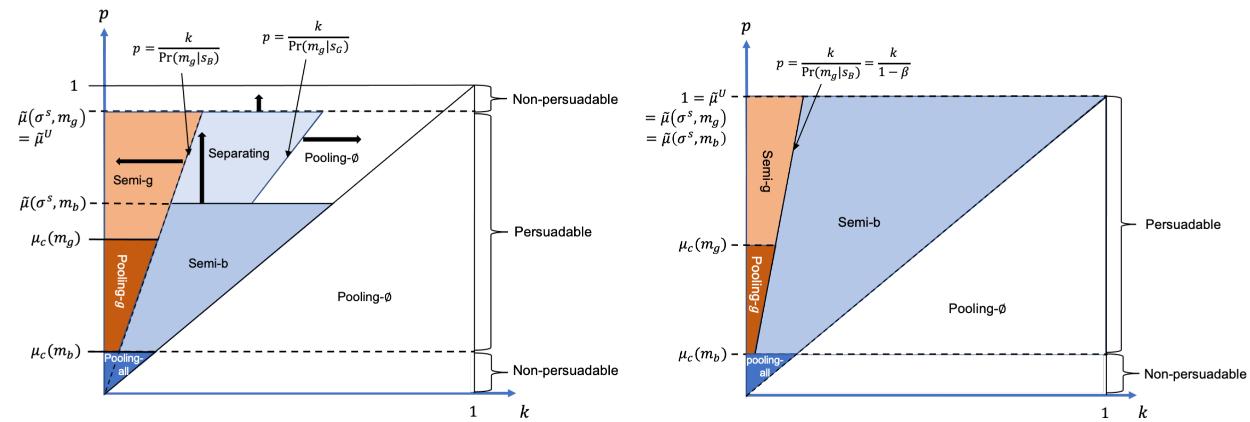

(a) Changes to the regions for each equilibrium as α increases (in the direction of red arrows)

(b) Equilibria in the limit $\alpha \to 1$

Figure 3: Equilibria of the game

equilibrium ('*semi-b*' type, in particular) increases, ultimately engulfing both separating and pooling equilibrium regions. In particular, the region for the separating equilibrium (of trapezoid shape) shrinks because its height diminishes as $\widetilde{\mu}(\sigma^s, m_g)$ increases and converges to $\widetilde{\mu}^U \to 1$. In other words, as the firm's targeting becomes more accurate, the firm's incentive problem worsens such that the separating equilibrium disappears. On the other hand, the scope of semi-separating equilibria in which the firm intentionally sends an ad to consumers with s_B increases, thus becoming more prevalent. Also, note that in the limit where $\alpha \to 1$, '*semi-b*' type semi-separating equilibrium still exists and is the unique equilibrium when $\mu_c(m_b) < p < \widetilde{\mu}^U$ and $p \cdot (1-\beta) < k < p$.

The results confirm that our key findings from the benchmark without the customer's informative signal are robust. Any semi-separating equilibrium exhibits the same qualitative properties. Namely, as α increases, the firm engages in less individually-targeted advertising.

Proposition 6. *In any semi-separating equilibrium, the firm engages in individually targeted advertising, i.e.,* $|\sigma(s_G) - \sigma(s_B)| > 0$. *As the firm's information becomes more precise, the firm engages in less targeted advertising, i.e.,* $|\sigma(s_G) - \sigma(s_B)|$ *decreases in* α.

The following limiting result where $\alpha \to 1$, is qualitatively similar to Corollary 1, focusing on the '*semi-b*' equilibrium. Under exogenous price, even though there are misaligned incentives between the firm and consumers, we show the consumer surplus can be positive.

Corollary 2. *As* $\alpha \to 1$, *if* $\mu_c(m_b) < p < \frac{k}{1-\beta}$ *and* $k < p$, *we have the following results:*

1. *The firm knowingly targets consumers with* s_B *with probability* $\frac{1-p}{p}\frac{\mu(1-\beta)}{(1-\mu)\beta}$. *Targeted consumers mistakenly buy the product with probability* $1-\beta+\beta(\frac{k}{p}-\frac{p-k}{p}\frac{\mu(1-\beta)}{(1-\mu)\beta})$ *and realize utility* $-p$*;*

2. *If* μ *is sufficiently large, the consumer surplus is positive for a sufficiently large* β*.*

Some consumers will end up making suboptimal purchasing decisions even with private information. In particular, even when the advertiser's prediction becomes perfectly accurate (i.e., $\alpha \to 1$), the advertiser still targets consumers with s_B with probability $\frac{1-p}{p}\frac{\mu(1-\beta)}{(1-\mu)\beta}$. Targeted consumers will sometimes buy the product and suffer a utility loss.[11] Nevertheless, the consumer surplus can be positive, especially when β is large enough and μ is not too small. When the consumer's private information becomes more accurate, the firm sends fewer ads to consumers with s_B. Moreover, consumers can avoid making wrong purchases by relying on private signals. So, the consumer surplus can be positive despite the incentives of the firm and the consumers being misaligned.

4.3 Personal Data Opt-out

In order to send targeted advertising, advertisers need to access consumers' personal data. Given the firm's advertising incentives, we consider consumers' privacy decisions to opt out of data collection and its welfare implications in this section. At the beginning of the game, consumers have an option to opt out of the data collection. If consumers opt out, the advertiser does not access consumers' personal data, effectively leading to $\alpha = 1/2$. Otherwise, the advertiser can predict according to the firm's prediction capability $\alpha > 1/2$. Consumers can opt out of the data collection for an infinitesimal cost.[12]

The order of the game is as follows: First, consumers choose whether to opt out of the data collection. Then, the game we analyzed previously is played as a subgame. If consumers opt out of data collection, only pooling equilibria can exist because the advertiser has no information on individual consumers' preferences. Borrowing from Proposition 4, we can describe the equilibrium under opt-out in the following Lemma:

[11] Similarly, if $\mu_c(m_g) \leq p < 1$ and $k < (1-\beta)\cdot p$, we have the semi-g equilibrium. The firm sends an ad to consumers with s_B with probability $\sigma^{\text{semi-g}}(s_B) \to \frac{1-p}{p}\frac{\mu\beta}{(1-\mu)(1-\beta)}$, which is increasing in β. Among these consumers, $\Pr(m_g|s_B)\delta_c(m_g) \to \frac{k}{p}$ fraction purchases and eventually realizes a negative utility.

[12] We capture the consumer's privacy concerns in a minimal fashion. In our model, consumers do not have any explicit disutility for giving up personal information or disutility from receiving ads. Nor is there any utility from allowing the data collection, as websites or apps often provide full service to a consumer upon agreeing to share data. Any incentive for opting in or out of data collection must come only from the transaction utility of the focal product.

Lemma 5. *Suppose consumers are persuadable ($\mu_c(m_b) < p \leq \widetilde{\mu}^U$), and consumers opts out of the firm's data collection. If $k < p \cdot \Pr(m_g)$ and $p \leq \mu_c(m_g)$, the advertiser mass advertises (i.e., the pooling-g) in equilibrium. Otherwise, there is no advertising (i.e., the pooling-$\varnothing$) in equilibrium.*

If the unit cost of advertising is cheap ($k \leq p \cdot \Pr(m_g|s_B)$), and the price is not too high ($p \leq \mu_c(m_g)$), the equilibrium is pooling with mass advertising. It holds under *both* cases whether or not consumers choose to opt out. Thus, given the small cost of opting out, consumers must stay with the default option and not opt-out.

Outside this parameter region, the equilibrium can differ depending on the consumer's privacy choice. If consumers opted out, only the no-advertising pooling equilibrium exists in which consumers can only expect a zero surplus. However, suppose consumers do not opt out of data collection. Then, they may secure a positive expected surplus outside the region for pooling-g equilibrium because the equilibrium can be either semi-separating or separating equilibrium. For example, under semi-b equilibrium, consumers with a good impression (m_g) buy the product with probability 1 and expect a positive surplus, whereas consumers with a bad impression (m_b) are indifferent between buying and not buying, thus expecting a zero surplus. Therefore, it follows that consumers never opt out of the data collection, allowing tracking of the consumer information.

Proposition 7. *Suppose the consumers are persuadable. When consumers have an option to opt out of data collection, they never opt out. Moreover, for an intermediate range of k (i.e., either semi-b or separating equilibrium region), consumer surplus is positive. This parameter region expands as α increases so that the consumer surplus is more likely to be positive.*

Proposition 7 shows that, when the price is exogenous, consumers have no incentive to withhold information from the firm. This is because the firm may not have sufficient information about consumers and, thus, respond by withholding advertising altogether, which strips away the consumers' opportunity to buy the product. Moreover, under exogenous pricing, consumers need not worry about being charged a higher price in case they opt-in and find a good match for the product. The proposition also shows that the consumer surplus can be positive under certain parameter regions, which a higher prediction accuracy enlarges.

5 Endogenous Pricing

In the previous sections, we assume that the price is exogenously given. As the algorithm's prediction accuracy changes, the advertiser's optimal price should also change. In this section, we allow the advertiser to optimally set the price at the very beginning of the game and, subsequently, the game we have analyzed previously be played as a subgame. A consumer observes the price in the ad if she receives one.[13] We focus on analyzing how the accuracy of prediction α affects the optimal price and consumer welfare.

5.1 Pricing equilibrium analysis

First, we investigate how profit changes with a marginal increase in price. Within the interior of each equilibrium region presented in Propositions 3, 4, and 5 and Figure 3(a), we can show that the advertiser's profit is continuously (weakly) increasing in price p. For example, in the interior of the pooling-g equilibrium, which exists for $k < p \cdot \Pr(m_g|s_B)$ and $\mu_c(m_b) < p \leq \mu_c(m_g)$, the players' strategies do not locally depend on the price. Therefore, the firm's profit must increase in the firm's price for $p < \mu_c(m_g)$. Also, in the interior region of the separating equilibrium, the firm's profit must increase in price as players' strategies do not locally depend on price. The remaining part is the interior regions of both semi-separating equilibria. In both types of semi-separating equilibrium, the firm is indifferent between sending an ad to consumers with s_B and not sending it. This implies that the firm's expected profit from these consumers must be zero. Hence, it suffices to compute the expected profit from consumers with s_G, which is continuously (weakly) increasing in price p (see Lemma A-1 in the Appendix for more details). Moreover, we show that the advertiser's profit is continuous everywhere in p except on some boundaries between different types of equilibrium where a shift in the equilibrium regime occurs (see Lemma A-2 in Appendix for more details).

We can then determine the optimal price by comparing the firm's profits under different prices. In general, the optimal pricing decision can be described by three thresholds on k.

Proposition 8. *Let* $\overline{\beta} = \frac{\sqrt{\mu^2-6\mu+5}-\mu-1}{2(1-2\mu)}$. *There exist thresholds* $\overline{k}_1 \leq \overline{k}_2 \leq \overline{k}_3$ *such that:*

1. *For* $k < \overline{k}_1$, *if* $\beta < \overline{\beta}$, *the optimal price is* $p^* = \mu_c(m_b)$ *under the the pooling-all equilibrium.*

[13] The firm chooses its price prior to observing the noisy signal about individual consumer types $t_i \in \{G, B\}$. Therefore, the price cannot provide any information about the consumer's t_i.

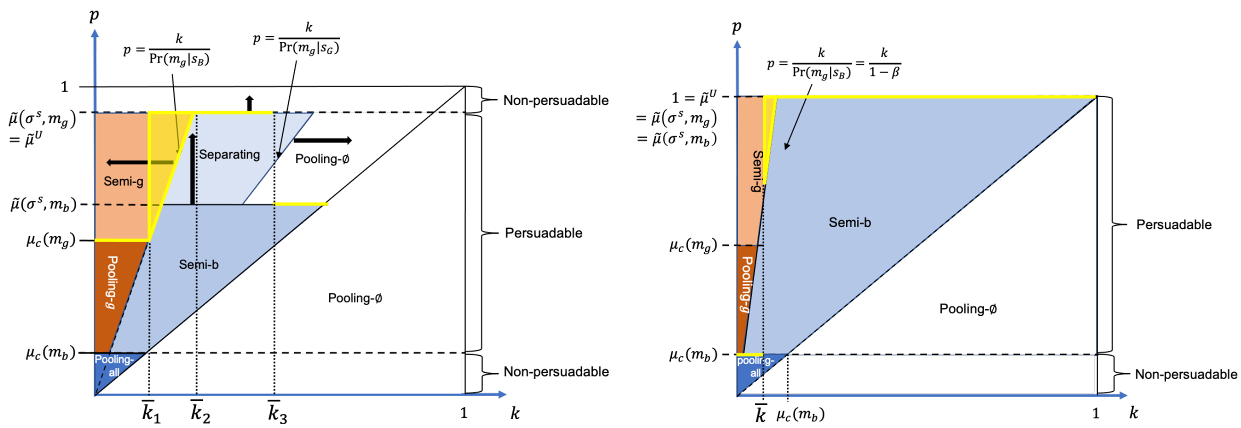

(a) Equilibrium Under Optimal Price ($\beta \geq \overline{\beta}$) (b) Optimal Price as $\alpha \to 1$ ($\beta < \overline{\beta}$)

Figure 4: Equilibrium prices (colored in yellow)

Otherwise (i.e., $\beta \geq \overline{\beta}$), the optimal price is $p^* = \mu_c(m_g)$ under the pooling-g equilibrium.

2. *If $\overline{k}_1 \leq k < \overline{k}_2$, any $p \in \left[\max\left\{\mu_c(m_g), \frac{k}{Pr(m_g|s_B)}\right\}, \widetilde{\mu}(\sigma^s, m_g)\right]$ is optimal under the semi-g equilibrium.*

3. *If $\overline{k}_2 \leq k < \overline{k}_3$, $p^* = \widetilde{\mu}(\sigma^s, m_g)$ is optimal under the separating equilibrium.*

4. *If $k \geq \overline{k}_3$, $p^* = \widetilde{\mu}(\sigma^s, m_b)$ is optimal under the semi-b equilibrium.*[14]

Figure 4(a) illustrates the optimal price for different values of k when $\beta \geq \overline{\beta}$. The optimal pricing strategy is divided into four regions of k where different advertising strategy arises as an equilibrium. If k is sufficiently small ($k \leq \overline{k}_1$), the unit cost of advertising is so inexpensive that the firm sends an ad to the entire market, resulting in a pooling equilibrium. If the consumer's private signal is sufficiently noisy ($\beta < \overline{\beta}$), the consumer's posterior beliefs are homogeneous enough that the firm charges a low price $p^* = \mu_c(m_b)$, and consumers buy the product irrespective of their private signals (i.e., pooling-all equilibrium). However, if $\beta \geq \overline{\beta}$, consumers' beliefs are heterogeneous enough that the firm charges a high price, and, in turn, only a consumer with a good private signal will make a purchase (i.e., pooling-g equilibrium). For larger values of k ($k \geq \overline{k}_1$), the firm engages in individually targeted advertising, where the equilibrium of the advertising game corresponds to a semi-separating or separating equilibrium. The details of the bounds ($\overline{k}_1, \overline{k}_2, \overline{k}_3$) for the optimal price and the corresponding equilibrium are provided in the appendix.

[14]If k is very large, i.e., $k \geq \overline{k}_4 = \max\{\overline{k}_3, \widetilde{\mu}(\sigma^s, m_b)\}$, the unique equilibrium is pooling-∅ equilibrium in which the firm does not send any ads. So, $p^* = \widetilde{\mu}(\sigma^s, m_b)$ is weakly optimal because transaction does not occur anyway.

When the price is exogenously fixed, the firm responds to a higher prediction accuracy by advertising to consumers that it predicts to have a bad match with a higher probability (Proposition 6). In contrast, the advertiser may not do so under endogenous pricing. Instead of targeting "wrong" customers, which dilutes its recommendation power, the firm can increase the price for "right" customers. For example, when $\overline{k}_2 \leq k < \overline{k}_3$, the firm can chooses a price $p = \widetilde{\mu}(\sigma^s, m_g)$ under separating equilibrium where the advertising strategy remains fully individually targeted such that $|\sigma(s_G) - \sigma(s_B)| = 1$. Also, even under semi-separating equilibrium, both $\widetilde{\mu}(\sigma^s, m_g)$ and $\widetilde{\mu}(\sigma^s, m_b)$ (the upper limits of prices which sustain the semi-separating equilibrium) increase in α. This, in turn, encourages the advertiser to reach consumers with a good match s_G with higher prices. Thus, it deploys a more individually targeted advertising strategy to improve the persuasive power of the targeted advertising, which again permits it to charge a higher price as prediction accuracy increases. Thus, advertising serves as a clearer recommendation for a consumer, but consumers may face a higher price as α increases under endogenous pricing. The following proposition states these implications for the firm's individually targeted advertising.

Proposition 9. *If the firm chooses its price optimally, as α increases, the firm engages in weakly more individually targeted advertising in the following sense:*

1. *Under the semi-b equilibrium, the firm engages in maximally targeted advertising where $|\sigma(s_G) - \sigma(s_B)| = 1$.*

2. *Under the semi-g equilibrium, both the upper and lower bound of the equilibrium price range become higher as α increases. Also, as p increases, the firm engages in more individually targeted advertising, i.e., $\partial\, |\sigma(s_G) - \sigma(s_B)| / \partial\, p \geq 0$.*

Note that even in semi-b equilibrium, where the firm's advertising strategy is fully individually targeting ($|\sigma(s_G) - \sigma(s_B)| = 1$), consumers still mix, i.e., $\delta_c^{*\text{semi-b}} \in (0, 1)$. So, the semi-separating equilibrium does not converge completely to the separating equilibrium, where $\delta_c^s = (\delta_c^s(m_g) = 1, \delta_c^s(m_b) = 0)$. Moreover, as the firm's information precision approaches 1, the region for each equilibrium simplifies as follows, and we show that the consumer surplus reduces to zero under a wide range of parameters.

Corollary 3. *As $\alpha \to 1$, there exists $\overline{k}$ such that*

1. *If $k < \overline{k}$, the optimal price is $p^* = \mu_c(m_b)$ for $\beta < \overline{\beta}$ with the pooling-all equilibrium, and $p^* = \mu_c(m_g)$ for $\beta \geq \overline{\beta}$ with the pooling-g equilibrium. (2) Otherwise (i.e.,, $\overline{k} \leq k$), any $p^* \in [\frac{k}{1-\beta}, 1)$ is optimal with the semi-g equilibrium if $k < 1 - \beta$, or $p^* = 1$ if $k \geq 1 - \beta$.*

2. *Moreover, consumer surplus approaches 0 except when $p^* = \mu_c(m_b)$ under pooling-all equilibrium (i.e., $k < \overline{k}$ and $\beta < \overline{\beta}$).*

Figure 4(b) depicts the limiting case as $\alpha \to 1$ for the case of $\beta < \overline{\beta}$. First, the optimal price must approach 1 for $k \geq 1 - \beta$ under the semi-b equilibrium, so the consumer welfare must approach 0. Also, when $\overline{\beta} \leq \beta$, the advertiser charges $p = \mu_c(m_g)$ for $k \leq \overline{k}$, where consumer welfare is also 0 under the "semi-g" equilibrium. Only when the equilibrium advertising strategy is pooling-all, where the consumer purchases irrespective of her own signal m and the price is $p = \mu_c(m_b)$, consumer surplus is positive. This pooling can arise only when the advertising cost is small ($k < \overline{k}$), and the consumer signal is not sufficiently informative ($\beta < \overline{\beta}$).

The welfare implications in Corollary 3 is in a stark contrast to those under the exogenous price case in Corollary 2. Under exogenous pricing, despite the firm's strategic mis-targeting, consumer welfare is strictly positive under a wide range of parameters (where semi-separating equilibria exist). However, if the firm chooses its price optimally, except in the region for the pooling-all equilibrium, the firm no longer "cheats" with its advertising. More specifically, both the firm's advertising strategy under both types of semi-separating equilibria converges to the fully individually targeting strategy $\sigma^s = (\sigma_G = 1, \sigma_B = 0)$.[15] Nevertheless, it raises its price to extract all the surplus. Therefore, consumers can be worse off when they have a more accurate signal ($\beta > \overline{\beta}$), and the consumer welfare goes to zero. More importantly, this generates an important implication for the firm if consumers make their own privacy choice to opt out of the firm's data collection, which is the focus of the next section.

5.2 Personal Data Opt-out

Same as in the exogenous pricing case, we allow consumers to choose whether to opt out of the firm's data collection when the price is the firm's endogenous choice. To gain analytical tractability, we focus on the limiting case as $\alpha \to 1$. Hence, if a consumer opts out, the firm cannot access the

[15] Again, even in the limit $\alpha \to 1$, consumers still mix. So, the semi-separating equilibria do not converge completely to the separating equilibrium.

consumer data and does not receive a signal about the consumer's match type, effectively $\alpha = 1/2$. Otherwise, the advertiser can predict with accuracy $\alpha \geq 1/2$, according to the firm's prediction capability. First, consumers choose whether or not to opt out of the firm's data collection. Then, the firm chooses the price, and the game from Section 5 is played. The firm charges two different prices, one observed by consumers who allow data collection, and the other observed by consumers who opt out of data collection.[16]

If consumers opt out of the firm's data collection, only pooling equilibria (either "pooling-g" or 'pooling-all") are possible because the advertiser does not have a signal on consumers' match type. Then, the optimal price for consumers who have opted out of data collection will be $\mu_c(m_g)$ or $\mu_c(m_b)$. Given an infinitesimal cost of opting out of the data collection, consumers opt out if and only if their expected surplus is strictly greater than the default option of allowing the data collection. From Corollary 3, as $\alpha \to 1$, the consumer surplus is zero except under the pooling-all equilibrium. Therefore, consumers opt out if and only if their surplus under the default is zero (i.e., pooling-g, semi-g, or semi-b equilibrium), and opting out results in the pooling-all equilibrium.

Proposition 10. *As $\alpha \to 1$, when consumers have an option to opt out of data collection, consumers opt out for some intermediate range of k: $k \in [\bar{k}, \mu_c(m_b)]$. However, as $\beta \to 1$, the region for opting out disappears.*

When the price is exogenous, we find that consumers *never* opt out (Proposition 7) but the consumer surplus is positive for an intermediate range of k (i.e., either semi-b or separating equilibrium region). However, the results are now reversed when the price is endogenous. Under the firm's endogenous pricing, the pooling-all equilibrium is the only equilibrium that grants a positive consumer surplus. If the firm uses consumer data for targeting, the scope of semi-separating equilibria expands when α goes to 1, as demonstrated in Figure 4(b). This reduces the scope for pooling equilibria, including the pooling-all equilibrium. Thus, the pooling-all equilibrium exists in a larger parameter region when the firm engages in non-targeted advertising under consumer data opt-out. In particular, when $k \in [\bar{k}, \mu_c(m_b)]$, consumers opt out of the data collection to be in the pooling-all equilibrium rather than be in the semi-separating equilibrium by allowing the firm to use their personal data for targeted advertising.

[16] A common reason consumers opt out of data collection is the belief that firms charge lower prices to anonymous consumers (Ichihashi, 2020).

Moreover, as the consumer's own signal becomes more accurate, $\beta \to 1$, $\mu_c(m_b) = \frac{(1-\beta)\mu}{(1-\beta)\mu+\beta(1-\mu)}$ converges to 0. Thus, the region under which consumers opt out disappears. In other words, whether the consumer has an incentive to opt out of data collection depends on the quality of the consumer's own signal. Also, it reveals an important managerial insight regarding consumer privacy choice and firm response. There can be a complementarity between the advertiser's prediction technology and information provision technology. As the advertiser's algorithm becomes more accurate, it should also consider providing enough information for consumers to make their own assessment of product fit to prevent consumers from opting out of data collection, which would render the advertiser's algorithm useless. For example, the advertiser could improve the accuracy of the consumer's own assessment of product fit (β) by providing more informative ads, more information upon clicking the ads, or making the consumer's information search easier.

6 Conclusion

This paper examined a firm's optimal targeting strategy when the firm can predict how much consumers might enjoy the product by using various information about consumers. Then, merely being targeted can act as the firm's implicit recommendation that can persuade consumers' purchasing decisions. However, the firm may have incentives to exploit the recommendation role of advertising by sending ads to the wrong consumers even though the advertiser does not believe that the product is a good fit for them. A higher prediction accuracy strengthens the recommendation role of a targeted ad but also worsens this incentive problem; the profit-maximizing firm's incentives are not necessarily aligned with those of the consumers. In this paper, we have analyzed under which conditions targeted advertising can act as an implicit recommendation and how the main economic forces are affected by the firm's prediction accuracy and endogenous pricing decision.

Under exogenous pricing, being targeted can be either a perfect (under the separating equilibrium) or noisy signal (under the semi-separating equilibria), except when the unit cost of advertising or the exogenous price is sufficiently small. Interestingly, as the firm's prediction becomes more accurate, the firm engages in less targeted advertising. Even when the firm's prediction capability is perfect, the firm knowingly targets a fraction of the wrong consumers. As a result, consumers receive ads for unfit products and make incorrect purchase decisions because they expect the firm's advertising is targeted individually, albeit imperfect. Despite the negative consequences of generating some consumers' wrong purchasing decisions, the consumer surplus can remain positive because the firm can better identify consumers with a good fit for the product under the exogenous pricing case.

Then, we show the firm's ability to choose its price optimally reverses these effects of the prediction accuracy. Under the endogenous pricing case, instead of diluting its targeted advertising, the firm can simply raise the price and focus on serving the right consumers. Therefore, the firm increasingly engages in targeted advertising in a wide range of parameter spaces. However, when the firm can perfectly predict each consumer's match for the product, the firm's ability to increase price drives the consumer surplus down to zero, except in a minuscule parameter region.

Given the implications for the consumer surplus, the paper analyzes the consumers' decision on whether or not to opt out of personal data collection. We show that under the exogenous price, consumers never opt out of data collection. This is because the absence of the firm's targeting prevents consumers from making the right purchasing decisions. However, under the endogenous price, the consumers find opting out of data collection more appealing because the firm's individual targeting often leads to a greater price and extraction of consumer surplus.

Our findings reveal useful insights regarding the prediction technologies and the firm's targeting strategies. Mistargeting is not only an outcome of firms' technological capabilities but can also of the advertiser's economic incentives. Given the misaligned incentives, the firm's better prediction capability does not necessarily translate into more accurate targeting. Moreover, it may discourage consumers from allowing the firm's data collection. These results can shed important implications for regulators that seek to protect consumer welfare in light of increasing data collection and targeted advertising. For example, what is the optimal data collection policy both from the perspectives of the firm and the social planner? Does giving the consumers a choice between data collection and no data collection options benefit the consumers, consistent with recent practices by major tech companies, such as Apple?[17] Also, what should be the consumer's default decision? Should consumers opt out of the data collection, or should opt in to the data collection? These are interesting and important questions from a managerial and regulatory perspective. However, given the current paper's focus and scope, we leave these questions to be explored in future research.

[17]https://www.nytimes.com/2021/04/26/technology/personaltech/apple-app-tracking-transparency.html.

Appendix

Proof of Lemma 1

The first part that $\widetilde{\mu}(\widetilde{\sigma}, m) \geq \mu_c(m)$ if and only if $\widetilde{\sigma}(s_G) \geq \widetilde{\sigma}(s_B)$ follows directly from Equation (2). Also, we have $\frac{\partial\,[\widetilde{\mu}-\mu_c]}{\partial\,\alpha} \geq 0$ if and only if $\frac{\partial\,\widetilde{\mu}}{\partial\,\alpha} = \frac{\mu_c(1-\mu_c)(\widetilde{\sigma}(s_G)^2-\widetilde{\sigma}(s_B)^2)}{\big(\mu_c(\alpha\cdot\widetilde{\sigma}(s_G)+(1-\alpha)\cdot\widetilde{\sigma}(s_B))+(1-\mu_c)(\alpha\cdot\widetilde{\sigma}(s_B)+(1-\alpha)\cdot\widetilde{\sigma}(s_G))\big)^2} \geq 0$ if and only if $\widetilde{\sigma}(s_G) \geq \widetilde{\sigma}(s_B)$. ■

Proof of Lemma 2

The proof is laid out in the main text two paragraphs above the statement. ■

Proof of Proposition 1

The steps of the analysis provided in the main text show most of the proof of Proposition 1. It only remains to show $\sigma^*(s_G) = \phi \cdot \sigma^*(s_B)$. It follows directly from Equation (3) such that $p \cdot \big(\widetilde{\sigma}(s_G)\big(\mu\alpha + (1-\mu)(1-\alpha)\big) + \widetilde{\sigma}(s_B)\big(\mu(1-\alpha) + (1-\mu)\alpha\big)\big) = \big(\widetilde{\sigma}(s_G) \cdot \mu\alpha + \widetilde{\sigma}(s_B) \cdot \mu(1-\alpha)\big)$. This simplifies to $\widetilde{\sigma}(s_G) = \phi \cdot \widetilde{\sigma}(s_B)$, where $\phi = \frac{(1-\mu)\alpha\cdot p-\mu(1-\alpha)(1-p)}{\mu\cdot\alpha(1-p)-(1-\mu)(1-\alpha)p}$. ■

Proof of Proposition 2

The proof follows from differentiating ϕ with respect to α: $\frac{\partial\,\phi}{\partial\,\alpha} = -\frac{(p-\mu)\big(\mu(1-p)+(1-\mu)p\big)}{\big(\mu\cdot\alpha(1-p)-(1-\mu)(1-\alpha)p\big)^2} \leq 0$ because $p \in (\mu, \widetilde{\mu}(\sigma^s)]$. ■

Proof of Corollary 1

The first part follows directly from taking the limit $\lim_{\alpha\to 1} \phi = \frac{(1-\mu)p}{\mu(1-p)}$, which is greater than 1 because $p \in (\mu, \widetilde{\mu}(\sigma^s)]$. In the benchmark, k/p fraction of consumers buy the product. ■

Proof of Lemma 3

Suppose there exists an equilibrium where the firm totally mixes, i.e., $0 < \sigma^*(s_B) \leq \sigma^*(s_G) < 1$. Given consumers' purchase strategy $\delta_c^*(m_g)$ and $\delta_c^*(m_b)$, the firm must be indifferent between sending an ad or not sending an ad for both consumers with s_G and s_B. The firm's indifference conditions for both cases are $p\cdot(\Pr(m_g|s_G)\cdot\delta_c^*(m_g)+\Pr(m_b|s_G)\cdot\delta_c^*(m_b))-k = 0$ and $p\cdot(\Pr(m_g|s_B)\cdot$

$\delta_c^*(m_g) + \Pr(m_b|s_B) \cdot \delta_c^*(m_b)) - k = 0$, respectively. It is straightforward to see that, however, the left-hand side of the former condition is greater than or equal to the left-hand side of the latter condition. This is because $\Pr(m_g|s_G) > \Pr(m_g|s_B)$ given that $\alpha, \beta > 1/2$. Moreover, as we will show, $\delta_c^*(m_g) \geq \delta_c^*(m_b)$. Therefore, the term in the former condition places a strictly greater weight on $\delta_c^*(m_g)$ (which is weakly greater than $\delta_c^*(m_b)$). This shows that the indifference condition cannot hold simultaneously unless $\delta_c^*(m_g) = \delta_c^*(m_b)$.

If $\delta_c^*(m_g) = \delta_c^*(m_b)$, the firm's expected profit is zero because the firm is indifferent between sending an ad or not. Also, for the firm to adopt a totally mixed strategy, consumers must also mix, i.e., $\delta_c^*(m_g) \in (0,1)$. A consumer's indifference condition is $\widetilde{\mu}(\widetilde{\sigma}, m) - p = 0$ for both $m = m_g$ and m_b as defined in Equation (1) and (2). In fact, it is easy to check that $\widetilde{\mu}(\widetilde{\sigma}, m_g) > \widetilde{\mu}(\widetilde{\sigma}, m_b)$ because given $\beta > 1/2$, a consumer's posterior is more positive if her own impression is better. So, if $\delta_c^*(m_b) \geq 0$, then $\delta_c^*(m_g) = 1$. Also, if $\delta_c^*(m_g) \in (0,1)$, then $\delta_c^*(m_b) = 0$ must hold (which proves Lemma 4). So, $\delta_c^*(m_g) = \delta_c^*(m_b)$ cannot hold unless $\delta_c^*(m_g) = \delta_c^*(m_b) = 0$, or $= 1$. Moreover, if $\sigma^*(s_B) > 0$, then $\sigma^*(s_G) = 1$. ■

Proof of Proposition 3, 4 and Lemma 4

The propositions are proved in the main text and the lemma in the proof of Lemma 3. ■

Proof of Proposition 5

(1) For *semi-b* case, where consumers with m_b is indifferent , the indifference condition for consumers with m_b, $\widetilde{\mu}(\sigma^{\text{semi}}, m_b) = \frac{\mu_c \cdot \alpha + \mu_c \cdot (1-\alpha) \cdot \sigma^{\text{semi-b}}(s_B)}{\mu_c \alpha + (1-\mu_c)(1-\alpha) + \sigma(s_B) \cdot \left(\mu_c(1-\alpha) + (1-\mu_c)\alpha\right)} = p$ must hold. This helps us to pin down the firm's mixing probability for a consumer with s_B is $\sigma^{\text{semi-b}}(s_B) = \frac{p(1-\mu_c)(1-\alpha) - (1-p) \cdot \mu_c \cdot \alpha}{(1-p)\mu_c(1-\alpha) - p(1-\mu_c)\alpha}$. where $\mu_c = \mu_c(m_b)$. Also, $\delta_c(m_g) = 1$ must hold. To ensure the existence of $\sigma^{\text{semi-b}}(s_B) \in (0,1)$, it is necessary and sufficient to have $\widetilde{\mu}(\sigma^p, m_b) = \mu_c(m_b) < p < \widetilde{\mu}(\sigma^s, m_b)$. Also, because the firm adopts $\sigma(s_B) \in (0,1)$, $\Pr(m_g|s_B) + \Pr(m_b|s_B) \cdot \delta_c^{\text{semi}-b}(m_b) = \frac{k}{p}$ must hold. That $\delta_c^{\text{semi}-b}(m_b) \in (0,1)$ holds implies that $0 < \delta_c^{\text{semi-b}}(m_b) = \frac{k/p - \Pr(m_g|s_B)}{\Pr(m_b|s_B)} < 1$ must satisfy, which leads to $\Pr(m_g|s_B) < k/p < 1$.
(2) For *semi-g case*, $\delta_c^{\text{semi-g}}(m_b) = 0$ and $\delta_c^{\text{semi-g}}(m_g) \in (0,1)$. Solving the indifference condition gives us $\delta_c^{\text{semi-g}}(m_g) = \frac{k}{p} \cdot \frac{1}{\Pr(m_g|s_B)}$, which is less than 1 if and only if $\frac{k}{p} < \Pr(m_g|s_B)$. The indifference condition of consumers with m_g is the same as that of consumers with m_b in the previous case,

except that their private prior is $\mu_c = \mu_c(m_g)$. To ensure that $\sigma^{\text{semi-g}}(s_B) \in (0,1)$ exists such that $\widetilde{\mu}(\sigma^{\text{semi-g}}(s_B), m_g) = p$, we need $\widetilde{\mu}(\sigma^p, m_g) = \mu_c(m_g) < p < \widetilde{\mu}^U = \widetilde{\mu}(\sigma^s, m_g)$. ■

Proof of Proposition 6

In the semi-g equilibrium, $\sigma(s_B) = \frac{p(1-\mu_c(m_g))(1-\alpha)-(1-p)\cdot\mu_c(m_g)\cdot\alpha}{(1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha}$. Differentiating it with respect to α, we have $\frac{p^2(1-\mu_c(m_g))^2-(1-p)^2\mu_c(m_g)^2}{\left((1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha\right)^2}$, which is ≥ 0 if and only if $(1-p)\mu_c(m_g) \leq p(1-\mu_c(m_g))$. This condition holds from the assumption that $\mu_c(m_g) \leq p$. In the semi-b equilibrium, $\sigma(s_B) = \frac{p(1-\mu_c(m_b))(1-\alpha)-(1-p)\cdot\mu_c(m_b)\cdot\alpha}{(1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha}$. Differentiating it with respect to α, we have $\frac{p^2(1-\mu_c(m_b))^2-(1-p)^2\mu_c(m_b)^2}{\left((1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha\right)^2}$, which is ≥ 0 if and only if $(1-p)\mu_c(m_b) \leq p(1-\mu_c(m_b))$. This condition holds from the assumption that $p \geq \mu_c(m_b)$. Since $\sigma(s_G) = 1$ in both semi-separating equilibria, this proves that $|\sigma(s_G) - \sigma(s_B)|$ decreases in α. ■

Proof of Corollary 2

In this region, we have the semi-b equilibrium in which the firm's strategy converges to $\sigma^*(m_g) = 1$ and $\sigma^*(s_B) = \frac{1-p}{p} \cdot \frac{\mu(1-\beta)}{(1-\mu)\beta}$. Also, among those targeted consumers, only consumers with m_b mix with $\delta_c^*(m_b) = 1 - \beta + \beta(\frac{k}{p} - \frac{p-k}{p} \cdot \frac{\mu(1-\beta)}{(1-\mu)\beta})$. Thus, a fraction of consumers of type $t_i = B$ who purchases the product and suffers a utility loss is $1 - \beta + \beta \cdot \delta_c^*(m_b) = 1 - \beta + \beta(\frac{k}{p} - \frac{p-k}{p} \cdot \frac{\mu(1-\beta)}{(1-\mu)\beta})$. The consumer surplus is $CS = \mu \cdot \left(\beta + (1-\beta) \cdot \delta_c^*(m_b)\right) \cdot (1-p) + (1-\mu) \cdot \left(1 - \beta + \beta \cdot \delta_c^*(m_b) \cdot (0-p)\right)$. At $\beta = 1$, $CS|_{\beta=1} = \mu(1-p) - (1-\mu)k$, which is ≥ 0 if and only if $\mu \geq \frac{k}{1+k-p}$. So, if μ is above a certain threshold, by continuity, the consumer surplus is positive for β close to 1. ■

Proof of Lemma 5

It follows directly by plugging $\alpha = 1/2$ into Proposition 4. ■

Proof of Proposition 7

The first part follows directly from Lemma 5, which shows that the consumer surplus from opting out of data collection is zero. So, consumers weakly prefer opting in. Second part is also straightforward because if $k \leq p \cdot \Pr(m_g|s_B)$, then either pooling-g or semi-g equilibrium exists, each of which has zero consumer surplus. Also, if $k > p \cdot \Pr(m_g|s_G)$ and $p > \widetilde{\mu}(\sigma^s, m_b)$, then pooling-∅ is the unique equilibrium, which also gives zero consumer surplus. On the other hand, if $p \cdot \Pr(m_g|s_B) < k \leq p \cdot \Pr(m_g|s_G)$, then either semi-b or separating equilibrium exists where consumer surplus is positive. Lastly, as α increases, the region for parameter k identified above expands, i.e., $\Pr(m_g|s_B)$ and $\Pr(m_g|s_G)$ decreases and increases in α, respectively. Also, the boundary for p, i.e., $\widetilde{\mu}(\sigma^s, m_b)$, increases in α. ■

Proof of Proposition 8

We prove the proposition in a few following steps. First, we show that the firm's expected profit under semi-b equilibrium is weakly increasing in p.

Lemma A-1. *Within the semi-b equilibrium, the firm's profit is increasing in p. However, within the semi-g equilibrium, it is invariant in p.*

Proof. In both types of semi-separating equilibria, the firm is indifferent between sending and not sending an ad to a consumer with s_B. This implies that the firm's expected profit from this consumer is zero. Thus, it suffices to compute the expected profit from consumers with s_G.

In the "semi-b" equilibrium, a consumer's purchase decision is described by $\delta_c^{\text{semi-b}}(m_g) = 1$ and $\delta_c^{\text{semi-b}}(m_b) = \frac{k}{p} - (1 - \frac{k}{p}) \cdot \frac{\Pr(m_g|s_B)}{\Pr(m_b|s_B)}$. Therefore, we can write the advertiser's profit: $\Pi^{\text{semi-b}}(p) = \Pr(s_G)\Big[\Pr(m_g|s_G) + \Pr(m_b|s_G) \cdot \delta_c^{\text{semi-b}}(m_b)\Big] \cdot p - \Pr(s_G) \cdot k = \Pr(s_G)\Big[\Pr(m_g|s_G) - \Pr(m_b|s_G) \cdot \frac{\Pr(m_g|s_B)}{\Pr(m_b|s_B)}\Big](p - k)$ which strictly increases in p due to a positive correlation between s and m.

In the "semi-g" equilibrium, a consumers' purchasing strategies are $\delta_c^{\text{semi-g}}(m_g) = \frac{k}{p} \cdot \frac{1}{\Pr(m_g|s_B)}$ and $\delta_c^{\text{semi-g}}(m_b) = 0$. So, we can write the advertiser's profit: $\Pi^{\text{semi-g}}(p) = \Pr(s_G) \cdot \Pr(m_g|s_G) \cdot \delta_c^{\text{semi-g}}(m_g) \cdot p - \Pr(s_G) \cdot k = \Pr(s_G) \cdot k \cdot \big(\frac{\Pr(m_g|s_G)}{\Pr(m_g|s_B)} - 1\big)$ which is invariant in p. Thus, the firm's profit for semi-separating equilibrium is continuously (weakly) increasing in p. □

Next, we show the firm's equilibrium profit is continuous in p except for some boundaries between different types of equilibrium.

Lemma A-2. *Equilibrium profit is continuous in p except when $p = \mu_c(m_b)$ and $k \in (p \cdot Pr(m_g|s_B), p)$, or $p = \widetilde{\mu}(\sigma^s, m_b)$ and $k/p \in (Pr(m_g|s_B), Pr(m_g|s_G))$, or $p = \mu_c(m_g)$ and $k < p \cdot Pr(m_g|s_B)$.*

Proof. Under exogenous pricing, multiple equilibria can exist on the boundary between different types of equilibrium. We describe here the set of equilibria on each boundary. We focus only for the case of α sufficiently high such that $\widetilde{\mu}(\sigma^s, m_b) > \mu_c(m_g)$.

Case 1. Boundary between Pooling-g and Semi-g. Consider $p = \mu_c(m_g)$ and $k < p \cdot \Pr(m_g|s_B)$. The limit of the pooling-g equilibrium as p approaches $\mu_c(m_g)$ from below is $\sigma^{\text{pool-}m_g} \to (\sigma(s_G) = 1, \sigma(s_B) = 1)$ and $\delta_c^{\text{pool-}m_g} \to (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 4. The limit of the semi-b equilibrium as p approaches $\mu_c(m_g)$ from above is $\sigma^{\text{semi-g}} \to (\sigma(s_G) = 1, \sigma(s_B) = 1)$ and $\delta_c^{\text{semi-g}} \to (\delta_c(m_g) = \frac{k}{\mu_c(m_g)} \cdot \frac{1}{\Pr(m_g|s_B)}, \delta_c(m_b) = 0)$ from Proposition 5. Note that $\delta_c^{\text{semi-g}}(m_g) < 1$ for $k < p \cdot \Pr(m_g|s_B)$. Thus, we have $\sigma = (\sigma(s_G) = 1, \sigma(s_B) = 1)$ and $\delta_c = (\delta_c(m_g), \delta_c(m_b) = 0)$ is an equilibrium for any $\delta_c(m_g) \in [\frac{k}{p} \cdot \frac{1}{\Pr(m_g|s_B)}, 1]$. Because the firm's profit increases in $\delta_c(m_g)$, it is maximized by the pooling-g equilibrium at this boundary. There is a discontinuous drop in the firm's profit as we move from the pooling-g to the semi-g equilibrium, as $\delta_c(m_g)$ drops discontinuously from 1 to $\frac{k}{p} \cdot \frac{1}{\Pr(m_g|s_B)}$.

Case 2. Boundary between Pooling-g and Semi-b. Consider $p \in [\mu_c(m_b), \mu_c(m_g)]$ and $p = \frac{k}{\Pr(m_g|s_B)}$. The limit of the pooling equilibrium is $\sigma^{\text{pool-}m_g} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$ and $\delta_c^{\text{pool-}m_g} = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 4. The limit of the semi-separating equilibrium is $\sigma^{\text{semi-b}} = \left(\sigma(s_G) = 1, \sigma(s_B) = \frac{p(1-\mu_c(m_b))(1-\alpha)-(1-p)\cdot\mu_c(m_b)\cdot\alpha}{(1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha}\right)$ and $\delta_c^{\text{semi-b}} = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 5. Note that $\sigma^{\text{semi-b}}(s_B)$ must be lower than 1 for consumers with a bad impression to be indifferent. Thus, we have that $\sigma = (\sigma(s_G) = 1, \sigma(s_B))$ and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ is an equilibrium for any $\sigma(s_B) \in [\frac{p(1-\mu_c(m_b))(1-\alpha)-(1-p)\cdot\mu_c(m_b)\cdot\alpha}{(1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha}, 1]$. The firm's profit is the same for all equilibria because the advertiser is indifferent to target consumers with s_B.

Case 3. Boundary between Pooling-g and Separating. Consider $p \in [\widetilde{\mu}(\sigma^s, m_b), \mu_c(m_g)]$ and $p = \frac{k}{\Pr(m_g|s_B)}$. The limit of the pooling equilibrium is $\sigma^{\text{pool-}m_g} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$ and $\delta_c^{\text{pool-}m_g} = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 4. The limit of the separating equilibrium is $\sigma^s = (\sigma(s_G) = 1, \sigma(s_B) = 0)$, and $\delta_c^s = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 3. Thus, we have that $\sigma = (\sigma(s_G) = 1, \sigma(s_B))$ and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ is an equilibrium for any $\sigma(s_B) \in [0, 1]$. The firm's profit is the same for all equilibria because the advertiser is indifferent to target consumers with s_B.

Case 4. Boundary between Semi-g and Semi-b. Consider $p \in [\mu_c(m_g), \widetilde{\mu}(\sigma^s, m_b)]$ and $p = \frac{k}{\Pr(m_g|s_B)}$. The limit of the semi-g equilibrium is $\sigma^{\text{semi-g}} = (\sigma(s_G) = 1, \sigma(s_B) = \frac{p(1-\mu_c(m_g))(1-\alpha)-(1-p)\cdot\mu_c(m_g)\cdot\alpha}{(1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha})$ and $\delta_c^{\text{semi-g}} = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 5. The limit of the semi-b equilibrium is $\sigma^{\text{semi-b}} = \left(\sigma(s_G) = 1, \sigma(s_B) = \frac{p(1-\mu_c(m_b))(1-\alpha)-(1-p)\cdot\mu_c(m_b)\cdot\alpha}{(1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha}\right)$ and $\delta_c^{\text{semi-b}} = (\delta_c(m_g) = 1, \delta_c(m_b) =$

0) from Proposition 5. Note that $0 < \sigma^{\text{semi-b}}(s_B) < \sigma^{\text{semi-g}}(s_B) < 1$ for consumers with a bad impression to be indifferent under semi-b and consumers with a good impression to be indifferent under semi-g. Thus, we have that $\sigma = (\sigma(s_G) = 1, \sigma(s_B))$ and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ is an equilibrium for any $\sigma(s_B) \in [\frac{p(1-\mu_c(m_g))(1-\alpha)-(1-p)\cdot\mu_c(m_g)\cdot\alpha}{(1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha}, \frac{p(1-\mu_c(m_b))(1-\alpha)-(1-p)\cdot\mu_c(m_b)\cdot\alpha}{(1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha}]$. The firm's profit is the same for all equilibria because the advertiser is indifferent to target consumers with s_B.

5. Boundary between Semi-g and Separating. Consider $p \geq \widetilde{\mu}(\sigma^s, m_b)$ and $p = \frac{k}{\Pr(m_g|s_B)}$. The limit of the separating equilibrium is $\sigma^s = (\sigma(s_G) = 1, \sigma(s_B) = 0)$, and $\delta_c^s = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 3. The limit of the semi-separating equilibrium is $\sigma(s_G) = 1$, and $\sigma(s_B) = \frac{p(1-\mu_c(m_g))(1-\alpha)-(1-p)\cdot\mu_c(m_g)\cdot\alpha}{(1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha}$, and $\delta_c^{\text{semi-g}} = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 5. Note that $\sigma^{\text{semi-g}}(s_B)$ must be greater than 0 for consumers with a good impression to be indifferent under semi-g. Thus, we have that $\sigma = (\sigma(s_G) = 1, \sigma(s_B))$ and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ is an equilibrium for any $\sigma(s_B) \in [0, \frac{p(1-\mu_c(m_g))(1-\alpha)-(1-p)\cdot\mu_c(m_g)\cdot\alpha}{(1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha}]$. The firm's profit is the same for all equilibria because the advertiser is indifferent to target consumers with s_B.

6. Boundary between Semi-b and Separating. Consider $p = \widetilde{\mu}(\sigma^s, m_b)$ and $k/p \in [\Pr(m_g|s_B), \Pr(m_g|s_G)]$. The limit of the separating equilibrium is $\sigma^s = (\sigma(s_G) = 1, \sigma(s_B) = 0)$, and $\delta_c^s = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 3. The limit of the semi-b equilibrium is $\sigma^{\text{semi-b}} = (\sigma(s_G) = 1, \sigma(s_B) = 0)$ and $\delta_c^{\text{semi-b}} = (\delta_c(m_g) = 1, \delta_c(m_b) = \frac{k/p-\Pr(m_g|s_B)}{\Pr(m_b|s_B)})$ from Proposition 5. Note that $\delta_c^{\text{semi-b}} \geq 0$ because $\Pr(m_g|s_B) \leq k/p \leq 1$. Thus, we have that $\sigma = (\sigma(s_G) = 1, \sigma(s_B) = 0)$ and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b))$ is an equilibrium for any $\delta_c(m_b) \in [0, \frac{k/p-\Pr(m_g|s_B)}{\Pr(m_b|s_B)}]$. The firm's profit strictly increases in $\delta_c(m_b)$. Also note that for $k/p \in (\Pr(m_g|s_B), \Pr(m_g|s_G)]$, there is a discontinuous drop in profit as we move from the semi-separating equilibrium to the separating equilibrium, as $\delta_c(m_b)$ drops discontinuously to 0.

7. Boundary between pooling-all and Pooling-g. Consider $p = \mu_c(m_b)$ and $k \leq p \cdot \Pr(m_g|s_B)$. The limit of the pooling-all equilibrium is $\sigma^{\text{pool}-all} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$, and $\delta_c^{\text{pool}-all} = (\delta_c(m_g) = 1, \delta_c(m_b) = 1)$ from Proposition 4. The limit of the pooling-g equilibrium is $\sigma^{\text{pool-}m_g} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$, and $\delta_c^{\text{pool-}m_g} = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$ from Proposition 4. Note that when $p = \mu_c(m_b)$ and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b) = 0)$, consumers with a bad impression are indifferent between buying and not buying. Thus, $\sigma = (\sigma(s_G) = 1, \sigma(s_B) = 1)$, and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b))$ is an equilibrium for any $\delta_c(m_b) \in [0, 1]$. Profit increases in $\delta_c(m_b)$

and is maximized at the pooling-all equilibrium. There is a discontinuous drop in profit as p increases above $\mu_c(m_b)$.

8. Boundary between pooling-all and Semi-b. Consider $p = \mu_c(m_b)$ and $k \in [p \cdot \Pr(m_g|s_B), p]$. The limit of the pooling-all equilibrium is $\sigma^{\text{pool}-all} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$, and $\delta_c^{\text{pool}-all} = (\delta_c(m_g) = 1, \delta_c(m_b) = 1)$ from Proposition 4. The limit of the semi-separating equilibrium is $\sigma^{\text{semi-b}} = (\sigma(s_G) = 1, \sigma(s_B) = 1)$ and $\delta_c^{\text{semi-b}} = (\delta_c(m_g) = 1, \delta_c(m_b) = \frac{k/p - \Pr(m_g|s_B)}{\Pr(m_b|S_B)})$ from Proposition 5. Note that $\delta_c(m_b) = \frac{k/p - \Pr(m_g|s_B)}{\Pr(m_b|S_B)} < 1$ for $k < p$. Thus, $\sigma = (\sigma(s_G) = 1, \sigma(s_B) = 1)$, and $\delta_c = (\delta_c(m_g) = 1, \delta_c(m_b))$ is an equilibrium for any $\delta_c(m_b) \in [\frac{k/p - \Pr(m_g|s_B)}{\Pr(m_b|S_B)}, 1]$. Profit increases in $\delta_c(m_b)$ and is maximized at the pooling-all equilibrium. There is a discontinuous drop in profit as p increases above $\mu_c(m_b)$.

9. Upper Boundary of Semi-g Consider $p = \widetilde{\mu}(\sigma^s, m_g)$ and $k \leq p \cdot \Pr(m_g|s_B)$. The limit of the separating equilibrium is $\sigma^s = (\sigma(s_G) = 1, \sigma(s_B) = 0)$, and $\delta_c^s = (\delta_c(m_g) = \frac{k}{p} \cdot \frac{1}{\Pr(m_g|s_B)}, \delta_c(m_b) = 0)$ from Proposition 5.

Combining the above analyses (cases 1-9) proves the lemma. □

Now, we prove Proposition 8 using those two lemmas.

First consider the choices between $p = \mu_c(m_b)$ and $p = \mu_c(m_g)$ in the range of $k \leq \min\{\mu_c(m_g) \cdot \Pr(m_g|s_B), \mu_c(m_b)\}$. Setting $p = \mu_c(m_b)$ is optimal if and only if: $\mu_c(m_b) - k \geq \Pr(m_g) \cdot \mu_c(m_g) - k$, or equivalently, $\frac{\mu(1-\beta)}{\mu(1-\beta)+(1-\mu)\beta} \geq \mu\beta$. Note that the right-hand side increases in β and the left-hand side decreases in β. Also note that the right-hand side is larger at $\beta = 1$ and the left-hand side is larger at $\beta = 1/2$. Thus, there exists $\overline{\beta} \in (1/2, 1)$ such that $p = \mu_c(m_b)$ is optimal if and only if $\beta \leq \overline{\beta}$. The closed-form solution for $\overline{\beta}$ is $\overline{\beta} = \frac{\sqrt{\mu^2 - 6\mu + 5} - \mu - 1}{2(1-2\mu)}$.

For a given β, there are four potential candidates for the optimal price:

Candidate 1 (when $\beta \leq \overline{\beta}$): $p = \mu_c(m_b)$ under the pooling-all equilibrium. It can only be optimal for $k < \mu_c(m_b)$. The firm's profit in this case is: $\Pi_1(k) = \Pi^{\text{pool-all}} = \mu_c(m_b) - k$.

Candidate 1 (when $\beta > \overline{\beta}$): $p = \mu_c(m_g)$ under the pooling-g equilibrium. It can only be optimal for $k < \mu_c(m_g) \cdot \Pr(m_g|s_B)$. The firm's profit in this case is: $\Pi_1(k) = \Pi_1(k) = \Pi^{\text{pool-}m_g}(p) = \Pr(m_g) \cdot \mu_c(m_g) - k$.

Candidate 2: any $p \in \left[\frac{k}{\Pr(m_g|s_B)}, \widetilde{\mu}(\sigma^s, m_g)\right]$ under the semi-g equilibrium. It can only be optimal for $\mu_c(m_g) \cdot \Pr(m_g|s_B) \leq k \leq \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B)$. The firm's profit in this case is:

$\Pi_2(k) = \Pi^{\text{semi-g}}(p) = \Pr(s_G) \cdot k \cdot \big(\frac{\Pr(m_g|s_G)}{\Pr(m_g|s_B)} - 1\big)$.

Candidate 3: $p = \widetilde{\mu}(\sigma^s, m_g)$ under the separating equilibrium. It can only be optimal for $\widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B) \leq k \leq \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_G)$. The firm's profit in this case is: $\Pi_3(k) = \Pi^s(p) = \Pr(m_g|s_G) \cdot \Pr(s_G) \cdot \widetilde{\mu}(\sigma^s, m_g) - \Pr(s_G) \cdot k$.

Candidate 4: $p = \widetilde{\mu}(\sigma^s, m_b)$ under the semi-b equilibrium. It can only be optimal for for $\widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_B) \leq k \leq \widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_G)$. The firm's profit in this case is: $\Pi_4(k) = \Pi^{\text{semi-b}}(p) = \Pr(s_G)\Big[\Pr(m_g|s_G) - \Pr(m_b|s_G) \cdot \frac{\Pr(m_g|s_B)}{\Pr(m_b|s_B)}\Big](\widetilde{\mu}(\sigma^s, m_b) - k)$.

Now, we show that given the model primitives μ, α and β, only one of $p = \mu_c(m_g)$ (under the pooling-g equilibrium) and $p = \mu_c(m_b)$ (under the pooling-all) can be the optimal price in the entire interval for $k \in [0, 1]$. First, if $\mu_c(m_b) < \mu_c(m_g) \cdot \Pr(m_g|s_B)$, then the region where the pooling-all equilibrium can exist is a subset of the region for the pooling-g equilibrium. Moreover, it implies that the profit is greater under the pooling-g equilibrium between the two pooling equilibria because $\mu_c(m_g) \cdot \Pr(m_g|s_B) \leq \mu_c(m_g) \cdot \Pr(m_g)$. Therefore, pooling-all equilibrium is dominated by pooling-g equilibrium in its entire region of existence.

Second, if $\mu_c(m_b) \geq \mu_c(m_g) \cdot \Pr(m_g|s_B)$, then there are two sub-cases. Suppose $\mu_c(m_b) \geq \mu_c(m_g) \cdot \Pr(m_g)$. Then, pooling-g equilibrium is dominated by pooling-all equilibrium in its entire region. The only remaining case is $\mu_c(m_g) \cdot \Pr(m_g|s_B) \leq \mu_c(m_b) \leq \mu_c(m_g) \cdot \Pr(m_g B)$. In this case, for $k \leq \mu_c(m_g) \cdot \Pr(m_g|s_B)$, where both pooling equilibria can exist, pooling-g equilibrium dominates. Between the two equilibria, for $k > \mu_c(m_g) \cdot \Pr(m_g|s_B)$, only the pooling-all equilibrium can exist. However, in this region, semi-g equilibrium (Candidate 2 defined above) dominates pooling-all equilibrium. This is because the profit under the pooling-g and semi-g equilibrium coincide precisely at $k = \mu_c(m_g) \cdot \Pr(m_g|s_B)$ and $p = \mu_c(m_g)$. Moreover, note that the profit of pooling-g equilibrium decreases in k at a higher rate than the profit of semi-g equilibrium. Therefore, for $k \geq \mu_c(m_g) \cdot \Pr(m_g|s_B)$, the profit of semi-g equilibrium is greater than the profit under the pooling-all equilibrium. This proves that only one of $p = \mu_c(m_g)$ (under the pooling-g equilibrium) and $p = \mu_c(m_b)$ (under the pooling-all) can be the optimal price. Next, note that regardless of β, we have $\frac{d\Pi_1}{dk} < \max\{\frac{d\Pi_2}{dk}, \frac{d\Pi_3}{dk}, \frac{d\Pi_4}{dk}\}$. This proves Proposition 8-1. Also, we can have $\overline{k}_1 < 0$, so that the statement in Proposition 8-1 becomes trivial.

Consider $k \in [\widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_B), \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B)]$. This is the overlapping region between candidate 2 and candidate 4. Note that $\Pi_2(k = \widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_B)) = \Pi_4(k = \widetilde{\mu}(\sigma^s, m_b) \cdot$

$\Pr(m_g|s_B))$ if $\widetilde{\mu}(\sigma^s, m_b) \geq \mu_c(m_g)$, and $\Pi_2(k = \widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_B)) > \Pi_4(k = \widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_B))$ if $\widetilde{\mu}(\sigma^s, m_b) < \mu_c(m_g)$. Thus, $\Pi_2(k = \widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_B)) \geq \Pi_4(k = \widetilde{\mu}(\sigma^s, m_b) \cdot \Pr(m_g|s_B))$ in this range. In addition, we have $\frac{d\Pi_2}{dk} > 0 > \frac{d\Pi_4}{dk}$. Thus, either candidate 1 or candidate 2 is optimal in this range. Let $\overline{k}_2 = \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B)$ if $\overline{k}_1 < \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B)$, and let $\overline{k}_2 = \overline{k}_1$ if $\overline{k}_1 \geq \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B)$. This proves Proposition 8-2. Also, if $\overline{k}_1 = \overline{k}_2$, the statement in Proposition 8-2 becomes trivial.

Finally, consider the choice between candidate 3 and candidate 4. Note that $\frac{d\Pi_3}{dk} < \frac{d\Pi_4}{dk}$ because $\left[\Pr(m_g|s_G) - \Pr(m_b|s_G) \cdot \frac{\Pr(m_g|s_B)}{\Pr(m_b|s_B)}\right] < 1$. Thus, Π_3 and Π_4 have a single crossing. Note that from the above we have $\Pi_3(k = \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B)) = \Pi_2(k = \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B)) > \widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B))$, and thus the crossing happens at some k above $\widetilde{\mu}(\sigma^s, m_g) \cdot \Pr(m_g|s_B))$. Letting $\overline{k}_3$ denote this point of crossing proves Proposition 8-3 and 4. Also, if $\overline{k}_3 > \widetilde{\mu}(\sigma^s, m_b)$, the statement in 8-4 is trivial. There is no advertising under any price, so all prices are weakly optimal. ■

Proof of Proposition 9

In semi-b equilibrium, plugging the optimal price $p^{*\text{semi-b}} = \widetilde{\mu}(\sigma^s, b) = \frac{\mu\cdot\alpha(1-\beta)}{\mu\cdot\alpha(1-\beta)+(1-\mu)(1-\alpha)\beta}$ into $\sigma^{\text{semi-b}}(s_B) = \frac{p(1-\mu_c(m_b))(1-\alpha)-(1-p)\cdot\mu_c(m_b)\cdot\alpha}{(1-p)\mu_c(m_b)(1-\alpha)-p(1-\mu_c(m_b))\alpha} = 0$. In semi-g equilibrium, it is straightforward to show that the upper bound $\widetilde{\mu}^U = \widetilde{\mu}(\sigma^s, m_g)$ increases in α. Moreover, the lower bound weakly increases in α because $k/\Pr(m_g|s_B)$ increases in α, whereas $\mu_c(m_g)$ is constant in α. Lastly, differentiating $\sigma(s_B)$ with respect to p, we get $-\frac{(2\alpha-1)\mu_c(m_g)(1-\mu_c(m_g))}{\left((1-p)\mu_c(m_g)(1-\alpha)-p(1-\mu_c(m_g))\alpha\right)^2} < 0$.

Proof of Corollary 3

Consider the limiting case $\alpha \to 1$ of the four optimal price candidates in the proof of Proposition 8.

Limit of Candidate 1 (when $\beta \leq \overline{\beta}$): $p = \mu_c(m_b)$. It can only be optimal for $k < \mu_c(m_b)$. Let $\widehat{\Pi}_1(k)$ denote the firm's profit in this case, we have: $\widehat{\Pi}_1(k) = \mu_c(m_b) - k$.

Limit of Candidate 1 (when $\beta > \overline{\beta}$): $p = \mu_c(m_g)$. It can only be optimal for $k < (1-\beta)\mu_c(m_g)$. Let $\Pi_1(k)$ denote the firm's profit in this case, we have: $\widehat{\Pi}_1(k) = \Pr(m_g) \cdot \mu_c(m_g) - k$.

Limit of Candidate 2: any $p \in \left[\max\left\{\mu_c(m_g), \frac{k}{1-\beta}\right\}, 1\right)$. It can only be optimal for $k \leq 1-\beta$. Let $\widehat{\Pi}_2(k)$ denote the firm's profit in this case, we have: $\widehat{\Pi}_2(k) = \mu \cdot k \cdot \frac{2\beta-1}{1-\beta}$.

Limit of Candidate 3: $p = 1$. It can only be optimal for $k \geq 1-\beta$. Let $\widehat{\Pi}_3(k)$ denote the firm's profit in this case, we have: $\widehat{\Pi}_3(k) = \mu \cdot (\beta - k)$.

Limit of Candidate 4: $p = 1$. It can only be optimal for for $k \geq 1-\beta$. Let $\widehat{\Pi}_4(k)$ denote the firm's profit in this case, we have $\widehat{\Pi}_4(k) = \mu \cdot \frac{2\beta-1}{\beta} \cdot (1-k)$.

Note that the limit of Candidate 3 and that of Candidate 4 have the same range ($k \geq 1-\beta$), and one can confirm that $\widehat{\Pi}_4(k) \geq \widehat{\Pi}_3(k)$ for $k \geq 1-\beta$, so we can ignore candidate 3 in the limit.

Because $\frac{d\widehat{\Pi}_1}{dk} < \frac{d\widehat{\Pi}_2}{dk}$ and $\frac{d\widehat{\Pi}_1}{dk} < \frac{d\widehat{\Pi}_3}{dk}$, if candidate 1 is optimal for k, then candidate 1 must be optimal for all $k' < k$. This proves the existence of $\overline{k}$.

If $\beta > \overline{\beta}$, then $\overline{k} = (1-\beta) \cdot \mu_c(m_g)$. For $\beta \leq \overline{\beta}$, to get the value of $\overline{k}$, we find the the value of k where $\widehat{\Pi}_1(k)$ crosses $\widehat{\Pi}_2(k)$ for $k \leq 1-\beta$ or $\widehat{\Pi}_3(k)$ for $k \geq 1-\beta$, which is $\overline{k} = \max\left\{\frac{(1-\beta)\cdot\mu_c(m_b)}{(1-\beta)+\mu\cdot(2\beta-1)}, \frac{\beta\cdot\mu_c(m_b)-\mu\cdot(2\beta-1)}{\beta-\mu\cdot(2\beta-1)}\right\}$. ■

Proof of Proposition 10

Consumers opt out if and only if consumer surplus becomes strictly higher by doing so. From Lemma 5, the optimal price under opt-out can only be $p = \mu_c(m_g)$ or $p = \mu_c(m_b)$ and consumer surplus is 0 under $p = \mu_c(m_g)$. From Corollary 3, consumer surplus is positive only when price is $p = \mu_c(m_b)$. Thus, combining Corollary 3 and Lemma 5 produces the range of k in the first part of the proposition. The second part is also trivial. When $\beta \to 1$, $\mu_c(m_b) = \frac{(1-\beta)\mu}{(1-\beta)\mu+\beta(1-\mu)} \to 0$. Thus, the consumer opt-out region disappears. ■

Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

References

Acquisti, A. and Varian, H. R. (2005). Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381.

Agarwal, N., Athey, S., and Yang, D. (2009). Skewed bidding in pay-per-action auctions for online advertising. *American Economic Review*, 99(2):441–47.

Agrawal, A., Gans, J., and Goldfarb, A. (2018). *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.

Anand, B. N. and Shachar, R. (2009). Targeted advertising as a signal. *Quantitative Marketing and Economics*, 7(3):237–266.

Anderson, S. P. and Renault, R. (2006). Advertising content. *American Economic Review*, 96(1):93–113.

Armstrong, M. and Zhou, J. (2022). Consumer information and the limits to competition. *American Economic Review*, 112(2):534–77.

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.

Bleier, A. and Eisenbeiss, M. (2015). Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science*, 34(5):669–688.

Chakraborty, A. and Harbaugh, R. (2010). Persuasion by cheap talk. *American Economic Review*, 100(5):2361–82.

Chen, J., Liu, D., and Whinston, A. B. (2009). Auctioning keywords in online search. *Journal of Marketing*, 73(4):125–141.

De Corniere, A. and De Nijs, R. (2016). Online advertising and privacy. *The RAND Journal of Economics*, 47(1):48–72.

Dukes, A. and Liu, L. (2015). Online shopping intermediaries: The strategic design of search environments. *Management Science*, 62(4):1064–1077.

Dzyabura, D. and Hauser, J. R. (2019). Recommending products when consumers learn their preference weights. *Marketing Science*, 38(3):417–441.

Goldfarb, A. and Tucker, C. (2011). Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404.

Ichihashi, S. (2020). Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–95.

Iyer, G., Soberman, D., and Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3):461–476.

Joshi, A., Bagherjeiran, A., and Ratnaparkhi, A. (2011). User demographic and behavioral targeting for content match advertising. In *Proceedings of the Fifth International Workshop on Data Mining and Audience Intelligence for Advertising (ADKDD 2011)*, pages 53–60. Citeseer.

Jullien, B. and Pavan, A. (2019). Information management and pricing in platform markets. *The Review of Economic Studies*, 86(4):1666–1703.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Li, X. and Xu, Z. (2022). Superior knowledge, price discrimination, and customer inspection. *Marketing Science*, forthcoming.

Mayzlin, D. and Shin, J. (2011). Uninformative advertising as an invitation to search. *Marketing Science*, 30(4):666–685.

Montes, R., Sand-Zantman, W., and Valletti, T. (2019). The value of personal information in online markets with endogenous privacy. *Management Science*, 65(3):1342–1362.

Rafieian, O. and Yoganarasimhan, H. (2021). Targeting and privacy in mobile advertising. *Marketing Science*, 40(2):193–218.

Shen, Q. and Villas-Boas, M. J. (2018). Behavior-based advertising. *Management Science*, 64(5):2047–2064.

Shin, J. (2005). The role of selling costs in signaling price image. *Journal of Marketing Research*, 42(3):302–312.

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.

Yoganarasimhan, H. (2020). Search personalization using machine learning. *Management Science*, 66(3):1045–1070.

Zhong, Z. Z. (2022). Platform search design: The roles of precision and price. *Marketing Science*, forthcoming.