# Consumer Purchase Journey, Targeted Advertising, and Privacy Choices

W. Jason Choi    Kinshuk Jerath    Miklos Sarvary*

June 2021

*W. Jason Choi (jason.choi@rutgers.edu) is Assistant Professor of Marketing at Rutgers Business School, Kinshuk Jerath (jerath@columbia.edu) is Professor of Marketing at Columbia Business School, and Miklos Sarvary (miklos.sarvary@columbia.edu) is the Carson Family Professor of Business at Columbia Business School.
The authors would like to thank Jacob Goldenberg, Oded Netzer, Martin Peitz, Amin Sayedi, Andrey Simonov, Hema Yoganarasimhan, and seminar participants at Columbia University, Rutgers University, University of Illinois Urbana-Champaign, Chinese University of Hong Kong, Singapore Management University, National University of Singapore, University of Texas Dallas, University of Illinois Chicago, Hong Kong University, University of Hawaii, and the 12th Workshop on the Economics of Advertising & Marketing for their valuable comments.

# Consumer Purchase Journey, Targeted Advertising, and Privacy Choices


## Abstract

Advertising to a consumer moves her along the purchase journey, and tracking the consumer's online activities enables the advertiser to target ads based on inferences about the consumer's purchase journey state. While targeted ads provide information that benefits consumers, many consumers dislike being tracked and, furthermore, repeated advertising may lead to ad wearout. We develop a theoretical framework to investigate, under the above considerations, the impact of regulations that endow consumers with the choice to opt in to or out of having their online activities tracked. A consumer's opt-in decision balances two aspects of privacy—its intrinsic value (value of privacy for its own sake) and its instrumental value (indirect impact of allowing tracking on utility due to the actions of other agents; in this case the advertiser and the ad network). We identify forces that can have countervailing effects on different components of a consumer's instrumental privacy value — ads provide product match information that benefits the consumer but repeated ads cause ad wearout that reduces her utility — and determine conditions under which consumers opt in to or out of tracking. We also determine conditions under which regulation would be necessary to induce firms to take actions to impact consumers' intrinsic privacy values (e.g., improving data security). Finally, we show that when advertisers set product prices endogenously, endowing consumers with privacy choices can lower consumers' utilities compared to the setting where all consumers are opted in by default.


**Keywords**: consumer purchase journey, ad wearout, consumer tracking, privacy regulations, intrinsic and instrumental privacy.

Before purchasing a product, consumers typically go through a "purchase journey." A canonical model of this is the Awareness-Interest-Desire-Action (AIDA) "purchase funnel" model; i.e., consumers first need to become Aware of a product, then their Interest needs to be generated, then they need to have Desire to consume the product, and finally they may take Action and purchase the product (Lavidge and Steiner, 1961; Kotler and Keller, 2012). Firms use advertising (in addition to other marketing communications) to build awareness of their products and to move consumers along the successive states of the purchase funnel (Kotler and Keller, 2012; Abhishek et al., 2017, 2018; Todri et al., 2020).

With advancements in data collection and machine learning, tracking consumers' online activities has enabled advertisers to make real-time inferences about consumers' purchase journey states. Firms can then target and retarget ads to consumers based on their inferred purchase journey states (Lambrecht and Tucker, 2013; Hoban and Bucklin, 2015; Johnson et al., 2017; Seiler and Yao, 2017; Sahni et al., 2019). Industry practitioners also advocate targeting based on consumers' "states in the decision journey" (Edelman, 2010), and providers such as Google allow advertisers to specify retargeting audiences based on such inferences (Google, 2020a).

However, imperfect targeting can lead to ad repetition from the same firm which elicits negative consumer response. This phenomenon is known as "ad wearout" or "ad annoyance." Aaker and Bruzzone (1985), Chen et al. (2016), Chae et al. (2019) and Todri et al. (2020) show that consumers' responses to ads become negative with repetition. Campbell and Keller (2003) find that consumer attitude toward an ad can become negative with repetition because "at higher levels of ad repetition, consumers may ... consider the inappropriateness of advertising tactics."

Furthermore, widespread adoption of consumer tracking has deepened consumers' concerns about their online privacy (McDonald and Cranor, 2010). For instance, 77% of the US Internet users indicate that they are "concerned about how tech/social media companies are using [their] online data ... for commercial purposes" (eMarketer, 2019a) and 68% of the US Internet users report feeling concerned about "companies displaying ads based on their data" (eMarketer, 2018). In response to the growing outcry from consumers and privacy advocates, advertising organizations and regulators have sought to curb privacy-infringing practices, such as online tracking. In May 2018, the European Union (EU) implemented the General Data Protection Regulation (GDPR) and, in January 2020, the State of California implemented the California Consumer Privacy Act (CCPA), following it up with the passing of the California Privacy Rights Act (CPRA) in November 2020. These regulations require that firms inform consumers what data will be collected and for what purposes, and that firms obtain explicit consent from consumers to use their data.[1,2] If consumers opt out of tracking, then advertisers cannot monitor consumers' online behavior. Consequently, advertisers' targeting capabilities are undermined and ad impressions could be wasted (e.g., repeated exposure to consumers who had already purchased) (Srinivasan, 2019). On the other hand, if consumers opt in to tracking, advertisers can target ads more effectively to specific audiences.

The impact of privacy regulations on the advertising industry is a topic of ongoing debate among practitioners, academics, and policymakers. On one hand, regulations are expected to limit advertisers' tracking capability, thereby reducing ad effectiveness (Goldfarb and Tucker, 2011; Aziz and Telang, 2016). On the other hand, there is evidence suggesting that despite consumers' stated aversion towards tracking, consumers appear not as reluctant to allow tracking in practice (known as the "privacy paradox;" (Norberg et al., 2007; Athey et al., 2017)). For example, Johnson et al. (2020) find that less than 0.26% of the US and EU consumers opt out of behavioral targeting in the AdChoices program started by Google in 2010, under which consumers could opt out of online behavioral advertising. de Matos and Adjerid (2020) find that, at a large European

[1]Before these regulations were in effect, it was difficult, if not impossible, for consumers to prevent being tracked online. For instance, firms used flash cookies to re-spawn deleted cookies (Stern, 2018; Angwin, 2010). The Ad Choices program, which enabled consumers to opt out of personalized ads (but not out of being tracked) was (and continues to be) riddled with loopholes (Sankin, 2021). Moreover, firms purchased personal data from third-party information vendors without consumers' consent. Such activities are regulated under the GDPR and the CCPA/CPRA laws.

[2]We note that opting out of tracking is different from opting out of personalized ads, both of which are different from ad blocking. For instance, the Ad Choices program enabled (to some extent) opting out of personalized ads but did not enable opting out of being tracked. Regarding ad blocking, first, ad blocking does not prevent tracking; i.e., it does not prevent data collection and usage by firms. Second, most ad blockers have a business model under which they act as gatekeepers to the consumers who use ad blockers, and allow certain advertisers who pay the ad blockers to show ads to consumers. In this case, consumers may indeed see personalized ads. In this paper, we focus on opting out of being tracked; if a consumer opts out of being tracked, this only implies that their data cannot be collected and used for targeting; however, this consumer still sees ads, though likely less personalized.

telecommunications provider, the opt-in rate increased from 53% to 64% once GDPR-compliant consent was elicited. These findings suggest that privacy regulations that endow consumers with privacy choices may not necessarily lower opt-in rates.

In this paper, we develop a game theory model to shed light on the above debate about the effects of privacy regulations. We investigate how endowing consumers with the choice of allowing firms to track their online activities, and (imperfectly) infer their purchase funnel states for ad targeting, affects firms' advertising strategies and consumers' opt-in decisions. We consider a two-period model in which consumers visit web pages, and each consumer creates one opportunity for an ad impression per period. Advertisers bid for ad impressions supplied by the ad network that sells ad inventory available at the web pages. A consumer's progression down the purchase funnel is driven by the product information she possesses. Advertising conveys product match information (e.g., Anderson and Renault, 2006) that transitions an uninformed consumer from her current funnel state to the next. Each consumer, based on her preferences for ad exposure, for the advertised product, and for privacy, chooses whether to allow the ad network and advertiser to track her online behavior.

We model privacy as a multi-dimensional construct consisting of intrinsic and instrumental values (Becker, 1980; Posner, 1981; Wathieu and Friedman, 2009; Farrell, 2012). The intrinsic value of privacy refers to the utility a consumer derives from protecting privacy for its own sake. Simply the act of sharing data may reduce the consumer's utility; this may be, for instance, due to the threat of a data breach that compromises the consumer's privacy.

The instrumental value of privacy stems from the indirect effects of the consumer's privacy choice on her utility. We focus on the indirect effects related to the consumer's online ad experiences: ads provide product match information that benefits the consumer, but if shown repetitively, cause ad wearout that reduces her utility. A consumer's purchase journey information influences the advertiser's ad strategy, which in turn impacts the consumer's utility through the effects on (a) ad repetition intensity and (b) the consumer's product match knowledge. Therefore, when making her opt-in choice, a consumer balances the two instrumental components — the informational value of ads and ad repetition wearout — and her intrinsic value of privacy.

Using this model, we address the following research questions: What are the implications of purchase journey tracking on the consumers' privacy utilities? How do key economic forces interact to shape the consumers' opt-in choices? How does consumer tracking impact ad revenue? And does endowing consumers with privacy choices always increase consumer welfare?

Our analysis yields a number of interesting insights, some of which we highlight here. First, we find that a consumer's choice to allow her purchase journey to be tracked or not has significant implications for her instrumental value of privacy. This effect is through the influence of consumer trackability on the advertiser's and the ad network's actions. Under certain conditions, opting in to tracking generates countervailing changes in the components of the consumer's instrumental value — fewer repeated ads reduce the disutility from ad wearout but this also reduces the match information that the consumer obtains, and vice versa.

Second, consumers may opt in to or out of tracking under different conditions. For instance, if the effectiveness with which ads convey product information is intermediate, consumers benefit from opting in because tracking reduces ad repetition. Interestingly, the precision with which consumers are tracked moderates this opt-in incentive. For instance, if consumers are sensitive to ad wearout, fewer consumers may opt in as their funnel state information is tracked with higher precision. The reason is that as tracking precision increases, consumers to whom product information was not successfully conveyed by early ads are more likely to be identified and targeted for repeat advertising.

Third, we uncover a novel mechanism through which the total ad revenue decreases in tracking precision. On one hand, more precise tracking of consumer information increases the efficiency with which consumers are matched with ads, thereby increasing the advertiser's willingness to pay for ads. On the other hand, as explained above, improvements in tracking precision may induce advertiser actions that decrease the consumers' instrumental values of privacy (e.g., higher intensity of repeated ads). This motivates more consumers to opt out as tracking precision increases which, in turn, reduces the total ad revenue as lack of consumer information dampens advertising efficiency. We delineate the conditions under which the latter negative effect dominates the former positive effect such that ad revenue decreases in tracking precision.

While in the main model we assume product price to be exogenous, in an extension, we allow product price to be set endogenously by the advertiser, in the presence of asymmetric information between the advertiser and the ad network. We find that the insights from the main model continue to hold. Interestingly, we also identify a welfare-diminishing effect of consumer privacy choices, driven by a high product price when it is set endogenously.

In another extension, we allow the advertiser and the ad network to take actions to impact the intrinsic value of privacy. For instance, the firms can improve data security, which would reduce the disutility felt by consumers on sharing data. We find that self-regulation may be feasible if the consumers' instrumental cost of giving up privacy is high, such that firms are incentivized to reduce the intrinsic cost of privacy on their own. On the other hand, if the consumers' instrumental benefit of giving up privacy is high, firms lack the incentive to take actions such as improving data security, and regulation may be needed to induce firms to take these actions.

Taken together, our paper provides a useful theoretical framework for analyzing the nuanced interaction between consumers, who exercise their privacy rights to balance the intrinsic and instrumental values of privacy, and advertisers, who seek to target ads to inform consumers about their products and ultimately induce them to purchase. We advance our understanding of consumers' privacy choices by elucidating their effects on the firms' strategies that indirectly impact consumers' utilities. We demonstrate that in a setting where consumers' purchase funnel states can be tracked by firms, multiple forces are at play and they can have countervailing effects on the consumers' utilities.

Another important contribution of our paper is that we assess the impact of privacy regulations on the advertising ecosystem. In particular, we show that compared to a benchmark in which all consumers are opted in by default, the endowment of consumer privacy choices can either increase or decrease the average intensity of ad repetition, depending on the effectiveness of ads. Furthermore, we find that privacy regulations may not always increase consumer welfare. If firms can adjust product prices endogenously, the endowment of privacy choices can elicit firms' price responses that hurt consumers.

In addition to the papers referenced above, our paper relates to the literature on targeted advertising and online privacy. Extant literature on targeted advertising studies various implications of targeting. For example, it examines the impact of targeting on ad supply, ad prices, ad strategies, ad intensity and adoption of ad avoidance tools (Esteban et al., 2001; Iyer et al., 2005; Athey and Gans, 2010; Bergemann and Bonatti, 2011; Johnson, 2013; Aziz and Telang, 2017; Shen and Villas-Boas, 2018). We extend the existing literature in a novel and important way by modeling the consumer purchase journey, which allows us to study funnel state-dependent ad effects. We show that modeling funnel considerations sheds light on an understudied link between cross-period ads in a setting where ad exposures endogenously create interim funnel-state heterogeneity. This has the effect of inducing heterogeneous consumers to make different ex ante decisions to opt in to or out of tracking. The value of modeling the purchase funnel is clear from the fact that, in our framework, not explicitly modeling the purchase funnel would lead to all consumers opting out of tracking.

We also contribute to the growing literature on online privacy. Research on price discrimination examines consumers' implicit privacy decisions, wherein consumers strategically time their purchase to control the disclosure of their preferences to the firm, thereby mitigating price discrimination (Taylor, 2004; Villas-Boas, 2004). Other papers investigate more explicit privacy decisions, wherein consumers take actions to control the information disclosed to firms (Acquisti and Varian, 2005; Conitzer et al., 2012; Montes et al., 2019; Ichihashi, 2020). de Cornière and de Nijs (2016) investigate the ad network's incentive to disclose consumer information to advertisers whereas, in our paper, the consumers control the flow of their personal information. The mechanisms driving our results are orthogonal to market thickness (Bergemann and Bonatti, 2011; Rafieian and Yoganarasimhan, 2020) and market structure (Campbell et al., 2015) as we abstract from advertiser competition in the main model.

D'Annunzio and Russo (2020) study a similar setting to ours where consumers can endogenously decide whether to be tracked or not. However, our paper is different in several important ways. First, we explicitly model consumer tracking along the purchase journey; i.e., advertisers track consumers' progressions through the purchase journey after a series of ad exposures, rather than tracking single- vs. multi-homing consumers across different publishers. Second, the consideration of consumers' transitions down the purchase funnel by virtue of previously shown ads gives rise to multi-period dynamics as advertisers consider retargeting some consumers due to interim heterogeneity, which we explicitly model; on the other hand, D'Annunzio and Russo (2020) consider a static setting. Third, while D'Annunzio and Russo (2020) assume that the number of ads a consumer is exposed to is independent of her privacy choice, we explicitly incorporate potential changes in advertising intensity, which substantially influences a consumer's instrumental value of privacy.

The rest of the paper is organized as follows. We first describe the main model and present the key results including the consumers' opt-in behaviors and the implications of endogenous privacy choice on the total ad revenue. We then analyze various extensions that assess the robustness of the main insights and also generate new insights. Finally, we summarize the key findings and conclude. Proofs for the main results are provided in the Appendix, while the proofs for other results and additional analyses are provided in the Web Appendix.

# *MODEL*

The game consists of three entities — consumers, an advertiser and an ad network — interacting over two time periods. In each time period, a consumer visits a web page and the ad network enables an advertiser to show an ad to the consumer. We describe the components of the model in detail below.

### *Consumers*

There is a unit mass of consumers, each of whom visits two web pages (both of which are covered by the ad network), one in each Period $t \in \{1, 2\}$.[3] A consumer is exposed to at most one ad for the advertiser's product per period from the page she visits. While consumers are aware of the advertiser's product, they are uncertain about product match. Product match is binary, and if a consumer realizes a product match, she derives positive utility on consumption which we normalize to 1; otherwise, she derives zero utility from the product. For all consumers, ex ante (i.e., before Period 1), the product matches with probability $\phi \in [0, 1]$ and does not match with probability $1 - \phi$; this common prior belief is shared by the consumers, the advertiser and the ad network. We label this initial state of consumers, in which they are uncertain about product match and have a prior $\phi$ for match, as the "top funnel state," or funnel state $T$.

Consumers can resolve match uncertainty based on information conveyed through advertising. A single ad conveys match information to consumers with a "success rate" of $\mu$: if consumers see an ad, they obtain match information with probability $\mu$, and do not obtain the information with probability $1 - \mu$.[4] Thus, $\mu$ can be interpreted as the effectiveness of ads in conveying product information.[5] Consumers to whom product match information is successfully conveyed transition to a state we label the "bottom funnel state," or funnel state $B$. If and when a consumer reaches funnel state $B$, the match is revealed to her (recall that she realizes a match with probability $\phi$ and does not realize a match with probability $1 - \phi$). We assume that while it is possible for a consumer to obtain match information on her own, this involves information search costs which are prohibitively high.

[3]We implicitly assume that consumers obtain a large positive utility per web page visit. This ensures that consumers visit two web pages in equilibrium, despite the potential negative ad effects, which will be discussed later in the text.

[4]We assume this to be an exogenous process, not in the consumer's control.

[5]Our results remain unaltered if we scale the ad effectiveness parameter $\mu$ to vary between smaller ranges. For instance, if we consider a sub-population of consumers who potentially respond to ads while the other consumers do not, then the unconditional effectiveness of ads shown to the whole population would be scaled down in proportion to the sub-population of responsive consumers. Therefore, ad effectiveness rates can be scaled to vary between, say, 0% and 5% (which are arguably closer to empirical estimates) without impacting our results. However, for expositional clarity, we use the formulation with effectiveness numbers varying between 0 and 1.

A consumer's expected utility consists of intrinsic and instrumental components:

$$u(x) = \underbrace{-\theta x}_{\text{intrinsic}} \underbrace{-\eta\,\mathbb{E}\left[q(x)\right] + \mathbb{E}\left[S(x)\right]}_{\text{instrumental}}. \tag{1}$$

In the intrinsic component of (1), $x$ denotes the consumer's privacy decision, which equals 1 if she opts in to tracking, and 0 if she opts out of tracking,[6] and $\theta$ denotes the consumer's intrinsic privacy cost; i.e., the disutility she incurs for allowing tracking. Privacy cost $\theta$ is distributed according to a cumulative distribution function $F$ on the support $\left[0, \overline{\theta}\right]$.[7]

In the instrumental component of (1), $\eta > 0$ is the ad wearout parameter (i.e., the disutility she incurs per repeated ad exposure), $q(x)$ is the total number of repeated ads she sees, and $S(x)$ is her product surplus. Given product price $p$, the consumer's expected product surplus is

$$\mathbb{E}[S(x)] = \begin{cases} \max\left[\mathbb{E}[v-p], 0\right] & \text{without match information,} \quad (2) \\ \mathbb{E}[\max\left[v-p, 0\right]] & \text{with match information,} \quad (3) \end{cases}$$

where the match utility $v$ equals 1 if she realizes a product match and 0 otherwise. In (2), the consumer does not know product match and thus makes the product purchase decision based on the prior distribution (i.e., match with probability $\phi$ and not match with probability $1-\phi$). In (3), the consumer makes her purchase decision after learning $v$. Note, by Jensen's inequality, (2)$\leq$(3); i.e., match information (weakly) increases the consumer's expected product surplus. This constitutes the informative value of advertising in our model and will be discussed further in the analysis. Depending on the consumer's privacy choice $x$, different numbers of ads may be shown to the consumer such that her product match knowledge, and hence $\mathbb{E}[S(x)]$, potentially changes with $x$. We assume that a consumer purchases at most one unit.

In sum, a consumer's opt-in decision depends on (i) her intrinsic valuation of privacy, and (ii) the intensity of repeated ads and product surplus she expects as a result of her privacy decision (see Figure 1).[8] These components constitute the intrinsic and instrumental aspects of privacy,

[6] We assume that (i) consumers make a one-time privacy decision, and (ii) privacy decision is binary (e.g., Conitzer et al., 2012; Montes et al., 2019). These assumptions keep the analysis simple without significantly changing qualitative insights.

[7] Heterogeneity in $\theta$ may stem from numerous factors, such as differences in consumers' perceptions of how secure their shared data is and differences in what consumers believe constitutes personal information (Acquisti et al., 2016).

[8] In an extension in the Web Appendix, we consider a scenario in which a fraction of consumers hold naïve beliefs

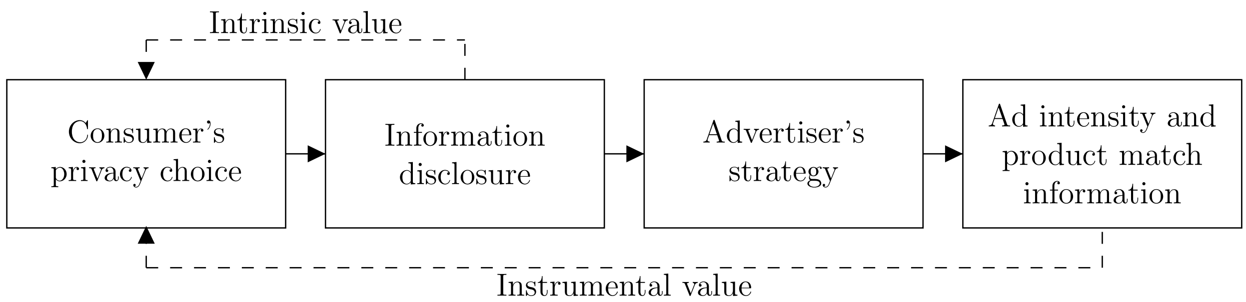


*Figure 1: Conceptual Diagram*

respectively.[9]

Before proceeding, we clarify two points about the utility function in (1). First, the unit weight of product surplus, relative to intrinsic cost and wearout disutility, is normalized to 1. Second, the instrumental components $q(\cdot)$ and $S(\cdot)$ depend not only on $x$, but also on ad effectiveness $\mu$, match probability $\phi$, and product price $p$. However, we suppress the latter three variables in denoting $q(\cdot)$ and $S(\cdot)$ to highlight the changes in the instrumental components with respect to a consumer's privacy choice $x$.

## *Advertiser*

Depending on whether consumers can be tracked or not, the advertiser can buy different types of ads. If tracking is not possible, then the advertiser can only buy untargeted impressions (e.g., ads displayed to all website visitors independent of their browsing histories). In particular, even if an ad is shown in Period 1 and redistributes consumers along the funnel states, the advertiser cannot target ads in Period 2 based on the consumers' inferred funnel states.

On the other hand, if tracking is possible, the ad network can sell and the advertiser can buy ad impressions based on inferences about the consumers' funnel states. The advertiser can specify the target audience such that ads are shown only to consumers who meet some pre-specified criteria that correlate with consumers' funnel states. While we assume purchase history to be

about the instrumental consequences of their privacy choices. We show that the trade-offs that naïve consumers face are qualitatively different from those that sophisticated consumers face.

[9] Consistent with the literature, the instrumental aspect of privacy captures the notion that an individual's utility depends on the actions of the entity that receives the individual's information (Becker, 1980; Posner, 1981; Farrell, 2012).

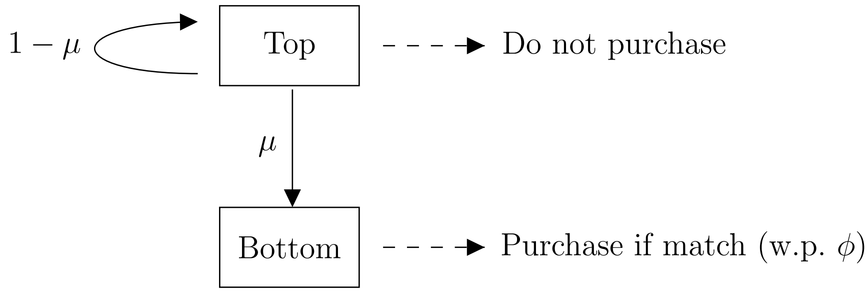


*Figure 2: Purchase Journey and Ad Effects*

perfectly observable for opt-in consumers, we assume the inference of their funnel states to be imperfect. This is because purchase is a specific, observable action taken by a consumer, whereas a consumer's funnel state is a latent, conceptual construct and there can be various confounding factors in inferring it. We define the accuracy of the funnel state signal $s \in \{T, B\}$ of an opt-in consumer who has not purchased as

$$\rho = \mathbb{P}\{s = f | f\} \in [1/2, 1] \,, \tag{4}$$

where $f \in \{T, B\}$ denotes the consumer's true funnel state. Based on the signal, a consumer will be a $T$-signal or a $B$-signal consumer; note that a $T$-signal consumer is actually a $T$-consumer with probability $\rho$ (and similarly for a $B$-signal consumer). We use terminologies "signal accuracy" and "tracking precision" interchangeably. Note that the funnel state signal is relevant only for consumers who have not yet purchased because purchases are observable and indicative of consumers being in state $B$. In each period, the advertiser decides, conditional on the purchase history and funnel state signal $s$, the bid amounts for each ad impression.

We assume that product price $p$ is exogenous and the advertiser's marginal production cost $c$ is in the interval $[\phi + \eta, 1]$.[10] This implies $p \geq \phi + \eta$, such that in funnel state $T$ (i.e., under product match uncertainty), consumers do not purchase the product.[11] For ease of exposition, we

[10] In an extension, we show that the qualitative insights from the main model are maintained under endogenous product pricing.

[11] To see this, the maximum expected utility a consumer obtains from purchasing the product under match uncertainty is $\mathbb{E}[v] - p = \phi - p$. If $p > \phi + \eta$, then $\phi - p$ is less than the minimum expected utility a consumer obtains from not purchasing under match uncertainty: $0 - \eta$ (zero utility from not purchasing in Period 1 and the maximum ad wearout disutility due to Period 2 ad repetition). Furthermore, $p \geq \eta$ ensures that opt-in consumers do not exploit purchase tracking by strategically purchasing the product in Period 1, despite not realizing a product match, to avoid seeing repeated ads in Period 2.

suppress the marginal cost notation; thus, whenever the advertiser's and ad network's profits are concerned, price $p$ should be interpreted as margin $p - c$, unless specified otherwise. Note that under $p \geq \phi + \eta$, consumers' expected product surplus in (2) and (3) simplifies to

$$\mathbb{E}[S(x)] = \begin{cases} 0 & \text{without match information,} \\ \phi(1-p) & \text{with match information.} \end{cases}$$

Figure 2 summarizes the consumer's funnel state-specific purchase decisions and the ad effects. The advertiser maximizes its net profit across two periods, which is the difference of its product sales margin and advertising costs.

### *Ad Network*

The ad network sells ad impressions to advertisers via second-price auctions.[12] It sets a reserve price in each Period $t \in \{1, 2\}$ for each of the different types of ad impressions sold (e.g., ads targeted to $B$-signal consumers). In our paper, this is equivalent to selling through a posted price; however, we choose the auction format in line with how the vast majority of display ads are sold.[13] The ad network maximizes its revenue from ad sales across two periods. We assume that if an ad slot is not sold to the advertiser, then it is sold to a mass of advertisers (e.g., via ad exchanges) at price $\underline{b}$.[14] We hereafter refer to the focal advertiser, who bids prior to the remnant inventory auction, simply as "the advertiser," and the focal advertiser's ad simply as "ad."

The timing of the game is as follows.

**Period 0:** Each consumer decides whether or not to opt in to being tracked.

**Period 1:** Ad network sets reserve prices for Period 1 ads for opt-in consumers and for opt-out consumers. Advertiser bids for ad impressions. If ads are shown, some consumers transition through the funnel, and product purchases are realized.

[12] Note that, due to the revenue equivalence principle, our results would not change if we considered first-price auctions, a mechanism to which some firms have recently transitioned (e.g., see Bigler (2019)).

[13] In 2019, over 83.5% of total display ad spend in US were transacted through real-time auctions (eMarketer, 2019b).

[14] The fixed clearing price for remnant ad inventory can be interpreted as an approximation to the asymptotic case with a large number of bidders with independent valuations distributed on a common support $\mathcal{S}$ and $\underline{b} = \max\{b : b \in \mathcal{S}\}$.

**Period 2:** Ad network sets reserve prices for targeted ads (based on the realized funnel state signal) for opt-in consumers, and untargeted ads for opt-out consumers. Advertiser bids for ad impressions. If ads are shown, some consumers transition through the funnel, and product purchases are realized.

We solve for the subgame-perfect equilibrium of the game. We assume that

$$\underline{b} \leq \mu\phi p \tag{5}$$

to rule out the degenerate case in which the advertiser's ad is not shown in either period (see Claim 1 in the Appendix). We also note that the publishers who own the web pages are treated as passive in the model. However, insofar as the ad network and the publishers share the same objective function of maximizing monetization by showing ads and they split these revenues on a proportional basis — which is typically the case (Jatain, 2019; Google, 2020b) — this is a reasonable assumption.

We consider various extensions of the main model to address simplifying assumptions that we make in the main model. Specifically, we analyze scenarios in which: (a) the advertiser is informationally advantaged over the ad network and sets product price endogenously; (b) ad wearout leads to reduced ad responsiveness; and (c) firms can take actions to influence the consumers' intrinsic costs of privacy. In the Web Appendix, we consider additional extensions where: (d) some consumers are naïve and hold erroneous beliefs about the impact of opting in to tracking; (e) multiple advertisers compete for the same ad slot; (f) consumers arrive in overlapping generations over an infinite horizon; (g) consumers are heterogeneous in their initial funnel states; (h) ad content can be personalized based on the inferred funnel states; and (i) the purchase funnel is degenerate. These extensions confirm the robustness of our main results and generate new insights.

# *ANALYSIS*

We first analyze the subgame for opt-out consumers and then the subgame for opt-in consumers. We describe the advertiser's and the ad network's strategies in each subgame and then characterize $\mathbb{E}[q(\cdot)]$, the expected intensity of repeated ads, and $\mathbb{E}[S(\cdot)]$, the expected product surplus. Recall from (1) that $\mathbb{E}[q(\cdot)]$ and $\mathbb{E}[S(\cdot)]$ are the two instrumental factors that determine, along with intrinsic privacy cost $\theta$, the consumers' privacy choices in Period 0. Per the assumption in (5), we restrict attention to Period 2 subgames in which the advertiser's ad is shown in Period 1 to all consumers. After solving the subgames for opt-out and opt-in consumers, we analyze consumers' opt-in choices.

## *Analysis for Opt-Out Consumers*

Consider the Period 2 subgame for opt-out consumers in which a Period 1 ad is shown. The Period 1 ad informs $\mu$ fraction of consumers about product match, thereby inducing them to move down to funnel state $B$, where $\phi$ fraction realize a product match and purchase. Consequently, the Period 1 ad creates interim heterogeneity as it redistributes consumers along the purchase funnel such that in Period 2, $1-\mu$ fraction are in state $T$ and $\mu(1-\phi)$ are in funnel state $B$ (while $\mu\phi$ fraction have already purchased).

The advertiser and the ad network cannot track opt-out consumers' funnel states nor their purchase histories. Therefore, they strategize based on prior distributions and expected ad effects. The advertiser's valuation, and hence its bid, for Period 2 ads shown to opt-out consumers is

$$b_0 = \mathbb{P}\{f = T\}\mu\phi p + \mathbb{P}\{f \neq T\} \cdot 0 = (1-\mu)\mu\phi p, \tag{6}$$

where $\mathbb{P}\{f = T\}$ is the probability that an opt-out consumer is in funnel state $T$ (and is equal to $1-\mu$), $\mu\phi$ is the probability of the ad inducing these $1-\mu$ fraction of $T$-consumers in Period 2 to realize product match and purchase, and $p$ is the advertiser's margin. Since $\mu(1-\phi)$ fraction of $B$-consumers and $\mu\phi$ fraction of consumers who have already purchased in Period 1 do not respond to advertising, the value of advertising to them is 0.

The ad network anticipates (6) and sets the reserve price for Period 2 ads accordingly. If $\underline{b} \leq$ (6), where $\underline{b}$ is the "outside bid," then the ad network's profit-maximizing reserve price is $b_0$; otherwise, it is $\underline{b}$. Therefore, the reserve price for Period 2 ads shown to opt-out consumers is $\max[\underline{b}, (1-\mu)\mu\phi p]$; the advertiser's Period 2 ad is shown if and only if $\underline{b} \leq (1-\mu)\mu\phi p$. In the following proposition, we summarize the implications of Period 2 advertising (with an ad shown to all consumers in Period 1) for the opt-out consumers' instrumental values.

**Proposition 1.** *Opt-out consumers' expected intensity of repeated ads and expected product surplus (across the two periods), respectively, are*

$$\mathbb{E}[q(0)] = \begin{cases} 1 & \text{if } \underline{\mu} < \mu \leq \overline{\mu}, \\ 0 & \text{otherwise}, \end{cases} \quad \text{and } \mathbb{E}[S(0)] = \mu\phi(1-p) \times \begin{cases} 2-\mu & \text{if } \underline{\mu} < \mu \leq \overline{\mu}, \\ 1 & \text{otherwise}, \end{cases} \tag{7}$$

*where*

$$\underline{\mu} = \left(1 - \sqrt{1 - 4\underline{b}/(\phi p)}\right)/2 \text{ and } \overline{\mu} = \left(1 + \sqrt{1 - 4\underline{b}/(\phi p)}\right)/2. \tag{8}$$

Proposition 1 shows that opt-out consumers see repeated ads if and only if ad effectiveness $\mu$ is intermediate (see Figure 3a). Intuitively, if $\mu$ is small, then the Period 2 ad is ineffective in informing consumers about product match and inducing purchase; therefore, the advertiser has low valuation for Period 2 ad such that its bid does not beat the reserve price. If $\mu$ is large, then the ad is also not shown because the first ad exposure *saturates*: a large fraction of consumers learn about product match and potentially purchase. Thus, a large $\mu$ diminishes the marginal value of Period 2 ads such that Period 2 ads are not shown. Overall, if $\mu$ is extreme, the ad is not shown. On the other hand, if $\mu$ is intermediate, then repeated ads are shown: Period 2 ads are effective enough in informing $T$-consumers and "pushing" them down to state $B$, but not too effective such that Period 1 ad saturates.

The implications for expected product surplus can be summarized by two effects. First, as illustrated by the dashed line in Figure 3b, the expected product surplus increases in $\mu$. This is because as ad effectiveness increases, ads are more likely to convey match information to consumers. The match information then allows consumers to reap the positive surplus upon match

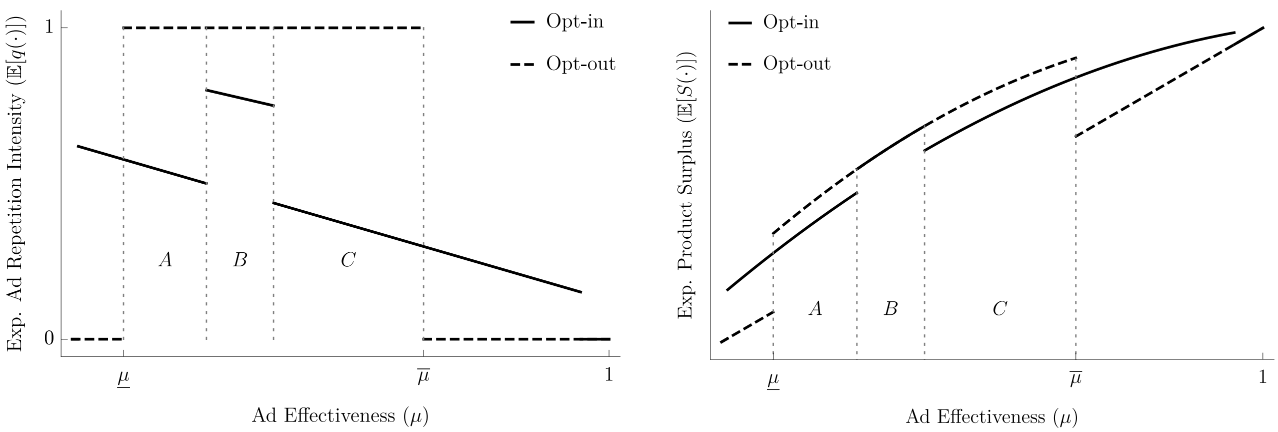


*(a) Ad Repetition Intensity* *(b) Consumer Product Surplus*

*Figure 3: Subgame Outcomes for Opt-Out and Opt-In Consumers*

and forego purchase upon mismatch (see (3)), which ultimately increases their expected product surplus. Second, the expected product surplus also changes discontinuously in $\mu$ due to repeat advertising. For intermediate $\mu$, consumers anticipate that even if the Period 1 ad exposure does not successfully convey match information, they will have another opportunity in Period 2 to obtain such information through repeat advertising.

## *Analysis for Opt-In Consumers*

Consider the Period 2 subgame for opt-in consumers in which a Period 1 ad is shown. In contrast to the case of opt-out consumers, the advertiser can condition its Period 2 bid, and the ad network its reserve price, on the opt-in consumers' purchase histories and inferred funnel states. Due to imperfect tracking, however, they can only condition on an imperfect signal $s \in \{T, B\}$ of the consumer's true funnel state $f \in \{T, B\}$, where the signal accuracy is given by (4).

Since only $T$-consumers respond to advertising with positive probability, the expected marginal value of advertising to an $s$-signal consumer is $\mathbb{P}\{f = T|s\}\mu\phi p$. By Bayes' rule, $\mathbb{P}\{f = T|s\} = \frac{\mathbb{P}\{s|f=T\}\mathbb{P}\{f=T\}}{\mathbb{P}\{s|f=T\}\mathbb{P}\{f=T\}+\mathbb{P}\{s|f=B\}\mathbb{P}\{f=B\}}$, so that the advertiser's valuation, and hence its Period 2 bid, for an $s$-signal consumer is

$$b_s = \mu\phi p \times \begin{cases} \frac{\rho(1-\mu)}{\rho(1-\mu)+(1-\rho)\mu(1-\phi)} & \text{if } s = T, \\ \frac{(1-\rho)(1-\mu)}{(1-\rho)(1-\mu)+\rho\mu(1-\phi)} & \text{if } s = B. \end{cases} \tag{9}$$

The advertiser bids 0 for consumers who already purchased, and the signal for them is irrelevant.

The ad network anticipates (9) and sets reserve price $\max[\underline{b}, b_s]$ for an $s$-signal consumer. This implies that the advertiser's Period 2 ad is shown to an $s$-signal consumer if and only if $\underline{b} \leq b_s$. We define two precision thresholds:

$$\rho_B = \frac{(1-\mu)(\mu\phi p - \underline{b})}{(1-\mu)(\mu\phi p - \underline{b}) + \mu(1-\phi)\underline{b}} \text{ and } \rho_T = \frac{\mu(1-\phi)\underline{b}}{(1-\mu)(\mu\phi p - \underline{b}) + \mu(1-\phi)\underline{b}}, \tag{10}$$

such that $\underline{b} \leq b_B \iff \rho \leq \rho_B$ and $\underline{b} \leq b_T \iff \rho_T \leq \rho$. We present the instrumental values for opt-in consumers in the following proposition.

**Proposition 2.** *Let $\rho_B$ and $\rho_T$ be as defined in (10). Opt-in consumers' expected intensity of repeated ads and expected product surplus (across the two periods), respectively, are*

$$\mathbb{E}[q(1)] = \begin{cases} 1-\mu\phi & \text{if } \underline{b} \leq \frac{1-\mu}{1-\mu\phi}\mu\phi p \text{ and } \rho \leq \rho_B, \\ \rho(1-\mu) + (1-\rho)\mu(1-\phi) & \text{if } \max[\rho_B, \rho_T] < \rho, \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

*and*

$$\mathbb{E}[S(1)] = \mu\phi(1-p) \times \begin{cases} 2-\mu & \text{if } \underline{b} \leq \frac{1-\mu}{1-\mu\phi}\mu\phi p \text{ and } \rho \leq \rho_B, \\ 1+(1-\mu)\rho & \text{if } \max[\rho_B, \rho_T] < \rho, \\ 1 & \text{otherwise.} \end{cases} \tag{12}$$

Similar to the case of opt-out consumers, the expected intensity of repeated ads for opt-in consumers is non-monotonic in ad effectiveness $\mu$ (as illustrated by the solid line in Figure 3a) — it may increase in $\mu$ because ads are shown to more types of opt-in consumers (e.g., both $T$- and $B$-signal opt-in consumers) as ad effectiveness increases, whereas it may decrease in $\mu$ due to the saturation effect, whereby Period 1 ad under large $\mu$ decreases the size of $T$-consumers in Period 2.

Next, we compare the instrumental values of opt-out and opt-in consumers. If $\mu$ is intermediate, opt-in consumers see fewer repeated ads than opt-out consumers due to reduced ad wastage: while the advertiser shows repeated ads to all opt-out consumers, the advertiser shows repeated ads only to a subset of opt-in consumers (i.e., either to all non-purchasers as in Region $B$ of Figure 3a, or to those who are inferred to be in state $T$ as in Regions $A$ and $C$ of Figure 3a). If $\mu$ is either small or large, then opt-in consumers see (weakly) more repeated ads than opt-out consumers as increased targeting efficiency for opt-in consumers induces the advertiser to show Period 2 ads (see $\mu \leq \underline{\mu}$ and $\mu > \overline{\mu}$ in Figure 3a).

The opt-in consumers' expected product surplus may be either greater or smaller than that of opt-out consumers (as illustrated by the solid line in Figure 3b). Enabling tracking can increase the expected product surplus because tracking leads to more efficient ad targeting: $T$-consumers, who benefit most from the match information conveyed through ads are exposed to repeated ads in regions where opt-out consumers would not have seen ads. On the other hand, $T$-consumers who are misinferred as $B$-consumers may not see repeated ads, which, if seen, would have aided their product match resolution, thereby increasing their expected product surplus. In this case, opt-in consumers receive *lower* product surplus than opt-out consumers due to mistargeting (see Regions $A$ and $C$ in Figure 3b).

Counter-intuitively, the mistargeting effect is most pronounced if tracking precision is sufficiently high. To see this, recall that the advertiser wants to show ads to $T$-consumers and not to $B$-consumers because match information is relevant only for $T$-consumers. Therefore, if tracking precision is high, the advertiser refines its targeting strategy such that only $T$-signal consumers see Period 2 ads. Due to imperfect tracking, however, $\mathbb{P}\{s = B | f = T\} = 1 - \rho$ fraction of $T$-consumers are misinferred as $B$-consumers, and hence do not see Period 2 ads under this refined targeting regime. While reduced ad intensity mitigates ad wearout, the missed ad exposures also hurt misinferred consumers because they miss valuable information that would have aided their purchase decisions.

Consumers compare the instrumental values across opt-out and opt-in subgames when making their privacy choices. Define the changes in the two components of the instrumental values of privacy from opting in to tracking as:

$$\Delta_q = \mathbb{E}[q(1)] - \mathbb{E}[q(0)] \text{ and } \Delta_S = \mathbb{E}[S(1)] - \mathbb{E}[S(0)]. \tag{13}$$

We summarize the instrumental trade-offs consumers face in the following proposition.

**Proposition 3.** *Let $\rho_B$ and $\rho_T$ be as defined in (10), and $\Delta_q$ and $\Delta_S$ as defined in (13).*

*A. If (a) $(1-\mu)\mu\phi p < \underline{b} \leq \frac{1-\mu}{1-\mu\phi}\mu\phi p$ or (b) $\frac{1-\mu}{1-\mu\phi}\mu\phi p < \underline{b}$ and $\rho_T < \rho$, then $\Delta_q > 0$ and $\Delta_S > 0$.*

*B. If $\underline{b} \leq (1-\mu)\mu\phi p$ and $\rho_B < \rho$, then $\Delta_q < 0$ and $\Delta_S < 0$.*

*C. If $\underline{b} \leq (1-\mu)\mu\phi p$ and $\rho \leq \rho_B$, then $\Delta_q < 0$ and $\Delta_S = 0$.*

*D. If $(1-\mu)\mu\phi p < \underline{b} \leq \frac{1-\mu}{1-\mu\phi}\mu\phi p$ and $\rho \leq \rho_T$, then $\Delta_q = 0$ and $\Delta_S = 0$.*

Proposition 3 shows that the direction of changes in instrumental privacy values from opting in to tracking is determined by a nuanced interplay of ad effectiveness, $\mu$, and tracking precision, $\rho$. For instance, if $\mu$ is intermediate and $\rho$ is large (Condition $B$ of Proposition 3 and Region $B$ of Figure 4), then the advertiser adopts a finer targeting strategy, whereby only $T$-signal consumers are shown repeat ads. This generates the mistargeting effect such that opt-in consumers see fewer repeated ads (positive instrumental value) but obtain lower expected product surplus (negative instrumental value). The net instrumental gains from opting in to tracking may be negative if $\eta|\Delta_q| < |\Delta_S|$. If $\mu$ is intermediate and $\rho$ small (Condition $C$ of Proposition 3 and Region $C$ of Figure 4), then only opt-in consumers who do not purchase see repeated ads, while all opt-out consumers, regardless of their purchase histories, see repeated ads. Therefore, opt-in consumers expect to see fewer repeated ads (positive instrumental value) without reducing their expected product surplus (zero instrumental value). In this case, the net instrumental gains from opting in is positive.

Overall, Proposition 3 highlights the complex implications of the consumer's privacy choice on her instrumental privacy utilities. Next, we examine how these instrumental trade-offs impact the consumers' privacy choices.

### *Consumer Privacy Choices and Their Impact*

We solve for the consumers' equilibrium privacy choices. We then discuss the implications of the privacy choices for the total ad revenue (defined as the sum of the ad network's and advertiser's

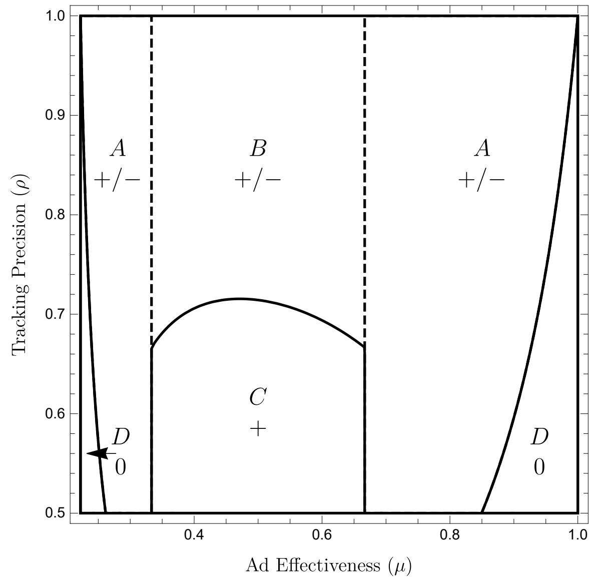


*Figure 4: Net Instrumental Gains from Opting In to Tracking ($-\eta\Delta_q + \Delta_S$)*

profits) and highlight an important effect of endogenous privacy choices on the relationship between ad revenue and tracking precision.

***Privacy choices.*** Consumers opt in to tracking if and only if the expected instrumental gains from opting in are positive and outweigh their intrinsic privacy cost, $\theta$. The key trade-off within the instrumental components is that consumers dislike seeing repeated ads due to wearout effects, but they value the match information conveyed through advertising that increases their expected product surplus. The set of consumers who opt in is $\mathcal{I} = \{\theta : u(0) \leq u(1)\} = \{\theta : \theta \leq -\eta\Delta_q + \Delta_S\}$. The total proportion of opt-in consumers is simply the probability measure of the set $\mathcal{I}$; i.e., $|\mathcal{I}| = F(-\eta\Delta_q + \Delta_S)$. Substituting the expressions for the instrumental values from Propositions 1 and 2, we obtain the equilibrium proportion of opt-in consumers as stated in the following proposition.

**Proposition 4.** *Let $b_B$ and $b_T$ be as defined in (9) and let*

$$\tilde{\theta} = \begin{cases} \eta\mu\phi & \text{if } \underline{b} \leq \underline{\xi}, \\ \eta(1-(1-\mu)\rho - \mu(1-\phi)(1-\rho)) - \mu\phi(1-p)(1-\mu)(1-\rho) & \text{if } \underline{\xi} < \underline{b} \leq (1-\mu)\mu\phi p, \\ -\eta(1-\mu\phi) + \mu\phi(1-p)(1-\mu) & \text{if } (1-\mu)\mu\phi p < \underline{b} \leq \overline{\xi}, \\ -\eta((1-\mu)\rho + \mu(1-\phi)(1-\rho)) + \mu\phi(1-p)(1-\mu)\rho & \text{if } \overline{\xi} < \underline{b} \leq b_T, \\ 0 & \text{if } b_T < \underline{b}, \end{cases} \tag{14}$$

*where $\underline{\xi} = \min[b_B, (1-\mu)\mu\phi p]$ and $\overline{\xi} = \max[b_B, (1-\mu)\mu\phi p]$. The proportion of opt-in consumers is*

$$|\mathcal{I}| = F(\tilde{\theta}). \tag{15}$$

Proposition 4 allows us to assess the impact of various model parameters on the consumers' opt-in decisions. Figure 5 illustrates the non-monotonic opt-in patterns with respect to $\rho$. The discontinuous changes in opt-in rates are attributable to the advertiser's shift in targeting strategies: for higher tracking precision, the advertiser adopts a finer targeting strategy such that only $T$-signal consumers see Period 2 ads. If $\underline{b}$ is small, this targeting refinement decreases the expected intensity of repeated ads as well as the expected product surplus. This implies a negative instrumental value from opting in for consumers who place more weight on product surplus than on ad wearout (i.e., small $\eta$); hence the discontinuous decrease in opt-in rate in Figure 5a. Conversely, the shift in targeting regime implies a positive instrumental value for consumers who care more about ad wearout (i.e., large $\eta$); hence the discontinuous increase in opt-in rate in Figure 5b.

The opt-in rate also changes continuously with tracking precision. In particular, if $\rho$ is large, $\mu$ moderately small, and $\eta$ large, then the proportion of opt-in consumers decreases in $\rho$ (see dashed line in Figure 5b). The intuition is the following: if $\rho$ is large, the advertiser shows repeated ads only to $T$-signal opt-in consumers in Period 2; if $\mu$ is small, most opt-in consumers remain in funnel state $T$ after Period 1 ad exposure (but $\mu$ is large enough that Period 2 advertising is worthwhile). Put together, as tracking precision increases, the advertiser receives more $T$-signals

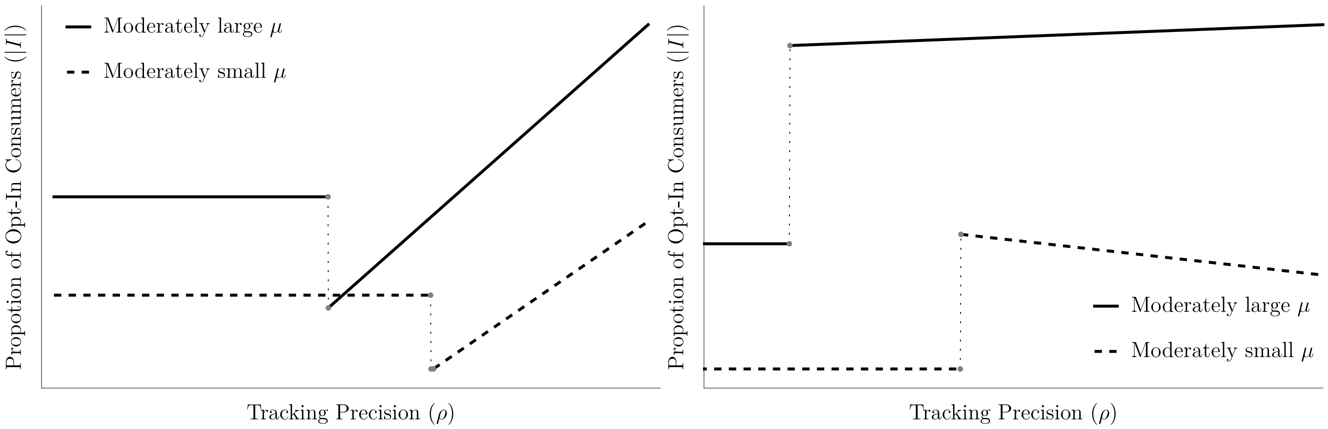


*(a) Low Ad Wearout ($\eta = 0.01$)* *(b) High Ad Wearout ($\eta = 0.25$)*

*Figure 5: Tracking Precision and Proportion of Opt-In Consumers;* $\underline{b} = 0.05, p = 0.75, \phi = 0.5, F(\theta) = \theta$

from the $T$-consumers such that more ads are shown in Period 2. Combined with consumers' high sensitivity to ad wearout, the negative instrumental value from increased ad repetition intensity induces more consumers to opt out. We summarize this result in the following corollary.

**Corollary 1.** *If* $\rho > \max[\rho_B, \rho_T]$, $\mu < \frac{1}{2-\phi}$, *and* $\eta > \frac{(1-p)(1-\mu)\mu\phi}{1-(2-\phi)\mu}$, *then* $\frac{\partial|\mathcal{I}|}{\partial\rho} < 0$.

Before proceeding, we note that in the Web Appendix we provide the analysis for a scenario in which the purchase funnel is degenerate; i.e., there is no state-wise progression towards purchase and there is no interim consumer heterogeneity created by advertising. In this scenario, we find that, in our framework, all consumers opt out of tracking, essentially because the instrumental value component of privacy gets muted. This highlights the importance of explicitly modeling the state-wise consumer purchase journey (i.e., the purchase funnel).

***Ad revenue.*** Another object of interest is the total ad revenue, $\Pi$, which we define as the sum of the ad network's and the advertiser's profits. In the main model, total ad revenue reduces to the ad network's profit because the ad network extracts all the surplus under information symmetry. The ad revenue can be written as

$$\Pi = |\mathcal{I}|\Pi_1 + (1 - |\mathcal{I}|)\Pi_0,$$

where $|\mathcal{I}|$ is the total proportion of opt-in consumers as defined in (15); $\Pi_1$ and $\Pi_0$ are ad revenues from opt-in and opt-out consumers, respectively, and are given by (see Claim 2 in the Appendix

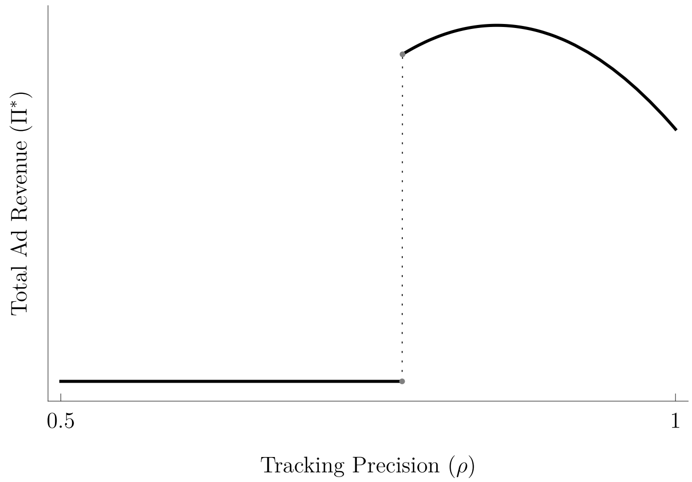


*Figure 6: Tracking Precision and Ad Revenue;* $\underline{b} = 0.05, p = 1, \phi = 0.5, \mu = 0.125, \eta = 0.1, F(\theta) = \theta$

for derivations)

$$\Pi_1 = \mu\phi p + (1 - \mu\phi)\left(\mathbb{P}\{s = T\}\max\left[\underline{b}, b_T\right] + \mathbb{P}\{s = B\}\max\left[\underline{b}, b_B\right]\right) + \mu\phi\underline{b}, \tag{16}$$

and

$$\Pi_0 = \mu\phi p + \max\left[\underline{b}, (1 - \mu)\mu\phi p\right]. \tag{17}$$

Tracking precision affects total ad revenue through two channels, the direct effect and the indirect effect. The direct effect follows from the fact that the advertiser's ad valuation increases with tracking precision; therefore, enhanced tracking precision (weakly) increases total ad revenue. On the other hand, higher tracking precision may indirectly reduce total ad revenue through its effect on the consumers' privacy choices. Per Corollary 1, this happens if consumers are sensitive to ad wearout (i.e., large $\eta$) and opt-in consumers expect to see more repeated ads under higher tracking precision. The negative instrumental value due to ad wearout incentivizes consumers to opt out for higher precision levels (see dashed line in Figure 5b). Since total ad revenue from opt-in consumers, $\Pi_1$, is weakly greater than the total ad revenue from opt-out consumers, $\Pi_0$ (see Claim 2 in the Appendix), total ad revenue decreases as consumers opt out. Overall, if the indirect opt-out effect dominates the positive direct effect, total ad revenue may decrease in tracking precision (see Figure 6). The following proposition lays out the exact condition for this.

**Proposition 5.** *Let $|\mathcal{I}|$ be as defined in (15), and consider the parameter range for which $\partial\Pi_1/\partial\rho$*

*and $\partial|\mathcal{I}|/\partial\rho$ are well-defined. The total ad revenue $|\mathcal{I}|\Pi_1 + (1-|\mathcal{I}|)\,\Pi_0$ decreases in tracking precision if and only if*

$$\underbrace{|\mathcal{I}|\frac{\partial\Pi_1}{\partial\rho}}_{\textit{direct effect}} + \underbrace{\frac{\partial|\mathcal{I}|}{\partial\rho}\,(\Pi_1 - \Pi_0)}_{\textit{indirect effect}} < 0.$$

Proposition 5 sheds light on a novel mechanism through which the potential to track consumers more accurately reduces the ad network's profit. Even if market thickness is fixed and the ad network extracts all of the advertiser's surplus, higher tracking precision may decrease the ad network's profit. In an environment where consumers are endowed with data privacy choices, increased tracking precision may induce consumers to opt out of tracking to mitigate ad wearout from repeated ads. As consumers opt out, the proportion of consumers in the population that can be targeted falls such that the ad network's total profit decreases.

An important managerial implication of Proposition 5 is that marketers should carefully assess the consumers' "opt-out elasticity" with respect to tracking precision. If consumers anticipate higher tracking precision to reduce the instrumental value of sharing their data, then consumers will withdraw from data sharing. In total, the risks of consumers opting out of tracking may outweigh the benefits of increased tracking precision (e.g., greater surplus due to efficiency gains) such that the total ad revenue decreases.

Thus far, we have examined the consumers' endogenous privacy choices and their implications for the total ad revenue. An important question for policymakers is: what is the impact of consumers' privacy choices on the ad ecosystem? Specifically, what are the implications of privacy regulations that empower consumers with data privacy rights on consumer surplus, the total ad revenue, and the advertising outcomes? We address these questions by comparing the equilibrium results from the previous section with the benchmark case in which all consumers are opted in by default. Notice that the all-opt-in benchmark is simply the subgame for opt-in consumers, which was presented in the analysis section. Therefore, we need only compare the equilibrium outcomes under endogenous privacy choices with the subgame outcomes for opt-in consumers. Based on this comparison, we address the above questions in the following proposition.

**Proposition 6.** *Compared to the benchmark in which all consumers are opted in to tracking, endowing consumers with privacy choices*

1. *weakly increases the expected intensity of repeated ads shown in the system and the expected product surplus consumers receive if $\underline{\mu} < \mu \leq \overline{\mu}$, and weakly decreases them otherwise;*
2. *weakly increases consumer surplus; and*
3. *weakly decreases total ad revenue.*

First, empowering consumers with privacy choices has mixed effects on the average intensity of repeated ads served in the system. If ad effectiveness is intermediate, then opt-out consumers always see repeated ads, while opt-in consumers expect to see fewer repeated ads due to improved targeting. Therefore, as privacy-conscious consumers opt out, more repeated ads are served. On the other hand, if ad effectiveness is either small or large, then opt-out consumers do not see repeated ads whereas opt-in consumers in top funnel states do due to more efficient targeting. Therefore, as consumers exercise their opt-out rights, fewer repeated ads are shown on average compared to the all-opt-in benchmark. Second, in our setting wherein consumers' privacy choices do not exert externalities on other consumers' utilities, privacy choices enable consumers to allow tracking selectively such that consumer surplus weakly increases. In an extension with endogenous product prices, we show that this is not necessarily true: consumers may be worse off having privacy choices than not. Finally, the endowment of privacy choices weakly reduces the total ad revenue because, as privacy-sensitive consumers opt out of tracking, the advertiser-consumer matching efficiency decreases.

The main model enabled us to develop an understanding of the factors that influence consumer opt-in choice, specifically, the intrinsic and instrumental values of privacy, and of the actions of the advertiser and the ad network. We now proceed to extensions of the main model.

# *EXTENSIONS*

In this section, we consider three extensions of the main model that confirm the robustness of our insights and also provide new insights. We analyze additional extensions in the Web Appendix.

## *Endogenous Product Price Under Information Asymmetry*

We investigate whether our main insights under exogenous product prices carry over in a setting where the advertiser adjusts price endogenously. We consider a stylized analogue of the purchase funnel in the main model. We assume an exogenous product consideration process, whereby consumers in funnel state $T$ do not consider the product, while consumers in funnel state $B$ consider the product with probability $\phi + \epsilon$ if the consumer is $H$-type, and with probability $\phi - \epsilon$ if the consumer is $L$-type, for some small $\epsilon > 0$. The mass of $H$- and $L$-type consumers is each 1/2. Upon consideration, consumers privately draw a product valuation according to a cumulative distribution function $Z$ on the support $\left[0, \overline{V}\right]$.

We assume that the advertiser can perfectly observe this consideration probability for opt-in consumers (i.e., either $\phi+\epsilon$ or $\phi-\epsilon$), while the ad network cannot. For instance, the advertiser and ad network may have access to the same consumer data, but the advertiser may be better equipped to translate the data to meaningful information (de Cornière and de Nijs, 2016). This information asymmetry is essential to analyzing endogenous product prices, for without the asymmetry, the advertiser's surplus would be completely extracted by the ad network and the advertiser would be indifferent between various possible product prices.

The advertiser sets product price before consumers make privacy choices, such that all consumers face the same product price. To obtain sharp insights, we assume that the opt-in consumers' funnel states can be perfectly inferred by the advertiser and the ad network (i.e., $\rho = 1$).

Finally, under the assumption that the advertiser's expected profit from opt-in consumers considering purchase, $(1 - Z(p))\,(p - c)$, is strictly quasi-concave in $p$, we show that if $\epsilon$ is sufficiently small, then the non-monotonic advertising outcomes from the main model, and thus, the qualitative insights from the main model, carry over.

Interestingly, we find that under endogenous product price, consumers may face different product prices between all-opt-in benchmark and the endogenous privacy choices regime and, under certain conditions, the price may be higher in the latter. To see why, note that the advertiser's profit from opt-in consumers is greater (see proof of Proposition 7 below) than that from opt-out consumers.[15] Therefore, the advertiser has incentive to increase the opt-in rate. Even though a high product price reduces the advertiser's expected profit (recall that $(1 - Z(p))\,(p - c)$ decreases for large $p$ due to strict quasi-concavity), a high product price could increase consumer opt-in rate through the following mechanism. By setting a high product price, the advertiser lowers the expected value of advertising (i.e., through reduced expected margin). This implies that opt-out consumers do not see repeated ads in Period 2, while opt-in consumers see repeated ads that provide valuable information (see Condition $A$ in Proposition 3). Overall, if consumers have low sensitivity to ad wearout, then opt-in consumers receive greater instrumental utility than opt-out consumers. This induces more consumers to opt in, thereby increasing the advertiser's expected profit. In summary, the intuition is that high product price can induce more consumers to opt in by serving as the advertiser's commitment device to provide greater instrumental utility to opt-in consumers than to opt-out consumers. The following proposition summarizes this finding.

**Proposition 7.** *Let $\hat{\epsilon}$, $\hat{\eta}$, $\underline{p}'$ and $\overline{p}'$ be as defined in the proof. Suppose $0 < \epsilon < \hat{\epsilon}$. The product price under the endogenous privacy choice regime is higher than that under the all-opt-in benchmark if and only if $\underline{p}' < c < \overline{p}'$ and $\eta < \hat{\eta}$.*

A high-level takeaway from Proposition 7 is that in a setting where consumers are endowed with privacy choices, firms may adjust their strategies to impact the opt-in rate. In our setting, the advertiser raises the product price to induce more consumers to opt in. Note that if this price increase is sufficiently large, consumers may be worse off under the endogenous privacy choice regime than in the setting where consumers are stripped of privacy choice and are opted in by default. We state this result in the following corollary.

[15] In fact, the latter is 0 because the advertiser and the ad network are equally uninformed about opt-out consumers, such that the ad network completely extracts the advertiser's surplus.

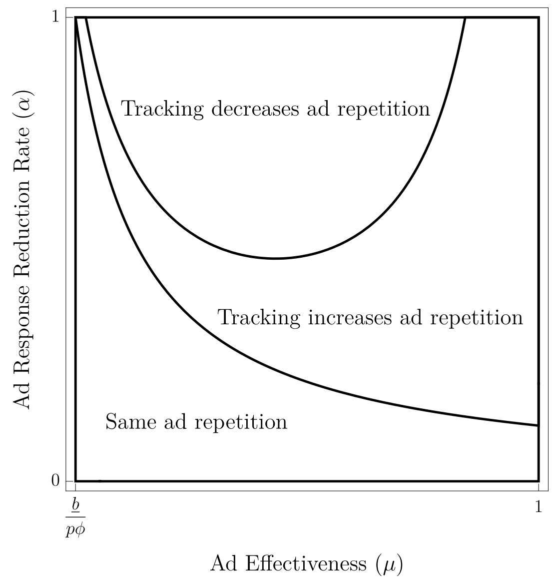


*Figure 7: Tracking and Intensity of Repeated Ads*

**Corollary 2.** *Let $p_{(i)}$, $\overline{p}'$, and $\sigma(p)$ be as defined in the proof. Consumers who opt in to tracking have lower utilities under the endogenous privacy choice regime than in the all-opt-in regime. Consumers who opt out of tracking have lower utilities under the endogenous privacy choice regime than in the all-opt-in regime if and only if $\eta < \left(\sigma\left(p_{(i)}\right)\mu\phi\left(2-\mu\right) - \sigma\left(\overline{p}'\right)\mu\phi\right)/(1-\mu)$.*

## *Ad Wearout Leading to Reduced Ad Responsiveness*

In the main model, we assumed that the wearout disutility that a consumer experiences from exposure to repeated ads is orthogonal to the consumer's product demand. In this section, we analyze a scenario in which ad wearout "spills over" to the consumer's product utility. This is motivated by empirical findings that marginal responsiveness to advertising may diminish in the number of ad exposures (e.g., Aaker and Bruzzone, 1985; Chen et al., 2016; Chae et al., 2019; Todri et al., 2020). Specifically, we assume that if a consumer sees a repeated ad, then her product match probability decreases by a factor $\alpha \in [0, 1]$. To illustrate, suppose a $T$-consumer sees an ad. If it is her first exposure, then her conversion probability is $\mu\phi$, where $\mu$ is the probability of the ad conveying product match information to the consumer, and $\phi$ is the match probability. On the other hand, if it is a repeat exposure, then her conversion probability reduces to $\alpha\mu\phi$.

Due to space considerations, we provide the analysis in the Web Appendix and summarize here how the results are impacted by $\alpha$. As depicted in Figure 7, the parameter region for which tracking decreases ad repetition becomes smaller as $\alpha$ decreases (i.e., ad wearout induces larger ad response reduction). The intuition is that as ad wearout reduces consumers' responsiveness to ads, the marginal value of repeat advertising decreases. Therefore, the region for which repeat ads are shown to opt-out consumers — which is equivalent to the region for which tracking reduces ad repetition — becomes smaller.

Compared to the main model, if $\alpha$ is large, then the qualitative insights carry over. If $\alpha$ is small, then reduction in ad responsiveness has two effects. First, consumers expect lower product surplus than in the main model because: (a) their match probabilities are reduced by a factor $\alpha$ for repeat ads, and (b) the advertiser shows fewer repeat ads due to diminished ad valuation. Second, within the extension model, consumers expect to see (weakly) more repeat ads from opting in to tracking than from opting out. Since repeat ads increase the probability that consumers receive match information, thereby increasing consumers' expected product surplus, consumers have a stronger incentive to opt in to tracking than they do in the main model.

### *Changing Intrinsic Value of Privacy*

A high-level insight that we obtain from the main model is that the firms' strategies affect consumers' privacy choices through the instrumental benefits and costs of allowing tracking. While the main model focuses on the implications of firms' strategies for the different, and possibly countervailing, components of the instrumental value of privacy, firms can also take actions to influence consumers' intrinsic values of privacy. For example, firms' improvements in data security to reduce privacy breaches (Ke and Sudhir, 2021) or firms' adoption of different default privacy frames (Lin, 2020) could shift consumers' perceptions of privacy, thereby changing their intrinsic privacy preferences. Moreover, the fact that 67% of US and Canadian consumers report that they would feel "comfortable sharing personal information with a company" under transparent data policy disclosures (Ipsos, 2019) lends further credence to the notion that consumers' intrinsic values of privacy can be influenced by firms' actions.

Based on these motivations, we extend the main model by allowing firms' strategies to shift consumers' intrinsic values of privacy. We assume that before consumers make privacy choices, the ad network can take actions that affect the intrinsic value of privacy; specifically, it can change consumers' privacy costs by magnitude $\Delta_\theta \in \mathbb{R}$, at cost

$$c\left(\Delta_\theta\right) = \frac{k}{2}\Delta_\theta^2 \times \begin{cases} 1 & \text{if } \Delta_\theta < 0, \\ \tau & \text{if } \Delta_\theta \geq 0, \end{cases} \tag{18}$$

for $\tau \in [0, 1)$. The cost function (18) captures the notions that: (a) the marginal cost of shifting preferences increases, and (b) reducing consumer's intrinsic privacy cost is more costly than increasing it. To ensure the existence of interior solutions, we assume that the density function $f$ of the intrinsic cost distribution $F$ is differentiable and has a bounded first-derivative; i.e., $\max_{\theta \in [0,\overline{\theta}]} |f'(\theta)| < k$ where $k$ is a cost parameter in (18).

Recall that the ad revenue from opt-in consumers is weakly higher than that from opt-out consumers (see Claim 2 in the Appendix). It follows that the ad network only benefits from increasing the opt-in rate; therefore, the ad network will incur the cost $c\left(\Delta_\theta\right)$ only to *reduce* $\theta$ (e.g., invest in data security or revamp disclosure policies). However, there are conditions under which the instrumental benefits of opting in are high enough that all consumers opt in without firm intervention, thereby obviating the need for the ad network's costly efforts to reduce consumers' intrinsic privacy costs. Based on the opt-in choices derived in the main model (see Proposition 4), we summarize the ad network's optimal reduction of consumers' privacy costs in the following proposition (see Web Appendix for analysis details).

**Proposition 8.** *Let $\tilde{\theta}$ be as defined in (14), and let $\Pi_1$ and $\Pi_0$ be as defined in (16) and (17), respectively. Then the ad network reduces the consumers' privacy costs by $\Delta_\theta^* = \max\left[0, \tilde{\Delta}_\theta\right]$, where $\tilde{\Delta}_\theta$ solves the first-order condition $f\left(\tilde{\theta} + \Delta_\theta\right)\left(\Pi_1 - \Pi_0\right) = k\Delta_\theta$.*

Proposition 8 provides important insights for regulators. For example, if $\tilde{\theta} > \overline{\theta}$ (i.e., all consumers opt in even if $\Delta_\theta = 0$), then we obtain the corner solution $\Delta_\theta^* = 0$. This means that if the instrumental value of opting in is sufficiently large, then the ad network has little incentive to invest in, say, data security to alleviate consumers' privacy concerns. In such cases, self-regulation to improve data security would be insufficient and regulatory intervention would be necessary. Overall, Proposition 8 demonstrates the significance of considering the complex effects of firms' strategies on both the instrumental and the intrinsic aspects of privacy, especially in designing regulation policy.

# *CONCLUSION*

Tracking consumers' Internet activities enables dynamic ad targeting as consumers traverse the purchase funnel for a product. While tracking may increase the efficiency with which ads are matched to consumers in different funnel states, it may also affect advertisers' ad strategies that indirectly impact consumers' utilities. In this setting, we study the impact of regulations (e.g., the GDPR) that, motivated by privacy concerns, endow consumers with the choice to have their online activity be tracked or not. In particular, we develop a framework to analyze consumers' opt-in/-out decisions — co-determined by the intrinsic and instrumental aspects of privacy — and their impact on the strategies and profits of ad networks and advertisers.

We obtain a number of insights from our analysis. A greater number of ads increases disutility from ad repetition wearout, but increases expected product surplus as ads provide match information that improves consumers' product purchase decisions. Therefore, depending on which of these countervailing effects is stronger, opting in to tracking can lead to greater or smaller instrumental value of privacy. These forces combine to produce counter-intuitive results. For instance, even though consumers opt in to enable their information to be tracked, under certain conditions fewer consumers may opt in as information is tracked with higher precision. This happens if ad effectiveness is relatively low and consumers are sensitive to ad wearout, because with higher tracking precision, consumers for whom the early ad exposures were "unsuccessful" are more likely to be identified and shown more ads. In turn, when this happens, ad revenue decreases in tracking precision because fewer consumers opt in, and without tracking data the efficiency of ads, and therefore willingness to pay for ads, decreases.

Furthermore, we find that in a setting where product price is set endogenously, endowing consumers with privacy choices can hurt consumers because the advertiser may set high product price to increase consumer opt-in rate. A higher price implies a lower expected product surplus for the consumer and therefore acts as a commitment device for showing fewer ads to the consumer. Finally, firms' actions such as increasing data security would impact the intrinsic value of privacy by reducing the perceived risk of sharing data. We find that firms may, under certain conditions, naturally have the incentive to take such actions; under other conditions, such actions may need to be mandated through regulations.

We acknowledge several limitations of our research. First, it would be interesting to examine a more active role of publishers. One could consider publishers acting as information gateways and study the forces that affect the publishers' incentives to disclose or withhold consumers' information to the ad network. Second, future research could relax the exogenous consumer website visit assumption and investigate how consumers adjust their website visit decisions in tandem with their tracking choices. Third, while we allow heterogeneity in consumers' ad responsiveness, we assume that ad effectiveness for top-funnel consumers is homogeneous (i.e., all equal to $\mu$). An interesting avenue for future research could be to study the implications of heterogeneous ad effects and the advertiser's incentive to potentially "screen" consumers based on this heterogeneity. Fourth, regulations such as the GDPR and the CCPA/CPRA are consistent with the so-called "rights and responsibilities" model. This model states that consumers should have certain privacy rights, such as being tracked only on opting in, and firms have certain responsibilities, such as protecting consumer data. The instrumental and intrinsic values of privacy in our model correspond to the "responsibilities" and the "rights," respectively. However, we do not consider all rights and responsibilities covered in recent regulations, such as the ability of consumers to have their data shared with themselves or transferred to other parties, and the requirement that firms maintain data in a format that is easy to understand at the time of such data operations. Future work can consider these other aspects of regulation as well.

# References


Aaker, D. A. and Bruzzone, D. E. (1985). Causes of irritation in advertising. *Journal of Marketing*, 49(2):47–57.

Abhishek, V., Despotakis, S., and Ravi, R. (2017). Multi-channel attribution: The blind spot of online advertising. *Working paper*. Irvine, CA.

Abhishek, V., Fader, P., and Hosanagar, K. (2018). Media exposure through the funnel: A model of multi-stage attribution. *Working paper*. Irvine, CA.

Acquisti, A., Taylor, C. R., and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2):442–92.

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

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

Angwin, J. (2010). The web's new gold mine: Your secrets. *The Wall Street Journal*.

Athey, S., Catalini, C., and Tucker, C. (2017). The digital privacy paradox: Small money, small costs, small talk. Technical report, National Bureau of Economic Research.

Athey, S. and Gans, J. S. (2010). The impact of targeting technology on advertising markets and media competition. *American Economic Review*, 100(2):608–13.

Aziz, A. and Telang, R. (2016). What is a digital cookie worth? *Working paper*. Vancouver, BC.

Aziz, A. and Telang, R. (2017). Irrelevant advertising: How more consumer tracking can make ads less effective. *Working paper*. Vancouver, BC.

Becker, G. S. (1980). Privacy and malfeasance: A comment. *Journal of Legal Studies*, 9(4):823–826.

Bergemann, D. and Bonatti, A. (2011). Targeting in advertising markets: implications for offline versus online media. *RAND Journal of Economics*, 42(3):417–443.

Bigler, J. (2019). "Rolling out first price auctions to Google Ad Manager partners," (accessed October 18, 2020). https://bit.ly/3aKm4K5.

Campbell, J., Goldfarb, A., and Tucker, C. (2015). Privacy regulation and market structure. *Journal of Economics & Management Strategy*, 24(1):47–73.

Campbell, M. and Keller, K. (2003). Brand familiarity and advertising repetition effects. *Journal of Consumer Research*, 30(2):292–304.

Chae, I., Bruno, H. A., and Feinberg, F. M. (2019). Wearout or weariness? measuring potential negative consequences of online ad volume and placement on website visits. *Journal of Marketing Research*, 56(1):57–75.

Chen, J., Yang, X., and Smith, R. E. (2016). The effects of creativity on advertising wear-in and wear-out. *Journal of the Academy of Marketing Science*, 44(3):334–349.

Conitzer, V., Taylor, C. R., and Wagman, L. (2012). Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science*, 31(2):277–292.

D'Annunzio, A. and Russo, A. (2020). Ad networks and consumer tracking. *Management Science*, 66(11):5040–58.

de Cornière, A. and de Nijs, R. (2016). Online advertising and privacy. *RAND Journal of Economics*, 47(650):48–72.

de Matos, M. G. and Adjerid, I. (2020). Consumer consent and firm targeting after gdpr: The case of a telecom provider in europe. *Working paper*. Catolica Lisbon School of Business and Economics.

Edelman, D. (2010). Branding in the digital age: you're spending your money in all the wrong places. *Harvard Business Review*, 88(12):62–69.

eMarketer (2018). "Level of Concern US Internet Users Have Toward Social Media Companies Displaying Ads Based on Their Data" (accessed May 3, 2019). https://bit.ly/3gKv7ig.

eMarketer (2019a). "Consumer Attitudes on Marketing 2019: Privacy Concerns Mount, and Ad Blocking Isn't Going Away" (accessed December 3, 2019). https://bit.ly/3aJk9pp.

eMarketer (2019b). "Programmatic Digital Display Ad Spending" (accessed December 3, 2019). https://bit.ly/32UXGBu.

Esteban, L., Gil, A., and Hernandez, J. M. (2001). Informative advertising and optimal targeting in a monopoly. *Journal of Industrial Economics*, 49(2):161–80.

Farrell, J. (2012). Can privacy be just another good? *Journal on Telecommunications and High Technology Law*, 10(2):251–64.

Goldfarb, A. and Tucker, C. E. (2011). Privacy regulation and online advertising. *Management science*, 57(1):57–71.

Google (2020a). "About remarketing audiences" (accessed October 18, 2020). https://bit.ly/3vs5XZU.

Google (2020b). "AdSense revenue share" (accessed October 18, 2020). https://bit.ly/3b2uWLB.

Heidhues, P. and Kőszegi, B. (2017). Naivete-based discrimination. *Quarterly Journal of Economics*, 132(2):1019–1054.

Hoban, P. R. and Bucklin, R. E. (2015). Effects of internet display advertising in the purchase funnel: Model-based insights from a randomized field experiment. *Journal of Marketing Research*, 52(3):375–93.

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

Ipsos (2019). "Global Citizens and Data Privacy" (accessed May 3, 2019). https://bit.ly/3gRrRBo.

Iyer, G., Soberman, D., and Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3):461–76.

Jatain, V. (2019). "24 Best Ad Networks For Publishers," (accessed October 18, 2020). https://bit.ly/3xvvoLJ.

Johnson, G., Lewis, R. A., and Nubbemeyer, E. (2017). The online display ad effectiveness funnel & carryover: Lessons from 432 field experiments. *Working paper*. Boston University.

Johnson, G. A., Shriver, S. K., and Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1):33–51.

Johnson, J. P. (2013). Targeted advertising and advertising avoidance. *RAND Journal of Economics*, 44(1):128–44.

Ke, T. T. and Sudhir, K. (2021). Privacy rights and data security: Gdpr and personal data driven markets. *Working paper*. Chinese University of Hong Kong.

Kotler, P. and Keller, K. L. (2012). *Marketing Management*. Pearson.

Lambrecht, A. and Tucker, C. E. (2013). When does retargeting work? information specificity in online advertising. *Journal of Marketing Research*, 50(5):561–76.

Lavidge, R. J. and Steiner, G. A. (1961). A model for predictive measurements of advertising effectiveness. *Journal of Marketing*, 25(6):59–62.

Lin, T. (2020). Valuing intrinsic and instrumental preferences for privacy. *Working paper*. Boston University.

McDonald, A. M. and Cranor, L. F. (2010). Americans' attitudes about internet behavioral advertising practices. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society*, pages 63–72. ACM.

Montes, R., Sand-Zantman, W., and Valletti, T. (2019). The value of personal information in online markets with endogenous privacy. *Management Science*, 65(3):1342–62.

Norberg, P. A., Horne, D. R., and Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126.

Posner, R. A. (1981). The economics of privacy. *American Economic Review*, 71(2):405–409.

Rafieian, O. and Yoganarasimhan, H. (2020). Targeting and privacy in mobile advertising. *Marketing Science*. Forthcoming.

Sahni, N. S., Narayanan, S., and Kalyanam, K. (2019). An experimental investigation of the effects of retargeted advertising: The role of frequency and timing. *Journal of Marketing Research*, 56(3):401–418.

Sankin, A. (2021). "I Tried to Use the Ad Tech Industry's Tool to Opt Out of Personalized Ads. Did It Work?" (accessed April 13, 2021). https://bit.ly/3vs6afE.

Seiler, S. and Yao, S. (2017). The impact of advertising along the conversion funnel. *Quantitative Marketing and Economics*, 15(3):241–278.

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

Srinivasan, R. (2019). "A privacy-safe approach to managing ad frequency," (accessed October 18, 2020). https://bit.ly/3vo7r77.

Stern, J. (2018). Facebook really is spying on you, just not through your phone's mic. *The Wall Street Journal*.

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics*, pages 631–650.

Intersoft Consulting (2021). "Article 12 GDPR: Transparent information, communication and modalities for the exercise of the rights of the data subject" (accessed March 16, 2021). https://gdpr-info.eu/art-12-gdpr/.

Todri, V., Ghose, A., and Singh, P. V. (2020). Trade-offs in online advertising: Advertising effectiveness and annoyance dynamics across the purchase funnel. *Information Systems Research*, 31(1):102–125.

Villas-Boas, J. M. (2004). Price cycles in markets with customer recognition. *RAND Journal of Economics*, pages 486–501.

Wathieu, L. and Friedman, A. (2009). An empirical approach to understanding privacy concerns. *Working Paper*. Washington, DC.

# *Notes*

# Appendix

## *Proofs*

*Proof of Proposition 1.* As explained in the text, Period 2 ad is shown if and only if $\underline{b} \leq (1-\mu)\mu\phi p$, which simplifies to $\underline{\mu} < \mu \leq \overline{\mu}$. Therefore, we obtain the expression for $\mathbb{E}[q(0)]$ as presented in the lemma.

The expected product surplus can be derived as follows: if a consumer sees ads in both periods, then she purchases with probability $\mu\phi$ in Period 1, and of the remaining $1 - \mu\phi$ consumers, $(1 - \mu)\mu\phi$ purchase due to Period 2 ad. Adding the two terms and multiplying the product surplus $1 - p$ conditional on purchase, we obtain the expression for expected product surplus if $\underline{\mu} < \mu \leq \overline{\mu}$. If only Period 1 ads are shown, then the total purchase probability is $\mu\phi$, from which the expression in the proposition follows. ■

*Proof of Proposition 2.* If $\underline{b} \leq b_B$, then the reserve price for each signal will be set just as high as the advertiser's bid; therefore, the advertiser's ad is shown to all non-purchasers in Period 2, regardless of their signals. Therefore, $\mathbb{E}[q(1)] = 1 - \mu\phi$ and

$$\mathbb{E}[S(1)] = (1 - p)\left(\mu\phi + (1 - \mu)\mu\phi\right) = \mu\phi(1 - p)(2 - \mu).$$

If $b_B < \underline{b} \leq b_T$, then the reserve price for $B$-signal consumers is $\underline{b}$, which exceeds the advertiser's bid for $B$-signal consumers. This implies that the advertiser's Period 2 ad is shown only to $T$-signal consumers. Since $s = T$ can arise from $T$-consumers being correctly inferred or $B$-consumers being misinferred, we obtain $\mathbb{E}[q(1)] = (1-\mu)\rho+\mu(1-\phi)(1-\rho)$ and $\mathbb{E}[S(1)] = (1-p)\left(\mu\phi + \rho(1 - \mu)\mu\phi\right) = \mu\phi(1 - p)(1 + (1 - \mu)\rho)$.

Finally, if $b_T < \underline{b}$, then the reserve prices for both signals exceed the advertiser's respective bids. Therefore, the advertiser's Period 2 ad is not shown, yielding $\mathbb{E}[q(1)] = 0$ and $\mathbb{E}[S(1)] = \mu\phi(1 - p)$. ■

*Proof of Proposition 3.* We characterize the conditions for each of the four regions stated in the proposition:

A. $\Delta_q > 0$ and $\Delta_S > 0$ iff $\mathbb{E}[q(0)] = 0$ and at least $T$-signal opt-in consumers see repeated ads. The first condition is $(1-\mu)\mu\phi p < \underline{b}$ and the second $\underline{b} \leq b_T$, which simplifies to $\rho > \rho_T$. But if $\underline{b} \leq \frac{1-\mu}{1-\mu\phi}\mu\phi p$, then $\rho_T < \frac{1}{2}$ so that the second condition is implied. Therefore, the two conditions can be simplified to either (a) $(1-\mu)\mu\phi p < \underline{b} \leq \frac{1-\mu}{1-\mu\phi}\mu\phi p$ or (b) $\frac{1-\mu}{1-\mu\phi}\mu\phi p < \underline{b}$ and $\rho_T < \rho$.

B. $\Delta_q < 0$ and $\Delta_S < 0$ iff $\mathbb{E}[q(0)] = 1$ and repeated ads are shown only to $T$-signal opt-in consumers. The first condition is $\underline{b} \leq (1-\mu)\mu\phi p$ and the second $b_B < \underline{b} \leq b_T$, which simplifies to $\rho_T < \rho$ and $\rho_B < \rho$. But $\rho > 1/2 \Rightarrow (1-\rho)(1-\phi) < \rho \iff \rho(1-\mu) + (1-\rho)\mu(1-\phi) < \rho \iff 1-\mu < \frac{\rho(1-\mu)}{\rho(1-\mu)+(1-\rho)\mu(1-\phi)} \iff (1-\mu)\mu\phi p < b_T \Rightarrow \underline{b} \leq b_T$. Therefore, the intersection of the two conditions simplifies to $(1-\mu)\mu\phi p < \underline{b}$ and $\rho > \rho_B$.

C. $\Delta_q < 0$ and $\Delta_S = 0$ iff $\mathbb{E}[q(0)] = 1$ and repeated ads are shown to all opt-in consumers. The first condition is $(1-\mu)\mu\phi p < \underline{b}$ and the second $\rho \leq \rho_B$ and $\rho > \rho_T$. As above, since the first condition implies $\rho > \rho_T$, the intersection of the two conditions simplifies to $(1-\mu)\mu\phi p < \underline{b}$ and $\rho \leq \rho_B$.

D. $\Delta_q = \Delta_S = 0$ iff $\mathbb{E}[q(0)] = 0$ and no opt-in consumers see repeated ads. The first condition is $(1-\mu)\mu\phi p < \underline{b}$ and the second $b_T \leq \underline{b} \iff \rho \leq \rho_T$.

■

*Proof of Proposition 4.* If $\underline{b} \leq \underline{\xi}$, then all opt-out consumers see repeated ads, and all opt-in non-purchasers see repeated ads. Since the expected product surplus is the same regardless of the consumers' privacy choices, the ad repetition intensity determines the instrumental value differential: $\mathbb{E}[q(0)] = 1$ and $\mathbb{E}[q(1)] = 1 - \mu\phi$. Therefore, the expected instrumental gain from opting in is $\eta\mu\phi$, which implies that all consumers with $\theta \leq \eta\mu\phi$ opt in.

If $b_B < \underline{b} \leq (1-\mu)\mu\phi p$, then only $T$-signal opt-in consumers see repeated ads while all opt-out consumers see repeated ads. Therefore, the expected instrumental gain from opting in is $\eta(1-(1-\mu)\rho - \mu(1-\phi)(1-\rho)) - \mu\phi(1-p)(1-\mu-(1-\mu)\rho) = \eta(1-(1-\mu)\rho - \mu(1-\phi)(1-\rho)) - \mu\phi(1-p)(1-\mu)(1-\rho)$

If $(1-\mu)\mu\phi p < \underline{b} \leq b_B$, then all opt-in non-purchasers see repeated ads while none of the opt-out consumers do. Therefore, the expected instrumental gain from opting in is $-\eta(1-\mu\phi)+$

$\mu\phi(1-p)(1-\mu)$.

If $\overline{\xi} < \underline{b} \leq b_T$, then only $T$-signal opt-in consumers see repeated ads while none of the opt-out consumers do. Therefore, the expected instrumental gain from opting in is $-\eta((1-\mu)\rho + \mu(1-\phi)(1-\rho)) + \mu\phi(1-p)(1-\mu)\rho$. ■

*Proof of Corollary 1.* Since $\frac{\partial|\mathcal{I}|}{\partial\rho} = 0$ for all $\rho$ except $\rho > \max[\rho_B, \rho_T]$, for $\frac{\partial|\mathcal{I}|}{\partial\rho} < 0$ we need $\rho > \max[\rho_B, \rho_T]$. In this region, we have $\frac{\partial|\mathcal{I}|}{\partial\rho} = -\eta\,(1-(2-\phi)\mu) + (1-p)(1-\mu)\mu\phi$, which is negative iff $1-(2-\phi)\mu > 0$ and $\eta > \frac{(1-p)(1-\mu)\mu\phi}{1-(2-\phi)\mu}$. ■

*Proof of Proposition 5.* This follows immediately from the chain rule and $\frac{\partial\Pi_0}{\partial\rho} = 0$. ■

*Proof of Proposition 6.* Let the instrumental objects be as defined in Propositions 1 and 2.

1. If $\underline{\mu} < \mu \leq \overline{\mu}$, then $\mathbb{E}[q(0)] > \mathbb{E}[q(1)]$ and $\mathbb{E}[S(0)] \geq \mathbb{E}[S(1)]$. Therefore, under endogenous privacy choices, the expected intensity of repeated ads is $|\mathcal{I}|\mathbb{E}[q(1)] + (1-|\mathcal{I}|)\mathbb{E}[q(0)] \geq \mathbb{E}[q(1)]$ for all $\mathcal{I} \geq 0$. Similarly, the expected product surplus under endogenous privacy choices is $|\mathcal{I}|\mathbb{E}[S(1)] + (1-|\mathcal{I}|)\mathbb{E}[S(0)] \geq \mathbb{E}[S(1)]$. If $\mu \leq \underline{\mu}$ or $\mu > \overline{\mu}$, then $\mathbb{E}[q(0)] = 0 \leq \mathbb{E}[q(1)]$ and $\mathbb{E}[S(0)] = (1-p)\mu\phi \leq \mathbb{E}[S(1)]$. Therefore, $|\mathcal{I}|\mathbb{E}[q(1)] + (1-|\mathcal{I}|)\mathbb{E}[q(0)] \leq \mathbb{E}[q(1)]$ and $|\mathcal{I}|\mathbb{E}[S(1)] + (1-|\mathcal{I}|)\mathbb{E}[S(0)] \leq \mathbb{E}[S(1)]$.
2. Each consumer has the option to stay opted in, or to opt out if doing so provides higher utility. Since consumers' privacy choices do not exert externalities on other consumers' utilities, privacy choices can never decrease consumers' utilities.
3. This follows immediately from Claim 2.

■

## *Statements and Proofs of Claims*

**Claim 1.** *Inequality (5) implies that the ad network's strategy of selling Period 1 ads to the advertiser (instead of to the external passive advertisers at price $\underline{b}$) weakly dominates its strategy of not selling Period 1 ads to the advertiser.*

*Proof of Claim 1.* Since the advertiser's Period 2 surplus is completely extracted by the ad network, the advertiser's Period 1 surplus is $\max[\mu\phi p - R, 0]$, where $R$ is the reserve price set by the ad network in Period 1. If the ad network sells the Period 1 ad to the advertiser at the highest possible price $R = \mu\phi p$, then the ad network's total profit would be *at least* $\mu\phi p + \underline{b}$, because it can sell Period 2 ads at a price at least as high as $\underline{b}$. On the other hand, if the ad network sells Period 1 ads to the external advertisers at price $\underline{b}$, then its maximum total profit is $\underline{b} + \mu\phi p$. This completes the proof. ■

**Claim 2.** *The ad network's expected profit from opt-in consumers,* $\Pi_1$*, is weakly greater than that from opt-out consumers,* $\Pi_0$.

*Proof of Claim 2.* In Period 1, the advertiser bids $\mu\phi p$, which is the advertiser's Period 1 value of advertising. Therefore, the ad network's Period 1 payoff is $\mu\phi p$ for both opt-in and opt-out consumers. Thus, to prove the claim, it suffices to compare the ad network's Period 2 payoffs for opt-in and opt-out consumers. The ad network's Period 2 payoff from opt-in consumers is $\Pi_1^{(2)} = (1-\mu\phi)\left(\mathbb{P}\{s=T\}\max\left[\underline{b}, b_T\right] + \mathbb{P}\{s=B\}\max\left[\underline{b}, b_B\right]\right) + \mu\phi\underline{b}$, where $b_T$ and $b_B$ are as defined in (9); its Period 2 profit from opt-out consumers is $\Pi_0^{(2)} = \max\left[\underline{b}, (1-\mu)\mu\phi p\right]$. Due to convexity of the maximum operator, we have

$$
\begin{aligned}
\Pi_1^{(2)} &\geq (1-\mu\phi)\max\left[\underline{b}, \mathbb{P}\{s=T\}b_T + \mathbb{P}\{s=B\}b_B\right] + \mu\phi\underline{b} \\
&= (1-\mu\phi)\max\left[\underline{b}, \tilde{\xi}\right] + \mu\phi\underline{b} \\
&= \max\left[(1-\mu\phi)\underline{b}, (1-\mu)\mu\phi p\right] + \mu\phi\underline{b} \\
&= \max\left[\underline{b}, (1-\mu)\mu\phi p + \mu\phi\underline{b}\right] \\
&\geq \max\left[\underline{b}, (1-\mu)\mu\phi p\right] = \Pi_0^{(2)},
\end{aligned}
$$

where $\tilde{\xi} = \left(\rho\frac{1-\mu}{1-\mu\phi} + (1-\rho)\frac{\mu(1-\phi)}{1-\mu\phi}\right)b_T + \left((1-\rho)\frac{1-\mu}{1-\mu\phi} + \rho\frac{\mu(1-\phi)}{1-\mu\phi}\right)b_B$. In sum, we obtain that $\Pi_1 = \mu\phi p + (1-\mu\phi)\left(\mathbb{P}\{s=T\}\max\left[\underline{b}, b_T\right] + \mathbb{P}\{s=B\}\max\left[\underline{b}, b_B\right]\right) + \mu\phi\underline{b}$ is weakly greater than $\Pi_0 = \mu\phi p + \max\left[\underline{b}, (1-\mu)\mu\phi p\right]$. ■

# Web Appendix

## *Proofs of Extension Results*

*Proof of Proposition 7.* Let $\phi_H = \phi + \epsilon$ and $\phi_L = \phi - \epsilon$. Note that if $p = c$ or $p = \overline{V}$, where $c$ is the advertiser's marginal production cost, then the advertiser's profit is 0. Therefore, for any $\epsilon > 0$, the advertiser will set price between $(c, \overline{p}(\epsilon)]$ where $\overline{p}(\epsilon) < \overline{V}$. Let $\overline{p} = \sup_{\epsilon>0} \overline{p}(\epsilon)$, and assume that

$$0 < \epsilon \leq \frac{1}{3}\left(\phi - \frac{\underline{b}}{\mu \min\left[\omega(c), \omega\left(\overline{p}\right)\right]}\right), \tag{WA1}$$

where $\omega(p) = (1 - Z(p))\,(p - c)$, which is assumed to be strictly quasi-concave.

We consider three Period 2 subgames in turn:

1. Period 2 subgame in which Period 1 ad is shown to all opt-in consumers: In this subgame, the value of advertising to $T$-consumers is either $\mu\phi_L\omega(p)$ or $\mu\phi_H\omega(p)$. Therefore, the ad network decides between two reserve prices: $R = \mu\phi_L\omega(p)$ and $R = \mu\phi_H\omega(p)$. If $R = \mu\phi_L\omega(p)$, then the ad network's Period 2 profit is $\mu\phi_L\omega(p)$, and if $R = \mu\phi_H\omega(p)$, then it is $\frac{1}{2}\mu\phi_H\omega(p) + \frac{1}{2}\underline{b}$. Due to (WA1), it follows that the ad network earns a higher profit from the former.

   Under this reserve price, the advertiser's Period 2 surplus is

$$(1-\mu)\left(\mu\phi_H\omega(p) - \mu\phi_L\omega(p)\right) = (1-\mu)\mu\omega(p)\left(\phi_H - \phi_L\right), \tag{WA2}$$

   if the consumer is $H$-type, and it is 0 if the consumer is $L$-type.

2. Period 2 subgame in which Period 1 ad is shown to only $H$-type opt-in consumers. In this subgame, there would be $\frac{1}{2}(1-\mu)$ mass of $H$-type $T$-consumers, $\frac{1}{2}\mu$ mass of $H$-type $B$-consumers, and $\frac{1}{2}$ mass of $L$-type $T$-consumers. The ad network's Period 2 profit is

$$\left(\frac{1}{2}(1-\mu) + \frac{1}{2}\right)\mu\phi_L\omega(p),$$

   if it sets reserve price $R = \mu\phi_L\omega(p)$, and it would be

$$\frac{1}{2}(1-\mu)\mu\phi_H\omega(p) + \frac{1}{2}\underline{b},$$

   if it sets reserve price $R = \mu\phi_H\omega(p)$. Due to (WA1), the former is larger than the latter;

therefore, the reserve price for $T$-consumers is set to $R = \mu\phi_L\omega(p)$. Under this reserve price, the advertiser's Period 2 surplus is equal to (WA2) if the consumer is $H$-type, and is equal to 0 if the consumer is $L$-type.

3. Period 2 subgame in which Period 1 ad is not shown to any opt-in consumers. Similar to the analysis above of the subgame in which Period 1 ads are shown to all opt-in consumers, (WA1) implies that the ad network sets reserve price $R = \mu\phi_L\omega(p)$. Furthermore, the advertiser's Period 2 surplus if the consumer is $H$-type is $\mu(\phi_H - \phi_L)\omega(p)$; if the consumer is $L$-type, it is 0.

With the Period 2 subgames at hand, we solve for the ad network's Period 1 reserve prices. In particular, the ad network sets reserve price in anticipation of the advertiser's expected Period 2 surplus which are derived above. The ad network considers three reserve prices: the highest reserve price at which the advertiser buys (a) ads regardless of consumer's type, (b) ads only if the consumer is $H$-type, and (c) no ads. The reserve prices under these three regimes, respectively, are

$$\begin{aligned}
R_1^a &= \mu\phi_L\omega(p), \\
R_1^b &= \mu\phi_H\omega(p) + (1-\mu)\mu\omega(p)(\phi_H - \phi_L), \\
R_1^c &= \infty.
\end{aligned}$$

The ad network's expected profit under each regime is

$$\begin{aligned}
\pi_N^a &= \mu\phi_L\omega(p) + (1-\mu)\mu\phi_L\omega(p) + \mu\underline{b}, \\
\pi_N^b &= \frac{1}{2}\mu\omega(p)\left(\phi_H + (1-\mu)(\phi_H - \phi_L)\right) + \frac{1}{2}\underline{b} + \frac{1}{2}(2-\mu)\mu\phi_L\omega(p) + \frac{1}{2}\mu\underline{b}, \\
\pi_N^c &= \mu\phi_L\omega(p),
\end{aligned}$$

where the second summands in $\pi_N^a$ and $\pi_N^b$ are obtained above from the ad network's maximum Period 2 profits under the respective subgames. Since $\pi_N^a \geq \pi_N^c$, the third regime never constitutes an equilibrium. We compare $\pi_N^a$ and $\pi_N^b$. $\frac{d}{d\epsilon}\pi_N^a - \pi_N^b = -\frac{1}{2}\mu\omega(p)(5-3\mu) < 0$ and $(\pi_N^a - \pi_N^b)|_{\epsilon=0} = \frac{1}{2}(1-\mu)(\mu\phi\omega(p) - \underline{b})$, which is positive by (WA1). Therefore, $\pi_N^a > \pi_N^b$ if and only if $\epsilon < \frac{(1-\mu)(\mu\phi\omega(p)-\underline{b})}{\mu\omega(p)(5-3\mu)}$.

We assume

$$\epsilon \leq \frac{1-\mu}{5-3\mu}\left(\phi - \frac{\underline{b}}{\mu \min[\omega(c), \omega(\overline{p})]}\right) \tag{WA3}$$

such that the ad network sets reserve price $R = R_1^a$. (Note that since $\frac{1-\mu}{5-3\mu} \leq \frac{1}{5}$, (WA3) is a stronger condition than (WA1).) Therefore, the advertiser's expected surplus is

$$\frac{1}{2}\left(\mu\omega(p)\phi_H - \mu\omega(p)\phi_L + (\text{WA2})\right) = \frac{1}{2}\left(\phi_H - \phi_L\right)\mu\left(2-\mu\right)\omega(p) = \epsilon\mu\left(2-\mu\right)\omega(p). \tag{WA4}$$

We now solve for the consumers' privacy choices. Since (WA1) implies that $\mu\phi_L\omega(p) \geq \underline{b}$, we obtain that all opt-in consumers see Period 1 ad and all opt-in consumers in funnel state $T$ see repeated ads in Period 2; therefore,

$$\mathbb{E}[q(1)]^{\text{ext}} = 1-\mu \text{ and } \mathbb{E}[S(1)]^{\text{ext}} = \sigma(p)\mathbb{E}_\phi\left[\mu\phi\left(2-\mu\right)\right] = \sigma(p)\mu\phi\left(2-\mu\right),$$

where

$$\sigma(p) = \int_p^{\overline{V}} v - p\, dZ(v) \tag{WA5}$$

and $\mathbb{E}_\phi[h(\phi)] = \frac{1}{2}\left(h(\phi-\epsilon) + h(\phi+\epsilon)\right)$ for any $h(\cdot)$.

Similarly, all opt-out consumers see Period 1 ad; they see repeated ads in Period 2 if and only if $(1-\mu)\mu\mathbb{E}_\phi[\phi]\omega(p) \geq \underline{b}$. Therefore,

$$\mathbb{E}[q(0)]^{\text{ext}} = \begin{cases} 1 & \text{if } (1-\mu)\mu\phi\omega(p) \geq \underline{b}, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\mathbb{E}[S(0)]^{\text{ext}} = \sigma(p)\mu\phi \times \begin{cases} 2-\mu & \text{if } (1-\mu)\mu\phi\omega(p) \geq \underline{b}, \\ 1 & \text{otherwise.} \end{cases}$$

We assume

$$\epsilon \leq \mu\phi, \tag{WA6}$$

so that $(1-\mu)\mu\phi\omega(p) \geq \underline{b} \Rightarrow \mu\phi_L\omega(p) \geq \underline{b}$. Therefore, there are two regions to consider: (i) $\underline{b} \leq (1-\mu)\mu\phi\omega(p)$, and (ii) $(1-\mu)\mu\phi\omega(p) < \underline{b} \leq \mu\phi_L\omega(p)$. In Region (i), consumers obtain the same product surplus, so their privacy choice is determined by the differences in expected intensity of repeated ads: a consumer opts in if and only if $u_\theta(1) \geq u_\theta(0) \Leftrightarrow -\eta\mathbb{E}[q(1)]^{\text{ext}} - \theta \geq -\eta\mathbb{E}[q(0)]^{\text{ext}} \Leftrightarrow \theta \leq \eta\mu$. In Region (ii), a consumer opts in if and only if $u_\theta(1) \geq u_\theta(0) \Leftrightarrow -\eta(1-\mu) + \sigma(p)\mu\phi\,(2-\mu) - \theta \geq \sigma(p)\mu\phi \Leftrightarrow \theta \leq (1-\mu)(\mu\phi\sigma(p) - \eta)$.

Combined with advertiser's total expected surplus from opt-in consumers in (WA4), we obtain the advertiser's total expected profit with respect to product price $p$:

$$\pi_A = \epsilon\mu(2-\mu)\omega(p) \times \begin{cases} F(\eta\mu), & \text{if } (1-\mu)\mu\phi\omega(p) > \underline{b}, \\ F\left((1-\mu)(\mu\phi\sigma(p)-\eta)\right) & \text{otherwise.} \end{cases} \tag{WA7}$$

Next, we determine the advertiser's optimal product price $p$. If $\frac{\underline{b}}{(1-\mu)\mu\phi} > \max_p \omega(p)$, then $p^* = p_{(ii)} = \operatorname{argmax}_p \omega(p)F\left((1-\mu)(\mu\phi\sigma(p)-\eta)\right)$. If $\frac{\underline{b}}{(1-\mu)\mu\phi} \leq \max_p \omega(p)$, then there are three candidate prices: $p_{(i)} = \operatorname{argmax}_p \omega(p)$, $\min\left[\underline{p}', p_{(ii)}\right]$, and $\max\left[\overline{p}', p_{(ii)}\right]$, where $\underline{p}' < \overline{p}'$ are the two roots of $(1-\mu)\mu\phi\omega(p) = \underline{b}$. Note that the third price candidate $\max\left[\overline{p}', p_{(ii)}\right]$ simplifies to $\overline{p}'$ because $\frac{d}{dp}F\left((1-\mu)(\mu\phi\sigma(p)-\eta)\right) \leq 0$ (because $\sigma(p)$, defined in (WA5) as consumer's expected product surplus, decreases in $p$) implies that the price that maximizes $\omega(p)F\left((1-\mu)(\mu\phi\sigma(p)-\eta)\right)$ is less than the price that maximizes $\omega(p)$. Therefore, $p_{(ii)} \leq p_{(i)}$ and, by concavity of $\omega(p)$, $p_{(i)} \leq \overline{p}'$.

We derive conditions under which the product price under endogenous consumer opt-in is strictly higher than that under the all-opt-in benchmark. Note that if all consumers are opted in by default, then the optimal product price is $p^*_{\text{all}} = p_{(i)}$. Since the third price candidate $\overline{p}'$ is the only price candidate that is greater than $p_{(i)}$, a necessary and sufficient condition is when $\overline{p}'$ maximizes (WA7), the advertiser's surplus under endogenous opt-in.

First, note that

$$\begin{aligned} \pi_A(p = \underline{p}') &= \epsilon\mu(2-\mu)\omega(\underline{p}')F\left((1-\mu)(\mu\phi\sigma(\underline{p}')-\eta)\right) \\ &= \epsilon\mu(2-\mu)\frac{\underline{b}}{(1-\mu)\mu}F\left((1-\mu)(\mu\phi\sigma(\underline{p}')-\eta)\right) \\ &\geq \epsilon\mu(2-\mu)\frac{\underline{b}}{(1-\mu)\mu}F\left((1-\mu)(\mu\phi\sigma(\overline{p}')-\eta)\right) \\ &= \pi_A(p = \overline{p}'), \end{aligned}$$

where the second and fourth equalities are due to the definitions of $\underline{p}'$ and $\overline{p}'$, and the inequality is due to consumer's expected surplus $\sigma(p)$ decreasing in $p$. Therefore, the second price candidate always weakly dominates the third. The second price candidate is ruled out if and only if marginal cost $c$ is greater than $\underline{p}'$.

If $c \geq \overline{p}$, then $c$ is the optimal price, both under endogenous opt-in and under all-opt-in default; therefore, we need $\underline{p}' < c < \overline{p}'$. Second, we need $\pi_A(p = \max[c, p_{(i)}]) < \pi_A(p = \overline{p}')$, where we accounted for the marginal cost in the maximizing price for the first regime. This inequality simplifies to

$$\epsilon\mu(2-\mu)\omega(\max[c, p_{(i)}])F(\eta\mu) < \epsilon\mu(2-\mu)\omega(\overline{p}')F((1-\mu)(\mu\phi\sigma(\overline{p}') - \eta)). \tag{WA8}$$

Since (a) left-hand side increase in $\eta$ while the right-hand side decreases in $\eta$, and (b) the inequality holds at $\eta = 0$ (where left-hand side reduces to 0 due to $F(\eta\mu)|_{\eta=0}= 0$ and right-hand side is positive), we obtain that there exists a unique $\hat{\eta} > 0$ such that (WA8) holds if and only if $\eta < \hat{\eta}$.

In sum, if $\epsilon < \hat{\epsilon} = \min\left[\frac{1-\mu}{5-3\mu}\left(\phi - \frac{b}{\mu \min[\omega(c), \omega(\overline{p})]}\right), \mu\phi\right]$, then the product price under endogenous opt-in is higher than that under all-opt-in default if and only if $\underline{p}' < c < \overline{p}'$ and $\eta < \hat{\eta}$. ■

*Proof of Corollary 2.* Let $\underline{p}'$, $\overline{p}'$, and $p_{(i)}$ be as defined in the proof of Proposition 7. Suppose $\underline{p}' < c < \overline{p}'$ and $\eta < \hat{\eta}$ such that the product price under all-opt-in default is $p_{(i)}$, which is lower than $\overline{p}'$, the product price under endogenous opt-in. In this case, all consumers who opt in to tracking under the endogenous privacy choices regime face the same advertising outcomes but higher product prices, such that they have strictly lower utilities than they do under the all-opt-in benchmark. Consumers who opt out of tracking under the endogenous privacy choices regime see fewer repeated ads (i.e., instrumental gain of $\eta(1-\mu)$) but obtain lower product surplus (i.e., instrumental loss of $\sigma(p_{(i)})\mu\phi\,(2-\mu) - \sigma(\overline{p}')\mu\phi$). Therefore, if

$$\eta < \frac{\sigma(p_{(i)})\mu\phi\,(2-\mu) - \sigma(\overline{p}')\mu\phi}{1-\mu},$$

then opt-out consumers are also worse off in the endogenous privacy choice regime than in the all-opt-in benchmark. ■

*Proof of Proposition 8.* Let $\tilde{\theta}$ be as defined in (14). If the ad network reduces consumers' privacy costs uniformly by $\Delta_\theta$, then all consumers with privacy cost smaller than $\tilde{\theta}'$ opt in to tracking, where

$$\tilde{\theta}' = \begin{cases} \tilde{\theta} + \Delta_\theta & \text{if } \underline{b} \leq b_T, \\ 0 & \text{if } b_T < \underline{b}. \end{cases}$$

Therefore, if $b_T < \underline{b}$, then there are no changes in the opt-in rate, which implies that the ad network will not invest in data security. On the other hand, if $\underline{b} \leq b_T$, then the ad network's profit increases by

$$\left(F\left(\tilde{\theta} + \Delta_\theta\right) - F\left(\tilde{\theta}\right)\right)(\Pi_1 - \Pi_0).$$

The ad network maximizes this profit increase against the cost of reducing the privacy cost by $\Delta_\theta$, which is $\frac{k}{2}\Delta_\theta^2$. The first-order condition simplifies to the condition stated in the proposition. This completes the proof. ■

## *Analysis for Ad Wearout Leading to Reduced Ad Responsiveness*

Let $\underline{b}' = \underline{b}/\alpha$, and let $\rho'_B$ and $\rho'_T$ be the same as (10) except $\underline{b}$ replaced with $\underline{b}'$.

- **Period 2 subgame for opt-out consumers where Period 1 ad is shown**: The value of advertising to opt-out consumers in Period 2 is $(1-\mu)\alpha\mu\phi p$; therefore, Period 2 ads are shown if and only if $\underline{b} \leq (1-\mu)\alpha\mu\phi p$. The advertiser's surplus is $\mu(1-\phi)\phi p$, and the ad network's profit is $\max[\underline{b}, (1-\mu)\alpha\mu\phi p]$
- **Period 2 subgame for opt-out consumers where Period 1 ad is not shown**: The value of advertising to opt-out consumers is $\mu\phi p$; therefore, ads are shown if and only if $\underline{b} \leq \mu\phi p$. The advertiser's surplus is 0, and the ad network's profit is $\max[\underline{b}, \mu\phi p] = \mu\phi p$ by (5).
- **Period 1**: the ad network can charge Period 1 ads as high as $\mu\phi p$ for a total profit of $\mu\phi p + \max[\underline{b}, (1-\mu)\alpha\mu\phi p]$, which is weakly greater than its profit from not selling to the advertiser in Period 1 for total profit of $\underline{b} + \mu\phi p$ due to (5). Therefore, the ad network is weakly better off selling Period 1 ads, which the advertiser buys.

We obtain

$$\mathbb{E}[q'(0)] = \begin{cases} 1 & \text{if } \underline{\mu}' < \mu \leq \overline{\mu}', \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\mathbb{E}[S'(0)] = \mu\phi(1-p) \times \begin{cases} 1 + (1-\mu)\alpha & \text{if } \underline{\mu}' < \mu \leq \overline{\mu}', \\ 1 & \text{otherwise,} \end{cases}$$

where

$$\underline{\mu}' = \left(1 - \sqrt{1 - 4\underline{b}'/(\phi p)}\right)/2 \text{ and } \overline{\mu}' = \left(1 + \sqrt{1 - 4\underline{b}'/(\phi p)}\right)/2. \tag{WA9}$$

- **Period 2 subgame for opt-in consumers where Period 1 ad is shown**: The value of advertising to $T$- and $B$-signal consumers is $b'_T = \mathbb{P}\{T|s = T\}\alpha\mu\phi p$ and $b'_B = \mathbb{P}\{T|s = B\}\alpha\mu\phi p$, respectively. Since $\mathbb{P}\{T|s = T\} \geq \mathbb{P}\{T|s = B\}$, Period 2 ads are shown to (a) all opt-in non-purchasers if $\underline{b} \leq b'_B$, (b) only $T$-signal consumers if $b'_B < \underline{b} \leq b'_T$, and (c) none of the opt-in consumers if $b'_T < \underline{b}$. The advertiser's Period 2 surplus is 0. The ad network's profit is $\mu\phi\underline{b} + (1-\mu\phi)\left(\mathbb{P}\{s = T\}\max\left[\underline{b}, b'_T\right] + \mathbb{P}\{s = B\}\max\left[\underline{b}, b'_B\right]\right)$.
- **Period 2 subgame for opt-in consumers where Period 1 ad is not shown**: The value of advertising is $\mu\phi p$; therefore, ads are shown if and only if $\underline{b} \leq \mu\phi p$. The advertiser's surplus is 0, and the ad network's profit is $\max\left[\underline{b}, \mu\phi p\right]$.
- **Period 1**: the ad network can charge Period 1 ads as high as $\mu\phi p$ for a total profit of $\mu\phi p + \mu\phi\underline{b} + (1-\mu\phi)\left(\mathbb{P}\{s = T\}\max\left[\underline{b}, b'_T\right] + \mathbb{P}\{s = B\}\max\left[\underline{b}, b'_B\right]\right)$, which is greater than its profit from not selling to the advertiser in Period 1 for total profit of $\underline{b} + \mu\phi p$. Therefore, ad network is weakly better off selling ads in Period 1 to the advertiser.

  We obtain the same expected ad repetition intensity as (11) except $\underline{b}$, $\rho_B$, and $\rho_T$ are replaced with $\underline{b}'$, $\rho'_B$, and $\rho'_T$, respectively.

  The opt-in consumers' expected product surplus is also obtained in a similar manner by replacing (12) with the new prime parameters defined above. In addition, the conversion probabilities from Period 2 advertising are reduced by $\alpha$ such that

$$\mathbb{E}[S'(1)] = \mu\phi(1-p) \times \begin{cases} 1 + (1-\mu)\alpha & \text{if } \underline{b}' \leq \frac{1-\mu}{1-\mu\phi}\mu\phi p \text{ and } \rho \leq \rho_B', \\ 1 + (1-\mu)\rho\alpha & \text{if } \max[\rho_B', \rho_T'] < \rho, \\ 1 & \text{otherwise.} \end{cases}$$

Overall, we see that the advertising repetition and product surplus outcomes are qualitatively the same as the main model. Repetition-induced reduction in consumers' responsiveness to advertising has two effects. First, it reduces ad effectiveness such that conditions for the advertiser's bid beating the reserve price become stricter (loosely, bid must beat $\underline{b}'$ instead of $\underline{b}$, where $\underline{b}' = \underline{b}/\alpha \geq \underline{b}$). Second, it reduces consumers' expected product surplus because, in this framework, ad wearout causes consumers to respond less to ads that may generate positive information value.

## *Analysis for Naïve Consumers*

In the main model, we assumed that consumers were sufficiently knowledgeable about the ramifications of their privacy choices — in particular, about how their privacy choices affect the ad network's and the advertiser's strategies. This analysis sheds meaningful light on the long-term implications of the various privacy regulations (e.g., the GDPR and CCPA) as the players in the ad ecosystem (in particular, the consumers) learn and adapt to the new regulatory landscape. Indeed, a key requirement of these regulations is that firms inform consumers "in a concise, transparent, intelligible and easily accessible form, using clear and plain language" about their privacy choices and how their data are processed (Intersoft Consulting, 2021). Nevertheless, one may wonder how the qualitative insights from the main model would change if consumers less than fully understand the downstream implications of their data sharing decisions. To that end, we extend the main model by analyzing a situation with a mix of "naïve" and "sophisticated" consumers.

While sophisticated consumers anticipate the downstream ad network and advertiser strategies (i.e., as the consumers from the main model do), naïve consumers hold prior beliefs about the ad network and the advertiser's strategies that do not always align with the actual outcomes. To formalize these ideas, we assume that, for some $\beta \in [0, 1]$, $\beta$ proportion of consumers are naïve, and $1 - \beta$ proportion are sophisticated. Naïve consumers have heterogeneous beliefs denoted by the parameter $\gamma$ distributed according to a distribution function $G$ on the support $[0, 1]$.

A $\gamma$-type naïve consumer holds the following belief: with probability $\gamma$, an opt-in consumer will see ads in both periods while opt-out consumers see no ads; and with probability $1-\gamma$, an opt-in consumer will see no ads while opt-out consumers will see ads in both periods. $\gamma < \frac{1}{2}$ implies the belief that opting in to tracking will lead to fewer repeated ads being seen, while $\gamma > \frac{1}{2}$ implies the belief that opting in to tracking will lead to more repeated ads being seen; both beliefs may be erroneous. For instance, $\gamma = 1$ corresponds to the erroneous belief that opting out of tracking will have the same effect as "unsubscribing" from all ads. Note that this formulation of consumer naïvete is consistent with the literature. For instance, the case of $\gamma = \frac{1}{2}$ subsumes the myopic consumer who disregards how their present actions impact their future utilities (Taylor, 2004). This is because if $\gamma = \frac{1}{2}$, then the consumer believes that any privacy choice has the same instrumental implications. This formulation is also consistent with consumers omitting (either intentionally or not) certain outcome variables that impact her total utility (Heidhues and Kőszegi, 2017).

Given this belief, the expected utility of a $\gamma$-type consumer with privacy cost $\theta$ and ad wearout parameter $\eta$ who opts in to tracking is $u_\gamma(1) = \gamma\left(\overline{S} - \eta\right) - \theta$, and the expected utility of the same consumer who opts out of tracking is $u_\gamma(0) = (1-\gamma)\left(\overline{S} - \eta\right)$, where

$$\overline{S} = (1-p)\mu\phi(2-\mu). \tag{WA10}$$

Therefore, a $\gamma$-type naïve consumer opts in to tracking if and only if $\theta \leq (2\gamma - 1)(\overline{S} - \eta)$. Based on this, we obtain the following proposition that summarizes the two conditions under which a naïve consumer opts in to tracking.

**Proposition 9.** *A $\gamma$-type naïve consumer opts in to tracking if and only if her privacy cost $\theta$ is sufficiently small and either (i) $\overline{S} < \eta$ and $\gamma \leq \frac{1}{2}$, or (ii) $\eta \leq \overline{S}$ and $\frac{1}{2} < \gamma$, where $\overline{S}$ is as defined in* (WA10).

The intuition behind Proposition 9 is the following. If naïve consumers believe that by opting in to tracking they are less likely to be exposed to repeated ads (i.e., $\gamma \leq 1/2$) and they are highly sensitive ad wearout, then they will opt in to tracking. Conversely, if they are more concerned about product surplus than ad wearout (i.e., small $\eta$) and believe that they will see more ads from opting in, then they will allow tracking.

Overall, the $\beta$ segment of naïve consumers make privacy choices according to Proposition 9 independent of the downstream implications for the ad network's and advertiser's strategies. As the naïve consumers' privacy choices do not impact the privacy choices of $1-\beta$ segment of sophisticated consumers, the qualitative insights from the main model remain the same insofar as $\beta$ is not too large.

## *Analysis for Advertising Competition*

In this section, we analyze a scenario with two symmetric advertisers who compete for a single ad slot in each period. We assume that both advertisers' product prices are exogenous and the same $p \leq 1$—by abstracting from the product pricing decisions, we can focus on the impact of advertising competition. In addition, we assume that opt-in consumers' funnel states can be perfectly inferred and that there are no external passive advertisers (whose maximum bid is $\underline{b}$ in the main model).

More formally, there are two advertisers by $j \in \{1, 2\}$ who offer perfect substitutes and compete for advertising inventory in each period. For ease of exposition, we hereafter denote Advertiser $j$ by $A_j$. We modify the consumer model by considering a tuple of funnel states $(f_1, f_2)$, where $f_j \in \{T_j, B_j\}$ denotes the consumer's funnel state for $A_j$. Each consumer starts from $(T_1, T_2)$, i.e., at the top of the two purchase funnels, one for $A_1$ and one for $A_2$. An ad exposure from $A_j$ moves the consumer down from funnel state $T_j$ to funnel state $B_j$ with probability $\mu$, and has no effect with probability $1-\mu$. Consistent with the main model, we assume that a consumer in funnel state $B_j$ realizes a product match with probability $\phi$. In contrast to the main model, if a consumer is in funnel states $(B_1, B_2)$, then she realizes a product match independently for each advertiser. Moreover, if she realizes a match for both advertisers' products, then she purchases randomly from one of the advertisers.

We focus on a symmetric equilibrium in which both advertisers bid the same amount in Period 1. We assume that a tie is broken randomly, and without loss of generality, assume that $A_1$ wins the

Period 1 auction.

**Period 2 subgame for opt-out consumers:**

Consider the Period 2 subgame for opt-out consumers in which $A_1$'s ad was shown in Period 1. In Period 2, there are $1-\mu$ consumers in $(T_1, T_2)$, $\mu(1-\phi)$ consumers in $(B_1, T_2)$, and $\mu\phi$ purchasers.

If $A_1$ wins the Period 2 ads, then $A_1$'s Period 2 payoff is $\pi_1(\text{win}) = (1-\mu)\mu\phi p$. If $A_2$ wins the Period 2 ads, then the $\mu$ fraction of consumers move from $T_2$ to $B_2$, while $1-\mu$ fraction stay in $T_2$. Thus, $A_1$'s Period 2 payoff is 0. Therefore, $A_1$'s bid is

$$b_1 = \pi_1(\text{win}) - \pi_1(\text{lose}) = (1-\mu)\mu\phi p. \tag{WA11}$$

Since $\mu\phi$ fraction of consumers buy $A_1$'s product after its Period 1 ad exposure, there are $1-\mu\phi$ non-purchasers in Period 2 that $A_2$ can potentially induce to purchase through ads. If $A_2$ wins the Period 2 ads, then $A_2$'s Period 2 payoff is $\pi_2(\text{win}) = (1-\mu\phi)\mu\phi p$. If $A_1$ wins the Period 2 ads, then $A_2$'s Period 2 payoff is $\pi_2(\text{lose}) = 0$. Therefore, $A_2$'s bid is

$$b_2 = \pi_2(\text{win}) - \pi_2(\text{lose}) = (1-\mu\phi)\mu\phi p. \tag{WA12}$$

Since (WA12) > (WA11), we obtain that for opt-out consumers, the advertiser that loses the Period 1 auction wins in Period 2. Therefore, $\mathbb{E}[q(0)] = 0$ and $\mathbb{E}[S(0)] = (1-p)\left(\mu\phi + (1-\mu\phi)\mu\phi\right) = (2-\mu\phi)\mu\phi(1-p)$.

**Period 2 subgame for opt-in consumers:**

In the subgame for opt-in consumers, there are two types consumers that advertisers bid for: consumers in funnel states $(T_1, T_2)$ and those in funnel states $(B_1, T_2)$. Both advertisers bid $\mu\phi p$ for the former group, so the advertisers tie and each wins with equal probability. On the other hand, for consumers in $(B_1, T_2)$, their bids diverge. To derive the bids, note that if $A_1$ bids nothing because these are consumers who did not realize a product match with $A_1$:

$$b_1(B_1, T_2) = 0. \tag{WA13}$$

If $A_2$ wins $(B_1, T_2)$, $A_2$'s payoff is

$$\pi_2(\text{win}(B_1, T_2)) = \mu\phi p,$$

and if $A_2$ loses, its payoff is 0. Therefore, $A_2$'s bid for $(B_1, T_2)$ is

$$b_2(B_1, T_2) = \mu\phi p. \tag{WA14}$$

Thus, (WA14) $\geq$ (WA13): the advertiser that lost in Period 1 wins the Period 2 auction for consumers who transitioned down Period 1 winner's funnel but did not purchase.

Taken together, we have $\mathbb{E}[q(1)] = \frac{1-\mu}{2}$ and $\mathbb{E}[S(1)] = (1-p)\left(\mu\phi + \mu(1-\phi)\mu\phi + (1-\mu)\mu\phi\right) = \mu\phi(1-p)\left(2-\mu\phi\right)$.

## *Analysis for Overlapping Generations Over Infinite Time Horizon*

In this section, we extend the time horizon of the main model to infinity and consider overlapping arrivals of consumers to test the robustness of the main results beyond the two-period framework. To that end, we assume that in each Period $t \in \{1, 2, ...\}$, $T$-consumers of mass $\frac{1}{2}$ join the system, visit two web pages, one in each period, and then exits the system. Therefore, in each period, there is a unit mass of consumers, half of which are "old" and half of which are "new." We restrict our analysis to the opt-out consumers and demonstrate that the non-monotonic pattern of ad repetition intensity carries over.

We solve for a Markov Perfect equilibrium. In any given period, the payoff-relevant state can be fully characterized by the mass of $T$-consumers, which we denote by $|T|$. There are two states: one in which the advertiser's ad was shown in the previous previous, and another in which it was not. In the former, $|T| = \frac{1}{2}(1-\mu) + \frac{1}{2}$, where the first term represents the "old" consumers to whom the match information was not successfully conveyed, and the second term the "new" consumers who join the system. In the latter, $|T| = 1$ because "old" consumers are not exposed to advertising such that all of them remain in state $T$.

Given these two states, the ad network considers three ad selling strategies.

**I. Induce advertiser to advertise every period.**

The incentive compatibility constraint to induce this outcome is $\frac{1}{2}(2-\mu)\mu\phi p - R + \delta V_1 \geq \delta V_0$, where $\delta \in [0, 1)$ is the discount factor, and $V_1$ and $V_0$ are the advertiser's continuation values in a state where the total mass of $T$-consumers are $\frac{1}{2}(2-\mu)$ and 1, respectively. In equilibrium, the ad network sets $R$ such that this constraint binds; i.e.,

$$\frac{1}{2}(2-\mu)\mu\phi p - R + \delta V_1 = \delta V_0. \tag{WA15}$$

By recursive reasoning, we obtain

$$V_0 = \max[0 + \delta V_0, \mu\phi p - R + \delta V_1] \tag{WA16}$$

and

$$V_1 = \frac{1}{2}(2-\mu)\mu\phi p - R + \delta V_1. \tag{WA17}$$

Suppose $\delta V_0 > \mu\phi p - R + \delta V_1$, such that, from (WA16), $V_0 = \delta V_0 \Rightarrow V_0 = 0$. Then, (WA15) and (WA17) imply $V_1 = 0$ and $R = \frac{1}{2}(2-\mu)\mu\phi p$. But then, $\mu\phi p - R + \delta V_1 = \frac{1}{2}\mu^2\phi p > 0$, a contradiction. Therefore, it must be that $\delta V_0 < \mu\phi p - R + \delta V_1$, such that, from (WA16), $V_0 = \mu\phi p - R + \delta V_1$. Substituting this into (WA15) and (WA17) and solving yields $R = \frac{1}{2}(2-(\delta+1)\mu)\mu\phi p$. The ad network's average per-period payoff is

$$\frac{1}{2}(2-(\delta+1)\mu)\mu\phi p. \tag{WA18}$$

**II. Induce advertiser to advertise in alternating periods.**

The incentive compatibility constraints to induce this outcome are

$$\mu\phi p - R + \delta V_1' \geq \delta V_0' \tag{WA19}$$

and

$$\delta V_0' \geq \frac{1}{2}(2-\mu)\mu\phi p - R + \delta V_1', \tag{WA20}$$

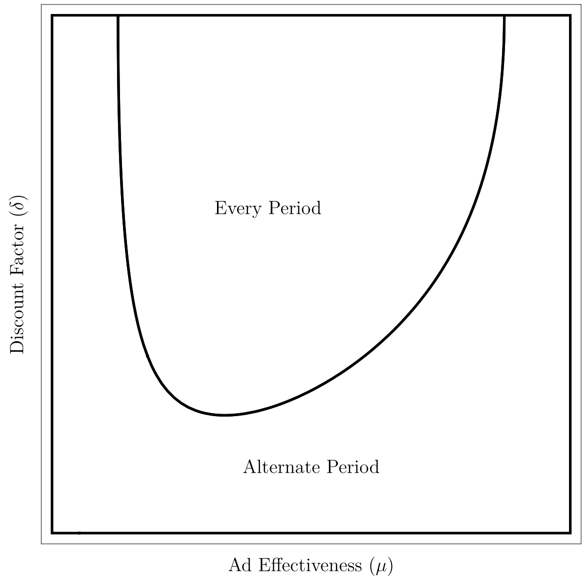


*Figure WA1: Overlapping Generations over Infinite Time Horizon: Advertising Strategy*

where $V_1'$ and $V_0'$ are the advertiser's continuation values in a state where the total mass of $T$-consumers is $\frac{1}{2}(2-\mu)$ and 1, respectively. Combining inequalities (WA19) and (WA20), we obtain

$$\frac{1}{2}(2-\mu)\mu\phi p+\delta(V_1'-V_0')<R<\mu\phi p+\delta(V_1'-V_0'); \tag{WA21}$$

therefore, in equilibrium, the ad network sets $R=\mu\phi p+\delta(V_1'-V_0')$. By recursive reasoning, we obtain $V_1'=\mu\phi p-R+\delta V_1'$ and $V_0'=\max[\delta V_0',\mu\phi p-R+\delta V_1']$. It can be shown that the system of equations yields a unique solution of $V_1'=V_0'=0$ and $R=\mu\phi p$. Therefore, the ad network's average per-period payoff is

$$(1-\delta)\left(\frac{1}{1-\delta^2}\mu\phi p+\frac{\delta}{1-\delta^2}\underline{b}\right)=\frac{1}{1+\delta}\mu\phi p+\frac{\delta}{1+\delta}\underline{b}. \tag{WA22}$$

**III. Induce advertiser to never advertise.**

In this case, the reserve price is set high enough that the advertiser never buys, and the ad network's per period payoff is simply $\underline{b}$.

By (5), the third strategy is dominated by the second. Therefore, the ad network chooses between the first and second strategy. Comparing (WA18) and (WA22), we obtain that for opt-out consumers, the advertiser's ad is shown in every period if ad effectiveness is intermediate, and in alternating periods otherwise (see Figure WA1). Specifically, the advertiser's ad is shown in every period if $\underline{\mu}'' < \mu \leq \overline{\mu}''$, and it is shown in alternating periods otherwise, where

$$\underline{\mu}'' = \frac{\delta p\phi - \sqrt{\delta p\phi\,(\delta p\phi - 2(\delta+1)^2\underline{b})}}{(\delta+1)^2 p\phi} \text{ and } \overline{\mu}'' = \frac{\delta p\phi + \sqrt{\delta p\phi\,(\delta p\phi - 2(\delta+1)^2\underline{b})}}{(\delta+1)^2 p\phi}. \tag{WA23}$$

These bounds are well-defined if and only if $\delta > \left(\phi p - 4\underline{b} - \sqrt{\phi p(\phi p - 8\underline{b})}\right)/4\underline{b}$, and when they are well-defined, the size of the interval $[\underline{\mu}'', \overline{\mu}'']$ increases in $\delta$; i.e., the parameter region for which ads are shown in every period increases in $\delta$. The intuition is the following: from the advertiser's perspective, advertising consecutively is worthwhile because the successive ad may induce consumers for whom the first ad was not successful to purchase. Thus, the value of this strategy is created in the successive period (as opposed to the present period). Therefore, as the advertiser values the future more heavily (i.e., larger $\delta$), this strategy of advertising consecutively becomes more valuable.

### *Analysis for General Initial Funnel State Distribution*

In the main model, we assumed that all consumers started their journey from funnel state $T$. While this simplifying assumption allowed us to derive basic insights cleanly, in practice, initial funnel states may vary across consumers. We capture this heterogeneity by allowing $\lambda$ proportion of consumers to start from funnel state $B$, and $1-\lambda$ from funnel state $T$, where $\lambda \in [0,1]$. Broadly, $\lambda$ can be interpreted as the advertiser's "brand knowledge:" the larger the $\lambda$, the greater the extent to which consumers possess match information pertaining to the advertiser's product. To avoid consumers' privacy choices potentially signaling their initial funnel states, we assume that consumers make privacy choices before learning their funnel states. We provide the analysis in the Web Appendix and discuss the key results here.

Since $B$-consumers are less responsive to advertising than $T$-consumers, we anticipate larger $\lambda$ to reduce the intensity of repeated advertising. As illustrated in Figure WA2, the ad repetition level plots confirm this intuition. Similar to the outcome from the extension where ad wearout reduces consumers' responsiveness to ads, the intensity of repeat advertising decreases as the value of advertising is diminished. However, the underlying mechanisms are different. In the other extension, the value of advertising decreased due to ad wearout negatively impacting consumers'

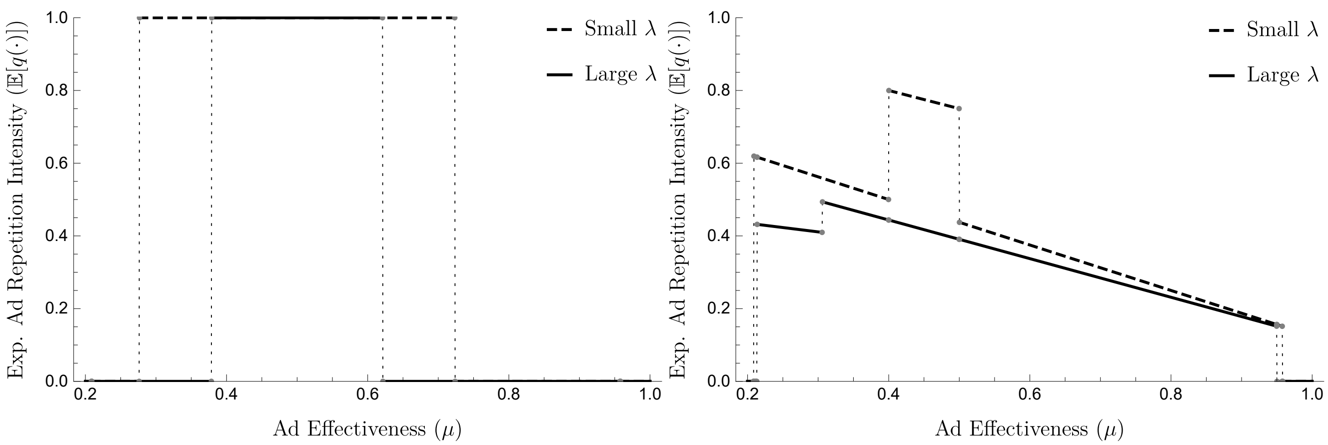


*(a) Opt-Out Consumers* *(b) Opt-In Consumers*

*Figure WA2: Intensity of Repeated Ads Under General Initial Distribution;* $\underline{b} = 0.05, p = \phi = \mu = 0.5, \rho = 0.75$

ad responsiveness; in this extension, the value of advertising falls because a fraction of consumers are already informationally saturated prior to ad exposures. Overall, in both extensions, the qualitative insights from the main model carry over with the moderating effect of diminished ad valuations.

We begin the analysis for opt-out consumers.

(a) **Period 2 subgame for opt-out consumers where Period 1 ad is shown**: the value of showing another ad in Period 2 is $(1-\lambda)(1-\mu)\mu\phi p$; therefore, repeated ads are shown if and only if $\underline{b} \leq (1-\lambda)(1-\mu)\mu\phi p$

(b) **Period 2 subgame for opt-out consumers where Period 1 ad is not shown**: the value of showing ad in Period 2 is $(1-\lambda)\mu\phi p$; therefore, repeated ads are shown if and only if $\underline{b} \leq (1-\lambda)\mu\phi p$

Consider the ad network's strategy in Period 1. The ad network can charge up to $(1-\lambda)\mu\phi p$. Therefore, selling Period 1 ad weakly dominates not selling Period 1 ad if and only if $(1-\lambda)\mu\phi p + \max[\underline{b}, (1-\lambda)(1-\mu)\mu\phi p] \geq \underline{b} + \max[\underline{b}, (1-\lambda)\mu\phi p]$, which is equivalent to $\underline{b} \leq (1-\lambda)\mu\phi p$.

Therefore, if $\underline{b} \leq (1-\lambda)(1-\mu)\mu\phi p$, then ads are shown in both periods; if $(1-\lambda)(1-\mu)\mu\phi p < \underline{b} \leq (1-\lambda)\mu\phi p$, then ads are shown only in Period 1; and otherwise, no ads are shown in either period.

We thus obtain

$$\mathbb{E}[q^{\text{gen}}(0)] = \begin{cases} 1 & \text{if } \underline{b} \leq (1-\lambda)(1-\mu)\mu\phi p, \\ 0 & \text{otherwise.} \end{cases}$$

As for expected product surplus, we distinguish between consumers who begin in funnel state $T$ and those who begin in funnel state $B$. For the former consumers, we have

$$\mathbb{E}[S_T^{\text{gen}}(0)] = \mu\phi(1-p) \times \begin{cases} 2-\mu & \text{if } \underline{b} \leq (1-\lambda)(1-\mu)\mu\phi p, \\ 1 & \text{if } (1-\lambda)(1-\mu)\mu\phi p < \underline{b} \leq (1-\lambda)\mu\phi p, \\ 0 & \text{otherwise,} \end{cases}$$

whereas for the latter consumers who begin in funnel state $B$, their product surplus is unaffected by ads. Since these consumers independently realize a product match with probability $\phi$, we have

$$\mathbb{E}[S_B^{\text{gen}}] = \phi(1-p). \tag{WA24}$$

Next, consider the opt-in consumers.

(a) **Period 2 subgame for opt-in consumers where Period 1 ad is shown to all non-purchasers**: there are $(1-\lambda)(1-\mu)$ $T$-consumers, $(1-\lambda)\mu(1-\phi)+\lambda(1-\phi)$ $B$-consumers, and $(1-\lambda)\mu\phi+\lambda\phi$ purchasers. Therefore, the value of advertising to $T$-signal consumers, and hence the bid for these consumers, is

$$b_T^a = \mathbb{P}\{f = T|s = T\}\mu\phi p = \frac{\rho(1-\lambda)(1-\mu)}{\rho(1-\lambda)(1-\mu)+(1-\rho)((1-\phi))((1-\lambda)\mu+\lambda)}\mu\phi p,$$

and similarly, the value of advertising to $B$-signal consumers, and hence the bid for these consumers, is

$$b_B^a = \mathbb{P}\{f = T|s = B\}\mu\phi p = \frac{(1-\rho)(1-\lambda)(1-\mu)}{(1-\rho)(1-\lambda)(1-\mu)+\rho((1-\phi))((1-\lambda)\mu+\lambda)}\mu\phi p.$$

Since $\mathbb{P}\{f = T|s = T\} \geq \mathbb{P}\{f = T|s = B\}$ (due to $\rho \geq 1/2$), there are three parametric regions to consider: $\underline{b} \leq b_B^a$, $b_B^a < \underline{b} \leq b_T^a$, and $b_T^a < \underline{b}$. In the first condition, ads are shown to all non-purchasers; in the second, ads are shown only to $T$-signal consumers, and in the third, no ads are shown to any opt-in consumers. Note that in the second condition, two groups of consumers see repeated ads: those who start off in funnel state $T$, do not purchase, and are correctly inferred as $T$-consumers (with probability $\rho$) and those who start off in funnel state $B$, do not purchase, and are misinferred as $T$-consumers (with probability $1-\rho$). Therefore, a consumer who starts in funnel state $T$ expects to see repeated ads at intensity

$$\mathbb{E}[q_T^a(1)] = \begin{cases} 1-\mu\phi & \text{if } \underline{b} \leq b_B, \\ (1-\mu)\rho + \mu(1-\phi)(1-\rho) & \text{if } b_B < \underline{b} \leq b_T, \\ 0 & \text{if } b_T < \underline{b}, \end{cases}$$

whereas a consumer who starts in funnel state $B$ expects to see repeated ads at intensity

$$\mathbb{E}[q_B^a(1)] = (1-\phi) \times \begin{cases} 1 & \text{if } \underline{b} \leq b_B, \\ 1-\rho & \text{if } b_B < \underline{b} \leq b_T, \\ 0 & \text{if } b_T < \underline{b}. \end{cases}$$

Similarly, the expected product surplus for a consumer who starts in funnel state $T$ is

$$\mathbb{E}[S_T^a(1)] = \mu\phi(1-p) \times \begin{cases} 2-\mu & \text{if } \underline{b} \leq b_B, \\ 1+(1-\mu)\rho & \text{if } b_B < \underline{b} \leq b_T, \\ 1 & \text{if } b_T < \underline{b}. \end{cases}$$

(b) **Period 2 subgame for opt-in consumers where Period 1 ad is shown to $T$-signal consumers**: the total number of $T$-signal consumers in Period 1 is $(1-\lambda)\rho + \lambda(1-\rho)$, so after Period 1 ads, there are $(1-\lambda)\rho(1-\mu) + (1-\lambda)(1-\rho)$ $T$-consumers (i.e., correctly inferred $T$-consumers who see ads but ads are not effective, and misinferred $T$-consumers who don't see ads), $(1-\lambda)\rho\mu(1-\phi) + \lambda(1-\phi)$ $B$-consumers (i.e., correctly inferred $T$-consumers who see ads and ads are effective but do not purchase, and $B$-consumers who do not realize a product match), and $(1-\lambda)\rho\mu\phi + \lambda\phi$ purchasers (i.e., correctly inferred $T$-consumers who purchase after seeing ads and $B$-consumers who realize a product match). The value of showing ads to $T$-signal consumers in Period 2, and hence the bids for these consumers, is

$$
\begin{aligned}
b_T^b &= \mathbb{P}\{f = T | s = T\}\mu\phi p \\
&= \frac{\rho\left((1-\lambda)\rho(1-\mu) + (1-\lambda)(1-\rho)\right)}{\rho\left((1-\lambda)\rho(1-\mu) + (1-\lambda)(1-\rho)\right) + (1-\rho)\left((1-\lambda)\rho\mu(1-\phi) + \lambda(1-\phi)\right)}\mu\phi p.
\end{aligned}
$$

Similarly we obtain

$$
\begin{aligned}
b_B^b &= \mathbb{P}\{f = T | s = B\}\mu\phi p \\
&= \frac{(1-\rho)\left((1-\lambda)\rho(1-\mu) + (1-\lambda)(1-\rho)\right)}{(1-\rho)\left((1-\lambda)\rho(1-\mu) + (1-\lambda)(1-\rho)\right) + \rho\left((1-\lambda)\rho\mu(1-\phi) + \lambda(1-\phi)\right)}\mu\phi p.
\end{aligned}
$$

Again, since $\mathbb{P}\{f = T|s = T\} \geq \mathbb{P}\{f = T|s = B\}$, there are three parametric regions to consider: $\underline{b} \leq b_B^b$, $b_B^b < \underline{b} \leq b_T^b$, and $b_T^b < \underline{b}$. A consumer who starts in funnel state $T$ expects to see repeated ads at intensity

$$
\mathbb{E}[q_T^b(1)] = \rho(1-\mu\phi) \times \begin{cases} 1 & \text{if } \underline{b} \leq b_B^b, \\ \rho & \text{if } b_B^b < \underline{b} \leq b_T^b, \\ 0 & \text{otherwise.} \end{cases}
$$

On the other hand, a consumer who starts in funnel state $B$ expects to see repeated ads at intensity

$$
\mathbb{E}[q_B^b(1)] = (1-\rho)(1-\phi) \times \begin{cases} 1 & \text{if } \underline{b} \leq b_B^b, \\ 1-\rho & \text{if } b_B^b < \underline{b} \leq b_T^b, \\ 0 & \text{otherwise.} \end{cases}
$$

The expected product surplus for a consumer who starts in funnel state $T$ is

$$
\mathbb{E}[S_T^b(1)] = \mu\phi(1-p) \times \begin{cases} \rho(2-\mu) + (1-\rho) & \text{if } \underline{b} \leq b_B^b, \\ \rho(1 + (1-\mu)\rho) + (1-\rho)\rho & \text{if } b_B^b < \underline{b} \leq b_T^b, \\ \rho & \text{otherwise.} \end{cases}
$$

(c) **Period 2 subgame for opt-in consumers where Period 1 ad is not shown**: in this case, there are $1-\lambda$ $T$-consumers, $\lambda(1-\phi)$ $B$-consumers, and $\lambda\phi$ purchasers. Therefore, the value of showing an ad to $T$-signal consumers, and hence the bid for these consumers, is

$$b_T^c = \mathbb{P}\{f = T | s = T\}\mu\phi p = \frac{\rho(1-\lambda)}{\rho(1-\lambda) + (1-\rho)\lambda(1-\phi)}\mu\phi p.$$

Similarly, for $B$-signal consumers, we obtain

$$b_B^c = \mathbb{P}\{f = T | s = B\}\mu\phi p = \frac{(1-\rho)(1-\lambda)}{(1-\rho)(1-\lambda) + \rho\lambda(1-\phi)}\mu\phi p.$$

Since Period 1 ads are not shown, we have

$$\mathbb{E}[q_T^c(1)] = \mathbb{E}[q_B^c(1)] = 0.$$

The expected product surplus for a consumer who starts in funnel state $T$ is

$$\mathbb{E}[S_T^c(1)] = \mu\phi(1-p) \times \begin{cases} 1 & \text{if } \underline{b} \le b_B^c, \\ \rho & \text{if } b_B^c < \underline{b} \le b_T^c, \\ 0 & \text{otherwise.} \end{cases}$$

With the Period 2 subgame outcomes at hand, we solve for the ad network and advertiser's Period 1 strategies. The value of advertising to $T$-signal consumers in Period 1, and hence the advertiser's bid for these consumers, is

$$b_T' = \mathbb{P}\{f = T | s = T\}\mu\phi p = \frac{\rho(1-\lambda)}{\rho(1-\lambda) + (1-\rho)\lambda}\mu\phi p.$$

Similarly, for $B$-signal consumers, we obtain

$$b_B' = \mathbb{P}\{f = T | s = B\}\mu\phi p = \frac{(1-\rho)(1-\lambda)}{(1-\rho)(1-\lambda) + \rho\lambda}\mu\phi p.$$

Since $\rho \ge 1/2$, we have $b_T' \ge b_B'$. Therefore, we obtain three cases: if $\underline{b} \le b_B'$, then the ad network sells ads for all opt-in consumers; if $b_B' < \underline{b} \le b_T'$, then the ad network only sells ads for $T$-signal opt-in consumers (while $B$-signal opt-in consumers' ads are sold instead to external advertisers at price $\underline{b}$); and if $b_T' < \underline{b}$, then all Period 1 ads are sold to external advertisers at price $\underline{b}$.

In sum, we have

$$\mathbb{E}[q^{\text{gen}}(1)] = \begin{cases} (1-\lambda)\mathbb{E}[q_T^a(1)] + \lambda\mathbb{E}[q_B^a(1)] & \text{if } \underline{b} \leq b'_B, \\ (1-\lambda)\mathbb{E}[q_T^b(1)] + \lambda\mathbb{E}[q_B^b(1)] & \text{if } b'_B < \underline{b} \leq b'_T, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\mathbb{E}[S^{\text{gen}}(1)] = \lambda\mathbb{E}[S_B] + (1-\lambda) \times \begin{cases} \mathbb{E}[S_T^a(1)] & \text{if } \underline{b} \leq b'_B, \\ \mathbb{E}[S_T^b(1)] & \text{if } b'_B < \underline{b} \leq b'_T, \\ \mathbb{E}[S_T^c(1)] & \text{otherwise.} \end{cases}$$

## *Analysis for Ad Content Personalization*

A simplifying assumption in the main model is that an ad shown to $B$-consumers by the same advertiser has no incremental effect. This is consistent with the informative view of advertising: consumers who are already aware and considering a product are informationally saturated such that they are unaffected by repeated ad exposures. However, it is conceivable that ad content may be tailored differently for repeated ads such that, despite being from the same advertiser, these ads may increase consumers' responsiveness to ads.

We present a parsimonious approach to modeling such effects. Specifically, we assume that for opt-in consumers, ad content can be tailored based on their inferred funnel states, such that repeat advertising tailored for $B$-consumers increases the $B$-consumers' product match probability by $\Delta_\phi > 0$. If $B$-consumers see repeat ads tailored for $T$-consumers, then their match probabilities do not change, and vice versa. Note that the main model is a special case of $\Delta_\phi \downarrow 0$. Therefore, unlike the main model, the advertiser has a positive valuation for ads shown to $B$-consumers. On the other hand, such personalization is not feasible for opt-out consumers, whose funnel states cannot be inferred beyond the prior.

In sum, the advertiser's Period 2 bid for opt-out consumers is the same as (6), while that for

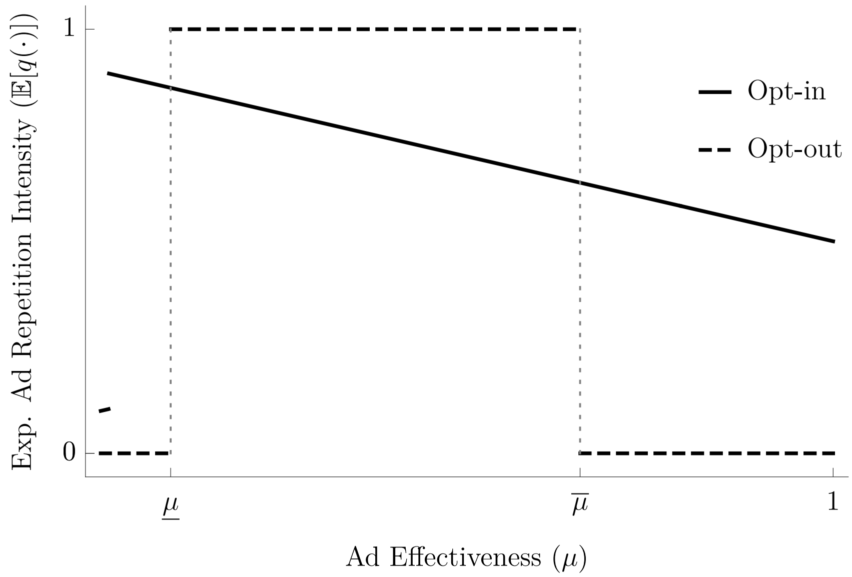


*Figure WA3: Ad Repetition Intensity with Ad Content Personalization*

opt-in consumers with signal $s$ is

$$\begin{aligned} b_s &= p \max\left[\mathbb{P}\{f = T|s\}\mu\phi, \mathbb{P}\{f = B|s\}\Delta_\phi\right] \\ &= p \times \begin{cases} \max\left[\frac{\rho(1-\mu)}{\rho(1-\mu)+(1-\rho)\mu(1-\phi)}\mu\phi, \frac{(1-\rho)\mu(1-\phi)}{\rho(1-\mu)+(1-\rho)\mu(1-\phi)}\Delta_\phi\right] & \text{if } s = T, \\ \max\left[\frac{(1-\rho)(1-\mu)}{(1-\rho)(1-\mu)+\rho\mu(1-\phi)}\mu\phi, \frac{\rho\mu(1-\phi)}{(1-\rho)(1-\mu)+\rho\mu(1-\phi)}\Delta_\phi\right] & \text{if } s = B. \end{cases} \end{aligned} \tag{WA25}$$

The maximum operators in (WA25) reflect the advertiser's ability to deliver different advertising content based on the inferred funnel state. Again, note that (WA25) converges to (9) as $\Delta_\phi \downarrow 0$.

The qualitative insights are preserved (compare Figures 3a and WA3) in that opt-out consumers expect to see fewer repeated ads when ad effectiveness is intermediate. Essentially, what changes with $\Delta_\phi > 0$ is that opt-in consumers expect to see higher intensity of repeated ads because the option to personalize ad content increases the advertiser's valuation for ads shown to opt-in consumers.

## *Discussion on Degenerate Purchase Funnel*

To better understand the role of the purchase funnel in the main model, we analyze a game *without* the purchase funnel. We show that with a degenerate purchase funnel, all consumers opt out of tracking.

To implement a degenerate purchase funnel, we assume that $\mu = 1$, such that *all* consumers transition to funnel state $B$ after the first ad exposure. This renders the second ad exposure wasteful, which implies that the advertising strategies would be the same for opt-in and opt-out consumers. In other words, the instrumental values of privacy are constant regardless of consumers' privacy choices, effectively rendering the instrumental values meaningless. In total, consumers make privacy choices solely based on their intrinsic values of privacy such that all consumers opt out of tracking.

On the other hand, as we see in the main model, the purchase funnel consideration induces some consumers to opt in to tracking as they trade off the cost of giving up privacy with the instrumental benefits of sharing data. In a setting with a non-degenerate purchase funnel structure, ad exposures endogenously create interim heterogeneity in consumers' funnel states. The advertiser then makes different inferences about opt-in and opt-out consumers, and thus implements different advertising strategies for opt-in and opt-out consumers. In particular, the instrumental values of privacy may be higher for opt-in consumers than for opt-out consumers such that consumers with sufficiently low privacy cost choose to opt in to tracking.