Customer Purchase Journey, Privacy Choices, and Advertising Strategies

W. Jason Choi Kinshuk Jerath Miklos Sarvary *

December 2019

*W. Jason Choi (jason.choi@columbia.edu) is a Ph.D. candidate, Kinshuk Jerath (jerath@columbia.edu) is Associate Professor of Marketing, Miklos Sarvary (miklos.sarvary@columbia.edu) is the Carson Family Professor of Business, all at the Columbia Business School, Columbia University, 3022 Broadway, New York, NY 10027. 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.

Customer Purchase Journey, Privacy Choices, and Advertising Strategies

Abstract

We investigate the impact of tracking consumers' Internet activities on the online advertising ecosystem in the presence of regulations that, motivated by privacy concerns, endow consumers with the choice to have their online activity be tracked or not. The consumers' strategic decisions to (dis)allow advertisers from tracking their activities depend on two aspects of privacy: its intrinsic value (protect privacy for its own sake) and its instrumental value (compromise privacy if doing so indirectly leads to some utility-enhancing outcome). This opt-in decision impacts the precision of inferences by advertisers about how far down a consumer is in the "purchase funnel" for a product by virtue of ads shown previously. We find that when ad effectiveness is intermediate, fewer ads are shown to opt-in consumers, who can be tracked and have their funnel stages inferred by advertisers, than to opt-out consumers, who cannot be tracked. In this case, consumers trade-off the benefit of seeing fewer ads by opting-in to tracking (positive instrumental value of privacy) with the disutility they feel from giving up their privacy (intrinsic cost of privacy). While privacy regulations generally increase consumer surplus, the implications for the ad network are mixed. Interestingly, the ad network's profit may be (i) higher under endogenous tracking than under full tracking, and (ii) lower as ads become more effective. We discuss managerial implications for advertisers as well as policy implications for regulators.

Keywords: customer purchase journey, online advertising, consumer tracking, data privacy, consumer privacy choice

1 Introduction

Advances in information technology have led to unprecedented levels of consumer tracking on the Internet (Lerner et al., 2016; Macbeth, 2017; Manjoo, 2019). According to Schelter et al. (2018), 355 third-party domains had installed trackers (e.g., cookies and web beacons) on over 90% of 41 million websites. Moreover, a recent study by Karaj et al. (2019) shows that 82% of the monitored web traffic had third-party scripts owned by Google, making it the largest third-party tracker by reach. These trackers allow firms to monitor not only which sites consumers visit, but also their browsing behavior such as whether the consumers interacted with the firms' ads (Roesner et al., 2012). Firms track consumers' online behavior for many reasons. Tracking helps firms (i) analyze site traffic and browsing patterns in order to deliver personalized content (Goldfarb and Tucker, 2011a; Bleier and Eisenbeiss, 2015), (ii) infer consumers' product preferences to inform pricing decisions (de Cornière and Nijs, 2016; Ichihashi, 2019; Montes et al., 2019; Taylor, 2004), and (iii) target ads to particular consumer segments (Bergemann and Bonatti, 2011; Iyer et al., 2005; Shen and Villas-Boas, 2018).

In particular, tracking helps advertisers to observe consumers' online browsing and purchase activities and infer their purchase journey stages. For instance, (Sahni et al., 2019) use consumer tracking data to infer whether a consumer is a "product viewer" or a "cart creator" and target ads accordingly. Google allows advertisers to use various tags to specify remarketing audiences based on such inferences.[1] Industry experts advocate that advertisers focus on targeting based on consumers' "stages in the decision journey" (Edelman, 2010). Empirical findings that ad effects, as measured by sales (Johnson et al., 2017; Lambrecht and Tucker, 2013; Seiler and Yao, 2017) or website return visits (Hoban and Bucklin, 2015; Sahni et al., 2019), vary widely across consumers' journey stages further highlight the importance of considering the purchase journey in developing advertising strategies (Todri et al., 2019).

While consumer tracking has benefited advertisers (Goldfarb and Tucker, 2011a; Johnson et al., 2019), its rapid expansion has deepened consumers' concerns about their online privacy (McDonald and Cranor, 2010). For instance, 77% of US Internet users indicate that they are "concerned about how tech/social media companies are using [their] online data ... for commercial purposes" (eMarketer, 2019a), 64% of UK Instagram users say "it's creepy how well online ads know me" (eMarketer, 2018a) and 68% of US Internet users report feeling concerned about "social media companies displaying ads based on their data" (eMarketer, 2018b).

In response to the growing outcry from consumers and privacy advocates, advertising organizations

[1] https://support.google.com/google-ads/answer/6335506 (accessed December 2019)

and regulators worldwide have sought to curb practices that potentially infringe on privacy, such as online tracking. Notably, in May 2018, the European Union (EU) enacted the General Data Protection Regulation (GDPR). Compared to its predecessors (e.g., Privacy and Electronic Communications Directive), the GDPR is considered the most stringent and comprehensive in terms of geographic and legislative scope.[2] Its hefty violation fines (maximum of $22.5 million and 4% of annual global turnover) are forcing even large firms like Google and Facebook to take compliance seriously.[3] The California Consumer Privacy Act (CCPA), a US analogue of the GDPR, is expected to go into effect in January 2020.[4]

One of the main tenets of the GDPR and the CCPA is the requirement that firms not only inform consumers what data will be collected for what purposes, but also obtain explicit affirmative consent to use their data. In other words, firms are not allowed to collect consumer data by default; consumers themselves must opt-in to their data being collected and processed by firms.[5] If consumers opt-out from tracking, then advertisers cannot monitor consumers' behavior across websites. Consequently, advertisers' targeting capabilities are drastically undermined and ad impressions could be potentially wasted (e.g., repeated exposure to consumers who had already purchased).[6] On the other hand, if consumers opt-in to tracking, advertisers can target ads to specific audiences based on a set of behavioral criteria (e.g., consumers who previously interacted with the ad but did not purchase).

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 (Aziz and Telang, 2016; Goldfarb and Tucker, 2011b). This may have contributed to the 50% decline in bids coming through sell-side ad platforms, and the 15% reduction in Google ad offerings via its ad exchange, after the GDPR went into effect (Kostov and Schechner, 2018). On the other hand, there is evidence suggesting that despite consumers' *stated*

[2]The regulation applies to all firms processing personal data of European subjects even if the firm operates outside of Europe. Personal data is defined as "any information relating to an identifiable person who can be directly or indirectly identified in particular by reference to an identifier" (`https://eugdpr.org/the-regulation/gdpr-faqs/`; accessed May 2019).

[3]In January 2019, Google was fined $57 million "for not properly disclosing to users how data is collected across its services ... to present personalized advertisements" (Satariano, 2019). Facebook revamped their privacy settings in compliance with the GDPR (`https://marketingland.com/what-marketers-need-to-know-about-facebooks-updated-business-tools-terms-238140`; accessed September 2019).

[4]See Future of Privacy Forum (2018) for a detailed comparison of the two regulations.

[5]While consumers were able to manually delete cookies even before the regulation, complete tracking prevention was extremely costly, if not impossible. For example, data collectors often used flash cookies technology to re-spawn cookies that were deleted by consumers (Stern, 2018; Angwin, 2010). Moreover, firms were able to purchase personal data from third-party information vendors without consumers' consent — such activities are now subject to GDPR enforcement.

[6]`https://www.blog.google/products/marketingplatform/360/privacy-safe-approach-managing-ad-frequency/` (accessed October 2019)

aversion towards tracking, they appear not as reluctant to allow tracking in practice.[7] For example, Johnson et al. (2019) find that less than 0.26% of US and EU consumers opt-out from behavioral targeting in the AdChoices program. Moreover, 67% of US and Canadian consumers report that they would feel "comfortable sharing personal information with a company" if it transparently discloses how their data will be used (Ipsos, 2019). These findings suggest that privacy regulations that endow consumers with the choice to being tracked may not necessarily result in low opt-in rates. In this respect, the net effect of privacy regulations may not be as detrimental to advertisers as they fear.

The discussion above on the advancements in tracking technology and shifts in industry regulations raises important questions for marketers and regulators alike. Does consumer tracking, which enables targeting based on a consumer's inferred purchase journey stage, lead to higher or lower levels of advertising intensity? How do advertising intensity and advertising effectiveness influence consumers' privacy choices of whether to allow being tracked? What are the implications of consumers' endogenous privacy choices on the ad network's profit? Which market participants benefit and lose from the regulation?

In this paper, we seek to shed light on these questions by developing a game theory model. We consider a two-period model in which consumers visit content pages, and each consumer creates one opportunity for an ad impression per period. An advertiser buys ad impressions from an ad network that sells ad inventory supplied by the content pages. Motivated from the discussion above, we assume that ad effects depend on the consumers' journey states represented by a "funnel" and that their purchase journey is influenced by advertising (Abhishek et al., 2017, 2018; Kotler and Armstrong, 2012). Based on their preferences for ad exposure and privacy, consumers choose whether to allow advertisers to track their online behavior.

Importantly, we model privacy as a multi-dimensional construct and assume that consumers jointly consider two aspects of privacy: its intrinsic value and its instrumental value (Becker, 1980; Posner, 1981; Farrell, 2012; Wathieu and Friedman, 2009). The intrinsic value of privacy refers to the utility consumers derive from protecting privacy for its own sake. We assume that consumers derive positive utility from protecting their privacy. On the other hand, the instrumental value of privacy stems from the indirect effects of the consumers' privacy choices (e.g., opting-in to tracking may affect the consumers' expected ad experiences). In contrast to the intrinsic aspect of privacy, the instrumental aspect of privacy may be either positive or negative depending on the firms' strategic responses.

[7] The overstatement of privacy concerns relative to revealed preferences is known as the "privacy paradox." See Norberg et al. (2007) and Athey et al. (2017) for details.

Our analysis yields a number of interesting insights. First, we find that, under certain conditions, consumers choose to opt-in to being-tracked because they expect to see fewer ads when advertisers can track them and infer their funnel states. In particular, this is the case if ad effectiveness is intermediate. Intuitively, the reason is the following. Consider an opt-out consumer who cannot be tracked. To this consumer, the advertiser shows ads in both periods if ad effectiveness is intermediate: ad effectiveness is high enough such that the first ad is worthwhile, and low enough that the first ad does not render the second ad wasteful. In contrast, for opt-in consumers, the advertiser shows a targeted ad in the second period only to selected consumer segments. Therefore, if ad effectiveness is intermediate, some consumers trade-off their costs from the instrumental and intrinsic aspects of privacy; i.e., they trade-off the benefit of seeing fewer ads by opting-in to tracking (positive instrumental value of privacy) with the disutility they feel from giving up their privacy (intrinsic cost of privacy).[8] Under other conditions, opting-in to tracking (weakly) increases the number of ads seen, in which case there is no such trade-off and both aspects of privacy utility induce consumers to opt-out from being tracked.

Second, we find that the consumers' opt-in decisions have important implications for the ad ecosystem. In particular, due to changes in consumers' opt-in behaviors, the ad network's profit may decrease in ad effectiveness, even though higher effectiveness implies higher purchase conversion probability. Intuitively, high ad effectiveness induces the saturation effect whereby the marginal value of successive ads is diminished by previously shown ads. This causes the advertiser to forego showing successive ads to opt-out consumers. In contrast, for opt-in consumers, enhanced targeting efficiency induces the advertiser to show successive ads. Thus, consumers expect to see fewer ads under no tracking, which incentivizes them to opt-out from tracking. As consumers opt-out, targeting efficiency falls, lowering ad valuations. Consequently, the ad network's profit can decrease as ads become more effective.

Third, privacy regulations increase consumer surplus and decrease the ad network's profit compared to a regime in which everyone can be tracked. Interestingly, however, if the advertiser is privately informed about ad valuations, consumers opting out of tracking may be a boon to the ad network. The ad network's inability to track opt-out consumers serves as a commitment mechanism that induces the ad network to sell untargeted ads that reach a larger consumer segment than targeted ads. This supply-side "market thickening" effect sometimes induces the advertiser to bid more aggressively for

[8]As firms comply with high standards of transparency enforced by privacy regulations, consumers will become not only aware of the privacy choices they are entitled to, but increasingly knowledgeable about the downstream consequences of their choices. For example, Figure 13 in Appendix A shows a sample privacy notice from Google shown to consumers in Europe. It describes the potential changes in ad intensity that could result from consumers' privacy choices. In practice, a significant fraction of ad slots indeed can be left unsold; for display ads, ad fill rates (i.e., the ratio of the number of ad slots that are available to get filled to the number that actually get filled) typically range from as low as 67% to 100% (Balseiro et al., 2014; Johnson et al., 2019).

opt-out consumers than for opt-in consumers. Therefore, the ad network's profit may be higher if some consumers exercise their privacy rights and opt-out from tracking.

We have a number of extensions that relax some of the simplifying assumptions in the main model. We consider cases with (i) information asymmetry between the advertiser and the ad network, (ii) multiple competing advertisers, (iii) imperfect signals of purchase histories for opt-in consumers, and (iv) an infinite time horizon where consumers arrive in overlapping generations. Overall, we find that the main insights are not affected in these extensions, while we obtain certain interesting new insights. For instance, we find from one of the extensions that more consumers may opt-in to being tracked if the signal about purchase behavior is less accurate; as a result, the ad network may be better off having lower signal precision.

In addition to being related to the papers referenced earlier, our paper contributes to two interrelated streams of research: 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, advertising strategies, ad intensity and adoption of ad avoidance tools (Athey and Gans, 2010; Aziz and Telang, 2017; Bergemann and Bonatti, 2011; Esteban et al., 2001; Iyer et al., 2005; Johnson, 2013; 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 creates a previously-unstudied link between the effectiveness of cross-period ads, which leads to novel insights pertaining to the impact of tracking on advertising strategies.

We also contribute to the growing literature on online privacy. Research on price-discrimination examines consumers' implicit privacy decisions, whereby 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, whereby consumers take (often costly) actions to control the amount of information disclosed to firms (Acquisti and Varian, 2005; Conitzer et al., 2012; Ichihashi, 2019; Montes et al., 2019). de Cornière and Nijs (2016) investigate the ad network's incentive to disclose consumer information to advertisers, whereas in our paper, the consumers exercise their privacy rights to decide the flow of their personal information. The mechanisms behind our results are orthogonal to market thickness (Bergemann and Bonatti, 2011; Rafieian and Yoganarasimhan, 2018) and market structure (Campbell et al., 2015) as we abstract from advertiser competition in the main model.

D'Annunzio and Russo (2019) 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 journey by virtue of previously shown ads gives rise to multi-period dynamics as advertisers consider retargeting consumers along the journey which we explicitly model. On the other hand, D'Annunzio and Russo (2019) consider a reduced-form effect of tracking in a static environment with a focus on publishers' decisions (e.g., ad capacity and outsourcing advertising to ad networks). Finally, while D'Annunzio and Russo (2019) assume that the number of ads a consumer is exposed to is independent of her privacy choice, we explicitly incorporate potential changes in ad intensity as an instrumental aspect of the consumer's privacy choice. In our model, the consumer's instrumental value of privacy is weighed against her intrinsic value of privacy.

At a higher level, our research (i) advances the understanding of the impact of tracking on the advertising ecosystem from a novel purchase journey perspective and (ii) contributes to the ongoing debate on online privacy regulations. Our findings suggest that assessing the impact of privacy regulations on the advertising industry is a complex issue. Nevertheless, we identify several robust theoretical insights that can inform various regulatory implications.

The rest of the paper is organized as follows. In Section 2, we describe the main model. In Section 3, we present the main results including the impact of tracking on advertising intensity, consumers' opt-in behavior, and the implications of endogenous privacy choice on the ad network's profit. In Section 4, we assess the robustness of the main insights by analyzing four extensions. In Section 5, we summarize the key results and conclude. All proofs are relegated to Section B of the appendix.

2 Model

The game consists of three players: consumers, an advertiser and an ad network. Consumers sequentially visit content pages where the ad network enables showing ads to them. The advertiser buys ad impressions from the ad network to reach consumers. Before we discuss each player's decisions and payoffs, we first explain a key feature of our model: the consumer purchase journey. We describe the relationship between advertising and consumers' progression down the purchase funnel.

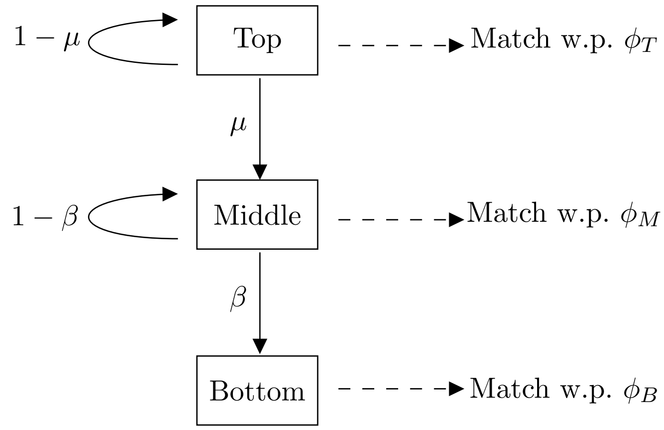

Figure 1: Purchase Journey and Ad Effects

Purchase Journey and Ad Effects

We consider a stylized purchase journey consisting of three distinct states labeled *top*, *middle*, and *bottom* (see Figure 1). For ease of exposition, we denote these states by T, M and B, respectively, and the consumers in the respective states by f-consumers for $f \in \{T, M, B\}$. We define funnel state $f \in \{T, M, B\}$ with a probability ϕ_f, which measures the likelihood of an f-consumer realizing a product match. In each time period, an f-consumer (who has not purchased yet) realizes a product match with probability (w.p.) ϕ_f, where $0 \leq \phi_T < \phi_M < \phi_B \leq 1$.[9] The three funnel states can be interpreted as follows: the top-funnel corresponds to "awareness" state, wherein the consumer is aware of the product's existence but is not seriously considering purchase; the mid-funnel corresponds to "consideration" or "interest" state, wherein the consumer is potentially considering purchase; and the bottom-funnel corresponds to even higher consideration and purchase interest by the consumer. We normalize ϕ_T and ϕ_B to 0 and 1, respectively.

If a consumer realizes a product match, she derives positive utility v from consuming the product; otherwise, she derives zero utility. In accordance with the empirical literature, we assume that ads affect consumers' likelihood of realizing a match with the advertised product (e.g., Johnson et al., 2016; Lee, 2002; Sahni, 2015; Shapiro et al., 1997; Xu et al., 2014); i.e., ads influence the consumers' progression through the funnel. Ads induce T-consumers to transition to funnel state M w.p. $\mu \in [0, 1]$, and have no effect w.p. $1 - \mu$. Similarly, ads induce M-consumers to transition to funnel state B w.p. $\beta \in [0, 1]$ and have no effect w.p. $1 - \beta$.

Note that our model specifications allow for flexible ad response curvatures using the funnel transition parameters μ, β, and mid-funnel match probability ϕ_M. For instance, Figure 2 depicts a convex ad

[9] Taylor (2004) calls this match probability the "intensity of taste for a particular class of goods" (pg. 635).

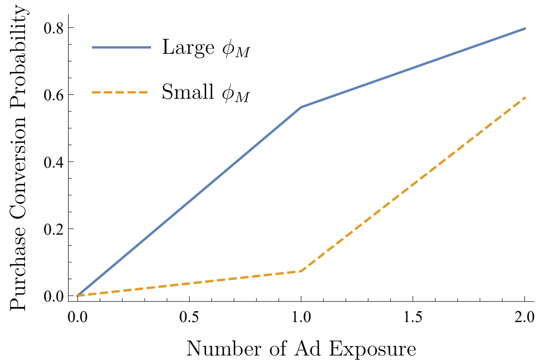

Figure 2: Ad Response Shapes

response curve with small ϕ_M, and a concave ad response curve with large ϕ_M. Advertising strategies, consumer choices and welfare outcomes will depend crucially on the curvature of the consumers' ad responses.

We also note that our results will remain completely unaltered if we scale the effectiveness numbers to vary between ranges different from 0 to 1. 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 numbers for the whole population of consumers can be scaled to vary between, say, 0 and 0.05 (i.e., effectiveness rates between 0% and 5%, which are arguably closer to empirical estimates), or any other range, rather than between 0 and 1 (i.e., effectiveness rates between 0% and 100%), without any impact on our results (see Section C of the appendix for more details). However, for model simplicity and expositional clarity, we use the formulation with effectiveness numbers varying between 0 and 1.

Consumers

A unit mass of consumers visit two content pages (both of which are in the ad network), one in each of Period $t \in \{1, 2\}$. Consumers are exposed to at most one ad impression per period from the page they visit. As described above, these ad exposures influence the consumers' progression through the purchase journey. For now, we assume that the initial state of newly arriving consumers is T. In other words, new consumers who visit content pages for the first time are *not* considering purchase. This helps deliver the main insights more cleanly. In Section 4.4, we relax this assumption by allowing some fraction of consumers to arrive in funnel state M.

Consumer utility consists of two components, product utility and privacy utility. The product consumption utility of an f-consumer (i.e., consumer in funnel state $f \in \{T, M, B\}$) is

$$u_{\text{prod}} = \tilde{v}_f - p,$$

where $\tilde{v}_f$ represents the stochastic match valuation, which equals v w.p. ϕ_f, and 0 w.p. $1 - \phi_f$, and p denotes the product price. If the consumer does not purchase, she derives the outside option utility 0. Without loss of generality, we normalize the match utility v to 1. Therefore, the consumer purchases if and only if she realizes a match and $p \leq 1$. Note that a consumer makes the purchase decision after realizing her match value. We assume that a consumer purchases at most one unit.

Next, we turn to privacy utility. We assume that consumers dislike being tracked, and that they are heterogeneous in their tracking disutility.[10] This disutility is captured by the privacy cost parameter θ, which has cumulative distribution function F. We assume that consumers can decide whether to opt-in or opt-out of being tracked. If a consumer opts-in, firms can track her identity and online browsing behavior (across content pages and across sessions) for targeting purposes. Later, we describe in more detail how tracking and targeting are implemented. The privacy utility of a consumer with privacy cost θ is

$$u_{\text{priv}}(x) = -\eta\, \tilde{q}(x) - \theta x, \tag{1}$$

where x denotes the consumer's privacy decision, which equals 1 if she opts-in, and 0 if she opts-out,[11] $\tilde{q}$ the total number of ads she expects to see, η the disutility she incurs per unit of ad impression (Johnson, 2013; de Cornière and Taylor, 2014), and θ the disutility she incurs for allowing tracking. In sum, consumers' privacy decisions are based on (i) the number of ads they anticipate to see as a result of their privacy decisions, and (ii) the extent to which they value privacy for its own sake. These two components constitute the instrumental and intrinsic aspects of privacy, respectively.

Advertiser

Depending on whether consumers can be tracked or not, the advertiser can buy different types of ads. If tracking is prohibited, then the consumers' identities cannot be matched across content page visits. In this case, the advertiser can only buy untargeted impressions (e.g., ads displayed to all website visitors

[10] Heterogeneity may stem from numerous factors such as differences in what consumers believe constitutes personal information (Acquisti et al., 2016) and differences in consumers' perception of their privacy control (Tucker, 2014).

[11] We clarify two implicit assumptions here. First, while consumers may modify their privacy decision at any time in practice, we assume consumers make a one-time privacy decision at the beginning of the game. Second, following the literature on endogenous privacy choices (e.g., Conitzer et al., 2012; Montes et al., 2019), we assume that privacy decision is binary. In practice, consumers may choose varying degrees of information disclosure. These assumptions keep the analysis simple without significantly changing the qualitative insights.

independent of their browsing histories). In particular, even if an ad is shown in Period 1 and shifts the distribution of consumers along the purchase journey, the advertiser is not able to target ads in Period 2 based on the funnel states.

On the other hand, if consumer tracking is allowed, the advertiser can buy ad impressions at the funnel-stage level. By installing tags on websites and embedding cookies on consumers' browsers, the advertiser can monitor the websites visited by the consumers, their browsing activity within the websites, and their purchase behavior.[12] Based on this information, the advertiser can specify the target audience such that their ads are shown only to consumers who meet some pre-specified criteria.[13] For example, the advertiser can target ads to consumers who are inferred to be in funnel state M and did not purchase. In each period, the advertiser decides which impressions to bid for and the respective bid amounts.

The advertiser also sets product price p_t in Period $t \in \{1, 2\}$. We normalize the marginal cost of the product to zero. Therefore, the advertiser's margin per conversion in Period t is p_t.

Ad Network

The ad network sells ad impressions to advertisers via second-price auctions.[14] It sets reserve price R_t^j in Period $t \in \{1, 2\}$, where j indexes the type of ad impression (e.g., ads targeted to M-consumers or untargeted ads for opt-out consumers). In our paper, this is equivalent to selling through a posted price; however, we choose the auction format to be consistent with how a vast majority of display ads are sold in the market.[15] The ad network maximizes its total profit across two periods, which consists of revenue and cost from ad impression sales. Costs may include operational costs associated with ad inventory management, as well as maintenance costs related to setting up ad auctions and delivering ads.[16] We denote this per-impression cost by $k \geq 0$.

Game Timing

The timing of the game is as follows.

[12]In Section 4.3, we a consider a general setting wherein consumers' purchase histories are observed *imperfectly*, and demonstrate that the main insights are preserved.

[13]For example, Facebook allows advertisers to target consumers "who engaged with any post or ad" or "clicked any call-to-action button" (`https://www.facebook.com/business/help/221146184973131?helpref=page_content`)

[14]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 `https://support.google.com/admanager/answer/9298211`).

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

[16]This model feature is motivated from our conversations with industry practitioners. In particular, they have indicated that selling ad inventory entails very significant operational costs associated with, for example, (i) storing, retrieving and relaying data to advertisers, and (ii) resolving the auction and announcing the outcome to all the bidders.

Period 0: Consumers decide whether or not to opt-in to being tracked.

Period 1: Ad network sets reserve prices for ads for opt-in consumers and ads for opt-out consumers. Advertiser sets product price and bids for ad impressions.

- If ads are shown, some consumers transition through funnel.
- Consumers make purchase decisions.

Period 2: Ad network sets reserve prices for targeted ads for opt-in consumers, and untargeted ads for opt-out consumers. Advertiser sets product price and bids for ad impressions.

- If ads are shown, some consumers transition through funnel.
- Consumers make purchase decisions.

We solve for the subgame-perfect equilibrium of the above game.

Before we proceed to the analysis, we note that, in Section 4, we consider a number of extensions to address several simplifying assumptions that we have made in the main model. We analyze scenarios in which (i) the advertiser has private information about its ad valuations, (ii) there are multiple, competing advertisers, (iii) the purchase histories of opt-in consumers are not perfectly observable, and (iv) consumers arrive in overlapping generations across an infinite time horizon. We also note that the publishers that own the content 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 commission basis (which is typically the case[17]), this is a reasonable assumption.

3 Analysis

First, note that the advertiser's product pricing decision is trivial and the optimal product price is always 1, which is the consumer's product utility on obtaining a match. Thus, the advertiser's margin per conversion is 1. Intuitively, if $p_t < 1$, then the advertiser leaves money on the table, and if $p_t > 1$, then no products are sold. Since $p_t^* = 1$ for $t \in \{1, 2\}$, the consumer purchases if and only if she realizes a match. This implies that the consumer's utility from product consumption is always zero whether or not she purchases. Therefore, when discussing consumer utility, we hereafter restrict attention to the privacy utility component.

To develop basic insights, we study the case of no tracking in Section 3.1 and full tracking in Section 3.2.

[17] For example, see https://support.google.com/adsense/answer/180195 and https://www.adpushup.com/blog/the-best-ad-networks-for-publishers/.

Then we discuss the main analysis with endogenous consumer tracking choice in Section 3.3.

3.1 No Tracking

In this section, we analyze the case in which consumers cannot be tracked; i.e., advertisers cannot distinguish consumers' funnel states nor their purchase histories. Our objective is to establish the baseline forces that determine the equilibrium advertising outcomes in the absence of consumer tracking. To solve for subgame perfect Nash equilibrium, we first analyze the advertiser's bidding problem in Period 2 and then proceed backwards. We assume that the advertiser plays weakly dominant bids, in the sense that the bidding strategies are not affected by "trembles" in various auction parameters, such as reserve prices and number of bidders.

Period 2

In Period 2, there are two possible subgames: one in which ads were shown in Period 1, and another in which they were not. We index the former Period 2 subgame with the subscript "2|ad" and the latter with "2|no ad." Consider the first subgame, in which ads were previously shown. The Period 2 distribution of consumers along the funnel can be characterized by three groups: (i) those who were not impacted by the first ad and remained in T, (ii) those who saw the ad, transitioned to M, and purchased, and (iii) those who transitioned to M but did not purchase. Given that the first ad induces interest w.p. μ, the first group is of size $1-\mu$. Since M-consumers realize a product match w.p. ϕ_M, the second group is of size $\mu\phi_M$. Finally, M-consumers do not purchase if they do not realize a product match; therefore, the third group is of size $\mu(1-\phi_M)$. While the advertiser knows this Period 2 distribution, it cannot identify which consumer belongs to which group in the absence of tracking.

To compute the advertiser's weakly dominant bid for the Period 2 untargeted ad, the advertiser compares its payoff when it wins vs. loses the ad auction. Let R_t denote the reserve price of untargeted ads in Period t.[18] If the advertiser bids b_2 in Period 2, its payoff is

$$\pi^A_{2|\text{ad}}(b_2) = \begin{cases} (1-\mu)\mu\phi_M + \mu(1-\phi_M)(\beta+(1-\beta)\phi_M) - R_2 & \text{if } b_2 \geq R_2, \\ \mu(1-\phi_M)\phi_M & \text{if } b_2 < R_2. \end{cases} \tag{2}$$

Consider the advertiser's payoff from winning the auction and displaying the ad, shown on the top row of (2). The first term denotes the conversion of T-consumers induced by Period 2 advertising: of the $1-\mu$ fraction of consumers who had not been affected by the Period 1 ad, μ fraction transition to funnel

[18] For ease of exposition, we suppress the ad type index j for the reserve price as only one type of ads (i.e., untargeted ads) is offered in the absence of tracking.

state M, of which ϕ_M fraction realize a match and purchase. Similarly, the second term denotes the conversion of M-consumers who had not converted in Period 1.

Note that if the advertiser bids below the reserve price and loses the auction, then its payoff is not 0 but $\mu(1-\phi_M)\phi_M$, as shown on the bottom row of (2). This is because even if no additional ads are shown in Period 2, the non-purchasers in funnel state M — who were pushed down from funnel state T after seeing the Period 1 ad — may realize a product match in Period 2 w.p. ϕ_M and purchase.

The payoffs of the second subgame in which ads were not shown in Period 1 can be analyzed in a similar manner. The following lemma states the subgame outcomes in Period 2.

Lemma 1 (Period 2 Bid and Reserve Price Without Tracking)**.**

- *Suppose the advertiser showed ads in Period 1. The advertiser's weakly dominant bid in Period 2 is* $b^*_{2|ad} = (1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta$*, and the ad network's optimal reserve price is* $R^*_{2|ad} = \max\left[k, b^*_{2|ad}\right]$.
- *Suppose the advertiser did not show ads in Period 1. The advertiser's weakly dominant bid in Period 2 is* $b^*_{2|no\ ad} = \mu\phi_M$*, and the ad network's optimal reserve price is* $R^*_{2|no\ ad} = \max\left[k, b^*_{2|no\ ad}\right]$.

The first part of Lemma 1 provides important preliminary insights into the conditions under which the advertiser buys *successive* ads in Period 2, conditional on having shown ads in Period 1. The advertiser buys successive ads if and only if $b^*_{2|\text{ad}} \geq R^*_{2|\text{ad}}$, which simplifies to

$$(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta \geq k. \tag{3}$$

That is, the marginal effectiveness of the successive ad, expressed on the left-hand side of (3), must be sufficiently large. Analyzing how this object changes with respect to the model primitives reveals two key determinants of a successive ad's marginal effectiveness.

The marginal effectiveness of the successive ad consists of two components: the marginal conversion of T-consumers (denoted by $(1-\mu)\mu\phi_M$) and the marginal conversion of M-consumers (denoted by $\mu(1-\phi_M)^2\beta$). It can be shown that $(1-\mu)\mu\phi_M+\mu(1-\phi_M)^2\beta$ decreases with respect to μ, the probability of an ad inducing interest, if and only if $\mu > \frac{\phi_M+\beta(1-\phi_M)^2}{2\phi_M}$. Moreover, the marginal effectiveness of the successive ad decreases with respect to ϕ_M, the product match probability of consumers in funnel state M, if and only if β, the probability of an ad inducing action, is greater than $\frac{1-\mu}{2(1-\phi_M)}$. The intuition for the first case is as follows. If μ is large, then the first ad exposure causes the Period 2 distribution of consumers to shift toward M. This implies a diminished role of successive ads in pushing T-consumers down to M in Period 2. Therefore, the marginal effectiveness of successive ads decreases in μ for large μ.

Consider the second case. If β is large, then the marginal effectiveness of successive ads is largely determined by their potential to convert M-consumers. Now, increasing ϕ_M has two effects. First, consumers are more likely to purchase after the first ad exposure such that there is a small segment of non-purchasers in funnel state M-consumers in Period 2. Second, if ϕ_M is large, those non-purchasers in funnel state M are likely to convert on their own without a successive ad exposure. Thus, increasing ϕ_M dampens the value of a successive ad.

Taken together, we see that an effective ad in Period 1 (i.e., large μ and ϕ_M) may diminish the marginal effectiveness of successive advertising in Period 2. We call this the *saturation effect*. It is visualized by the concave ad response curve for large ϕ_M in Figure 2.

Period 1

The reserve prices $R^*_{2|\cdot}$ from Lemma 1 imply that the advertiser's Period 2 payoff is $\mu(1-\phi_M)\phi_M$ if the advertiser shows ads in Period 1, and 0 otherwise. Taking this into account, the advertiser's problem in Period 1 is to determine the bid b_1 that maximizes

$$\pi_1^A(b_1) = \begin{cases} \mu\phi_M - R_1 + \mu(1-\phi_M)\phi_M & \text{if } b_1 \geq R_1, \\ 0 & \text{if } b_1 < R_1, \end{cases}$$

where R_1 is the reserve price for untargeted ads in Period 1. The following lemma states the advertiser's weakly dominant bid and the ad network's optimal reserve price.

Lemma 2 (Period 1 Bid and Reserve Price Without Tracking)**.** *Let* $x^+ \equiv \max[x, 0]$*. The advertiser's weakly dominant bid in Period 1 is* $b_1^* = \mu(2-\phi_M)\phi_M$*, and the ad network's optimal reserve price is*

$$R_1^* = \max\left[k + (\mu\phi_M - k)^+ - \left(b^*_{2|ad} - k\right)^+,\ b_1^*\right]. \tag{4}$$

We see from (4) that the ad network sometimes sets the Period 1 reserve price below the marginal cost k, even if that implies the ad network earns a negative payoff in Period 1. This occurs when showing successive untargeted ads in Period 2 is highly valuable for the advertiser; i.e., when $b^*_{2|\text{ad}} = (1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta$ is high. Intuitively, by setting a low reserve price in Period 1, the ad network helps the advertiser display ads, thereby creating an opportunity to extract greater surplus from the advertiser in Period 2. This ad pricing strategy can be viewed as the ad network capitalizing on the *convexity* of the ad response curve.

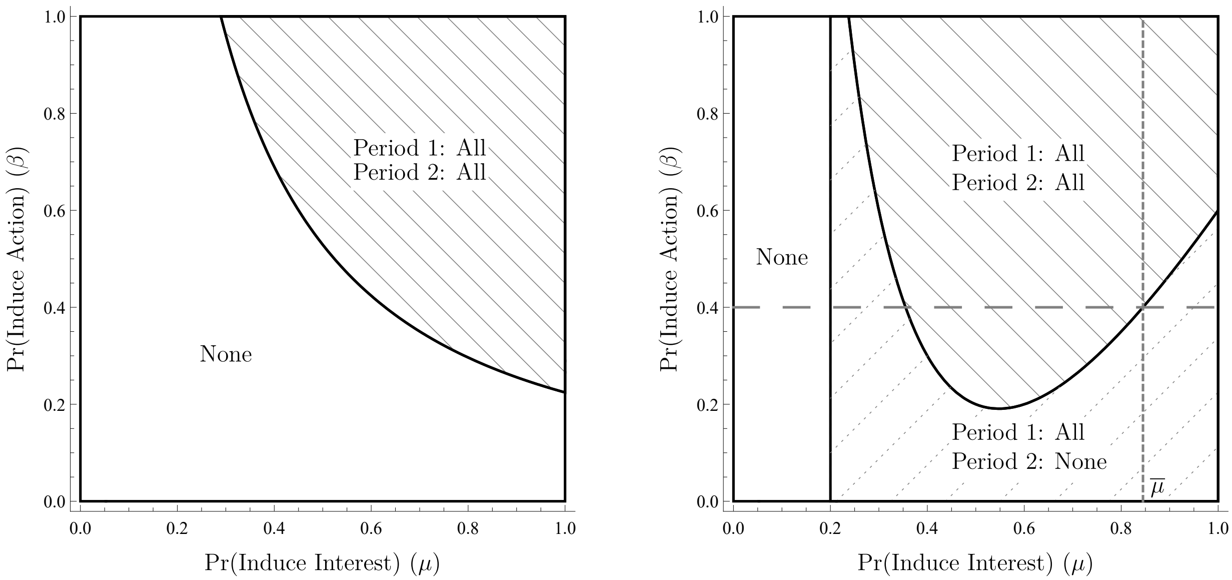

(a) Small $\phi_M (= 0.05)$ *(b) Large $\phi_M (= 0.5)$*

Figure 3: Ad Audiences with No Tracking; $k = 0.15$

Advertising Strategy Without Tracking

Given the equilibrium reserve prices and bids, we now characterize the conditions under which the advertiser buys ads in (i) both periods, (ii) only Period 1, and (iii) neither period. The following proposition summarizes the equilibrium advertising strategy across two periods.

Proposition 1 (Advertising Without Tracking)**.** *Suppose the advertiser cannot track consumers. For thresholds $\tilde{\beta}, \overline{\beta}, \underline{\beta}, \overline{\mu}$, and $\underline{\mu}$ defined in the proof, the equilibrium advertising strategy is as follows.*

- *Suppose $\phi_M < 1 - \sqrt{(\mu - k)^+/\mu}$. The advertiser buys ads in both periods if $\beta \geq \tilde{\beta}$, and does not buy any ads otherwise.*
- *Suppose $\phi_M \geq 1 - \sqrt{(\mu - k)^+/\mu}$.*
  - *if (i) $\beta \geq \overline{\beta}$ and $\mu \geq \underline{\mu}$, or (ii) $\underline{\beta} \leq \beta < \overline{\beta}$ and $\underline{\mu} \leq \mu < \overline{\mu}$, then advertiser buys ads in both periods;*
  - *if (i) $\beta < \underline{\beta}$ or $\mu \leq \underline{\mu}$, or (ii) $\underline{\beta} \leq \beta < \overline{\beta}$ and $\mu \geq \overline{\mu}$, then advertiser buys ads only in Period 1.*

Consider the case of small ϕ_M, depicted in Figure 3a, where the advertiser either buys ads in neither period or in both periods. This "all-or-nothing" pattern emerges when the ad response curve is convex. Specifically, if ϕ_M is small, the first ad exposure does little in terms of increasing the conversion probability. Thus, the advertiser does not find it worthwhile to advertise *only* in Period 1. However, if β is sufficiently large, a successive ad is highly likely to bring M-consumers in Period 2 down to B and induce purchase. This increase in effectiveness of successive ads compensates for the low effectiveness of the ads shown in Period 1. Thus, if ϕ_M is small, the advertiser either buys ads in neither period or in both.

In contrast, if ϕ_M is large, the total ad intensity across Periods 1 and 2 is more nuanced. In particular, the total ad intensity may be non-monotonic in μ (see Figure 3b). To understand the changes in ad intensity as μ increases, consider the cross-section of the plot represented by the large dashed line in Figure 3b for fixed $\beta = 0.4$. Along this line, as μ increases from 0 to $\overline{\mu}$ (denoted by the small dashed line), the ad intensity increases from 0 to 1 to 2 due to increasing effectiveness of ads. Past the $\overline{\mu}$ threshold, however, observe that Period 2 ads are foregone. This is because the combination of large ϕ_M and μ implies a high purchase conversion after Period 1 ads are shown. This diminishes the value of successive ads (i.e., Period 1 ads saturate) and consequently the intensity of ads shown in Period 2 is reduced.

In total, these results highlight the significance of considering the purchase journey in the analysis of advertising strategies, *even when* there is no trackability. In particular, modeling the funnel sheds light on how ads may influence consumer distribution along different funnel states. This distribution determines the marginal effectiveness of successive ads, which in turn affects ad buying decisions across time.

3.2 Full Tracking

We now analyze how the ability to track consumers affects the advertiser's strategy and the ad network's profit.

Advertising Strategy With Tracking

Without tracking, the advertiser was restricted to buying untargeted ads. With consumer tracking, however, the advertiser can target ads along two dimensions — consumers' positions in the purchase funnel and their product purchase histories. Specifically, in Period 2, it can target (i) T-consumers and (ii) M-consumers who did not purchase in Period 1 . In the Period 2 subgame for which ads were shown in Period 1, the advertiser's bidding strategy is a pair of bids $\left(b_2^T, b_2^M\right)$ for ad impressions associated with consumer segments (i) and (ii). Given reserve prices R_2^T and R_2^M for segments (i) and (ii), respectively,

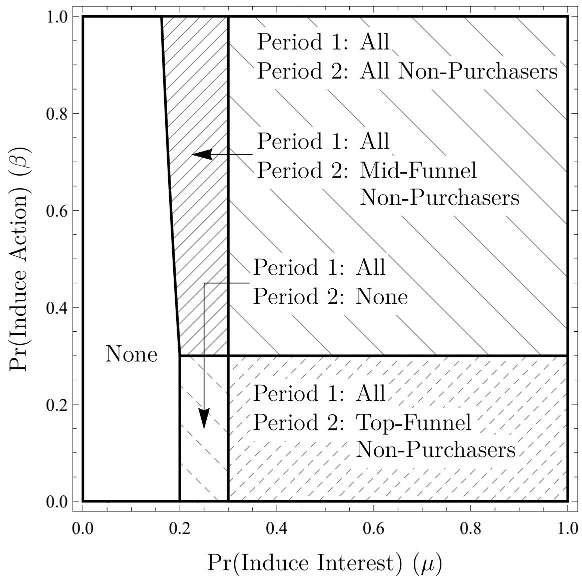

Figure 4: Ad Audiences With Tracking; $\phi_M = 0.5, k = 0.15$

the advertiser's expected payoff in Period 2 is

$$\pi^A_{2|\text{ad}}\left(b_2^T, b_2^M\right) = \begin{cases} (1-\mu)\left(\mu\phi_M - R_2^T\right) + \mu(1-\phi_M)\left(\beta + (1-\beta)\phi_M - R_2^M\right) & \text{if } b_2^T \geq R_2^T, b_2^M \geq R_2^M, \\ (1-\mu)(\mu\phi_M - R_2^T) + \mu(1-\phi_M)\phi_M & \text{if } b_2^T \geq R_2^T, b_2^M < R_2^M, \\ \mu(1-\phi_M)(\beta + (1-\beta)\phi_M - R_2^M) & \text{if } b_2^T < R_2^T, b_2^M \geq R_2^M, \\ \mu(1-\phi_M)\phi_M & \text{if } b_2^T < R_2^T, b_2^M < R_2^M. \end{cases}$$

The logic of solving for an equilibrium is similar to the case without no tracking, although more complex. Due to space considerations we relegate the full backwards induction analysis to the appendix. The following proposition characterizes the advertiser's strategies in the presence of consumer tracking.

Proposition 2 (Advertising With Tracking)**.** *Suppose the advertiser can track consumers along the purchase funnel. Let* $\tilde{\mu} = k\left(\phi_M(2-\phi_M) + (1-\phi_M)(\beta(1-\phi_M) - k)^+\right)^{-1}$*. The advertiser's equilibrium advertising strategy is as follows:*

- *if* $\mu \leq \tilde{\mu}$*, then do not buy any ads in either period;*
- *if* $\mu > \tilde{\mu}$*, then show ads to all consumers in Period 1. Furthermore,*
  - *if* $\mu > \frac{k}{\phi_M}$*, then in Period 2, buy ads targeted to* T*-consumers, and*
  - *if* $\beta > \frac{k}{1-\phi_M}$*, then in Period 2, buy ads targeted to* M*-consumers who did not purchase.*

Figure 4 depicts the advertising strategies in the presence of consumer tracking. Before we discuss this figure in detail, we note that a comparison with Figure 3 reveals that the advertising strategy

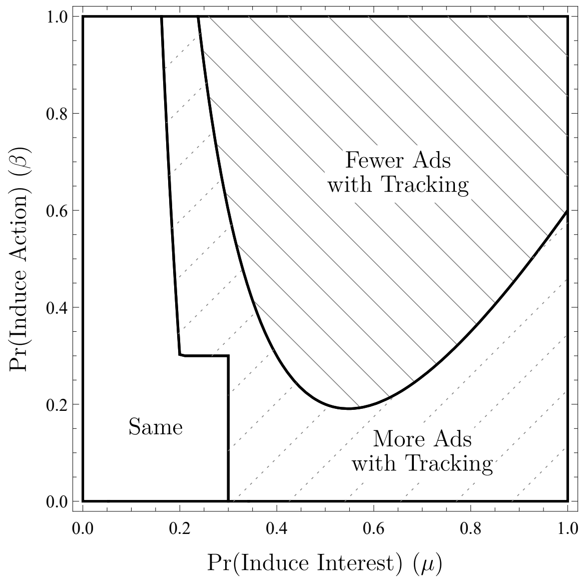

Figure 5: Ad Intensity With and Without Tracking; $\phi_M = 0.5, k = 0.15$

with tracking is significantly different from the advertising strategy without tracking. Proposition 2 shows that when the probability of the ad inducing interest is high, the advertiser adopts a *reach strategy*, whereby it targets successive ads to T-consumers who are not considering purchase. When the probability of the ad increasing purchase propensity is high, then the advertiser adopts a *frequency strategy*, whereby successive ads are shown to M-consumers who are already considering purchase.

Overall, Proposition 2 suggests that advertisers should be cognizant of the nuanced ad effects in relation to the consumers' journey down the funnel. As illustrated in Figure 4, the combination of consumer trackability and funnel considerations gives rise to various conditions under which one variant of advertising strategy is more profitable than another (e.g., reach vs. frequency).

Comparing No Tracking and Tracking Outcomes

A question of central interest that we can answer using this model is: how does consumer tracking impact overall advertising intensity? Our results show that the effect of tracking on ad intensity is nuanced (see Figure 5). If the funnel transition probabilities μ and β are small, then ad effectiveness is so low that no ads are shown to any consumers regardless of consumer trackability; hence, tracking may not change ad intensity. Otherwise, if either μ or β is sufficiently large, tracking may either increase or decrease the total intensity of ads.

When ad effectiveness is low, the average effectiveness of untargeted ads is low. With trackability, the advertiser can identify and bid for high-valuation impressions such that more ads are shown with tracking. On the other hand, consider the case when ad effectiveness is intermediate. Without tracking, untargeted ads are shown in both periods. However, some of the ads in Period 2 are wasted because they are shown to consumers who already purchased, which does not happen with tracking. Thus, tracking allows the advertiser to reduce spending on wasteful ad impressions, resulting in lower ad intensity with tracking than without.

While the above findings resonate with those from Esteban et al. (2001) and Iyer et al. (2005), our results are different from these papers at high levels of ad effectiveness. Specifically, when ad effectiveness is high, we find that the ad intensity differential *reverses*: more ads are shown under tracking than without. The intuition is that without tracking, high ad effectiveness dampens the value of successive ads, such that only first period ads are shown. On the other hand, when consumers can be tracked along the purchase journey, ads are targeted to (i) consumers who were not impacted by the first ad and stayed in top-funnel and/or (ii) consumers who moved down to mid-funnel but did not purchase. Therefore, more ads are shown under tracking than without. Put differently, the interdependence between the marginal effectiveness of cross-period ads stemming from funnel considerations reverses the ad intensity differential for high levels of ad effectiveness.

In sum, we establish a non-monotonic relationship between the impact of tracking on ad intensities and ad effectiveness, which is proxied by the funnel transition parameter μ when the match parameter ϕ_M is fixed. We summarize these results in the following proposition.

Proposition 3 (Ad Intensity)**.** *Consumer funnel tracking either increases or decreases the total ad intensity compared to the no tracking case. Specifically, for thresholds $\tilde{\beta}, \overline{\beta}, \underline{\beta}, \overline{\mu}$, and $\underline{\mu}$ defined in the proof,*

- *if (i) $\mu < \tilde{\mu}$ and $\beta < \tilde{\beta}$ or (ii) $\frac{k}{\phi_M(2-\phi_M)} < \mu \leq \frac{k}{\phi_M}$ and $\beta \leq \frac{k}{1-\phi_M}$, then the ad intensities are the same;*
- *if $\phi_M > 1 - \sqrt{(\mu - k)^+/\mu}$ and either (i) $\beta > \overline{\beta}$ and $\mu > \underline{\mu}$, or (ii) $\underline{\beta} \leq \beta < \overline{\beta}$ and $\underline{\mu} \leq \mu < \overline{\mu}$, then tracking reduces ad intensity;*
- *otherwise, tracking increases ad intensity.*

Tracking and Ad Network Profit

How does consumer tracking impact the ad network's profit? We find that tracking weakly increases the ad network's profit. Intuitively, consumer tracking endows the advertiser with more information on which the advertiser can condition its bid. In this case, the ad network can selectively supply ad impressions most highly valued by the advertiser while foregoing unprofitable ones, thereby raising its profit compared to the regime without consumer tracking. We state this result as a proposition.

Proposition 4 (Consumer Tracking and Ad Network Profit)**.** *Consumer funnel tracking weakly increases the ad network's profit.*

It is important to note that the tracking-induced improvement in the ad network's profit hinges on the assumption that the ad network is as knowledgeable about ad valuations as the advertiser. In Section 4.1, we show that the result of Proposition 4 does not always carry over to a setting where the advertiser has private information about ad valuations. Surprisingly, information asymmetry between the ad network and the advertiser may result in consumer tracking *lowering* the ad network's profit.

3.3 Endogenous Tracking Choice

In the preceding analysis, we examined two distinct cases in which the advertiser was either not able to track any consumers or able to track all consumers. In this section, we investigate the impact of endowing consumers with the choice to be tracked. That is, we analyze how consumers exercise their right to choose whether to allow tracking or not and how this decision affects the ad ecosystem. As discussed in the introduction, the analysis is largely motivated by the recent enactment of data privacy regulations that mandate affirmative consumer consent prior to acquiring and processing consumer data.

Consumers' Opt-In Behaviors

We first characterize the consumers' equilibrium privacy choices given the advertising outcomes under tracking and no tracking. Recall from the consumer privacy utility formulation in (1) that consumers dislike seeing ads (instrumental aspect) and also dislike being tracked (intrinsic aspect). This implies that consumers will choose to incur the privacy cost from opting-in to being tracked *only if* they expect to see fewer ads from doing so; i.e., only if the positive instrumental value outweighs the intrinsic cost of privacy. The following proposition summarizes the consumer opt-in behavior.

Proposition 5 (Consumer Opt-In Behavior)**.** *Let $q(0)$ and $q(1)$ denote the total number of ads consumers are exposed to when they opt-out and opt-in, respectively. The proportion of consumers who opt-in to being tracked can be non-monotonic in the funnel transition probability μ. In particular, for thresholds $\overline{\beta}, \underline{\beta}, \overline{\mu}$, and $\underline{\mu}$ defined in the proof, if either (i) $\underline{\beta} < \beta \leq \overline{\beta}$ and $\underline{\mu} < \mu < \overline{\mu}$, or (ii) $\beta > \overline{\beta}$ and $\mu > \underline{\mu}$, then $F\left(\eta\left(q(0) - q(1)\right)\right)$ consumers opt-in; otherwise, all consumers opt-out.*

The consumer opt-in pattern is driven by two forces. The first force stems from the changes in number of ads shown to opt-out consumers in Period 2. Recall that in the absence of tracking, for small ϕ_M, as μ increases, the advertiser's strategy sometimes changes from not advertising in either period to any consumer to advertising in both periods to all consumers (see Proposition 1). This pattern emerges

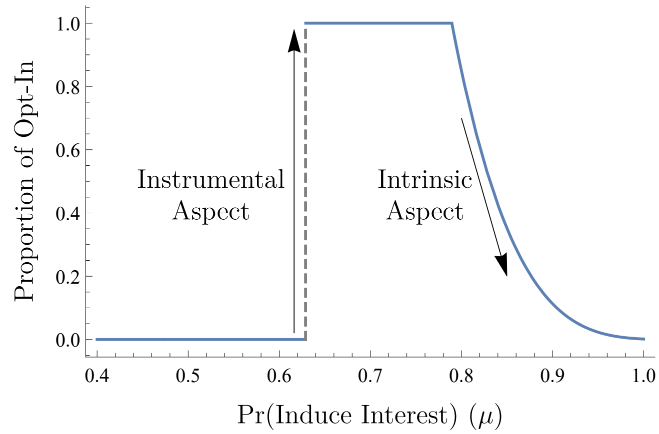

Figure 6: Proportion of Opt-In Consumers; $\phi_M = 0.05, \beta = 0.4, F(\theta) = \theta^4, \eta = 4, k = 0.15$

from the convexity of the ad response curve: while advertising only once is never profitable, showing successive ads might be. Thus, as consumers expect high ad intensity for large μ when they cannot be tracked, consumers are incentivized to opt-in to being tracked in order to see fewer ads. Figure 6 depicts the increase in opt-in rate as μ increases past the threshold $\mu \approx 0.63$, after which the advertiser shows ads to all consumers in both periods under no tracking.

The second force relates to the advertiser's targeting regime for opt-in consumers in Period 2. To illustrate, suppose μ and β are large. In this case, the advertiser adopts the frequency strategy in Period 2 such that successive ads are targeted to M-consumers (see Proposition 2). Now, as the probability of an ad inducing interest increases, consumers are more likely to transition to funnel state M after the first ad exposure, and hence become targets of Period 2 advertising. This dampens consumers' incentives to opting-in to being tracked. Figure 6 shows the associated decline in opt-in rate as μ increases in the neighborhood of $\mu \approx 0.85$.

Ad Network Profit

Next, we analyze how the consumers' opt-in behaviors characterized above impact the ad network's profit. Interestingly, we find that consumers' opt-in choices lead to non-monotonicities in the ad network's profit with respect to μ. In particular, under certain conditions, the ad network's profit *decreases* in μ, even though larger μ implies higher purchase conversion on average. To understand this, recall from Proposition 4 that the ad network's profit is (weakly) lower under no tracking than under tracking. Since larger μ may result in more consumers opting-out from being tracked (see Proposition 5), it follows that the ad network's profit may decrease in μ.

As described above, the change in the number of opt-in consumers can arise from two distinct forces. First, even if consumers expect to see fewer ads under tracking and thus opt-in to being tracked, if

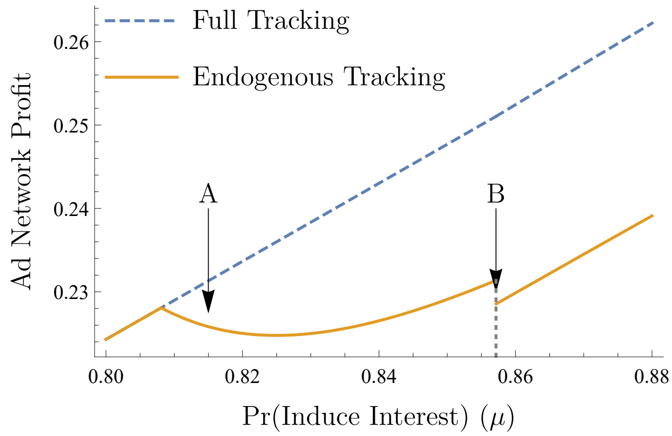

Figure 7: Ad Network Profit; $\phi_M = 0.175, \beta = 0.4, \eta = 3, k = 0.15, F(\theta) = \theta^{16}$

the ad intensity under tracking increases with μ, then less consumers choose to opt-in as μ increases. This decline in opt-in rate induces a continuous decrease in the ad network's profit, as illustrated in the region marked A in Figure 7. Second, consumers also consider the instrumental aspect of privacy: if consumers expect to see fewer ads without tracking, no consumer chooses to opt-in. This leads to discrete jump in the ad network's profit, as shown in the region marked B in Figure 7. The following proposition summarizes this finding.

Proposition 6 (Equilibrium Ad Network Profit)**.** *Let* $q^*(0)$ *and* $q^*(1)$ *denote the equilibrium ad intensity without and with tracking, respectively, and let* $\underline{\mu}'$ *and* $\overline{\mu}'$ *be as defined in the proof. Suppose consumers' privacy costs* θ *are uniformly distributed on* $[0, 1]$*. Under endogenous tracking, the ad network's profit decreases in* μ *if and only if either*

- $q^*(0) = 2$*,* $q^*(1) = 1 + \mu(1 - \phi_M)$*,* $\eta(1 - \mu(1 - \phi_M)) < 1$*, and either* $\mu < \underline{\mu}'$ *or* $\mu \geq \overline{\mu}'$*, or*
- $\phi_M \geq 1 - \sqrt{(\mu - k)^+/\mu}$ *and* $\underline{\beta} \leq \beta < \overline{\beta}$*.*

What does this mean for the ad network? Conventional wisdom suggests that the ad network would be better off if ads were more effective: ads that yield high purchase conversion are associated with high valuations, which allows the ad network to sell ad slots at higher prices. Proposition 6 provides a countervailing argument. Privacy regulations that allow consumers to choose between being tracked or not may result in more consumers opting-out from being tracked for higher levels of ad effectiveness, in particular if higher ad effectiveness implies more ads being shown to opt-in consumers. In this case, consumers choose to opt-out from tracking, thereby undermining targeting efficiency. This means that ad slots may be sold at lower prices, lowering the ad network's profit.

Consumer Surplus

One of the main objectives of privacy regulations is to protect consumers. Consistent with intuition, we find that giving consumers the choice to being tracked weakly improves consumer surplus, compared to the full tracking benchmark. Intuitively, the regulations allow consumers to make privacy decisions such that their individual surplus is maximized. And since their decision does not impose externalities on other consumers, net consumer surplus weakly increases. We state this as a proposition.

Proposition 7 (Consumer Surplus)**.** *Privacy regulations that allow consumers to choose whether to be tracked or not increase overall consumer surplus compared to the full tracking case.*

In sum, our analysis provides three important takeaways. First, consumers sometimes choose to opt-in to being tracked. Specifically, this happens if the effectiveness of the ads in inducing product interest, μ, is intermediate, where consumers trade-off the benefit of seeing fewer ads (positive instrumental value of privacy) with the disutility of compromising their privacy (intrinsic cost of privacy). Second, the ad network's profit when consumers can choose to be tracked or not may decrease in μ, even if larger μ implies higher purchase conversion. The intuition is that more consumers may choose to opt-out from being tracked as μ increases; this lowers targeting efficiency, which ultimately reduces the ad network's profit. Finally, consumer surplus always increases and the ad network's profit always decreases when consumers have the choice of being tracked or not. As we show in the next section, however, allowing the advertiser to have private information about their ad valuations may reverse the latter result. That is, the ad network's profit may be higher under endogenous tracking than under full tracking.

4 Extensions

In this section, we explore four extensions that demonstrate the robustness of the qualitative insights obtained from the main model. We also describe some additional insights that emerge from relaxing the assumptions from the main model.

4.1 Information Asymmetry about Ad Valuation

In the main model, we assumed that the advertiser's ad valuation is known by the ad network. Consequently, the ad network sets the reserve price such that the advertiser's surplus is extracted fully. In this information asymmetry extension, we allow the advertiser's ad valuation to be private information. In practice, ad valuations may not be fully known to the ad network for several reasons. First, the ad network may not perfectly observe all of the consumers' interactions with the advertiser (e.g., offline

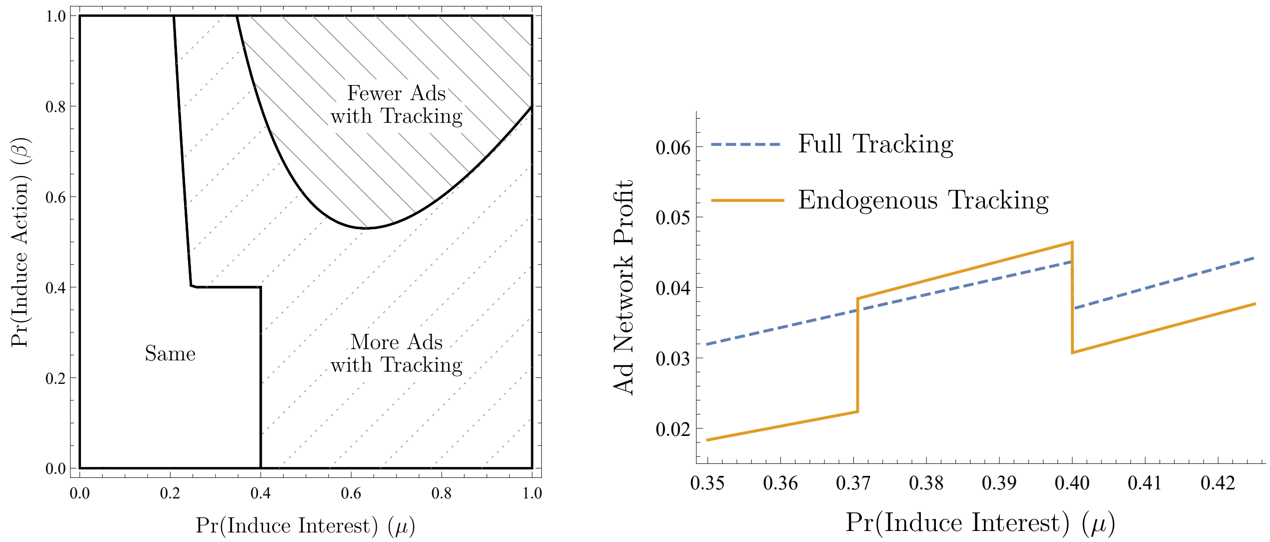

(a) Ad Intensity Differential; $\phi_M = 0.5, k = 0.05$

(b) Ad Network Profit; $\phi_M = \eta = 0.5, \beta = 0.9, k = 0.05, F(\theta) = \theta$

Figure 8: Ad Intensity and Ad Network Profit Under Information Asymmetry

interactions in the advertiser's physical store) that may inform the advertiser about consumer valuation. Second, even if the ad network possessed similar same levels of information as the advertiser, it may have less ability to infer consumer's willingness to pay for an advertiser's product (de Cornière and Nijs, 2016).

To that end, we assume that in each period, the advertiser's ad valuation is drawn independently from Uniform$[0, 1]$ and is known privately to the advertiser. In contrast to the main model, in this case the advertiser earns positive surplus. This adds interesting dynamics to the model because in Period 1 the advertiser will anticipate how the Period 1 outcome affects its Period 2 payoff and will bid accordingly.

While the patterns of ad intensity differential between the tracking and no tracking cases are qualitatively unaffected by the advertiser's bidding dynamics (see Figure 8a), the result pertaining to the ad network's profit is sometimes reversed. In particular, we find that under certain conditions, the ad network may benefit from regulations that allow consumers to endogenously choose to being tracked (see Figure 8b). This occurs when β is sufficiently high and μ intermediate.

The intuition is as follows. With tracking, the ad network has the option to sell select ad impressions targeted to a subgroup of consumers (e.g., M-consumers who did not purchase in Period 1). While such selective ad sales help the ad network to efficiently extract surplus from the advertiser in Period 2, they hurt the advertiser by limiting the size of consumer segments reached. However, as privacy regulations induce some privacy-conscious consumers to opt-out from tracking, the ad network's targetability is reduced. Thus, instead of selling targeted ads, the ad network sells untargeted ads that reach a larger consumer base. We call this the supply-side "market thickening" effect. Untargeted ads that reach more consumers are more profitable for the advertiser than, say, ads targeted to M-consumers. Thus, in anticipation of higher Period 2 payoffs for opt-out consumer segments, the advertiser bids more aggressively in Period 1 under endogenous tracking than under full tracking. This ultimately leads to higher ad network profit.

In sum, our analysis sheds light on a novel role of privacy regulations that allow endogenous tracking choices: regulations can serve as a commitment device for the ad network to sell more ad impressions.[19] This in turn better aligns the incentives of the ad network and the advertiser. Under certain conditions, privacy regulations can lead to higher profits for both parties compared to the full tracking benchmark.

4.2 Competing Advertisers

We extend the main model by considering two competing advertisers indexed by $i \in \{1, 2\}$. Consumers are heterogeneous in their product preferences: λ proportion of consumers are "loyal" to Advertiser 1, and $1 - \lambda$ proportion to Advertiser 2. A consumer transitions down the purchase funnel according to specifications of the main model (see Section 2) only if she sees an ad from the advertiser to which she is loyal. Without loss of generality, we assume that $\lambda > \frac{1}{2}$; i.e., Advertiser 1 is the dominant brand.[20]

We assume that opting-in to tracking reveals not only the consumers' funnel states and purchase histories, but also their *ex ante* product preferences; i.e, whether consumers are loyal to Advertiser 1 or 2. For example, tracking consumers' past visits to and browsing patterns within advertisers' websites may help advertisers infer consumers' product preferences.[21] The rest of the model specifications remain unchanged. Note that the extension model reduces to the main model if $\lambda = 1$.

We find that the qualitative insights from the main model carry over for a large range of parameters (see Figure 9). The results diverge if and only if either (a) the effectiveness of ads shown to T-consumers is high (i.e., large μ) or (b) product preference heterogeneity is sufficiently large (i.e., λ close to 0.5). The intuition is the following. Advertiser 2 has high incentive to advertise in Period 2 after Advertiser 1's ads has been shown in Period 1 if Advertiser 2 knows either (a) that its first ad exposure in Period 2 will be highly effective, or (b) that there is a large group of loyal consumers who will respond to Advertiser 2's ad. While in the main model ad slots would have been left unfilled in Period 2 due to the saturation

[19] Without such a commitment device, the advertiser would anticipate the ad network to sell only select profitable impressions in Period 2. Thus, the advertiser would expect to gain little surplus in Period 2, such that the ad network would not be able to charge high prices for Period 1 ads.

[20] In the case of symmetry $\left(\lambda = \frac{1}{2}\right)$, the only difference is that the first period outcome is randomly determined between the two advertisers; otherwise, the logic of the analysis is equivalent to the asymmetry case.

[21] `https://www.digitaltrends.com/computing/how-do-advertisers-track-you-online-we-found-out/`

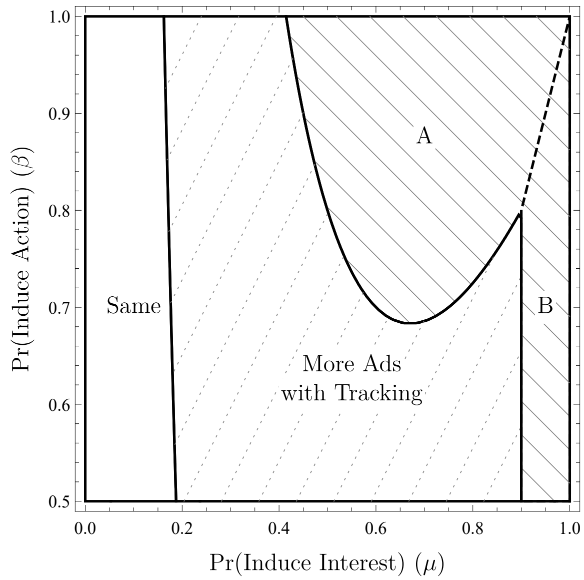

Figure 9: Competing Advertisers: Ad Intensity With and Without Tracking; $\phi_M = 0.5, k = 0.15, \lambda = 0.66$

effect, in this case the ad slots are filled by Advertiser 2 (Region B in Figure 9). Finally, to avoid seeing ads in both periods, consumers opt-in to tracking for this parameter range. We summarize this finding in the following proposition.

Proposition 8 (Competing Advertisers)**.** *The advertising intensity differential between tracking and no tracking regimes carries over from the main model if either μ is not too large or λ is not too small. Otherwise, opt-out consumers are exposed to ads in both periods: from Advertiser 1 in Period 1, and from Advertiser 2 in Period 2.*

4.3 Imperfect Observability of Purchase History

In the main model, we assumed that for an opt-in consumer, both her funnel state $f \in \{T, M, B\}$ and her purchase history were perfectly observable by the ad network. While this assumption helped us deliver the main insights clearly, it may not always reflect the information flow in practice. For example, firms may not be able to perfectly merge a consumer's identity across different website sessions, or they may not be able to match the online identities of consumers who search online but purchase offline. In this section, we relax the assumption that the purchase histories of opt-in consumers are perfectly observable. We make the following assumption which nests the main model as a special case: when an opt-in consumer arrives at a content page, the ad network and the advertiser can infer the consumer's funnel state f perfectly, but her purchase history only imperfectly. In particular, the ad network and the advertiser receive an imperfect signal about the consumer's purchase history.

To that end, let r denote the binary variable which equals 1 if the consumer purchased the advertised

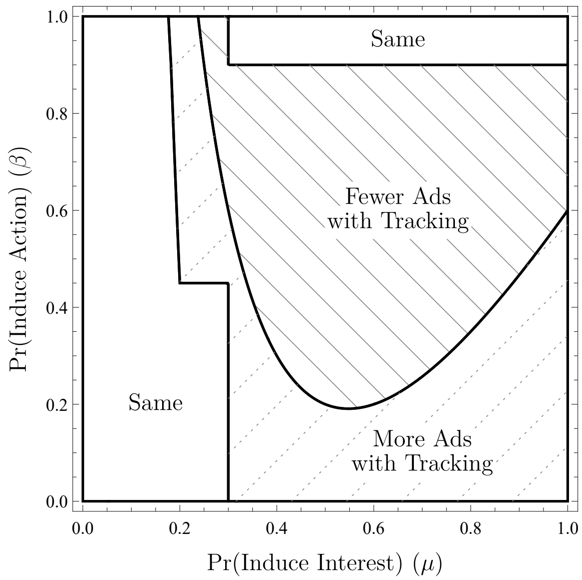

Figure 10: Ad Intensity With and Without Tracking Under Imperfect Purchase Observability; $\phi_M = 0.5, k = 0.15, \rho = \frac{2}{3}$

product, and 0 otherwise. Let S denote the binary signal that the ad network and the advertiser receive. The signal's accuracy is parametrized by $\rho \in \left[\frac{1}{2}, 1\right]$ where

$$\rho = \mathbb{P}\{S = r|r\}$$

for $r \in \{0, 1\}$. Note that this extension collapses to the base model when $\rho = 1$.

The main departure from the main model is the Period 2 subgame for opt-in consumers when an ad was shown in Period 1. With perfect observability, the ad network forewent selling ad impressions for consumers who had already purchased. In contrast, under imperfect purchase observability, the ad network decides ad supplies based on the signals it receives. In Period 2, the ad network infers that $1 - \mu$ fraction of consumers remain in funnel state T, and μ fraction move down to funnel state M. Of the μ fraction of M-consumers, $\phi_M \rho + (1 - \phi_M)(1 - \rho)$ are associated with the signal that the consumers already purchased the product advertised in Period 1, and $\phi_M(1 - \rho) + (1 - \phi_M)\rho$ are associated with the signal that the consumers did not purchase.

We find that under certain conditions, imperfect purchase observability changes the relative ad intensity between opt-in vs. opt-out consumers. Specifically, when the purchase signal accuracy ρ is low, and ads are highly effective (i.e., μ and β are large), the reduced targetability associated with imperfect observability leads to ads being shown in both periods to all opt-in consumers.

The intuition is that if the ad network is uncertain about a consumer's purchase history, then even if it receives a signal that the consumer had already purchased, the ad network deems the false positive probability to be sufficiently high that it puts up the ad impression for sale. Consequently, untargeted ads are shown to all opt-in consumers: those who are in state T, those in state M associated with "purchased" signal, and those in state M associated with "not purchased" signal. This implies that the ad intensities across opt-in and opt-out consumers are the same for large μ and β (see top-right corner of Figure 10), whereas fewer ads are shown to opt-in consumers with perfect observability (see Figure 5).

Next, we conduct comparative statics with respect to the signal accuracy ρ. What is the relationship between consumers' opt-in choices and the accuracy of their purchase history signals? Interestingly, we find that under certain conditions, more consumers choose to opt-in to tracking when purchase signals are less accurate. As the purchase signals become more accurate, more ads may be sold under tracking. Thus, the expectation of more intensive targeted ads motivates consumers with high privacy cost to opt-out from tracking for higher levels of purchase signal accuracy. Furthermore, the lower opt-in rate for high levels of signal accuracy may lower the ad network's profit due to a decline in targeting efficiency. The following proposition summarizes these findings.

Proposition 9. *Suppose* $\frac{\mu}{\mu-k}k < \beta$*,* $\max\left[\frac{k}{\mu}, 1-\sqrt{\frac{k}{\beta}}\right] < \phi_M < 1-\frac{k}{\beta}$ *and* $q^*(0) = 2$ *(i.e., ads are shown in both periods to all opt-out consumers). Then, the number of consumers opting-in to tracking decreases with the accuracy of the purchase signal. Consequently, the ad network's profit may decrease in the signal accuracy.*

Proposition 9 sheds light on a strategic force that goes against the lay intuition that accurate information about consumers is beneficial for the ad network. Distinct from the "market thinning" effect (Levin and Milgrom, 2010; Bergemann and Bonatti, 2011), less accurate signals about consumers' behaviors may induce consumers to expect fewer targeted ads should they opt-in to tracking, increasing the appeal of opting-in. And as more consumers opt-in to tracking, the targeting efficiency of the the ad network increases, thereby increasing its profit.

4.4 Infinite Horizon with Heterogeneous Overlapping Consumer Generations

We extend the game from the main model along two dimensions. First, we relax the assumption that all newly arriving consumers are at funnel state T. In particular, we allow $\sigma \in [0, 1]$ proportion of newly arriving consumers to be in funnel state M, and $1-\sigma$ in T. Broadly, σ can be interpreted as the advertiser's "brand strength:" the higher the σ, the greater the extent to which the advertiser's product is *a priori* known and considered by consumers. Second, we extend the game horizon from two-period

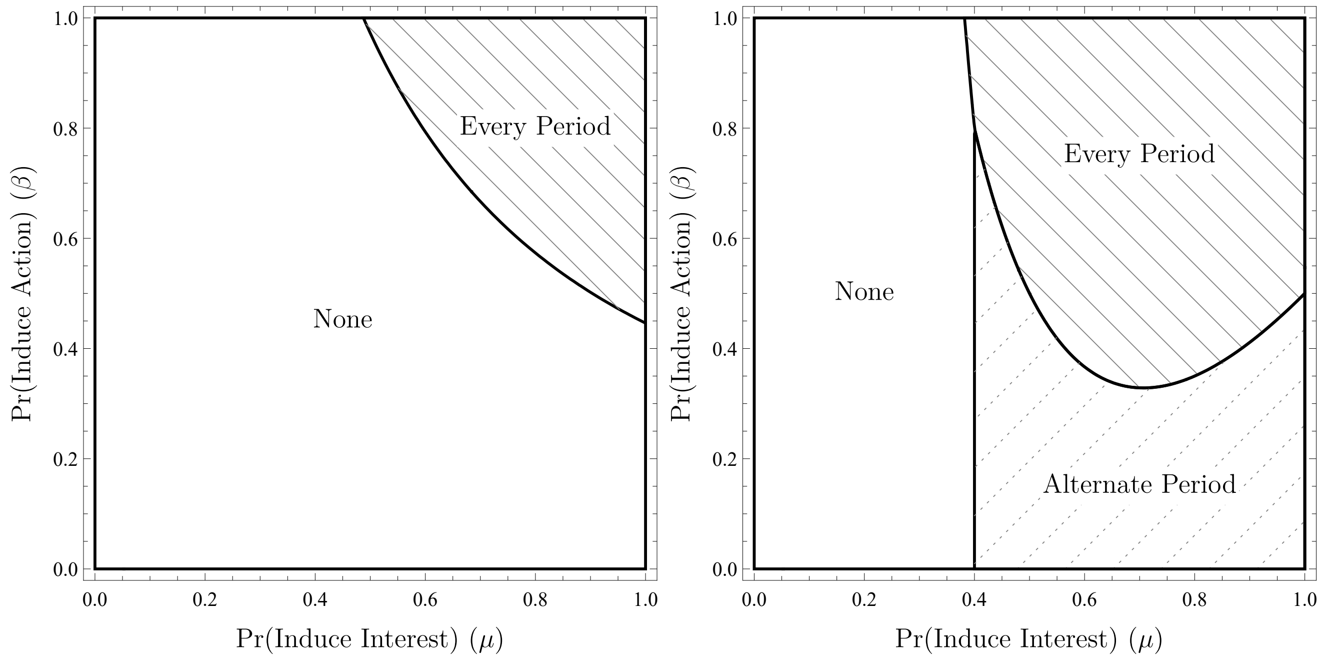

(a) $\phi_M = 0.05$ *(b)* $\phi_M = 0.5$

Figure 11: Advertising Strategy Without Tracking; $k = 0.25, \sigma = 0, \delta = 1 - 10^{-6}$

to infinite-period. In each period, a unit mass of consumers — σ mass of M-consumers and $1 - \sigma$ mass of T-consumers — arrive and live for two periods. Thus, in any given period, there are overlapping generations of consumers. The rest of the specifications remain the same as the main model.

We solve for a Markov-perfect equilibrium (MPE) wherein an advertiser's strategy depends only on the payoff-relevant state in that period. The ad network compares the total discounted profit (with discount factor $\delta \in [0, 1)$) obtained from inducing the different advertising outcomes, and then chooses whichever yields the highest profit.

Due to space considerations, we relegate the MPE derivation to Section D of the appendix. Here, we highlight how the qualitative insights obtained here compare to that of the main model. First, Figure 11 shows that if consumer heterogeneity is muted (i.e., $\sigma = 0$) and the discount factor is close to 1, the equilibrium outcomes closely mirror that of the two-period model (see Figure 3). In particular, the advertiser shows ads to all consumers in both periods if and only if β is sufficiently large and μ is intermediate. The underlying mechanism revolves around the saturation effect, and the qualitative insights remain essentially the same.

Second, we examine how the insights from the main model are moderated by two new parameters: the "brand strength" parameter σ, and the discount factor δ. As illustrated in Figure 12, we find that as either σ increases or δ decreases, the parametric region where the advertiser shows ads to all consumers

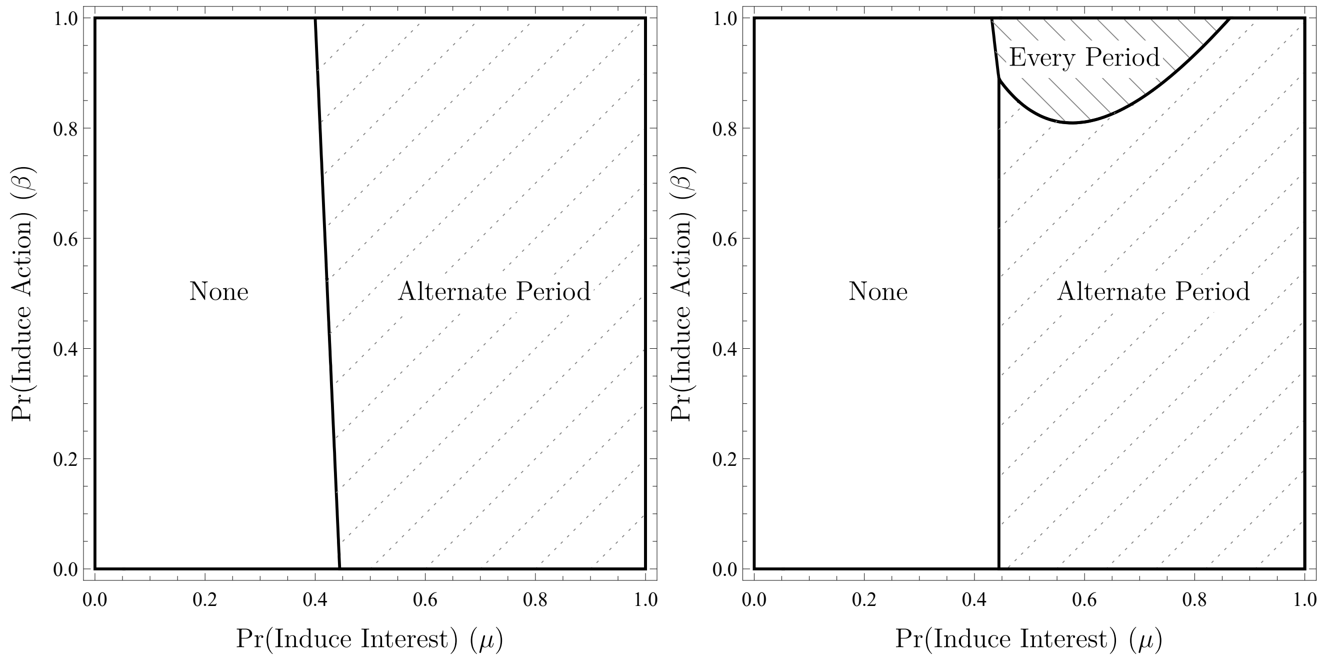

(a) $\sigma = 0.1, \delta = 1 - 10^{-6}$

(b) $\sigma = 0, \delta = 0.5$

Figure 12: Advertising Strategy Without Tracking; $\phi_M = 0.5, k = 0.25$

in both periods becomes smaller. The intuition is the following. As σ increases, a larger portion of newly arriving consumers are already mid-way down the funnel in state M. This accentuates the saturation effect: since many newly arriving consumers are already in the consideration phase and will likely convert without additional ad exposures, the value of a successive ad diminishes (see Figure 12a). On the other hand, as δ decreases, the advertiser places smaller weight on the value of a successive ad, whose payoff materializes in the future. Therefore, the incentive to buy successive ads decreases, even if convex ad response curves may have otherwise justified showing successive ads (see Figure 12b).

5 Conclusion

In this paper, we study the impact of tracking consumers' Internet activities on the online advertising ecosystem, and the impact of regulations that, motivated by privacy concerns, endow consumers with the choice to have their online activity be tracked or not (e.g., the GDPR). In particular, we model the consumer "purchase journey" and analyze the impact of consumers' opt-in decisions — co-determined by the intrinsic and instrumental aspects of privacy — on the strategies and profits of advertisers and ad networks.

Among others, we establish the following insights from the analysis. First, when given a choice, some consumers will choose to opt-in to tracking because they expect to see fewer ads when advertisers can track them and infer their funnel stages. Specifically when ad effectiveness is intermediate, ad-averse consumers opt-in to tracking, thereby trading-off the benefit of seeing fewer ads (positive instrumental value of privacy) with the disutility they feel from giving up their privacy (intrinsic cost of privacy). Second, consumers' opt-in behaviors have important implications for the ad ecosystem. For example, due to changes in consumers' privacy choices, the ad network's profit may decrease in ad effectiveness, even though higher ad effectiveness implies higher purchase conversion. Finally, we show that privacy regulations improve overall consumer surplus and reduce the ad network's profit. Interestingly, however, if the advertiser has private information about ad valuations, privacy regulations may increase the ad network's profit as well. Intuitively, as privacy-conscious opt-out from tracking, the ad network commits to selling untargeted ads that reach a larger consumer segment than targeted ads. This in turn incentivizes the advertiser to bid more aggressively for opt-out consumers, resulting in higher ad network profit compared to the regime in which all consumers can be tracked.

The results obtained in this paper provide important managerial insights for marketers and regulators alike. Our findings suggest that under certain conditions, the ad network and the advertiser could both earn higher profits if the ad network can credibly commit to *not* track consumers. Privacy regulations that allow consumer tracking only under affirmative consent can thus serve as a commitment device that helps the advertiser and the ad network "coordinate" in a mutually profitable manner. Furthermore, our results underscore the need for regulators to consider nuanced approaches to data privacy regulations that are based on various market conditions such as the accuracy of signals pertaining to consumers' online behavior, the degree of information asymmetry, consumer disutility for ads, their value of privacy, and the average effectiveness of ads.

Our research generates a number of interesting hypotheses that could be empirically tested. For instance, for ads with intermediate levels of effectiveness, our results suggest that compared to the full-tracking regime prior to the enforcement of privacy regulations, the ad fill rates are likely to increase when consumers have a choice to be tracked or not as more untargeted ads are shown to consumers who choose to opt-out from tracking. In a similar vein, our analysis predicts that ad prices will fall with the advent of privacy regulations due to declines in targeting efficiency associated with consumers opting-out. It would be interesting to investigate these hypotheses across different product categories that are associated with different levels of average ad effectiveness.

We acknowledge several limitations of the paper. First, our model does not account for flexible product pricing decisions by the advertiser because consumers' product utilities assume a "binary" functional form. While this assumption allowed us to focus on the advertising strategies, it would be interesting to consider a finitely elastic product demand that would allow for richer pricing strategies. Second, we implicitly assumed that the advertiser shows the same ad content to all consumers. In practice, advertisers may tailor their messages to consumers in different stages of the journey, insofar as consumers can be tracked (e.g., entice consumers lower down the funnel with price promotions). Thus, another fruitful avenue for future research would be to investigate how personalized ads for opt-in consumers may impact the funnel-transition probabilities. Finally, it would be interesting to examine a more active role of publishers. For example, 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.

References

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.

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*, pages 1–47. Vancouver, BC.

Balseiro, S. R., Feldman, J., Mirrokni, V., and Muthukrishnan, S. (2014). Yield optimization of display advertising with ad exchange. *Management Science*, 60(12):2886–2907.

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

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

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

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

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. (2019). Ad Networks and Consumer Tracking. *Management Science*.

de Cornière, A. and Nijs, R. D. (2016). Online Advertising and Privacy. *The RAND Journal of Economics*, 47(650):48–72.

de Cornière, A. and Taylor, G. (2014). Integration and search engine bias. *The RAND Journal of Economics*, 45(3):576–97.

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

eMarketer (2018a). Attitudes toward personalized advertising according to uk instagram users. `https://www.emarketer.com/Chart/Attitudes-Toward-Personalized-Advertising-According-UK-Instagram-Users-June-2018-of-respondents/224451`. (Accessed May 2019).

eMarketer (2018b). Level of concern us internet users have toward social media companies displaying ads based on their data. `https://www.emarketer.com/Chart/Level-of-Concern-US-Internet-Users-Have-Toward-Social-Media-Companies-Displaying-Ads-Based-on-Their-Data-March-2018-of-respondents/218053`. (Accessed May 2019).

eMarketer (2019a). Consumer attitudes on marketing 2019: Privacy concerns mount, and ad blocking isn't going away, aug 2019. `https://content-na1.emarketer.com/consumer-attitudes-on-marketing-2019`. Accessed Dec 2019.

eMarketer (2019b). Programmatic digital display ad spending, nov 2019. `https://content-na1.emarketer.com/programmatic-digital-display-ad-spending`. Accessed Dec 2019.

Esteban, L., Gil, A., and Hernandez, J. M. (2001). Informative Advertising and Optimal Targeting in a Monopoly. *The 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.

Future of Privacy Forum (2018). Comparing privacy laws: Gdpr v. ccpa. `https://fpf.org/wp-content/uploads/2018/11/GDPR_CCPA_Comparison-Guide.pdf`. (Accessed June 2019).

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

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

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. (2019). Online privacy and information disclosure by consumers. *Working paper*. Canada.

Ipsos (2019). Global citizens and data privacy. https://www.ipsos.com/sites/default/files/ct/news/documents/2019-01/ipsos-wef_-_global_consumer_views_on_data_privacy_-_2019-01-25-final.pptx_lecture_seule_0.pdf. (Accessed May 2019).

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

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

Johnson, G., Shriver, S., and Du, S. (2019). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Working paper*. Boston, MA.

Johnson, G. A., Lewis, R. A., and Nubbemeyer, E. I. (2016). The online display ad effectiveness funnel and carry-over: A meta-study of ghost ad experiments. *Working paper*, pages 1–28. Boston, MA.

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

Karaj, A., Macbeth, S., Berson, R., and Pujol, J. M. (2019). Who tracks me: Shedding light on the opaque world of online tracking. *arXiv preprint arXiv:1804.08959*.

Kostov, N. and Schechner, S. (2018). Google emerges as early winner from europes new data privacy law. *Wall Street Journal*.

Kotler, P. and Armstrong, G. (2012). *Principles of Marketing*. Prentice Hall.

Lambrecht, A. and Tucker, C. E. (2013). When does Retargeting Work? Information Specificity in Online Advertising. *Journal of Marketing Research*, 50(5):561–76.

Lee, A. Y. (2002). Effects of implicit memory on memory-based versus stimulus-based brand choice. *Journal of Marketing Research*, 39(4):440–54.

Lerner, A., Simpson, A. K., Kohno, T., and Roesner, F. (2016). Internet jones and the raiders of the lost trackers: An archaeological study of web tracking from 1996 to 2016. In *25th USENIX Security Symposium*.

Levin, J. and Milgrom, P. (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review*, 100(2):603–07.

Macbeth, S. (2017). Tracking the trackers: Analysing the global tracking landscape with ghostrank. Technical report, Technical report, Ghostery.

Manjoo, F. (2019). I visited 47 websites. hundreds of trackers followed me. *The New York Times*.

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. *The American economic review*, 71(2):405–409.

Rafieian, O. and Yoganarasimhan, H. (2018). Targeting and Privacy in Mobile Advertising. *Working paper*. Seattle, WA.

Roesner, F., Kohno, T., and Wetherall, D. (2012). Detecting and defending against third-party tracking on the web. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 12–12. USENIX Association.

Sahni, N. S. (2015). Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quantitative Marketing and Economics*, 13(3):203–247.

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.

Satariano, A. (2019). Google is fined $57 million under europe's data privacy law. *The New York Times*.

Schelter, S., Kunegis, J., et al. (2018). On the ubiquity of web tracking: Insights from a billion-page web crawl. *The Journal of Web Science*, 4(4):53–66.

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

Shapiro, S., MacInnis, D. J., and Heckler, S. E. (1997). The Effects of Incidental Ad Exposure on the Formation of Consideration Sets. *Journal of Consumer Research*, 24(1):94–104.

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

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

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

Todri, V., Ghose, A., and Singh, P. V. (2019). Trade-offs in online advertising: Advertising effectiveness and annoyance dynamics across the purchase funnel. *Forthcoming, Information Systems Research*.

Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. *Journal of marketing research*, 51(5):546–562.

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.

Xu, L., Duan, J. A., and Whinston, A. (2014). Path to Purchase: A Mutually Exciting Point Process Model for Online Advertising and Conversion. *Management Science*, 60(6):1392–412.

Appendix

A Sample Privacy Notice

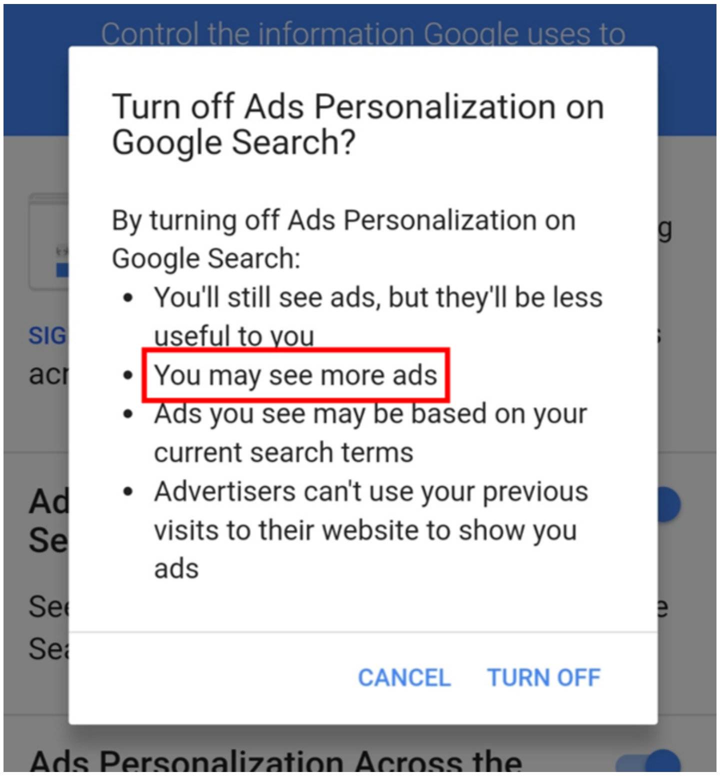

Figure 13: Google's Privacy Notice in Europe (July 2019)

B Proofs

Statement and Proof of Claim 1

Claim 1. *Suppose a player's payoff from bidding b in an auction parametrized by tuple (x, y, z, p) is*

$$\pi(b) = \begin{cases} x - y\,p & \text{if } b \geq p, \\ z & \text{if } b < p, \end{cases}$$

where $y > 0$ and $p > 0$. Then the player's weakly dominant bid (i.e., robust to any p) is $b^* = (x - z)^+/y$.

Proof. First, if $z \geq x$, then winning leads to strictly lower profit than losing. Therefore, the optimal bid is to lose for any p; hence $b^* = 0$.

Second, suppose $x > z$. We show that there is no strictly dominant deviation strategy for $b^* = (x-z)^+/y$. To that end, consider a deviation b' that is strictly less than $\frac{x-z}{y}$. Then for $p \in \left(b', \frac{x-z}{y}\right)$, we have $\pi(b') = z = x - y\left(\frac{x-z}{y}\right) < x - yp = \pi(b^*)$. For all other ranges of p, the two strategies yield the same payoff. Therefore, b^* weakly dominates b'.

Next, consider another deviation b'' that is strictly greater than $\frac{x-z}{y}$. Then for $p \in \left(\frac{x-z}{y}, b''\right)$, we have $\pi(b') = x - yp < x - y\left(\frac{x-z}{y}\right) = z = \pi(b^*)$. Again, the two strategies yield the same payoff for all other ranges of p. This completes the proof. ■

Statement and Proof of Claim 2

Claim 2. *Let $f(x) = \max[x, 0]$ for all $x \in \mathbb{R}$. Then $f(x) + f(y) \geq f(x+y)$ for all $x, y \in \mathbb{R}$.*

Proof.

$$\frac{1}{2}\left(f(x) + f(y)\right) = f\left(\frac{x}{2}\right) + f\left(\frac{y}{2}\right) \geq f\left(\frac{x}{2} + \frac{y}{2}\right) = \frac{1}{2}f(x+y)$$

where the equalities are due to linearity and inequality due to convexity. ■

Proof of Lemma 1

Proof. Consider the first subgame wherein the advertiser had shown ads in Period 1. Following Claim 1, the advertiser's weakly dominant bid (against any reserve price R_2) in Period 2 is

$$\begin{aligned} b^*_{2|\text{ad}} &= (1-\mu)\mu\phi_M + \mu(1-\phi_M)(\beta + (1-\beta)\phi_M) - \mu(1-\phi_M)\phi_M \\ &= (1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta. \end{aligned}$$

Similarly, the advertiser's weakly dominant Period 2 bid in the second subgame, wherein it did not advertise in Period 1, is $b^*_{2|\text{no ad}} = \mu\phi_M$. For each of the Period 2 subgames described above, the ad network sets R_2 as high as b^*_2, provided it is larger than k. Thus, we obtain the optimal Period 2 reserve prices $R^*_{2|\text{ad}} = \max\left[k, b^*_{2|\text{ad}}\right]$ and $R^*_{2|\text{no ad}} = \max\left[k, b^*_{2|\text{no ad}}\right]$. ■

Proof of Lemma 2

Proof. The advertiser's weakly dominant bid b^*_1 in Period 1 follows directly from Claim 1. For the ad network's optimal reserve price, consider its Period 1 payoff:

$$\pi^N(R_1) = \begin{cases} R_1 - k + (b^*_{2|\text{ad}} - k)^+ & \text{if } R_1 \leq b^*_1, \\ 0 + (b^*_{2|\text{no ad}} - k)^+ & \text{otherwise.} \end{cases}$$

It follows that $R_1^* = b_1^*$ if $b_1^* - k + (b_{2|\text{ad}}^* - k)^+ \geq (b_{2|\text{no ad}}^* - k)^+$, and $R_1^* \in (b_1^*, \infty)$ otherwise. The reserve price stated in the lemma satisfies this property. ■

Proof of Proposition 1

Proof. Given the reserve prices derived above, the ad network's profit in Period 2 if ads were shown in Period 1 is $(1-\mu)\mu\phi_M + \mu(1-\phi_M)(\beta+(1-\beta)\phi_M) - \mu(1-\phi_M)\phi_M - k = (1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k$, if the ad network sells Period 2 ads, and 0 otherwise. Therefore, the ad network's Period 2 profit given ads were shown in Period 1 is $\big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+$. Similarly, if ads were not shown in Period 1, then the ad network's Period 2 profit is $(\mu\phi_M - k)^+$.

Thus, the ad network's total profit from setting reserve price R_1 in Period 1 is

$$\pi^N(R_1) = \begin{cases} R_1 - k + \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+ & \text{if } R_1 \leq b_1^*, \\ 0 + (\mu\phi_M - k)^+ & \text{if } R_1 > b_1^*, \end{cases}$$

from which we obtain

$$R_1^* = \begin{cases} b_1^* & \text{if } b_1^* - k + \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+ \geq (\mu\phi_M - k)^+, \\ (b_1^*, \infty) & \text{otherwise.} \end{cases}$$

Since R_1^* can be any number greater than b_1^* when $b_1^* < k + (\mu\phi_M - k)^+ - \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+$, we can write

$$R_1^* = \max\Big[k + (\mu\phi_M - k)^+ - \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+,\ \mu\phi_M + \mu(1-\phi_M)\phi_M\Big].$$

Next, we derive the conditions under which the advertiser's weakly dominant bids exceed the optimal reserve prices set by the ad network.

Ads Shown Only in Period 2

We first show that showing ads only in Period 2 is never an equilibrium outcome. Towards a contradiction, suppose the conditions for such an equilibrium hold; i.e., $b_1^* < R_1^*$ and $\mu\phi_M - k \geq 0$. But $\mu\phi_M - k \geq 0$ implies that $R_1^* = \max\Big[k + (\mu\phi_M - k) - \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+,\ \mu\phi_M + \mu(1-\phi_M)\phi_M\Big]$, which simplifies to $\max\Big[\mu\phi_M - \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+,\ \mu\phi_M + \mu(1-\phi_M)\phi_M\Big]$. This is strictly greater than $b_1^* = \mu\phi_M + \mu(1-\phi_M)\phi_M$ if and only if $\mu\phi_M + \mu(1-\phi_M)\phi_M < \mu\phi_M - \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+$, which is equivalent to $\mu(1-\phi_M)\phi_M < -\big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+$. Since the left-hand side is strictly positive while the right-hand side is non-positive, this inequality never holds. A contradiction.

Ads Shown in Periods 1 and 2

The advertiser buys untargeted ads in both periods if and only if $b_1^* \geq R_1^*$ and $b_{2|\text{ad}}^* \geq R_{2|\text{ad}}^*$, which are equivalent to

$$\mu\phi_M + \mu(1-\phi_M)\phi_M \geq k + (\mu\phi_M - k)^+ - \big((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k\big)^+ \tag{5}$$

and

$$(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta \geq k, \tag{6}$$

respectively. Note that (6) implies that (5) simplifies to

$$(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta - k \geq k + (\mu\phi_M - k)^+ - (\mu(2-\phi_M)\phi_M)\,, \tag{7}$$

which can be re-arranged in terms of β as

$$\beta \geq \tilde{\beta} \equiv \frac{2k + (\mu\phi_M - k)^+ - \mu(3-\mu-\phi_M)\phi_M}{\mu(1-\phi_M)^2}.$$

The intersection of conditions (6) and (7) simplifies to (6) if $\mu > \frac{k}{\phi_M(2-\phi_M)}$, and to (7) if $\mu < \frac{k}{\phi_M(2-\phi_M)}$. These branching conditions in turn can be re-written as $\phi_M > 1 - \frac{\sqrt{(\mu-k)^+}}{\sqrt{\mu}}$ and $\phi_M < 1 - \frac{\sqrt{(\mu-k)^+}}{\sqrt{\mu}}$, respectively.

Ads Shown Only in Period 1

The advertiser buys only Period 1 ads if and only if $b_1^* \geq R_1^*$ and $b_{2|\text{ad}}^* < R_{2|\text{ad}}^*$. But if the second condition holds, the first simplifies to

$$\mu\phi_M - k + \mu(1-\phi_M)\phi_M \geq (\mu\phi_M - k)^+\,, \tag{8}$$

which holds if $\mu\phi_M \geq k$. If $\mu\phi_M < k$, then (8) simplifies to $k \leq \mu\phi_M(2-\phi_M)$. This last inequality can be re-arranged in terms of ϕ_M as $\phi_M \geq 1 - \frac{\sqrt{(\mu-k)^+}}{\sqrt{\mu}}$. In total, the intersection of the two conditions simplifies to $\phi_M \geq 1 - \frac{\sqrt{\mu(\mu-k)^+}}{\mu}$ and $(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta < k$. The latter condition can be simplified using its concavity. To that end, let

$$\underline{\mu} = \frac{\phi_M + \beta(1-\phi_M)^2 - \sqrt{\left(\beta(1-\phi_M)^2 + \phi_M\right)^2 - 4k\phi_M}}{2\phi_M},$$

$$\overline{\mu} = \frac{\phi_M + \beta(1-\phi_M)^2 + \sqrt{\left(\beta(1-\phi_M)^2\right)^2 - 4k\phi_M} + \phi_M}{2\phi_M}$$

be the two roots of $(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta = k$. The larger root $\overline{\mu}$ is greater than 1 for all β greater than $\overline{\beta} = \frac{k}{(1-\phi_M)^2}$, and the roots do not exist for all β smaller than $\underline{\beta} = \frac{2\sqrt{k\phi_M} - \phi_M}{(1-\phi_M)^2}$. Algebraic manipulations yield the conditions stated in the proposition.

■

Proof of Proposition 2

Proof. We derive the equilibrium strategies for two subgames: one in which the advertiser showed its ad in Period 1, and the other in which it did not.

First, consider the advertiser's Period 2 bidding problem when it has shown ads in Period 1. Let R_2^i be the reserve prices for impression type $i \in \{T, M, TM\}$. Impression type T (M) denotes the impression for which the consumer is in funnel state T (M), and TM denotes the impression for which the consumer is in either funnel state T or M.

Suppose the advertiser submits bid b_2^i for impression type i. The advertiser's payoff is

$$\pi_{2|\text{ad}}^A\left(b_2^T, b_2^M\right) = \begin{cases} (1-\mu)\left(\mu\phi_M - R_2^T\right) + \mu(1-\phi_M)\left(\beta + (1-\beta)\phi_M - R_2^M\right) & \text{if } b_2^T \geq R_2^T, b_2^M \geq R_2^M, \\ (1-\mu)(\mu\phi_M - R_2^T) + \mu(1-\phi_M)\phi_M & \text{if } b_2^T \geq R_2^T, b_2^M < R_2^M, \\ \mu(1-\phi_M)(\beta + (1-\beta)\phi_M - R_2^M) & \text{if } b_2^T < R_2^T, b_2^M \geq R_2^M, \\ \mu(1-\phi_M)\phi_M & \text{if } b_2^T < R_2^T, b_2^M < R_2^M. \end{cases}$$

Whether $b_2^M \geq R_2^M$ or $b_2^M < R_2^M$, the weakly dominant bid for the T-impression is $b_2^{T^*} = \mu\phi_M$. And regardless of b_2^T, the weakly dominant bid for the M-impression is $b_2^{M^*} = \beta(1-\phi_M)$.

Next, consider the advertiser's payoff from bidding for TM:

$$\pi_{2|\text{ad}}^A\left(b_2^{TM}\right) = \begin{cases} (1-\mu)\mu\phi_M + \mu(1-\phi_M)\left(\beta + (1-\beta)\phi_M\right) - \left((1-\mu) + \mu(1-\phi_M)\right) R_2^{TM} & \text{if } b_2^{TM} \geq R_2^{TM}, \\ \mu(1-\phi_M)\phi_M & \text{if } b_2^{TM} < R_2^{TM}. \end{cases}$$

It follows that $b_2^{TM^*} = \frac{(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta}{(1-\mu)+\mu(1-\phi_M)}$.

The ad network anticipates $b_2^{i^*}$ for $i \in \{T, M, TM\}$ and sets R_2^i that maximizes its Period 2 profit. There are four candidates that the ad network considers:

$$\left(R_2^T, R_2^M, R_2^{TM}\right) = \begin{cases} (\max[k, \mu\phi_M], \infty, \infty) & \text{induces } T\text{-ad sales}, \\ (\infty, \max[k, \beta(1-\phi_M)], \infty) & \text{induces } M\text{-ad sales}, \\ (\max[k, \mu\phi_M], \max[k, \beta(1-\phi_M)], \infty) & \text{induces } T\text{- and } M\text{-ad sales}, \\ \left(\infty, \infty, \max\left[k, \frac{(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta}{(1-\mu)+\mu(1-\phi_M)}\right]\right) & \text{induces } TM\text{-ad sales}. \end{cases}$$

If the ad network chooses the first candidate, then only Period 2 impressions for consumers in funnel state T are potentially sold. Since the size of T-consumers in Period 2 is $1-\mu$, this strategy yields ad network profit $(1-\mu)(\mu\phi_M - k)^+$. Similarly, the second candidate yields profit $\mu(1-\phi_M)(\beta(1-\phi_M)-k)^+$, the third $(1-\mu)(\mu\phi_M - k)^+ + \mu(1-\phi_M)(\beta(1-\phi_M)-k)^+$, and the fourth $((1-\mu)(\mu\phi_M - k) + \mu(1-\phi_M)(\beta(1-\phi_M)-k))^+$. From Claim 2, it follows that the third candidate $\left(R_2^T, R_2^M, R_2^{TM}\right) = (\max[k, \mu\phi_M], \max[k, \beta(1-\phi_M)], \infty)$ yields the highest payoff. Therefore, provided ads are shown in Period 1, ads are shown to T-consumers in Period 2 if and only if $\mu\phi_M \geq k$, and ads are shown to M-consumers if and only if $\beta(1-\phi_M) \geq k$.

Next, consider the second subgame wherein the advertiser did not show ads in Period 1. Then in Period 2, the advertiser's payoff from bidding b_2, given reserve price R_2 is

$$\pi^A_{2|\text{no ad}}(b_2) = \begin{cases} \mu\phi_M - R_2 & \text{if } b_2 \geq R_2 \\ 0 & \text{if } b_2 < R_2. \end{cases}$$

By similar reasoning as above, it follows that $b_2^* = \mu\phi_M$ and $R_2^* = \max[k, \mu\phi_M]$. The ad network's Period 2 payoff in this subgame is $(\mu\phi_M - k)^+$.

With the subgame results at hand, we can solve for the Period 1 game. The advertiser's total payoff from bidding b_1 in Period 1, given reserve price R_1, is

$$\pi^A(b_1) = \begin{cases} \mu\phi_M - R_1 + \mu(1-\phi_M)\phi_M & \text{if } b_1 \geq R_1, \\ 0 & \text{if } b_1 < R_1, \end{cases}$$

where the term $\mu(1-\phi_M)\phi_M$ represents the advertiser's Period 2 payoff when it shows ads in Period 1. Claim 1 implies that the advertiser's weakly dominant bid is $b_1^* = \mu\phi_M + \mu(1-\phi_M)\phi_M$. The ad network anticipates this and sets the reserve price as high as b_1^*, provided $R_1 - k + (1-\mu)(\mu\phi_M - k)^+ + \mu(1-\phi_M)(\beta(1-\phi_M)-k)^+ \geq (\mu\phi_M - k)^+$; i.e.,

$$R_1^* = \max\left[k - \left((1-\mu)(\mu\phi_M - k)^+ + \mu(1-\phi_M)(\beta(1-\phi_M)-k)^+\right) + (\mu\phi_M - k)^+, b_1^*\right].$$

Therefore, Period 1 ads are shown if and only if $b_1^* \geq R_1^*$, which is equivalent to

$$\mu\phi_M + \mu(1-\phi_M)\phi_M \geq k - \left((1-\mu)(\mu\phi_M - k)^+ + \mu(1-\phi_M)(\beta(1-\phi_M)-k)^+\right) + (\mu\phi_M - k)^+. \quad (9)$$

Suppose $\mu\phi_M \geq k$. Then (9) simplifies to $\mu(1-\phi_M)\phi_M \geq -\left((1-\mu)(\mu\phi_M - k) + \mu(1-\phi_M)(\beta(1-\phi_M)-k)^+\right)$, which is true. Suppose $\mu\phi_M < k$. Then (9) simplifies to $\mu \geq \tilde{\mu} \equiv k\left(\phi_M(2-\phi_M) + (1-\phi_M)(\beta(1-\phi_M)-k)^+\right)^{-1}$. Thus, Period 1 ads are shown if and only if either $\mu \geq \frac{k}{\phi_M}$ or $\tilde{\mu} \leq \mu < \frac{k}{\phi_M}$. Since $\frac{k}{\phi_M} \geq \tilde{\mu} \Leftrightarrow \phi_M \leq \phi_M(2-\phi_M) + (1-\phi_M)(\beta(1-\phi_M)-k)^+ \Leftarrow \phi_M \leq \phi_M(2-\phi_M) \Leftrightarrow 1 \leq 2-\phi_M$, which is true for all $\phi_M \in [0,1]$, we obtain that Period 1 ads are shown if and only if $\mu \geq \tilde{\mu}$. ■

Proof of Proposition 3

Proof. Let $q^*(0)$ and $q^*(1)$ denote the equilibrium ad intensities without and with tracking, respectively. We being the proof with three observations. First, note that $q^*(1) < 2$ because with tracking, ads are not shown to consumers who had already purchased. Second, if $q^*(0) > 0$, then $q^*(1) > 0$. To see this, suppose that $q^*(0) = 1$. Then under tracking, the ad network can replicate this no-tracking payoff by showing ads only in Period 1. Similarly, if $q^*(0) = 2$, then under tracking, the ad network can generate a weakly higher profit by showing ads to all consumers except those who already purchased. In either case, the ad network's profit under tracking when it shows ads is higher than not showing any ads, because $q^*(0) > 0$ implies showing ads generates positive surplus. Therefore, $q^*(0) > 0$ implies $q^*(1) > 0$.

Put together, we obtain that $q^*(0) > q^*(1)$ if and only if $q^*(0) = 2$. The condition for $q^*(0) = 2$ is given in Proposition 1. Moreover, $q^*(0) = q^*(1)$ if and only if either $q^*(0) = q^*(1) = 0$ or $q^*(0) = q^*(1) = 1$. The ad intensities are both zero if and only if $\mu < \tilde{\mu}$ (such that $q^*(1) = 0$) and $\beta < \tilde{\beta}$ and $\phi_M < 1 - \frac{\sqrt{(\mu-k)^+}}{\sqrt{\mu}}$ (such that $q^*(0) = 0$). But $\mu < \tilde{\mu}$ implies $\phi_M < 1 - \frac{\sqrt{(\mu-k)^+}}{\sqrt{\mu}}$, so the condition for $q^*(0) = q^*(1) = 0$ simplifies to $\mu < \tilde{\mu}$ and $\beta < \tilde{\beta}$.

Next, we derive the conditions under which the ad intensities are 1 in either tracking scenario. First, note that if $q^*(1) = 1$, then $q^*(0) < 2$. This is because $q^*(1) = 1$ implies that not showing ads in Period 2 under tracking is better than showing. And since showing ads in Period 2 with tracking yields weakly higher profit than showing ads in Period 2 without tracking, we obtain by transitivity that without tracking, not showing ads in Period 2 is more profitable than showing ads. Therefore, the condition $q^*(1) = 1$ and $q^*(0) = 1$ are jointly satisfied if and only if $\tilde{\mu} < \mu \leq \frac{k}{\phi_M}$ (such that $q^*(1) = 1$) and $\mu > \frac{k}{\phi_M(2-\phi_M)}$ (such that $q^*(0)$ is either 1 or 2). In total, $q^*(0) = q^*(1) = 1$ if and only if $\frac{k}{\phi_M(2-\phi_M)} < \mu \leq \frac{k}{\phi_M}$ and $\beta \leq \frac{k}{1-\phi_M}$. ■

Proof of Proposition 4

Proof. If ads are not shown in Period 1, then the ad network's Period 2 payoffs with and without tracking are the same at $(\mu\phi_M - k)^+$. On the other hand, if ads are shown in Period 1, then the Period 2 subgame under tracking yields the following ad network payoff $\pi^N_{2|\text{ad}} = (1-\mu)(\mu\phi_M - k)^+ + \mu(1-\phi_M)(\beta(1-\phi_M) - k)^+$. The ad network's payoff under no tracking is $\pi^N_{2|\text{no ad}} = ((1-\mu)\mu\phi_M + \mu(1-\phi_M)\beta(1-\phi_M) - k)^+$.

But we have

$$
\begin{aligned}
\pi^N_{2|\text{no ad}} &= ((1-\mu)\mu\phi_M + \mu(1-\phi_M)\beta(1-\phi_M) - k)^+ \\
&\leq ((1-\mu)(\mu\phi_M - k) + \mu(1-\phi_M)(\beta(1-\phi_M) - k))^+ \\
&\leq (1-\mu)(\mu\phi_M - k)^+ + \mu(1-\phi_M)(\beta(1-\phi_M) - k)^+ \\
&= \pi^N_{2|\text{ad}}.
\end{aligned}
$$

Finally, in Period 1, the ad network faces the same problem with and without tracking, except that it anticipates a higher Period 2 payoff with tracking if ads are shown in Period 1. Therefore, the total profit is weakly greater with tracking than without. ■

Proof of Proposition 5

Proof. Consumers opt-in to tracking only if $q^*(0) > q^*(1)$. But recall from from Proposition 3 that $q^*(0) > q^*(1)$ if and only if $q^*(0) = 2$. Therefore, the necessary condition for opting-in is $q^*(0) = 2$. The sufficient condition is that the consumer's privacy cost is low enough that the benefit of seeing fewer ads outweighs the privacy cost of opting-in. The marginal consumer is the consumer with cost $\min[1, \tilde{\theta}]$ such that $-\eta q^*(1) - \tilde{\theta} = -\eta q^*(0)$. ■

Proof of Proposition 6

Proof. Ad network's profit can decrease in μ due to two and only two reasons: (a) higher μ implies lower opt-in rate such that ad network profit decreases towards the opt-out profit, which is lower than opt-in profit, and (b) large μ implies higher ad intensity under tracking such that consumers opt-out.

The first part occurs if and only if $q^*(0) = 2$ and $q^*(1) = 1+\mu(1-\phi_M)$; i.e., under tracking, ads are only shown to M-consumers. If ads were shown to T-consumers as well, $q^*(1)$ would decrease in μ such that opt-in rate increases with μ. The opt-in rate is $F(\eta(2-(1+\mu(1-\phi_M)))) = F(\eta(1-\mu(1-\phi_M)))$, which decreases in μ if and only if $\eta(1-\mu(1-\phi_M)) \in (0,1)$. For the uniform distribution $F(\theta) = \theta$, the ad network's profit is

$$
\begin{aligned}
\pi_N = {} & \mu\phi_M + \mu(1-\phi_M)(\beta + (1-\beta)\phi_M) - (1+\mu(1-\phi_M))k \\
& + (1-\eta(1-\mu(1-\phi_M)))\left((1-\mu)\mu\phi_M - (1-(1-\mu)\phi_M)k\right).
\end{aligned} \tag{10}
$$

We want to find the conditions under which (10) decreases in μ. Note that

$$
\frac{\partial \pi_N}{\partial \mu} = \beta + 2\eta k(\mu-1) - \phi_M\left(2\beta + \eta\left(4k\mu - 2k + 3\mu^2 - 4\mu + 1\right) + 2\mu - 3\right) + \phi_M^2(\beta + \eta\mu(2k + 3\mu - 2) - 1).
$$

Since the second derivative of the above is $-6(1-\phi_M)\phi_M < 0$, we have that the derivative is concave in μ. Therefore, (10) is decreasing in μ for $\mu < \underline{\mu}'$ and $\mu > \overline{\mu}'$ where the thresholds are respectively given by the two roots of (10) in increasing order.

The second part follows from Proposition 5: if $\phi_M > 1 - \sqrt{(\mu-k)^+/\mu}$ and $\underline{\beta} \leq \beta < \overline{\beta}$, then for $\mu = \overline{\mu}^-$, consumers opt-in, and for $\mu = \overline{\mu}^+$, consumers opt-out. The efficiency loss associated with the increase in opt-out rate creates downward jump in the ad network's profit (cf. Proposition 4). ■

Proof of Proposition 7

Proof. Denote by $q(1)$ and $q(0)$ the total expected ad intensity with and without tracking, respectively. Furthermore, denote by $CS(1)$ and $CS(e)$ the total consumer surplus with full and endogenous tracking, respectively. Let $\tilde{\theta} = \max[0, \min[1, \eta(q(0) - q(1))]]$. Then the result follows from

$$CS(e) = \int_0^{\tilde{\theta}} -\eta q^*(1) - \theta \, dF + \int_{\tilde{\theta}}^1 -\eta q^*(0) \, dF \geq \int_0^{\tilde{\theta}} -\eta q^*(1) - \theta \, dF + \int_{\tilde{\theta}}^1 -\eta q^*(1) - \theta \, dF = CS(1).$$

■

Proof of Proposition 8

Proof. We first show that opting-out of tracking does not signal the consumer's types. Let ρ_i and ρ_j denote advertiser i and advertiser j's beliefs, respectively, that the consumer behind the opt-out impression is type i. By Bayes' rule, the beliefs must satisfy

$$\rho_i = \frac{\lambda S_i(\rho_i, \rho_j)}{\lambda S_i(\rho_i, \rho_j) + (1-\lambda) S_j(\rho_i, \rho_j)},$$

where $S_i(\rho_i, \rho_j)$ denotes the mass of type i consumers who choose to opt-out given advertisers' beliefs ρ. But a type i consumer will opt-out if and only if

$$-\theta - \eta \, q_i(1) < -\eta \, q_i(0; \rho_i, \rho_j),$$

where $q_i(1)$ and $q_i(0; \rho_i, \rho_j)$ is the total number of ads a type i consumer expects to see if she opts-in and -out, respectively. But $q_i(1)$ is independent of consumer's type i because if a consumer opts-in to tracking, the number of ads she expects to see depends only on the parameters μ, β and ϕ_M. Similarly, $q_i(0; \rho_i, \rho_j)$ is independent of consumer's type i because by definition, advertisers cannot base their strategies on consumers' types if they opt-out. Therefore, we obtain

$$S_i(\rho_i, \rho_j) = |\{\theta : -\theta - \eta \, q_i(1) < -\eta \, q_i(0; \rho_i, \rho_j)\}| \equiv S(\rho_i, \rho_j),$$

which implies

$$\rho_i^* = \frac{\lambda S(\rho_i^*, \rho_j^*)}{\lambda S(\rho_i^*, \rho_j^*) + (1-\lambda) S(\rho_i^*, \rho_j^*)} = \lambda.$$

Next, we derive the conditions under which the advertising outcomes diverge from the single-advertiser main model. Since advertiser i has more loyal consumers, the only new outcome that is possible is the following: in the opt-out market, advertiser i advertises in Period 1 and then advertiser j advertises in Period 2. This occurs if and only if the following three conditions hold:

1. advertiser j's Period 2 bid, conditional on advertiser i's ad begin shown in Period 1, (a) exceeds that of advertiser i and (b) is greater than or equal to the reserve price,
2. advertiser i's bid in Period 1 exceeds the reserve price, and
3. the ad network's profit is higher selling Period 1 ads that not selling them.

Condition 1(a) is equivalent to $(1-\lambda)\mu\phi_M > \lambda\left((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta\right)$. But the difference $(1-\lambda)\mu\phi_M - \lambda\left((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta\right)$ is convex with with respect to μ with two roots 0 and $-\frac{1}{\lambda} + \beta\left(\phi_M + \frac{1}{\phi_M} - 2\right) + 2$. And since $\lambda > \frac{1}{2}$ implies $-\frac{1}{\lambda} + \beta\left(\phi_M + \frac{1}{\phi_M} - 2\right) + 2 > \beta\left(\phi_M + \frac{1}{\phi_M} - 2\right) + 2 - 2 = \beta\left(\phi_M + \frac{1}{\phi_M} - 2\right) > 0$, we obtain that Condition 1(a) simplifies to $\mu > -\frac{1}{\lambda} + \beta\left(\phi_M + \frac{1}{\phi_M} - 2\right) + 2$.

Condition 1(b) is equivalent to $(1-\lambda)\mu\phi_M \geq k$, which, combined with $\lambda > \frac{1}{2}$, implies $\lambda\mu\phi_M \geq k$; this in turn implies Condition 2. Finally, Condition 3, provided Conditions 1 and 2, is equivalent to $\lambda\mu\phi_M + \lambda\mu(1-\phi_M)\phi_M - k + (1-\lambda)\mu\phi_M - k > \lambda\mu\phi_M - k$. This simplifies to $(1-\lambda)\mu\phi_M - k + \lambda\mu(1-\phi_M)\phi_M > 0$, which is implied by Condition 1(b): $(1-\lambda)\mu\phi_M \geq k \iff \lambda < 1 - \frac{k}{\mu\phi_M}$.

In sum, the conjunction of Conditions 1 through 3 simplify to $\mu > -\frac{1}{\lambda} + \beta\left(\phi_M + \frac{1}{\phi_M} - 2\right) + 2 \equiv \tilde{\mu}$ and $\lambda < 1 - \frac{k}{\mu\phi_M} \equiv \tilde{\lambda}$. ■

Proof of Proposition 9

Proof. Next, it suffices to characterize the conditions under which (i) ads are shown to all opt-out consumers in both periods, and (ii) the ad intensity for opt-in consumers increases with signal accuracy ρ. If both conditions hold, then fewer ads are shown under tracking, and more consumers opt-out from tracking as ρ increases. The first condition is derived from Proposition 1. For the second condition, we begin by characterizing the ad network's ad supply decisions for opt-in consumers with imperfect purchase observability. For expositional ease, denote by N- and P-impressions the impressions associated with "not purchased" and "purchased" signals, respectively.

Given the advertiser's weakly dominant bids for T-, N-, and P-impressions, the ad network's profits from selling each type of impressions are $\mu\phi_M - k$, $\frac{(1-\phi_M)(1-\rho)}{\phi_M\rho + (1-\phi_M)(1-\rho)}\beta(1-\phi_M) - k$, and $\frac{(1-\phi_M)\rho}{\phi_M(1-\rho) + (1-\phi_M)\rho}\beta(1-$

$\phi_M) - k$, respectively. The ad network sells whichever ad impressions yield positive profit.

Note that it is never profitable for the ad network to sell P-impressions but not N-impressions. The reason is that the fact that signals are at least partially informative imply N-impressions are valued more by the advertiser than P-impressions are.

First, note that $\frac{(1-\phi_M)(1-\rho)}{\phi_M\rho+(1-\phi_M)(1-\rho)} \leq \frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho}$ for all $\frac{1}{2} \leq \rho \leq 1$, because $\frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho} - \frac{(1-\phi_M)(1-\rho)}{\phi_M\rho+(1-\phi_M)(1-\rho)} = \frac{(2\rho-1)(\phi_M-1)\phi_M}{((2\rho-1)\phi_M-\rho)(-\rho+(2\rho-1)\phi_M+1)}$ and the latter term's sign is equivalent to that of $\frac{1-\phi_M}{\rho+(1-2\rho)\phi_M}$. Now, the denominator $\rho+(1-2\rho)\phi_M$ is always positive because it is a linear function of ρ and is positive at each endpoint $\rho = \frac{1}{2}$ and $\rho = 1$.

Second, note that $\frac{\partial}{\partial\rho}\frac{(1-\phi_M)(1-\rho)}{\phi_M\rho+(1-\phi_M)(1-\rho)} = -\frac{(1-\phi_M)\phi_M}{(-\rho+(2\rho-1)\phi_M+1)^2} < 0$ and $\frac{\partial}{\partial\rho}\frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho} = \frac{(1-\phi_M)\phi_M}{(\rho-2\rho\phi_M+\phi_M)^2} > 0$.

Third, since the bounds $\frac{(1-\phi_M)(1-\rho)}{\phi_M\rho+(1-\phi_M)(1-\rho)}\beta(1-\phi_M)$ and $\frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho}\beta(1-\phi_M)$ coincide at $\rho = \frac{1}{2}$, at which point the bounds equal $\beta(1-\phi_M)^2$, we obtain the following:

1. If $\mu\phi_M \geq k$, and $k > \beta(1-\phi_M)^2$, then as ρ increases from $\frac{1}{2}$ to 1, the regime changes from $k > \frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho}\beta(1-\phi_M)$ to $\frac{(1-\phi_M)(1-\rho)}{\phi_M\rho+(1-\phi_M)(1-\rho)}\beta(1-\phi_M) < k < \frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho}\beta(1-\phi_M)$; i.e., only T-impressions are shown for low ρ, and then N-impressions are also shown for high ρ.
2. If $\mu\phi_M < k$, and $k > \beta(1-\phi_M)^2$, then as ρ increases from $\frac{1}{2}$ to 1, the regime changes from $k > \frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho}\beta(1-\phi_M)$ to $\frac{(1-\phi_M)(1-\rho)}{\phi_M\rho+(1-\phi_M)(1-\rho)}\beta(1-\phi_M) < k < \frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho}\beta(1-\phi_M)$; i.e., no impressions are shown for low ρ, and then N-impressions are shown for high ρ.

This constitutes the second condition (i.e., the ad intensity for opt-in consumers increases with signal accuracy ρ).

However, the conditions $\mu\phi_M < k$ and $k > \beta(1-\phi_M)^2$ cannot hold jointly with $q^*(0) = 2$, which requires $(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta > k$. To see this, it suffices to show that $(1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta \leq \max[\beta(1-\phi_M)^2, \mu\phi_M]$, which holds because the left-hand side of the inequality is a linear combination of $\mu\phi_M$ and $\beta(1-\phi_M)^2$, so it must be smaller than the larger of $\mu\phi_M$ and $\beta(1-\phi_M)^2$. Therefore, for the conditions (i) and (ii) above to hold simultaneously, it must be that $\mu\phi_M \geq k$, and $k > \beta(1-\phi_M)^2$.

Finally, to ensure that the threshold of ρ past which the advertising regime changes from fewer to more advertising is between $\frac{1}{2}$ and 1, we must bound k by the largest value attained by the upper bound $\frac{(1-\phi_M)\rho}{\phi_M(1-\rho)+(1-\phi_M)\rho}\beta(1-\phi_M)$, which occurs at $\rho = 1$ and equals $\beta(1-\phi_M)$.

Re-arranging the condition $\beta(1-\phi_M)^2 < k < \min[\mu\phi_M, \beta(1-\phi_M)]$ with respect to ϕ_M and β yields the conditions in the proposition. ■

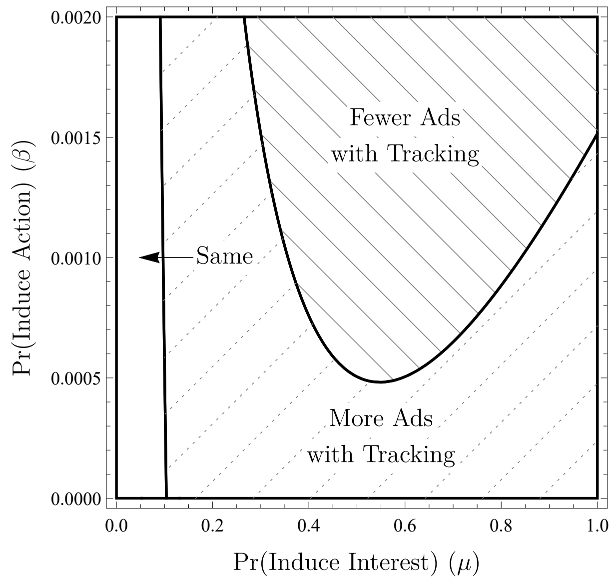

Figure 14: Ad Intensity with Parameter Scaling; $\alpha = 0.001, \phi_M = 0.005, k = 1.5 \times 10^{-6}$

C Parameter Scaling

We demonstrate the robustness of our main insights to smaller values of advertising effectiveness. To that end, suppose there exist two consumer segments: a potentially responsive segment and a non-responsive segment, whose sizes are given by α and $1 - \alpha$, respectively, for some small $\alpha \in (0, 1)$. We assume that the potentially responsive consumers respond to ads in the manner described in the main model, while the non-responsive consumers always ignore ads; i.e., they never respond to ads.

Without consumer tracking, the advertiser cannot distinguish between these segments, while with tracking, it can. Therefore, the ad intensity under no tracking is

- 2 if $\alpha\left(\mu\phi_M + \mu(1-\phi_M)\phi_M\right) - k + \alpha\left((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta\right) - k \geq (\alpha\mu\phi_M - k)^+$ and $\alpha\left((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta\right) - k \geq 0$,
- 1 if $\alpha\left(\mu\phi_M + \mu(1-\phi_M)\phi_M\right) - k \geq (\alpha\mu\phi_M - k)^+$ and $\alpha\left((1-\mu)\mu\phi_M + \mu(1-\phi_M)^2\beta\right) - k < 0$, and
- 0 otherwise.

Similarly, the ad intensity under tracking is

- $1 + \alpha(1-\mu)\mathbb{I}_{\{\mu\phi_M - k \geq 0\}} + \alpha\mu(1-\phi_M)\mathbb{I}_{\{\beta(1-\phi_M)-k\geq 0\}}$ if $\alpha\left(\mu\phi_M + \mu(1-\phi_M)\phi_M\right) - k + \alpha\left((1-\mu)(\mu\phi_M - k)^+ + \mu(1-\phi_M)(\beta(1-\phi_M) - k)^+\right) \geq 0$, and
- 0 otherwise.

Note that if we let $k' = k/\alpha$, then the ad intensity under no tracking is equivalent to the main model with ad cost k', and the ad intensity under tracking is either 0, 1 or between 1 and 2 under the same conditions as the main model. Thus, the conditions for the ad intensity differential are preserved from the main model.

As illustrated in Figure 14, we can replicate the advertising intensity differential patterns of the main model (Figure 5) for small values of α, ϕ_M, and k. Since the main insights rest on the ad intensity differential pattern, this suffices to show that the insights are robust to parameter scaling.

D Markov-Perfect Equilibrium

For any given Period t, define "old generation" as the mass of consumers who arrived in Period $t-1$, and "new generation" as those who arrive in Period t. In our setting, the payoff-relevant states can be fully characterized by the distribution of old generation non-converters in funnel states T and M. Consider the no tracking case where the advertiser cannot target ads based on the consumers' funnel states, nor their purchase history. Let λ_f^{old} denote the proportion of old-generation non-converters in funnel state $f \in \{T, M\}$. There are two possible states in each period: one in which the advertiser showed ads in the previous period, and another in which it did not show ads in the previous period.

To elaborate, suppose the advertiser showed ads in Period $t-1$. The old generations in Period $t-1$ (i.e., those who arrived in Period $t-2$) leave by Period t because consumers only live for two periods. Therefore, these consumers are irrelevant in the analysis of determining the successive distribution of old generation non-converters in Period t. Of the $1-\sigma$ T-consumers who arrived in Period $t-1$, $1-\mu$ fraction are not influenced by the ad and stay in T, μ fraction transition to M, of which ϕ_M convert and $1-\phi_M$ do not. Moreover, of the σ M-consumers who joined in Period $t-1$, $1-\beta$ stay in M, and still $1-\phi_M$ of those $\sigma(1-\beta)$ M-consumers do not convert. Therefore, the distribution of old generation non-converters in Period t would be

$$\left(\lambda_T^{\text{old}}, \lambda_M^{\text{old}}\right) = \left(\, (1-\sigma)(1-\mu),\, ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)\,\right).$$

We label this state as λ_1, where the subscript 1 indicates the advertiser showed ads in the previous period.

On the other hand, suppose the advertiser did not show ads in Period $t-1$. Without any ad exposures, the $1-\sigma$ T-consumers who arrived in Period $t-1$ would all remain in T by Period t. However, ϕ_M fraction of σ M-consumers convert and $1-\phi_M$ fraction remain in M in Period t. Therefore, the distribution of old-generation non-converters in this case is

$$\left(\lambda_T^{\text{old}}, \lambda_M^{\text{old}}\right) = \left(\, 1-\sigma,\, \sigma(1-\phi_M)\,\right).$$

We label this state as λ_0, where the subscript 0 indicates the advertiser did not show ads in the previous period.

Given two states and two possible advertising strategies at each state (i.e., advertise or not advertise), there are four Markov-perfect equilibrium (MPE) candidates: (i) always advertise regardless of the state; (ii) advertise only when the state is λ_0, which is equivalent to "pulse advertising" (i.e., alternate advertising with a single-period break in between; (iii) advertise only when the state is λ_1, which is effectively equivalent to (i); and (iv) never advertise. We compare the ad network's profits for the respective strategies.

I. Always advertise

For always advertising to be MPE, the advertiser's payoff from buying untargeted ads in Period t, given the state is either $\lambda_1 \equiv ((1-\sigma)(1-\mu), ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M))$ or $\lambda_0 \equiv (1-\sigma, \sigma(1-\phi_M))$, should be greater than that from not buying:

$$\begin{aligned}&(1-\sigma)(1-\mu)\mu\phi_M + ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)(\beta + (1-\beta)\phi_M) + (1-\sigma)\mu\phi_M - 2R_1 + \delta V_1 \\ &\quad \geq ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)\phi_M + \delta V_0,\end{aligned}$$

and

$$2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)(\beta + (1-\beta)\phi_M) - 2R_0 + \delta V_1 \geq \sigma(1-\phi_M)\phi_M + \sigma\phi_M + \delta V_0,$$

where V_1 is the continuation value from having shown ads in the previous stage, and V_0 is the continuation value from not having shown any ads in the previous stage. In equilibrium, the ad network will set reserve prices R_1 and R_0 such that these conditions bind; otherwise, it leaves money on the table. Therefore, from the second condition, we obtain

$$2R_0{*} = 2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)\beta(1-\phi_M) + \delta(V_1 - V_0).$$

But if the second condition holds, it must be that the continuation value from not showing ads is the continuation value from showing ads, such that

$$V_0 = 2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)(\beta + (1-\beta)\phi_M) - 2R_0^* + \delta V_1.$$

Then, substituting R_0^* yields $V_0 = \sigma(1-\phi_M)\phi_M + \sigma\phi_M + \delta V_0$, which in turn implies

$$V_0 = \frac{\sigma(2-\phi_M)\phi_M}{1-\delta}.$$

Similarly, after substituting $V_0 = \frac{\sigma(2-\phi_M)\phi_M}{1-\delta}$ into the first condition and letting it bind, we obtain

$$\begin{aligned}2R_1^* = &(1-\sigma)(1-\mu)\mu\phi_M + ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)(\beta + (1-\beta)\phi_M) \\ &+ (1-\sigma)\mu\phi_M + \delta V_1 - (((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)\phi_M + \delta V_0),\end{aligned}$$

which simplifies to

$$2R_1^* = (1-\sigma)(2-\mu)\mu\phi_M + ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)^2\beta + \delta V_1 - \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta}.$$

Since the continuation value of having shown ads is the continuation value from showing ads, we obtain

$$V_1 = (1-\sigma)(1-\mu)\mu\phi_M + ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)(\beta + (1-\beta)\phi_M) + (1-\sigma)\mu\phi_M - 2R_1^* + \delta V_1$$

which, upon substitution of R_1^* yields $V_1 = ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)\phi_M + \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta}$. Therefore,

$$\begin{aligned} 2R_1^* &= (1-\sigma)(2-\mu)\mu\phi_M + ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)^2\beta \\ &\quad + \delta\left(((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)\phi_M + \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta}\right) - \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta}. \end{aligned}$$

Since this strategy induces the state to be perpetually λ_1, the ad network's total profit is

$$\begin{aligned} \pi_N^I = \frac{1}{1-\delta}\Bigg(&(1-\sigma)(2-\mu)\mu\phi_M + ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)^2\beta \\ &+ \delta\left(((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)\phi_M + \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta}\right) - \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta} - 2k\Bigg). \end{aligned}$$

II. Advertise Only When State is $\lambda_0 = (1-\sigma, \sigma(1-\phi_M))$

For this pulsing strategy to be MPE, we need advertiser's payoff to be higher buying ads given $(1-\sigma, \sigma(1-\phi_M))$, and not buying ads given $((1-\sigma)(1-\mu), ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M))$, which respectively translate to:

$$2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)(\beta + (1-\beta)\phi_M) - 2R_0 + \delta V_1 \geq (\sigma(1-\phi_M)+\sigma)\phi_M + \delta V_0$$

and

$$(((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)+\sigma)\phi_M + \delta V_0 \geq 2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)(\beta + (1-\beta)\phi_M) - 2R_1 + \delta V_1.$$

The ad network sets $R_1 = \infty$ (such that no ads are bought at state λ_1) and $2R_0^* = 2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)(1-\phi_M)\beta + \delta(V_1 - V_0)$. This implies $V_0 = \sigma(2-\phi_M)\phi_M + \delta V_0$, which means $V_0 = \frac{\sigma(2-\phi_M)\phi_M}{1-\delta}$. Similarly, since $V_1 = (((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)+\sigma)\phi_M + \delta V_0$, we have

$$\begin{aligned} 2R_0^* &= 2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)(1-\phi_M)\beta \\ &\quad + \delta\left(\left((((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)+\sigma)\phi_M + \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta}\right) - \frac{\sigma(2-\phi_M)\phi_M}{1-\delta}\right). \end{aligned}$$

Since this strategy yields alternating states, the ad network's profit under this strategy is

$$\begin{aligned} \pi_N^{II} = \frac{1}{1-\delta^2}\Bigg(&2(1-\sigma)\mu\phi_M + (\sigma(1-\phi_M)+\sigma)(1-\phi_M)\beta \\ &+ \delta\left(\left((((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)+\sigma)\phi_M + \delta\frac{\sigma(2-\phi_M)\phi_M}{1-\delta}\right) - \frac{\sigma(2-\phi_M)\phi_M}{1-\delta}\right) - 2k\Bigg). \end{aligned}$$

III. Advertise Only When State is $\lambda_1 = (\,(1-\sigma)(1-\mu),\, ((1-\sigma)\mu + \sigma(1-\beta))(1-\phi_M)\,)$

Same as strategy I: set $R_0^* = \infty$ and the rest follows.

IV. Never advertise

This strategy yields 0 payoff.

Finally, comparing the payoffs π_N^I, π_N^{II} and 0 yield the presented equilibrium regions.