

D'Annunzio, Anna; Russo, Antonio

Working Paper

Ad Networks, Consumer Tracking, and Privacy

CESifo Working Paper, No. 6667

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: D'Annunzio, Anna; Russo, Antonio (2017) : Ad Networks, Consumer Tracking, and Privacy, CESifo Working Paper, No. 6667, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/171131>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Ad Networks, Consumer Tracking, and Privacy

Anna D'Annunzio, Antonio Russo

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Ad Networks, Consumer Tracking, and Privacy

Abstract

We study the relation between ad networks, consumer privacy and the online advertising market. We consider two publishers that can outsource their ad inventories to an ad network, in a market where consumers and advertisers endogenously multi-home. Differently from publishers, the ad network tracks consumers across websites, limiting wasteful repetition of ads. However, its tracking capability depends on consumer privacy-related choices (e.g., accepting third-party cookies). We show that tracking may increase or decrease the provision of ads, depending on its effect on expected advertising returns and on how audience sizes respond to ad quantities. When they decide whether to allow tracking, consumers exert a positive externality on advertisers. If tracking reduces the provision of ads, there is also a positive indirect externality on consumers. Hence, there may be too little tracking in equilibrium, even from consumers' perspective. We evaluate several privacy policies, including direct regulatory interventions and the creation of markets for the right to track consumers. Finally, we characterize the conditions such that outsourcing to the ad network expands the provision of ads compared to the case where publishers compete directly for advertisers.

JEL-Codes: D430, D620, L820, M370.

Keywords: advertising, ad network, internet, tracking, multi-homing, privacy.

Anna D'Annunzio
CSEF, University of Naples Federico II
Naples / Italy
dannunzio.anna@gmail.com

Antonio Russo
ETH Zurich
Zurich / Switzerland
russo@kof.ethz.ch

This version: September 2017

We thank Stefan Behringer, Alberto Bennardo, Emilio Calvano, Grazia Cecere, Jan Cauwer, Lapo Filistrucchi, Øystein Foros, Bruno Jullien, Hans Jarle Kind, Tore Nilssen, Jerome Pouyet, Alessia Russo, Marianne Verdier, Ying Lei Tho, Liad Wagman, and the participants at Telenor Research seminar, NHH seminar, CSEF seminar, CRESSE 2017, the 8th Searle Center Conference on Internet Commerce and Innovation, ICT Telecom-ParisTech conference 2017, EUIMedia Conference 2017, CESifo Area Conference on Applied Microeconomics 2017, ADRES conference 2017, the 6th EIEF-UNIBO-IGIER Bocconi workshop on IO, the 2016 Economics of IP, Software and Internet Conference in Toulouse, EARIE 2016 and the 2016 Media Economics Workshop in Zurich, for useful comments and discussions. All errors are ours alone. Part of this research was carried out while Anna D'Annunzio was employed by Telenor. The views expressed in this paper are those of the authors only and not those of Telenor.

1 Introduction

An increasingly large share of advertising takes place online: the Interactive Advertising Bureau estimates that the US online advertising market was worth about 90% of the TV market in 2015 (IAB, 2015). Several features of this market distinguish it from traditional ones. In this paper, we focus on two aspects which have received little attention in previous research. First, digital firms use new technologies to provide advertising, such that the effectiveness of online advertising hinges on their ability to monitor consumers and on the extent to which the latter protect their own personal information. Given the large amount of data that firms can collect, consumers and regulators are increasingly concerned about the potential consequences of reduced privacy.¹ Second, most digital publishers outsource their ad inventory to advertising networks, such as Google’s AdSense.² Ad networks typically have better technologies to track users online than publishers do. Our objective is to provide a theory of the online advertising market that incorporates ad networks, data-driven advertising technologies, and consumers’ choices regarding privacy.

One of the key factors behind the relevance of ad networks is that consumers commonly consult several online contents in a short time frame.³ This behavior reduces the reach (i.e. the expected number of informed consumers) of ad campaigns, because the same ad can hit a consumer several times on different publishers, being wasted with the consumer’s attention. Although each publisher can monitor a consumer on its own web-pages, and thus potentially avoid internal repetitions, it can hardly do so when the consumer visits other websites (Calvano and Jullien, 2012, Athey et al., 2016). Furthermore, whereas publishers are unique gatekeepers to the attention of exclusive viewers, competition for advertisers reduces the revenue generated from multi-homers. The combination of these factors implies that publishers’ ad inventories lose value as the share of multi-homers increases (Ambrus et al. 2016). Ad networks can address these problems in two ways. First, they centralize the sale of advertising space on different publishers. Second, they use data-driven technologies to improve the reach of ad campaigns. Specifically, in this paper we concentrate on *tracking* technologies, that

¹Turow et al. (2009) find that 84% of U.S. respondents say they do not want advertising tailored on their behavior on websites they have visited before (see also Tucker, 2014). A recent Gallup survey shows that 61% of respondents argued that online behavior tracking is unjustifiable (Morales, 2010). Acquisti et al. (2016) provide an overview of online privacy regulation in the EU and the US.

²Roesner et al. (2012) report that 89% of the (Alexa) top 500 sites include at least one cross-site tracker and 40% are tracked by Google’s ad network. The latter reaches 94% of total Internet audience (Comscore, 2016).

³Comscore (2016) reports that 94% of total internet audience visited Google and 80% visited Facebook in February 2016.

allow an ad network to follow consumers' activity online across multiple websites. For instance, by using third-party cookies, ad networks can keep track of which ads a consumer has already been exposed to, avoiding wasteful repetition.⁴

However, following a consumer usually entails storing substantial information on her behavior and preferences. Therefore, consumers may perceive tracking as a violation of their privacy and take measures to avoid it, at least partially (see Goldfarb and Tucker, 2012). For example, consumers can adjust their browser settings to block third-party cookies or use software to anonymize their browsing activity.⁵

The issues outlined above raise several interesting questions. How do data-driven advertising technologies affect market outcomes? What is the impact of consumers' privacy-related decisions on market outcomes and welfare? Should regulators restrict tracking? Should they assign property rights on personal data? How do ad networks affect market outcomes compared to the case where publishers compete? Our paper provides an analytical framework to tackle these questions.

We consider a model with two ad-financed online publishers, an ad network, and homogeneous advertisers. Consumers and advertisers possibly multi-home and ads reduce consumers' utility from browsing the Internet. We first study the scenario where publishers outsource their inventories to the ad network. We account for the fact that the ad network's ability to track consumers depends on the extent to which they protect their privacy, as discussed above. When consumers block tracking, they reduce the ability of the ad network to avoid cross-outlet repetition of impressions. Therefore, the number of consumers blocking tracking affects advertising revenues and, in turn, the equilibrium quantity of ads.

Our first finding is that the provision of online advertising may increase or decrease with tracking, depending on two forces that stem from the two-sided nature of digital publishers. On the one hand, tracking increases the marginal revenue from ads hitting multi-homers, which increases the return from advertising. However, tracking also increases the revenue from each multi-homer and, hence, the

⁴Tracking may also be useful when advertisers intend to show ads in a given sequence (see, for instance, https://support.google.com/dfp_premium/answer/1665531?hl=en) or reach with a personalized offer a consumer who has shown interest in a generic ad. Ad networks also perform other functions, including matching ads with consumers based on information regarding their interests (targeting).

⁵According to a 2011 USA Today.com article, Ghostery, one of the most common block-tracking tools, is "being downloaded by 140000 new users each month, with total downloads doubling to 4.5 million in the past 12 months". Furthermore, "software firm Abine projects the number of Internet users in North America using anti-tracking tools will be 28.1 million by the end of 2012, up from 17.2 million today". See <https://usatoday30.usatoday.com/tech/news/story/2011-12-29/internet-privacy/52274608/1>.

implicit cost of additional ads, because the quantity of multi-homers shrinks when advertising levels go up. Therefore, tracking reduces the equilibrium quantity of ads if the second effect dominates. We argue that this outcome is more likely, for instance, when users' preferences for the two publishers are negatively correlated, because, all else given, the share of multi-homing viewers tends to be smaller, but multi-homers constitute a larger component of the demand response to changes in advertising quantities on a publisher. To the best of our knowledge, these findings uncover a novel link between data-driven advertising technologies and online advertising levels.

We then compare consumers' and publishers' equilibrium choices, respectively, on tracking and advertising provision with optimal ones. We begin by comparing the competitive equilibrium to the first-best allocation, whereby these quantities are chosen to maximize welfare. We find that too few consumers allow tracking, because they ignore the positive effect on advertising revenues. Next, we consider the second-best scenario where tracking is chosen to maximize welfare, while advertising quantities are unregulated. We find that there may be too little or too much tracking in equilibrium, depending on two externalities. First, as mentioned, consumers ignore the impact of their choice on advertising surplus. Second, each consumer neglects the fact that her choice of allowing tracking affects the equilibrium quantity of ads and, in turn, the utility consumers get from browsing content online. If tracking reduces the provision of ads, both externalities are positive. Thus, there is too little tracking in equilibrium. Interestingly, this conclusion holds even from the perspective of consumers. On the other hand, if tracking increases the provision of ads, the externality on consumers is negative. Therefore, whether there is too much or too little tracking is ambiguous. Finally, we find that, compared to the second-best, advertising provision is excessive. The intuition is that firms neglect the disutility from ads suffered by consumers.

Recently, regulators in several countries have proposed new rules that protect consumer privacy by either reducing the costs of blocking tracking or forcing the advertising industry to adopt less intrusive practices.⁶ These policies are designed to reduce consumers' disutility from privacy losses. However, our findings suggest that they may have negative indirect consequences. In particular, as we discuss in Section 6.1, if tracking reduces the provision of ads, restricting it might penalize both the advertising industry and consumers. By contrast, if advertising levels increase with tracking,

⁶For example, the Federal Trade Commission proposed in 2012 a “do not track” mechanism, aimed at empowering consumers to avoid online tracking (see https://en.wikipedia.org/wiki/Do_Not_Track).

reducing the cost of blocking benefits consumers, but reduces advertising surplus.

As an alternative to direct regulatory intervention, some scholars and regulators have considered the adoption of market-based approaches to ensure an efficient level of exploitation of personal data (Acquisti et al., 2016). We provide some insights on the implications of these approaches, under different allocations of property rights (see Section 6.2). Specifically, we consider the case where the ad network is entitled to track consumers, but it offers an opt-out option in exchange for a fee. We also consider the symmetric case where the ad network has to pay the consumers who opt-in to tracking. We find that, depending on how property rights are allocated, either consumers or the advertising industry are better off compared to the status-quo. However, in neither case the distortions discussed above are fully corrected, due to the ad network’s market power and to the presence of other externality that consumers exert on each other.

To complete the analysis of the effects of ad networks on the online advertising market, in Section 7 we compare the case where publishers outsource and that where they compete directly for advertisers. Outsourcing can either restrict or expand the provision of ads, depending on two effects. The first is due to cross-outlet tracking, as discussed above. Secondly, the ad network centralizes the sale of ads. This implies that, whereas competing publishers can only extract the incremental revenue that advertisers make on overlapping consumers, the ad network captures all of this revenue. Hence, the ad network internalizes the effects of changes in the size and composition of audiences on both publishers. Therefore, it takes into account that, when the provision of advertising on a publisher increases, some multi-homers become single-homers on the other publisher. By contrast, these single-homers are lost from the perspective of a publisher competing head-to-head with the other. Due to these countervailing forces, the centralization of ad inventories may either restrict or expand the provision of ads in equilibrium. Overall, the welfare implications are generally ambiguous.

Throughout most of the analysis, we consider privacy as a “final good” (Farrell, 2012), in the sense that consumers care for it per se.⁷ However, in Section 8 we consider the case where tracking affects the overall relevance of the ads a consumer is exposed to. This extension implies that consumers’

⁷Farrell (2012) argues that privacy is both a final good (“Consumers care about privacy in part for its own sake. [...] Some consumers, and most consumers some of the time, don’t care at all; others care a lot.”) and an intermediate good (“Consumers also care about privacy in a more instrumental way. For instance, loss of privacy could identify a consumer as having a high willingness to pay for something, which can lead to being charged higher prices if the competitive and other conditions for price discrimination are present.”). See Lohr (2010) on the economic impact of privacy concerns.

privacy-related decisions are intertwined with their choice of publishers. Our main results are not fundamentally affected.

The remainder of this paper is organized as follows. Section 2 presents a brief literature review. In Section 3 we describe the baseline model. In Section 4 we study the case where publishers can outsource their inventories to the ad network. Section 5 compares equilibrium and optimal allocations. Section 6 discusses policy implications, including direct regulatory intervention and market-based approaches. Section 7 considers the case where publishers compete directly. Section 8 briefly studies privacy as an intermediate good. Section 9 concludes. Proofs and additional material are in Appendix. Furthermore, an Online Appendix contains formal derivations of the results in Section 8 and a microfoundation for the advertising market.

2 Literature

Our paper belongs to a recent literature that studies advertising in two-sided markets where consumers multi-home. See, e.g., Anderson et al. (2016) and Ambrus et al. (2016).⁸ Following this literature, we account for the fact that repeated ad impressions are partially wasteful, a particularly prominent issue when platforms' users multi-home. However, we depart from previous papers by modeling the ad network. Athey et al. (2016) consider the implications of publishers' inability to track multi-homers, in a framework with heterogeneous advertisers and exogenous audience sizes. They find that advertisers sort in equilibrium, as only those with the highest value of reaching consumers multi-home. Although we ignore advertiser heterogeneity, we allow for endogenous audience sizes. Furthermore, we link the ad network's tracking capability to consumer choices.

Previous literature has also studied the implications of data-driven technologies for online advertising. However, the literature has focused primarily on targeting (see, e.g., Bergemann and Bonatti, 2011). In this paper, we concentrate instead on tracking of consumers across outlets. These technologies make advertising more effective, but in a different way: the former improves the match between ads and consumers, whereas the latter improves the reach of an advertising campaign. A related literature studies targeting and advertising avoidance. Johnson (2013) shows that there may be too little ad-avoidance from a social standpoint. The reason is that consumers ignore that block-

⁸Earlier papers include Anderson and Coate (2005) and Peitz and Valletti (2008), who propose models where consumers single-home.

ing ads discourages advertising, benefiting other consumers who do not block. We instead focus on the decision whether to allow tracking, that has different welfare implications. Moreover, we model platforms (publishers and the ad network) bringing together consumers and advertisers. Hann et al. (2008) study consumer efforts to conceal themselves and to deflect marketing. They find that consumer surplus is higher with deflection, because it allows to discard solicitations rather than to divert them toward other consumers. Aziz and Telang (2017) study how access to third-party information affects competition between advertisers and the division of surplus. They find that access to this information can increase competition among advertisers. Hence, advertisers may be worse off and consumers better off when the former access third-party information.

Our paper also contributes to the growing literature on the economics of privacy (see Acquisti et al., 2016, for an exhaustive survey). We are primarily interested in the link between consumer incentives to protect their privacy and the provision of ads on competing publishers (Tucker, 2012). In this sense, our approach is complementary to previous studies that focus on the incentives of firms to gather and exploit personal information (Bergemann and Bonatti, 2015, de Cornière and de Nijs, 2016, Levin and Milgrom, 2010). Goldfarb and Tucker (2011) analyze how the introduction of Privacy Directive (2002/58/EC) in the European Union, which restricted data collection for targeting purposes, decreased the effectiveness of advertising. Campbell et al. (2015) study the effect of privacy regulation on market structure, finding that forced consent-gathering from consumers has negative effects on smaller firms.

A sizable literature studies privacy and price discrimination. Conitzer et al. (2012) study a setting where firms set personalized prices and consumers choose whether to stay anonymous. Belleflamme and Vergote (2016) study a monopolist that tracks consumers to price discriminate and allow consumers to acquire hiding technologies. Montes et al. (2015) study a similar case in a duopoly setting. de Cornière and de Nijs (2016) a platform’s incentives to reveal personal information, accounting for the consequences on prices in the product market. Differently from them, we study the effect of consumers’ privacy choices on the advertising market.

Finally, our paper relates to the literature on common distributors in media markets. Kind et al. (2016) consider a TV distributor to which channels can delegate the choice over viewer subscription prices. George and Hogendorn (2012) study news aggregators. Differently from them, we consider




Figure 1: Illustration of the market structure.

an aggregator on the advertising side of the market. de Cornière and Taylor (2014) study a search engine that directs consumers to two publishers. They allow one publisher to integrate (at various degrees) with the search engine. We consider instead coordination of both publishers induced by the ad network.

3 Model setup

We consider a market with two publishers, indexed by $i = 1, 2$, and an advertising network, denoted by AN (see Figure 1 for an illustration). For the first part of the analysis, we focus on the scenario where publishers outsource their inventories to the ad network (Section 4). Later on, we study the case where the ad network is unavailable, thus publishers compete directly for advertisers (Section 7).

Publishers and the ad network. Publishers provide free content to consumers and sell their attention to advertisers. We denote by q_i , $i = 1, 2$ the advertising level on publisher i , i.e. the quantity of impressions that the publisher exposes each consumer to.⁹ We sometimes also refer to q_i as publisher i ’s “advertising capacity”. We denote publisher i ’s profit as π_i , whereas the ad network’s profit is denoted as π_{AN} .

Publishers can either sell impressions directly or outsource this function to the ad network. When a publisher outsources, the ad network sells the publisher’s advertising inventory and retains the ensuing revenues, in exchange for a fee. We characterize the contract between publishers and the ad network in Section 4.2 and that with the advertisers in Section 4.2.

Whereas publishers are effective at monitoring the behavior of a consumer on their own website, they have limited ability to do so when a consumer visits other websites (Athey et al., 2016). Instead,

⁹We implicitly assume that all consumers visit the same number of web-pages on a publisher, and the layout of each page (including the quantity of ad space) does not vary with whom is browsing. Hence, the quantity of impressions per consumer is invariant.

by operating on multiple websites, ad networks are able to track consumers across outlets. As we clarify below, this feature has an important effect on the revenues earned by selling ads.

Consumers. There is a unit mass of consumers, who visit either one, both or no publisher. A consumer visits publisher i if and only if she receives a non-negative utility from doing so. We denote by D_{12} the quantity of multi-homing consumers, D_i the quantity of single-homers on publisher i , and D_0 the quantity of consumers visiting no publisher. Following Ambrus et al. (2016), we specify demands as follows

$$\begin{aligned} D_{12} &= Pr \{u_1 - \delta q_1 \geq 0, u_2 - \delta q_2 \geq 0\}, \\ D_i &= Pr \{u_i - \delta q_i \geq 0, u_j - \delta q_j < 0\}, \quad i, j = 1, 2 \text{ and } i \neq j, \\ D_0 &= 1 - D_1 - D_2 - D_{12}, \end{aligned} \tag{1}$$

where u_i is the reservation utility for content on publisher i . We assume that $(u_1, u_2) \sim g(u_1, u_2)$, a bivariate joint distribution with smooth density. The term δq_i captures the disutility generated by ads when visiting publisher i , where δ is a positive parameter.¹⁰

The demand system in (1) has the following notable properties. When q_i increases, publisher i loses multi-homers, who become single-homers on j , i.e. $\frac{\partial D_{12}}{\partial q_i} < 0$ and $\frac{\partial D_j}{\partial q_i} > 0$. Hence, the composition of j 's audience changes with q_i . However, j 's total demand does not, because $\frac{\partial (D_{12} + D_j)}{\partial q_j} = 0$.

As anticipated, firms are able to track consumers, where tracking involves monitoring consumers' browsing and collecting possibly sensitive information. Consequently, some consumers may perceive tracking as an invasion of their privacy. To capture this aspect, we assume consumers have a disutility θ from being tracked, which is distributed according to the c.d.f. $F(\cdot)$ on the interval $[0, \bar{\theta}]$, with a density function $f(\cdot)$. For simplicity, we assume that θ is distributed independently of consumers' preferences for content. To protect their privacy, consumers can take some actions that block tracking, e.g. deleting third-party cookies and installing browser plug-ins that anonymize one's browsing history. These actions come at a cost $c > 0$ for the individual. We denote by $\tilde{\theta}$ the consumer who is indifferent whether to allow tracking and by β the quantity of consumers that allows it, that is $\beta = F(\tilde{\theta})$.¹¹

¹⁰We assume that demands satisfy the necessary conditions for existence and uniqueness of an equilibrium with interior solutions. See Vives (2000) and Ambrus et al. (2016) for a discussion of these conditions.

¹¹We assume the decision to protect one's privacy is made independently of whether a consumer visits both, one

Total consumer surplus is given by $CS_c - CS_p$, where CS_c is the surplus from content and CS_p is the aggregate disutility from privacy losses, defined respectively as

$$CS_c = \int_{\delta q_1}^{\infty} \int_0^{\delta q_2} (u_1 - \delta q_1) h(u_1, u_2) du_2 du_1 + \int_0^{\delta q_1} \int_{\delta q_2}^{\infty} (u_2 - \delta q_2) h(u_1, u_2) du_2 du_1 \\ + \int_{\delta q_1}^{\infty} \int_{\delta q_2}^{\infty} (u_1 - \delta q_1 + u_2 - \delta q_2) h(u_1, u_2) du_2 du_1, \quad (2)$$

$$CS_p = \int_0^{\bar{\theta}} \theta f(\theta) d\theta + (1 - \beta) c. \quad (3)$$

Advertisers. There is a unit mass of advertisers. We assume advertising is informative and, following again Ambrus et al. (2016), we decompose the advertisers' expected revenue from informing consumers in several components, each related to a subset of viewers. These revenues depend on the probability that a consumer is informed about the advertiser's product on one or both publishers, as well as on the return from informing a consumer. We assume this return is the same for all advertisers and that advertisers capture the entire surplus generated when a consumer is informed (see e.g., Anderson and Coate, 2005, and Crampes et al. 2009).

Consider the case where an advertiser buys m_i impressions (per consumer) only from publisher i . We denote by $r_i(m_i)$ (resp. $\hat{r}_i(m_i)$) the expected revenue the advertiser obtains from a single-homing (resp. a multi-homing) consumer.¹² Thus, a single-homing advertiser's total revenue is

$$r_i(m_i) D_i + \hat{r}_i(m_i) D_{12}, \quad i = 1, 2. \quad (4)$$

To capture the fact that there are diminishing returns to advertising, we assume that r_i and \hat{r}_i are increasing and concave functions of m_i .¹³

or none of the publishers we model. This assumption captures the fact that many privacy-related choices (e.g., choosing a browser with different privacy features, changing cookies settings or installing anti-tracking software) are not made each time a consumer decides whether to visit a publisher, but rather on a "once-and-for-all" basis. Note that this decision is also relevant for "zero-homers" because they may visit other websites (we implicitly assume that $\bar{\theta}$ is not large enough to induce a consumer to completely stop using the Internet). Relaxing these assumptions would complicate the exposition without adding insights to our analysis. We consider some of the interactions between the decision whether to allow tracking and the utility from browsing content in Section 8.

¹²The difference between the advertising revenues generated by sending ads to single- and multi-homers on a publisher may stem, for instance, from the fact that multi-homers split their attention between several publishers, or from the different amount of time they spend browsing content.

¹³Diminishing returns may arise because ads are more likely to reach already informed consumers as the size of an advertising campaign increases. Furthermore, marginal impressions may fall on consumers that are less likely to be interested in a product than infra-marginal ones. See, e.g., Bagwell (2007) and Renault (2015).

Assume now the advertiser multi-homes, buying m_i impressions on each publisher. Its total revenue is

$$r_1(m_1)D_1 + r_2(m_2)D_2 + r_{12}^\sigma(m_1, m_2, \beta)D_{12}, \quad \sigma = \{nT, T\}. \quad (5)$$

The expected revenue from a single-homer is again $r_i(m_i)$, because she is exposed only to impressions on publisher i . The term $r_{12}^\sigma(m_1, m_2, \beta)$ is the expected revenue from a multi-homer, which is increasing and concave in m_i , $i = 1, 2$. We assume that the relation $\hat{r}_i \leq r_{12}^\sigma \leq \hat{r}_1 + \hat{r}_2$ holds. We assume that $\frac{\partial^2 r_{12}^\sigma}{\partial m_i \partial m_j} \leq 0$, for $i \neq j$ and $i, j = 1, 2$. This assumption captures the fact that ads on different publishers are imperfect substitutes, because a multi-homer may receive an impression that she has already registered while visiting another outlet (Ambrus et al., 2016, de Cornière and Taylor, 2014).

The revenue from a multi-homer $r_{12}^\sigma(m_1, m_2, \beta)$ is affected by cross-outlet tracking. Specifically, the superscript h denotes whether tracking applies ($\sigma = T$) or not ($\sigma = nT$). In the former case, r_{12}^T depends on the quantity of consumers that allow tracking, i.e. β . This formulation is consistent with the fact that the ad network can avoid wasteful impressions when it is able to follow consumers across publishers: upon identifying a consumer already informed on a publisher (for instance, by observing that she has already clicked on a given ad), the ad network can expose the consumer to a different message on the other publisher (e.g. an ad about a related product, possibly from the same firm). Hence, tracking improves the reach of an advertising campaign (for a given volume of impressions).¹⁴ Formally, tracking has the following properties. First, it increases the expected revenue from each multi-homer, i.e. $\frac{\partial r_{12}^T}{\partial \beta} > 0$. Second, it increases the value of a marginal impression on a multi-homer, i.e. $\frac{\partial^2 r_{12}^T}{\partial m_i \partial \beta} > 0$. Third, it reduces the loss in the value of a marginal impression on a multi-homer on one publisher as the number of impressions on the other publisher increases, i.e. $\frac{\partial^3 r_{12}^T}{\partial m_i \partial m_j \partial \beta} > 0$. That is, publishers become less substitutable for advertisers as tracking becomes more pervasive. Overall, these assumptions imply the following relations: $r_{12}^T > r_{12}^{nT}$, $\frac{\partial r_{12}^T}{\partial m_i} > \frac{\partial r_{12}^{nT}}{\partial m_i}$, and $\frac{\partial^2 r_{12}^{nT}}{\partial m_i \partial m_j} < \frac{\partial^2 r_{12}^T}{\partial m_i \partial m_j} \leq 0$ for any (m_1, m_2) .¹⁵

Note that our formulation captures the fact that the ability of the ad network to track consumers

¹⁴Tracking may also be useful for re-targeting purposes, i.e. proposing an ad containing a specific offer on a certain product to a consumer who has previously shown interest in the product. Another example where across outlet tracking is useful is the case of ad campaigns that send messages in a sequence to tell a brand story (see, e.g., https://scontent-arn2-1.xx.fbcdn.net/t39.2365-6/10333119_1457635661160496_168768318_n.pdf).

¹⁵We provide a microfoundation for the advertising revenue functions in the Online Appendix.

depends on whether they choose to protect their privacy. In reality, of course, this ability also depends on the effectiveness of the ad network’s technology. However, because we are primarily interested in the link between consumer privacy-related decisions and the advertising market, we ignore this aspect for the moment. We return to this point in Section 6.2.

Finally, observe that we assume revenues from single-homers are unaffected by cross-outlet tracking. In other words, the ad network can monitor the behavior of consumers within a given publisher as effectively as the publisher itself. This assumption is due to the fact that, in reality, publishers are quite effective at monitoring consumers on their own websites, but far inferior to the ad network at monitoring consumers on other outlets (Athey et al., 2016). Also, note that, to focus on the role of the ad network, we assume that internal tracking is not affected by consumers’ privacy decisions. This hypothesis is consistent with the observation that publishers typically use first-party cookies to monitor consumer behavior (on multiple visits) within their own websites. Blocking these cookies usually compromises the basic functionality of a website.¹⁶

Social welfare. We define welfare as the sum of consumer surplus, advertising surplus and publishers’ and ad network’s profits. Because payments collected by the latter are transfers from other players, they cancel out in the welfare formula, which boils down to

$$W = CS_c - CS_p + AS, \tag{6}$$

where the last term is total advertising surplus:

$$AS \equiv \sum_i r_i D_i + r_{12}^\sigma D_{12}. \tag{7}$$

Timing. The game proceeds as follows. In stage 1, *AN* offers a transfer to each publisher in exchange for its advertising inventory. Publishers accept or refuse. At stage 2 publishers choose the advertising capacity q_i . Consumers choose whether to block tracking at the same stage. At stage

¹⁶Cookies are pieces of code placed on a device or browser while a consumer visits a website. Typically, cookies store personalized user settings for that website, along with other information, and can be used to identify her during later visits. First-party cookies are issued by the website the consumer is visiting, whereas third-party cookies are generally issued by an ad network with the objective of monitoring the consumer across several websites. We assume that publishers share with the ad network the information on consumers (stored in first-party cookies) that block third-party tracking.

3, consumers observe the advertising capacity (q_1, q_2) on each website and decide which to join, if any. At stage 4, either the ad network or the publishers sell impressions to advertisers (depending on whether ad inventories have been outsourced). The advertisers decide which contract to accept, if any. Consumers then get exposed to ads and all payoffs are realized.

4 Solving the model

In this section, we illustrate the equilibrium of the model in the case where publishers outsource their inventory to the ad network. Note that, because consumer demands (stage 3) are fully described by (1), in the following we focus on the solution of the other stages of the game. We adopt Subgame-Perfect Nash Equilibrium as the solution concept and solve the model by backward induction.

4.1 Stage 4

Advertisers buy impressions from either the ad network or the publishers, depending on the outcome of the previous stages. Assume that none of the publishers outsourced to the ad network. Each publisher then offers a menu of contracts to each advertiser, specifying an advertising quantity m_i in exchange for a payment p_i . The choice of m_i is characterized as follows. First, because advertising revenues are increasing in m_i (see (4) and (5)), the publisher fills all the available advertising capacity.¹⁷ Furthermore, due to diminishing returns, the publisher splits the available capacity equally across all advertisers. Thus, in equilibrium each publisher offers the same contract to all advertisers, such that $m_i = q_i$.

To see how the payment p_i is determined, we need to evaluate the advertiser's payoff when rejecting the offer. If the advertiser places ads only on publisher j , its revenue is given by (4) net of the payment p_j . By contrast, if it also accepts publisher i 's offer, it gets (5) net of p_i and p_j (note that $\sigma = nT$ in this scenario, because the ad-network is not present on either publisher's website). Hence, each publisher sets

$$p_i = r_i(q_i) D_i(q_i, q_j) + (r_{12}^{nT}(q_1, q_2) - \hat{r}_j(q_j)) D_{12}(q_1, q_2), \quad i, j = 1, 2 \text{ and } i \neq j. \quad (8)$$

¹⁷In the spirit of Kreps and Scheinkman (1983), we assume the publisher gets a large negative payoff when exceeding the available capacity.

Because all advertisers accept this contract, each publisher obtains a profit $\pi_i = p_i$. In words, this payment equals the advertisers' incremental revenue from placing q_i impressions per consumer on publisher i . Publishers are therefore unable to extract the whole surplus generated by advertising on their platform because there are non-exclusive consumers. Furthermore, the revenue generated from multi-homers depends on how many impressions are placed on the other publisher.¹⁸

The outcome of the subgame that takes place when only publisher i outsources to AN is identical to that described above. The only difference is that the ad network sells the available impressions q_i to advertisers. Note that because AN is not present on both publishers, it cannot track consumers across outlets.

Consider now the subgame where both publishers outsourced to the ad network. The latter offers a menu of contracts to each advertiser, specifying a pair (m_1, m_2) and a total payment p_{AN} . Following the same logic as above, the ad network fills the available capacity on each publisher and divides it equally across all advertisers. Hence, in equilibrium it offers a single contract, specifying $m_i = q_i$, $i = 1, 2$. However, AN is able to extract the entire advertiser surplus, because it is the unique gateway to the audience of both publishers.¹⁹ Thus, we have

$$p_{AN} = AS = r_1(q_1) D_1(q_1, q_2) + r_2(q_2) D_2(q_1, q_2) + r_{12}^T(q_1, q_2, \beta) D_{12}(q_1, q_2). \quad (9)$$

Observe that, because it is present on both outlets, AN can track consumers across them. Given $r_{12}^T \geq r_{12}^{nT}$, tracking raises the advertisers' willingness-to-pay by increasing the revenue they obtain from multi-homers. As anticipated, the size of this effect depends on the ad network's tracking capability, captured by β . Note also that, given that it captures all the advertising surplus, the ad network internalizes the effects that impressions on a publisher produce on advertisers' willingness-to-pay for impressions on the other. We summarize these findings as follows

¹⁸We provide further discussion of expression (8) in Section 6, where we compare the equilibrium where publishers compete directly to that where they outsource to the ad network.

¹⁹The formulation of contracts between publishers, or the ad network, and advertisers follows Ambrus et al. (2016). We adopt it to avoid inessential complications, although other formulations are possible that do not change the analysis substantially. For instance, suppose each impression is sold separately via a second-price auction. Specifically, the auction's object is an opportunity to impress a user on a publisher, given information on the number of impressions the user will be exposed to on that publisher. In the case the ad network sells the impression, the auction also specifies the quantity of impressions on the other publisher. In equilibrium, each advertiser would bid the surplus gained with the impression (conditional on the quantity of impressions acquired). As a result, the equilibrium allocation of impressions would be the same as with our current formulation, i.e. such that $m_i = q_i$, $i = 1, 2$. Furthermore, the revenue publishers collect from advertisers equals (8). Similarly, the revenue of the ad network equals (9).

Lemma 1.

- *When no publishers outsource their inventory to the ad network, publishers offer a contract to advertisers specifying an advertising quantity $m_i = q_i$ and a payment p_i described in (8). All advertisers accept. The same outcome takes place when only one publisher outsources.*
- *When both publishers outsource, the ad network offers a contract to advertisers specifying $m_i = q_i$, $i = 1, 2$ and a payment p_{AN} as described in (9). All advertisers accept.*

One remark is in order before proceeding. First, the comparison of (8) and (9) suggests that outsourcing is Pareto-efficient from the perspective of the ad network and the publishers, for two reasons. First, the ad network centralizes the sale of ad impressions on different publishers, implying that the ad network is able to extract a larger share of the advertising surplus. Second, it is able to track consumers across outlets, reducing the waste of impressions.

4.2 Stages 2 and 1

In this section, we describe the contracting stage between the ad network and the publishers, and how the latter subsequently choose the respective advertising levels q_i . Next, we describe the consumers' decision whether to block tracking and, finally, we characterize the SPNE of the game.

At stage 1, AN simultaneously makes a take-it-or-leave-it offer to each publisher. In exchange for the publisher's ad inventory (and the ensuing revenue), the contract specifies a schedule of transfers $x_i(q_i)$ contingent on the advertising capacity q_i that the publisher chooses in Stage 2. Furthermore, the contract is contingent on the outcome of the negotiation with the other publisher. Specifically, the contract with publisher i is void if the other publisher rejects the ad network's offer, if any. This feature captures in the simplest possible way the fact that the ad network can make an advantageous offer to each publisher, while still making a positive profit, only if it manages the publishers' *joint* inventories. Indeed, as we have shown in Section 4.1, when a single publisher outsources, the ad network cannot make more profit from advertising than the publisher itself.²⁰

²⁰The assumption of contracts being contingent on the outcome of the negotiation with other parties is rather common in the literature on multi-agent relationships (Riley, 1979, and Bernheim and Whinston, 1985). By making the contract contingent, we take into account that the ad network and the publisher may want to renegotiate their contract when the other publisher does not outsource. Note also that our specification is consistent with evidence suggesting that ad networks try to influence the provision of advertising on the websites they serve (see,

Each publisher accepts the ad network's offer if the transfer is at least as large as its outside option, which we denote by π_i^o . Given the contract we just specified, π_i^o equals the revenue publishers get when competing directly, without the ad network, as described in Lemma 1. Letting q_i^C denote the advertising capacities in the equilibrium of this subgame, we have²¹

$$\pi_i^o = r_i(q_i^C) D_i(q_i^C, q_j^C) + (r_{12}^{nT}(q_i^C, q_j^C) - \hat{r}_j(q_j^C)) D_{12}(q_i^C, q_j^C). \quad (10)$$

It is fairly easy to see that both publishers outsource in equilibrium. The intuition is that, as we argued at the end of the previous section, outsourcing is Pareto-efficient from the perspective of the publishers. Thus, the ad network can provide each of them with a transfer at least as large as π_i^o , while making a positive profit.

Conditional on outsourcing to AN , in Stage 2 each publisher chooses q_i maximizing $x_i(q_i)$. It follows that the ad network can implement the q_i it desires by appropriately designing these schedules. Let the capacities that maximize the ad network's profit when managing the inventories of both publishers be q_i^* (to be determined below). To ensure participation by the publisher, $x_i(q_i^*) \geq \pi_i^o$ must hold. Given that $\pi_{AN} = p_{AN} - x_1(q_1^*) - x_2(q_2^*)$ when both publishers outsource, the constraint is binding in equilibrium. Hence, we have $x_i(q_i^*) = \pi_i^o$. Summarizing, AN 's profit is

$$\pi_{AN} = p_{AN} - \pi_1^o - \pi_2^o. \quad (11)$$

It remains to determine the equilibrium advertising levels (q_1^*, q_2^*) . These quantities maximize (11). Note from (10) that π_i^o is independent of the advertising capacities chosen on the equilibrium path. Therefore, the quantities (q_1^*, q_2^*) maximize p_{AN} . We have established that this revenue is equal to the entire surplus generated by advertising, AS (see (9)). Therefore, using the property that $\frac{\partial D_j}{\partial q_i} = -\frac{\partial D_{12}}{\partial q_i}$, the equilibrium capacities (q_1^*, q_2^*) satisfy the following first-order conditions:

$$\frac{\partial p_{AN}}{\partial q_i} = \frac{\partial AS}{\partial q_i} = \left[D_i \frac{\partial r_i}{\partial q_i} + D_{12} \frac{\partial r_{12}^T}{\partial q_i} \right] + \left[\frac{\partial D_i}{\partial q_i} r_i + \frac{\partial D_{12}}{\partial q_i} (r_{12}^T - r_j) \right] = 0, \quad i = 1, 2. \quad (12)$$

e.g., https://support.google.com/adsense/answer/17958?ref_topic=3001714&rd=1). However, our results are robust to alternative formulations. For instance, we could assume that the transfer to the publishers is given by $x_i(q_i)q_i(D_i + D_{12})$, where $x_i(q_i)$ a *per-impression* price contingent on q_i . Alternatively, we could assume the ad network and the publishers bargain directly over q_i . Both formulations were considered in a previous version of this paper, leading to identical results (details available from the authors).

²¹We characterize q_i^C in Section 7, where we study the case where publishers cannot outsource to the ad network.

The equilibrium quantity q_i^* is the result of the following trade-off: increasing the provision of advertising on one publisher generates more revenue from the infra-marginal consumers (first term in square brackets), but reduces the quantity of consumers that visit that publisher (second term in brackets). This expression also indicates that AN internalizes the impact of a change in the advertising quantity of one outlet on the revenue generated by the other. In particular, the net effect due to the loss of multi-homers is $r_{12}^T - r_j$, because those who stop browsing i 's website become single-homers on publisher j .

Finally, we consider the decision of consumers whether to block tracking. A consumer chooses to block if the disutility from being tracked is bigger than the cost of blocking, that is, if $\theta > c$. Therefore, denoting by β^* the share of consumers who at equilibrium allow tracking, we have

$$\beta^* = F(c). \quad (13)$$

Observe that this decision has an impact on the advertising market: by blocking tracking, a consumer prevents the ad network from following her across publishers. Therefore, the tracking effectiveness goes down, and so does AN 's revenue from multi-homers. We summarize our findings in the following

Lemma 2. *The unique SPNE of the game is such that:*

- *both publishers outsource their inventories to the ad network. The advertising quantities (q_1^*, q_2^*) maximize p_{AN} defined in (9), which coincides with total advertising surplus, AS ;*
- *the share of consumers that allow tracking is β^* , as defined in (13).*

4.3 The effect of tracking on the provision of advertising

We now investigate the effect of the ad network's tracking capability on the advertising market. Specifically, we investigate how a change in β affects the advertising quantities, q_i . Determining this effect is generally rather complex, because it depends on how changes in the quantity of impressions on one publisher affect the marginal revenues generated on the other. Nevertheless, assuming publishers are symmetric, it is possible to establish the following result:

Proposition 1. *Assume that publishers are symmetric. When they outsource to the ad network, the provision of advertising on each publisher increases with the extent of tracking if and only if*

$$D_{12} \frac{\partial^2 r_{12}^T}{\partial q_i \partial \beta} + \frac{\partial D_{12}}{\partial q_i} \frac{\partial r_{12}^T}{\partial \beta} > 0. \quad (14)$$

Interestingly, the equilibrium advertising quantities q_i^* do not necessarily increase with the ad network's tracking ability. To understand this result, consider that tracking has two main effects. On the one hand, it raises the revenue from infra-marginal multi-homers, because it reduces the likelihood of repeated impressions ($\frac{\partial^2 r_{12}^T}{\partial q_i \partial \beta} > 0$). This effect implies, all else given, a stronger incentive to expand the provision of ads. On the other hand, r_{12}^T increases with tracking ($\frac{\partial r_{12}^T}{\partial \beta} > 0$), because less repetition raises the value of impressions on multi-homers. Hence, given $\frac{\partial D_{12}}{\partial q_i} < 0$, the cost of losing the marginal multi-homers becomes more relevant with tracking. Intuitively, when the second effect dominates, the ad network prefers a lower advertising quantity as its tracking ability increases. Inspection of (14) suggests that this outcome is more likely when consumer preferences for the publishers are negatively correlated. The reason is that, because consumers tend to visit either one publisher or the other (but few visit both) one can expect D_{12} to be relatively small compared to the absolute value of $\frac{\partial D_{12}}{\partial q_i}$. Multi-homers are a relatively small share of publishers' audiences but a relatively large component of the change in demand triggered by an increase in the provision of ads. On the other hand, when preferences for publishers are positively correlated, more consumers tend to multi-home. In this scenario, (14) suggests that tracking should bring to higher advertising quantities. Moreover, when the disutility from advertising δ is high, we can again expect D_{12} to be relatively small compared to the absolute value of $\frac{\partial D_{12}}{\partial q_i}$, meaning that when consumers strongly dislike ads the ad network may decrease the level of ads when the tracking effectiveness increases.

Proposition 1 also suggests that, all else given, the effect of tracking on the provision of ads depends on its effect on the value of infra-marginal impressions on multi-homers $\frac{\partial r_{12}^T}{\partial \beta}$ relatively to the value of marginal impressions $\frac{\partial^2 r_{12}^T}{\partial q_i \partial \beta}$. Firms advertising mainstream products typically want to maximize the reach of their campaigns. In this case, we can expect that advertising revenue from marginal multi-homers increases more in the tracking capability of the ad network than the revenue from infra-marginal ones. The opposite may hold for advertisers that want to run ad campaigns on

a specific group of consumers. Tracking may bring added value to these advertisers because of the ability to control how the brand story is presented or to show to consumers different products they may be interested in.

To our knowledge, Proposition 1 identifies a novel effect of data-driven technologies on the online advertising market. This effect stems from the two-sided nature of digital publishers. Specifically, the impact of across-outlet tracking depends on its effect on advertising returns on multi-homing consumers and on how consumer demand for digital publishers changes with advertising quantities. Previous papers have studied the effects of related technologies on the provision of ads. Athey and Gans (2010) find that it is optimal for publishers to cut their supply of ad spaces when targeting is introduced, because the publisher is more effective at reaching the intended recipient of the ad. Johnson (2013) finds that the level of advertising increases with the precision of targeting if the value of the marginal ad is higher than the value of an ad on a random consumer. However, these results are distinctive of targeting technologies. Our model focuses instead on the impact of tracking on the effectiveness of ad campaigns that involve multiple publishers. Furthermore, previous papers do not account for competition for consumers, because they either model consumer demands as exogenous (Athey and Gans, 2010) or ignore publishers altogether (Johnson, 2013). Proposition 1 highlights the importance of these factors in determining the effect of data-driven technology on the provision of online advertising.

5 Advertising and tracking: first- and second-best

The objective of this section is to analyze how the equilibrium level of tracking and provision of advertising deviate from the optimal ones. Intuition suggests that there are several externalities related to tracking and advertising choices. First, consumers do not consider the effect of their choice on the overall effectiveness of tracking, β . This oversight implies a direct externality on the advertising industry, because tracking affects the value that advertisers get by reaching multi-homers. Furthermore, there is an indirect externality on consumers because of the effect on advertising quantities, which we identified in Section 4.3. In addition, publishers (and the ad network) also exert an externality when deciding on the provision of ads because they ignore the reduction in utility that consumers sustain. The presence of these externalities suggests that, generally, the market

equilibrium does not coincide with the welfare optimum. Hence, some regulation may be desirable. To streamline the exposition, however, we postpone the discussion of the policy implications of our findings to Section 6.1.

5.1 First-best

As a first step, we consider the first-best allocation, such that a welfare-maximizing regulator sets the quantity of consumers allowing tracking, as well as the advertising quantities. This is a rather unrealistic setting, because regulators hardly intervene so extensively on the online advertising market. Nevertheless, this scenario serves as a useful benchmark for the analysis that follows.

Using (6), we find that the first-best allocation, denoted $(\beta^{FB}, q_1^{FB}, q_2^{FB})$, is such that the following conditions are satisfied:

$$\begin{aligned}\frac{\partial W}{\partial \beta} &= -\frac{\partial CS_p}{\partial \beta} + \frac{\partial AS}{\partial \beta} = 0 \\ \frac{\partial W}{\partial q_i} &= \frac{\partial CS_c}{\partial q_i} + \frac{\partial AS}{\partial q_i} = 0, \quad i = 1, 2.\end{aligned}\tag{15}$$

The first expression suggests that the quantity of consumers that allows tracking has two effects on welfare. The first is on consumers' disutility from privacy losses: $\frac{\partial CS_p}{\partial \beta} = \tilde{\theta} - c$.²² Because $\tilde{\theta} = c$, this difference is zero when evaluated in equilibrium. The second effect is on advertising surplus. This effect is strictly positive, because, for any pair (q_1, q_2) , tracking raises the effectiveness of ads on multi-homers. Formally, starting from (7), we have $\frac{\partial r_{12}^T}{\partial \beta} D_{12} > 0$. This effect is not internalized by consumers, because they do not care for firms' profits. Hence, compared to the first-best, the equilibrium has too little tracking: $\beta^{FB} > \beta^*$.

The comparison between advertising quantities (q_1^{FB}, q_2^{FB}) and (q_1^*, q_2^*) is not straightforward. First, when deciding the advertising quantities, firms do not consider the direct effect on consumer surplus, i.e. $\frac{\partial CS_c}{\partial q_i} < 0$. However, the regulator does take it into account, which reduces the first-best advertising quantities compared to the equilibrium ones. Nevertheless, because there is too little tracking in equilibrium, the quantity of advertising might be smaller than the first-best one. The reason is that, as Proposition 1 shows, the equilibrium quantity of ads increases with tracking under certain conditions. Hence, the two effects may go in different directions, making the comparison

²²The derivative $\frac{\partial CS_p}{\partial \beta}$ is obtained by defining $\tilde{\theta}(\beta) \equiv F^{-1}(\beta)$, which implies that $\tilde{\theta}'(\beta) = \frac{1}{f(\tilde{\theta})}$. Using (3), we obtain $\frac{\partial CS_p}{\partial \beta} = \tilde{\theta} - c$.

ambiguous. We summarize in the following:

Proposition 2. *Compared to the first-best, there is too little tracking in equilibrium. The comparison with the equilibrium advertising quantity is ambiguous.*

5.2 Second-best level of tracking

We now consider a second-best scenario where the regulator maximizes welfare with respect to β , and the publishers are free to set the quantity of advertising. This is a more realistic setup than the previous one, because typically regulators do not impose limits on the provision of advertising in digital outlets, but they do intervene on privacy policy.

To find the second-best allocation, we maximize (6) with respect to β , denoting the solution as β^{SB} . The advertising quantities (q_1^*, q_2^*) are determined in equilibrium, as shown in Lemma 2. Because the latter satisfy $\frac{\partial AS}{\partial q_i} = 0$ for $i = 1, 2$ (see (12)), β^{SB} is such that:

$$\frac{\partial W}{\partial \beta} = -\frac{\partial CS_p}{\partial \beta} + \frac{\partial AS}{\partial \beta} + \sum_{i=1,2} \frac{\partial CS_c}{\partial q_i} \frac{\partial q_i^*}{\partial \beta} = 0. \quad (16)$$

This expression indicates that there are three effects of tracking on welfare. The first is the direct effect on consumers' net losses from privacy, CS_p . As discussed above, this effect is zero at the margin when evaluated in equilibrium, so $\frac{\partial CS_p}{\partial \beta} = 0$. The last two terms capture the effects of tracking that work through the advertising market. First, we have an impact on advertiser surplus, AS . As discussed above, this effect is positive ($\frac{\partial AS}{\partial \beta} > 0$) because r_{12}^T increases in β . Second, there is an indirect effect on consumers, captured by the last term in (16). As the extent of tracking changes, so does the equilibrium quantity of ads and, thus, the surplus that consumers get from browsing content, CS_c . Consumers ignore both of these effects when deciding whether to allow tracking. Broadly speaking, while the first externality seems to be widely recognized by policymakers, the second one has apparently received less attention in the policy debate.

Observe that, whereas the externality on advertisers is positive, the externality on consumers can go either way. The sign depends on the effect of tracking on advertising quantities, $\frac{\partial q_i^*}{\partial \beta}$. As we argued in Proposition 1, the sign of this derivative hinges on the impact of tracking on the advertising revenue from marginal and inframarginal multi-homers. Because $\frac{\partial CS_c}{\partial q_i} < 0$, if $\frac{\partial q_i^*}{\partial \beta} < 0$

consumers experience a lower disutility from advertising when browsing the publishers. By contrast, if $\frac{\partial q_i^*}{\partial \beta} > 0$, consumers are exposed to more ads and thus get lower utility from content when β increases. Summing up, the comparison among the second-best level of tracking and the equilibrium one depends on how tracking changes the quantity of ads. If $\frac{\partial q_i^*}{\partial \beta} < 0$, the externality on consumers is positive, like that on advertisers. Hence, there is too little tracking in equilibrium: $\beta^{SB} > \beta^*$. Otherwise, if $\frac{\partial q_i^*}{\partial \beta} > 0$, the comparison is ambiguous, because the externality on consumers is negative.

Proposition 3. *Compared to the second-best optimum, there is too little tracking in equilibrium if $\frac{\partial q_i^*}{\partial \beta} < 0$, $i = 1, 2$; otherwise, the comparison is ambiguous.*

Observe that, if the effect of tracking on advertising quantities is negative, there may be too little tracking in equilibrium even if one disregards the profits of the advertising industry. That is, consumers themselves may be better off with more tracking.

We note that Proposition 3 differs from the results in Johnson (2013), who studies ad-blocking. He finds that, from a second-best point of view, there is too little blocking of ads when firms' profits from advertising are relatively small (see his Proposition 4). The reason is that blocking produces a negative externality on advertisers, but a positive one on consumers (because it unambiguously leads to lower provision of advertising). Hence, when firm's profits from advertising are small, the second externality is more relevant. In our model, consumers who block tracking produce a negative externality on advertisers as well. However, the externality on consumers may be positive or negative, depending on the effect on advertising quantities.

5.3 Second-best level of advertising

To complete the analysis of the second-best, we now assume the regulator controls the level of advertising, and consumers are free to choose whether to allow tracking. Formally, the regulator maximizes welfare with respect to (q_1, q_2) , anticipating that the share of consumers allowing tracking is β^* (see (13)). The second best advertising levels solve

$$\frac{\partial W}{\partial q_i} = \frac{\partial CS_c}{\partial q_i} + \frac{\partial AS}{\partial q_i} = 0, \quad i = 1, 2. \quad (17)$$

Evaluating this expression at the equilibrium quantities (q_1^*, q_2^*) , we find that, because $\frac{\partial AS}{\partial q_i} = 0$ and $\frac{\partial CS_c}{\partial q_i} < 0$ for $i = 1, 2$, the equilibrium level of advertising exceeds the second-best one. The intuition is straightforward: publishers do not internalize the effect of ads on inframarginal consumers.

Proposition 4. *Compared to the second-best optimum, there is too much advertising in equilibrium.*

6 Policy implications

This section discusses the policy implications of our previous findings. We focus on two main themes. First, we consider how the collection of consumers data for advertising purposes should be regulated, for instance in terms of whether regulators should encourage or discourage consumers to protect their own personal information. Secondly, we focus on the applicability of market-oriented solutions to achieve an efficient level of exploitation of consumer data.

6.1 Privacy regulation and consumer consent

The results in Section 5.2 have direct implications for the regulation of online privacy. The reason is that, as anticipated, online publishers typically do not face any regulation on advertising quantities (e.g., advertising caps), but a vibrant policy debate exists concerning how to regulate tracking and empowering consumers to protect their personal information. For example, several regulators have evaluated the possibility of reducing the cost of avoiding tracking (Tucker, 2014), e.g. by mandating the inclusion of tools that facilitate the deletion of third-party cookies in Internet browsers.²³ Our analysis suggests that these measures are justified only in certain circumstances. We have shown that, when the level of advertising decreases with the extent of tracking, both advertiser surplus and consumer surplus from content increase with tracking quantity (see Proposition 3). In this case, the regulator should actually encourage tracking, even if it only cares for consumer surplus.

To formalize the discussion, we endogenize the cost of blocking tracking c . Assume that a regulator chooses c to maximize total welfare. Therefore, the optimal c is such that (we remind the

²³The Federal Trade Commission (FTC) recently considered a Do Not Track mechanism that would persistently record on a consumer's device her choice not to be tracked (see <https://www.ftc.gov/news-events/media-resources/protecting-consumer-privacy/do-not-track>). Also, several European Data Protection Authorities encourage the development of “do-not-track” tools allowing consumers to easily manage cookies (see http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2016/wp240_en.pdf).

reader that $\frac{\partial CS_p}{\partial \beta} = 0$ at equilibrium)

$$\frac{\partial W}{\partial c} = \frac{\partial \beta^*}{\partial c} \left(\frac{\partial AS}{\partial \beta} + \sum_{i=1,2} \frac{\partial CS_c}{\partial q_i} \frac{\partial q_i^*}{\partial \beta} \right) + \frac{\partial CS_p}{\partial c} = 0. \quad (18)$$

Intuitively, the share of tracked consumers increases with the cost of blocking c , i.e. $\frac{\partial \beta^*}{\partial c} > 0$. The last term in expression (18) is negative, capturing the increased cost for those who block. A sufficient condition for the first term to be positive is that $\frac{\partial q_i^*}{\partial \beta} < 0$. Hence, when advertising quantities decrease with β , the regulator wants to set a higher c to encourage tracking. Again, this conclusion applies also if the regulator only cares for consumer surplus.²⁴

Our results also address the implications of opt-in and opt-out options for collecting consumer consent regarding the use of personal data. Previous studies have shown that consumers have a tendency to stick to default settings, even when these settings are suboptimal (Acquisti et al., 2016). Hence, opt-out policies have the effect of increasing the perceived cost for consumers to block tracking. Thus, a legitimate concern is that ad networks may take advantage of this tendency by placing third-party cookies on consumers' browsers, without asking for their consent. Our results suggest that, when promoting tracking is desirable, opt-out policies for collecting consent may actually increase welfare. On the other hand, when tracking is excessive, regulators should promote the adoption of no-tracking as default settings and support the implementation of opt-in policies.²⁵

6.2 Control rights on personal data and market-based solutions

As an alternative to direct regulatory intervention, some scholars have proposed market-based approaches relying on the “propertization” of personal information, starting from a clear assignment of property rights, to tackle the issues associated with privacy protection (Samuleson, 2000, Schwartz, 2004). However, economic theory suggests that there are several issues impeding the formation of a well-functioning market for personal data (Acquisti et al., 2016, Hermalin and Katz, 2006). These

²⁴To encourage tracking in a way that minimizes consumers' privacy losses, a regulator may focus on decreasing the disutility that users perceive when tracked, e.g. by discouraging the adoption of the most intrusive tracking technologies, or fostering the provision of more transparent information to consumers. A related intervention could be to limit the timespan during which firms retain consumer data and requiring data to be anonymized and/or aggregated. Chiou and Tucker (2014) provide evidence that changes in retention policy did not change the quality of search engine results. Finally, the regulator may introduce a tax (or subsidy) on those who block (or allow) tracking.

²⁵Similar policy implications can be drawn regarding hardware and software providers that are integrated in the advertising industry (e.g., Google, Microsoft and Apple). Exploiting consumers' inertia, these providers may design their products with tracking set as the default option.

include the fact that information is dispersed among multiple players (including publishers and ad networks), which increases transaction costs. Furthermore, consumers themselves are imperfectly informed, preventing them from determining the value of their own data. In addition, there are spillovers between the market for personal data and the advertising market.²⁶ An extensive analysis capturing all these issues is outside the scope of this paper. However, to provide some insights on the desirability of markets for personal information, in this section we use the model to evaluate the implications of letting consumers and the ad network trade the right (not) to be tracked in exchange for a payment. Furthermore, we study the consequences of allocating the right to control tracking to either side of the market.

To set the stage, we assume the ad network has full control on which consumers are tracked, the latter being unable to opt-out. This rather extreme scenario is a useful benchmark, because it allows to clearly characterize the ad network's incentives concerning the adoption of technologies that improve tracking, given the effort that consumers make to block it.²⁷ Formally, we consider the case where publishers outsource their ad inventories, but assume that, at stage 2, the ad network sets the extent of tracking, β . Building on the analysis in Section 4, the ad network maximizes p_{AN} with respect to q_i , $i = 1, 2$, and β . Because $\frac{\partial p_{AN}}{\partial \beta} = \frac{\partial AS}{\partial \beta} = D_{12} \frac{\partial r_{12}^T}{\partial \beta} > 0$, the ad network tracks as many consumers as possible. Hence, tracking is maximal in equilibrium: $\beta = 1$. Since the ad network ignores the privacy losses of consumers, giving it full control over consumers' data leads to an excessive amount of tracking, compared to both the first and second best outcomes.

Assume now the regulator still assigns the right to control tracking to the ad network, but requires it to provide an opt-out option to consumers in exchange for a payment y . In this scenario, a consumer allows tracking if and only if $\theta - c \leq y$. At stage 1, the ad network sets y to maximize $p_{AN} + (1 - \beta)y$, because it collects this payment from the share $(1 - \beta)$ of consumers who block. Therefore, in equilibrium we have (recall that $p_{AN} = AS$)

$$y = \frac{1 - \beta}{\frac{\partial \beta}{\partial y}} + \frac{\partial AS}{\partial \beta}, \quad (19)$$

²⁶Given these hurdles, implementing a solution based on the well-known theorem by Coase (1960) may be difficult. There may of course be also ethical and moral reasons for avoiding the creation of a unified market for data. For instance, the EU General Data Protection Regulation (2016/679) deliberately ignores transferable ownership rights for personal data, on the basis that privacy is a human right that cannot be alienated.

²⁷Some of the tracking technologies available to websites and ad networks are quite hard to avoid for consumers (e.g. device fingerprinting). See Hoofnagle et al. (2012) for a review.

where all quantities are evaluated at the equilibrium ad quantities (q_1^*, q_2^*) . The first term on the right hand side of this expression is a monopoly mark-up (where $\frac{\partial \beta}{\partial y} > 0$). The second term represents the externality on advertisers that a consumer generates when blocking. The ad network internalizes it because it captures the whole advertising surplus. However, expression (19) suggests that forcing the ad network to set a price for opting out of tracking is not sufficient to implement the second best optimum. Indeed, when we evaluate the derivative $\frac{\partial W}{\partial \beta}$ in (16) at equilibrium, we find²⁸

$$\frac{\partial W}{\partial \beta} = -\frac{1-\beta}{\frac{\partial \beta}{\partial y}} + \sum_{i=1,2} \frac{\partial CS_c}{\partial q_i} \frac{\partial q_i^*}{\partial \beta}. \quad (20)$$

The second best is generally not implemented, for two reasons. First, the ad network has market power, so it incorporates a mark up in the opt-out price. Hence, it induces more consumers to accept tracking than it would if it simply had them internalize the externality on advertising surplus. Second, the ad network ignores the externality consumers impose on each other due to the change in advertising quantities (see Section 5.2). In fact, welfare might even decrease with respect to the equilibrium with no opt-out price: although advertising surplus is strictly higher (because more consumers allow tracking), consumers suffer higher privacy losses.²⁹

Next, we consider the symmetric scenario, where the regulator assigns full control rights to consumers, but the ad network can pay an opt-in fee k to those who allow tracking. To start, we assume the ad network does not know consumers' taste for privacy. Now, a consumer allows tracking if and only if $\theta - c \leq k$. The ad network sets k to maximize $AS - \beta k$. Therefore, noting that $\frac{\partial \beta}{\partial k} > 0$ and assuming that $k \geq 0$, we have

$$k = \max \left[\frac{\partial AS}{\partial \beta} - \frac{\beta}{\frac{\partial \beta}{\partial k}}, 0 \right], \quad (21)$$

where ad quantities are evaluated at the equilibrium (q_1^*, q_2^*) . Expression (21) shows that the price the ad network is willing to pay to track a consumer is smaller than the external effect on advertising surplus. The reason is that, while the ad network benefits from more tracking, it sustains a monetary cost k per each consumer tracked. Intuitively, the condition in (21) indicates that, if the effect of

²⁸The last equality follows from the fact that payments from consumers to the ad network cancel out in welfare and that $\frac{\partial CS_p}{\partial \beta} = \theta - c + y = 0$ in equilibrium.

²⁹Consumer net losses from privacy, $-CS_p - (1-\beta)y$, necessarily increase: some consumers sustain higher disutility, while those who opt out have to pay an extra price to the ad network.

tracking on advertising surplus is not large enough, the ad network is better off not paying any fee. Hence, the market for the right to track consumers does not deliver the (second-best) optimum. In this scenario, when we evaluate $\frac{\partial W}{\partial \beta}$ at equilibrium, we find that (assuming $k > 0$)

$$\frac{\partial W}{\partial \beta} = \frac{\beta}{\frac{\partial \beta}{\partial k}} + \sum_{i=1,2} \frac{\partial CS_c}{\partial q_i} \frac{\partial q_i^*}{\partial \beta} = 0. \quad (22)$$

Because the first term on the right hand side is positive, too few consumers allow tracking. The intuition is that the ad network cares for its monetary expenditure, on top of the externality on advertising surplus. However, as long as the opt-in fee is positive, advertising surplus is higher than in the status-quo (because more consumers allow tracking). In addition, the net consumer loss from privacy, $-CS_p + \beta k$, is smaller, because those who allow tracking are more than compensated by the fee. Hence, consumers are strictly better off as long as the external effect on advertising quantities is either negative, i.e. $\frac{\partial q_i}{\partial \beta} < 0$, or positive but relatively small.

As a final remark, note that we assumed that the ad network cannot observe consumers' disutility from privacy losses, θ . Although this information may be inferred from user behavior, introducing this aspect in the model would not fundamentally change the results. To see this, consider an hypothetical scenario where θ is observable and the ad network can set personalized prices to those who allow tracking. Although the ad network would induce more consumers to allow tracking than in the scenario considered above (because it can pay less to the inframarginal ones), the optimum would still not be reached because the ad network ignores the externality on consumers.

7 The ad network's impact on the advertising market

To further develop our understanding of the influence of ad networks on the online advertising market, now we consider the case where AN is inactive. Hence, publishers compete directly not only for consumers, but also for advertisers. We then compare the equilibrium in this scenario with that characterized in Section 4, whereby the ad network manages both publishers' inventories.

7.1 The equilibrium when publishers compete directly

As a first step, we characterize the SPNE when publishers compete directly. The sequence of events is the same as in the baseline model, except that Stage 1 is void. Note also that, because there is no cross-outlet tracking by assumption (given the absence of the ad network), the consumers' blocking decision at Stage 2 is irrelevant.

Recall that, as we found in Section 4.1, publishers offer a single contract to advertisers, specifying the payment $p_i^C = r_i(q_i) D_i(q_i, q_j) + (r_{12}^{nT}(q_1, q_2) - \hat{r}_j(q_j)) D_{12}(q_1, q_2)$ (described in (8)), in exchange for the advertising quantity $m_i = q_i$. In equilibrium, all advertisers accept and, therefore, each publisher earns $\pi_i^C = p_i^C$. Observe from (8) that multi-homers create an interdependence between the advertising revenues of the two publishers: the higher the revenue from reaching a multi-homer only on publisher j , the less advertisers are willing to pay for impressions on i .

Consider now the choice of advertising capacities q_i , which maximize π_i^C . Using (8), the system of first-order conditions is

$$\frac{\partial \pi_i^C}{\partial q_i} = \left[\frac{\partial r_i}{\partial q_i} D_i + \frac{\partial r_{12}^{nT}}{\partial q_i} D_{12} \right] + \left[r_i \frac{\partial D_i}{\partial q_i} + (r_{12}^{nT} - \hat{r}_j) \frac{\partial D_{12}}{\partial q_i} \right] = 0, \quad i, j = 1, 2, \quad i \neq j. \quad (23)$$

As q_i increases, the marginal revenue from infra-marginal consumers goes up (the first term in brackets of (23) is positive), but consumer demand for the publisher goes down (the second term in brackets is negative). The profit-maximizing q_i optimizes the trade-off between these two effects. Observe that losing a multi-homer entails a loss to publisher i equal to the incremental value $r_{12}^{nT} - \hat{r}_j$ that an advertiser gets from reaching such a consumer, given it already buys q_j impressions (per consumer) from the other publisher. We summarize these results in the following

Lemma 2. *When publishers compete without the ad network, the SPNE is such that publisher $i = 1, 2$ offers a single contract to all advertisers, whereby $m_i = q_i$ and p_i^C is given by (8). Thus, the publisher's profit is $\pi_i^C = p_i^C$. Advertising capacities (q_1^C, q_2^C) satisfy (23).*

7.2 Comparison: direct competition vs. ad network

We are now in a position to compare the equilibria where the ad network is active and where it is not. For the sake of exposition, we perform this comparison in steps. To begin, we assume AN is

unable to track consumers across outlets (i.e., $\beta = 0$). This assumption allows us to focus squarely on one of the functions of the ad network, namely that of centralizing the sale of ads on multiple outlets. By comparing (23) and (12) when $\beta = 0$, we obtain

Lemma 3. *If the ad network is unable to track consumers, the equilibrium provision of ads exceeds that in the case where publishers compete directly if and only if*

$$r_i(q_i) > \hat{r}_i(q_i), \quad i = 1, 2, \quad (24)$$

where all functions are evaluated in (q_1^C, q_2^C) .

The result in Lemma 3 suggests that, in the absence of tracking, the incentives of the ad network when choosing the provision of ads differ from those of the publishers only due to the effect on the revenues from marginal multi-homers. As q_i increases, some multi-homers become single-homers on j . Because the ad network sells impressions on both publishers, this switch produces a net loss equal to $r_{12}^{nT} - r_j$ (see (12)). By contrast, when publishers compete head-to-head, this switch imposes on publisher i a loss equal to the incremental value of a multi-homer, that is, $r_{12}^{nT} - \hat{r}_j$ (see (23)). If the former loss is smaller, i.e. if $r_j > \hat{r}_j$, the ad network perceives a larger opportunity cost of increasing the advertising quantity than each publisher does. Consequently, it supplies more ads than when publishers compete directly. In the following, we refer to this effect on the equilibrium advertising levels as the “joint control” effect.

It is useful to note that, because the ad network centralizes the sale of the ads, outsourcing generates similar effects as a merger of the two publishers. Because the ad network is able to extract the whole advertising surplus, its incentives when deciding on the advertising quantities are the same as those of a monopolist that owns both outlets jointly.³⁰

Let us now turn to the effect of cross-outlet tracking. To isolate this effect, we make use of (12) and compare the case where the ad network is able to track consumers ($\beta > 0$) to that where it is not ($\beta = 0$). Let $(\bar{q}_1^*, \bar{q}_2^*)$ be the advertising quantities that satisfy (12) with $\beta = 0$. We get:

³⁰Indeed, Lemma 3 is identical to Proposition 2 in Ambrus et al. (2016), who analyze the effects of a merger of media outlets on the advertising market.

Lemma 4. *The provision of ads when the ad network is able to track consumers exceeds that when it is unable to do so if and only if*

$$D_{12} \left(\frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i} \right) > -\frac{\partial D_{12}}{\partial q_i} (r_{12}^T - r_{12}^{nT}), \quad i = 1, 2, \quad (25)$$

where all functions are evaluated in $(\bar{q}_1^*, \bar{q}_2^*)$.

As we have seen in Section 4.3, tracking has two countervailing effects: the revenue from additional ads on infra-marginal multi-homers increases (left hand side of (25)), but so does the opportunity cost of losing multi-homers (right hand side of (25)). In line with Proposition 1, the equilibrium advertising levels decrease with tracking if and only if the second effect dominates. In the following, we refer to the effect captured by (25) as the “tracking effect”.

Of course, the fully-fledged comparison between the equilibria with and without *AN* depends on both the joint control and the tracking effects. Using (12) and (23), we get:

Proposition 5. *The provision of ads when publishers outsource to the ad network exceeds that in the case where publishers compete directly if and only if*

$$D_{12} \left(\frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i} \right) > -\frac{\partial D_{12}}{\partial q_i} (r_{12}^T - r_{12}^{nT} + \hat{r}_j - r_j), \quad i, j = 1, 2, \quad i \neq j, \quad (26)$$

where all functions are evaluated in (q_1^C, q_2^C) .

Inequality (26) puts together the two effects identified previously. The left hand-side captures the extra revenue from infra-marginal multi-homers that comes from tracking. The right hand side captures the difference in forgone revenues from marginal multi-homers, which are due to both tracking, i.e. $r_{12}^T - r_{12}^{nT}$, and joint control, i.e. $\hat{r}_j - r_j$.

In sum, Proposition 5 indicates that the advertising network affects the provision of advertising through its superior ability to track consumers across outlets and through its centralizing role. Because these effects may go in different directions, it is difficult to generally establish whether the provision of advertising increases with the ad network. Nevertheless, our framework identifies the conditions under which it increases or decreases. For instance, if the quantity of multi-homers is relatively large, and tracking strongly reduces the waste of marginal impressions, we can expect

the ad network to bring to an expansion in the volume of ads. By contrast, when the quantity of multi-homers is relatively small, we expect the ad network to have a restrictive effect.

7.3 Welfare analysis

To conclude this section, we compare welfare with outsourcing and when publishers compete for advertisers. When deciding the level of advertising, the ad network maximizes total advertiser surplus. However, neither the ad network nor the publishers take consumer surplus into account. Hence, compared to the second-best, advertising is overprovided when it is outsourced to the ad network, whereas it may be over- or under-provided when publishers compete head-to-head. As indicated by Proposition 5, publishers may set advertising levels higher or lower than the ad network does. Hence, the effect of the ad network on consumer surplus from content, CS_c , is positive when $q_i^* < q_i^C$. Nonetheless, one may argue that the disutility from privacy losses CS_p increases when advertising is outsourced to the ad network, because tracking by the ad network may be perceived by consumers as more invasive than that by publishers. Thus, even if $q_i^* < q_i^C$, consumer surplus and social welfare may decrease when advertising is outsourced. Summing up, the desirability of the ad network is generally ambiguous.

8 Privacy as an intermediate good

We have so far assumed that consumers care about privacy for its own sake, in the sense that tracking has no implications for the utility they obtain from browsing content. Thus, we modeled privacy as a “final good” (Farrell, 2012). In reality, the (dis)utility from ads may depend on whether consumers allow tracking, for example, because the ad network may cap the frequency of impressions on multi-homers more effectively with tracking. Hence, although consumers may care about privacy for its own sake (this is what the parameter θ captures), they may be more willing to accept tracking if it improves their online experience overall. To capture these aspects, we now extend the model to consider privacy as an “intermediate good”. Specifically, we assume the disutility from advertising depends on tracking: when a multi-homer allows tracking, her disutility from ads is reduced by a quantity z per impression.³¹ To avoid repetition, in this section we discuss informally how this

³¹In a broader interpretation, z could also capture the fact that tracking brings more relevant ads to the consumer.

change affects our main results. We relegate the formal analysis to the Online Appendix.

A first consequence of letting tracking affect the disutility from ads is that consumer privacy and content choices are intertwined. In particular, consumers with a lower disutility from tracking are more likely to multi-home, because tracking reduces the disutility from advertising and, with it, the marginal cost of visiting a publisher. Hence, for given reservation utilities u_1 and u_2 , an individual allowing tracking is more likely to visit both publishers. In turn, the choice whether to allow tracking also depends on whether a consumer multi-homes. Specifically, a consumer always (resp., never) allows tracking when θ is small (resp., large) enough. However, there exists a set of intermediate values of θ such that a consumer allows tracking only if she visits both publishers. The intuition is that the consumer benefits from the effect of tracking on the disutility from advertising only when multi-homing. Finally, we find that the quantity of consumers that allows tracking increases with the provision of advertising. The reason is that tracking becomes relatively more beneficial to consumers when the quantity of ads increases.

Given these complications, it is not surprising that the consumer demand system becomes quite involved. For illustrative purposes, we here represent it (see Figure 2) for a specific distribution of preferences, with horizontally differentiated publishers located at the extremes of the Hotelling line. Assume that $u_1 = u - tx$ and $u_2 = u - t(1 - x)$, with $x \sim U[0, 1]$, $u > 0$ and $t > 0$. Assume also that $\theta \sim U[0, \bar{\theta}]$.³²

Despite the increased complexity, our main findings do not qualitatively change. The reason is that, although the derivatives of D_1 , D_2 and D_{12} with respect to q_1 and q_2 differ with respect to the baseline model, the formulas of the first-order conditions (12) and (23), from which we derived the effects of tracking on advertising quantities, do not change. Therefore, Proposition 1 still holds. Furthermore, the comparison of ad quantities with the case where publishers compete directly is as in Proposition 5. Finally, the effects of changes in β on welfare also have the same structure as in expressions (15) and (17). Thus, Propositions 2 and 4 are qualitatively confirmed.

³²We assume parameter values are such that the market is covered.




Figure 2: Representation of demands in the Hotelling case. In this scenario, we have $\bar{x}_B = 1 - (u - \delta q_2)/t$, $\bar{\bar{x}}_B = (u - \delta q_1)/t$, $\bar{x}_{NB} = 1 - (u - \delta q_2 - z(q_1 + q_2))/t$ and $\bar{\bar{x}}_B = (u - \delta q_1 - z(q_1 + q_2))/t$. The subscript B (NB) stands for “block” (“not block”).

9 Concluding remarks

We have studied the link between consumer privacy choices, tracking and the advertising market. We have found that the provision of advertising may decrease with the tracking effectiveness, identifying a testable condition for this to occur. We have also identified several externalities linked to consumers’ decisions to protect their privacy. In particular, we have shown that there is a direct negative externality on the advertising industry, because allowing tracking makes advertising more effective. There is also an indirect externality on consumers, because tracking affects the quantity of ads consumers are exposed to. The latter externality may be positive or negative, depending on the effect of tracking on the provision of advertising. Assuming consumers dislike advertising, this result implies that tracking may be too low in equilibrium, even if the regulator’s objective is to maximize consumer surplus. We have explored several possible policy interventions to restore efficiency, including opt-in and opt-out regulation and the assignment of property rights on consumer personal data.

We also have investigated how an ad network affects the online advertising market, focusing on its role in centralizing the sale of ad inventories from several competing outlets and its ability to track multi-homing consumers across them. We have identified several testable conditions that characterize the ad network’s effect on advertising quantities, relating to the quantity of multi-

homing consumers, to how tracking increases the revenue generated from multi-homers and to the elasticity of consumer demands with respect to advertising quantities. Finally, we have considered the case where privacy is an intermediate good, confirming the qualitative results of the main model.

Our model is based on several assumptions that one may relax in future work. First, one could introduce consumer heterogeneity in product preferences and targeting of advertisements by the digital outlets. Second, one might consider competition among large and small publishers, which provide high and low values ad spaces, in order to understand which type of publishers chooses to outsource ad spaces and which inventories are outsourced.

References

- [1] Acquisti, A., C. Taylor and L. Wagman, 2016. The Economics of Privacy. *Journal of Economic Literature*, 54, 442-492.
- [2] Ambrus, A., E. Calvano, and M. Reisinger, 2016. Either or Both Competition: a Two-Sided Theory of Advertising With Overlapping Viewerships. *American Economic Journal: Microeconomics*, 8, 189-222.
- [3] Anderson, S.P. and S. Coate, 2005. Market Provision of Broadcasting: A Welfare Analysis. *Review of Economic Studies*, 72, 947-972.
- [4] Anderson, S.P., Ø. Foros and Kind, H. J., 2016. Competition for Advertisers in Media Markets. *Economic Journal*, forthcoming.
- [5] Athey, S., E. Calvano and J.S. Gans, 2016. The Impact of Consumer Multi-Homing on Advertising Markets and Media Competition. *Management Science*, forthcoming.
- [6] Athey, S., and J.S. Gans, 2010. The Impact of Targeting Technology on Advertising Markets and Media Competition. *The American Economic Review P&P*, 100, 608-613.
- [7] Aziz, A. and R. Telang, 2017. Consumer Tracking and Competition in Targeted Advertising. Available at SSRN: <https://ssrn.com/abstract=2993653>.
- [8] Bagwell, K., 2007. The Economic Analysis of Advertising, in Mark Armstrong and Rob Porter (eds.). *Handbook of Industrial Organization*, Vol. 3, Amsterdam: North-Holland.

- [9] Belleflamme, P. and W. Vergote, 2016. Monopoly price discrimination and privacy: The hidden cost of hiding. *Economics Letters* 149, 141–144
- [10] Bergemann, D. and A. Bonatti, 2011. Targeting in Advertising Markets: Implications for Offline versus Online Media. *RAND Journal of Economics*, 42, 417-443
- [11] Bergemann, D. and A. Bonatti, 2015. Selling Cookies. *American Economic Journal: Microeconomics*, 7, 259-294.
- [12] Bernheim D. and M.D. Whinston, 1985. Common Marketing Agency as a Device for Facilitating Collusion. *RAND Journal of Economics*, 16, 269-281
- [13] Calvano, E. and B. Jullien, 2012. Issues in online advertising and competition policy: a two-sided market perspective, *Recent Advances in the Analysis of Competition Policy and Regulation*. Edited by J.E. Harrington and Y. Katsoulacos, Edward Elgar Publishing.
- [14] Campbell, J., A. Goldfarb, and C. Tucker, 2015. Privacy Regulation and Market Structure. *Journal of Economics & Management Strategy*, 24, 47-73.
- [15] Chiou, L. and C. Tucker, 2014. Search Engines and Data Retention: Implications for Privacy and Antitrust, MIT Sloan School Working Paper 5094-14.
- [16] Coase, R. H., 1960. The Problem of Social Cost. *Journal of Law and Economics*, 3, 1-44.
- [17] Comscore 2016. Top 50 Digital Rankings. Available at <http://www.comscore.com/Insights/Rankings/comScore-Ranks-the-Top-50-US-Digital-Media-Properties-for-February-2016>
- [18] Conitzer, V., C. R. Taylor, and L. Wagman, 2012. Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases. *Marketing Science*, 31, 277-92.
- [19] Crampes, C., C. Haritchabalet, and B. Jullien, 2009. Advertising, Competition, and Entry in Media Industries. *Journal of Industrial Economics*, 57, 7-31.
- [20] de Cornière, A. and R. de Nijs, 2016. Online advertising and privacy. *RAND Journal of Economics*, 47, 48-72.

- [21] de Cornière, A. and G. Taylor, 2014. Integration and search engine bias. *RAND Journal of Economics*, 45, 576-597.
- [22] Farrell, J., 2012. Can privacy be just another good? *Journal on Telecommunications and High Technology Law*, 10, 251-264.
- [23] George, L.M. and C. Hogendorn, 2012. Aggregators, search and the economics of new media institutions. *Information Economics and Policy*, 24, 40-51.
- [24] Goldfarb, A. and C. E. Tucker, 2011. Privacy regulation and online advertising. *Management Science* 57, 57-71.
- [25] Goldfarb, A., and C. E. Tucker, 2012. Shifts in Privacy Concerns. *American Economic Review*, 102, 349-53.
- [26] Hann, I.-H., K. L. Hui, S.-Y. T. Lee, and I.P.L. Png, 2008. Consumer Privacy and Marketing Avoidance: A Static Model. *Management Science*, 54, 1094-1103.
- [27] Hermalin, B. E., and M. L. Katz, 2006. Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy. *Quantitative Marketing and Economics* 4, 209-39.
- [28] Hoofnagle C. J., A. Soltani, N. Good and D.J. Wambach, 2012. Behavioral Advertising: The Offer you Cannot Refuse. *Harvard Law and Policy Review*, 6, 273-296
- [29] Interactive Advertising Bureau (IAB), 2015. Internet advertising revenue report. Available at: http://www.iab.com/wp-content/uploads/2016/04/IAB_Internet_Advertising_Revenue_Report_FY_2015-final.pdf
- [30] Johnson, J.P., 2013. Targeted advertising and advertising avoidance. *RAND Journal of Economics*, 44, 128-144.
- [31] Kind, H. J., T. Nilssen, and L. Sørgard, 2016. Inter-Firm Price Coordination in a Two-Sided Market. *International Journal of Industrial Organization*, 44, 101-112.
- [32] Kreps, D.M. and J.A. Scheinkman, 1983. Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes. *The Bell Journal of Economics*, 14, 326-337.

- [33] Levin, J., and P. Milgrom, 2010. Online Advertising: Heterogeneity and Con action in Market Design. *American Economic Review*, 100, 603-07.
- [34] Lohr, S., 2010. Privacy Concerns Limit Online Ads, Study Says, *New York Times*.
- [35] Montes, R., W. Sand-Zantman and T.M. Valletti, 2015. The Value of Personal Information in Markets with Endogenous Privacy. CEIS Working Paper No. 352.
- [36] Morales, L., 2010. US Internet users ready to limit tracking of ads. Available at: <http://www.gallup.com/poll/145337/internet-users-ready-limit-onlinetracking/>
- [37] Peitz, M. and T.M. Valletti, 2008. Content and advertising in the media: Pay-TV versus free-to-air. *International Journal of Industrial Organization*, 26, 949-965.
- [38] Renault, R., 2015. Advertising in markets. In Anderson, S., Stromberg, D., and J. Waldfogel (eds.). *Handbook of Media Economics*. Amsterdam: Elsevier.
- [39] Riley, J., 1979. Informational Equilibria. *Econometrica*, 47, 331-359.
- [40] Roesner, F., T. Kohno and D. Wetherall, 2012. Detecting and Defending Against Third-Party Tracking on the Web. NSDI 2012.
- [41] Samuelson, P. 2000. Privacy as intellectual property. *Stanford Law Review* 52 (1125).
- [42] Schwartz, P. 2004. Property, Privacy, and Personal Data. *Harvard Law Review*, 117, 2056-2128.
- [43] Tucker, C. E., 2012. The economics of advertising and privacy. *International Journal of Industrial Organization* 30, 326–329.
- [44] Tucker, C. E., 2014. Social Networks, Personalized Advertising, and Privacy Controls. *Journal of Marketing Research*, 51, 546-562.
- [45] Turow, J., J. King, C.J. Hoofnagle, A. Bleakley and M. Hennessy, 2009. Americans Reject Tailored Advertising and Three Activities that Enable It. Mimeo, Berkeley
- [46] Vives, X., 2000. *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press: Cambridge, Massachusetts.

A Proofs of propositions and lemmas

A.1 Proof of Lemma 1

We begin by proving that the ad network offers a single contract. Assume the ad network offers two contracts, $C_i^1 \equiv (p_{AN}^1, m_i^1, m_j^1)$ and $C_i^2 \equiv (p_{AN}^2, m_i^2, m_j^2)$, and that some advertisers accept C_i^1 , while others accept C_i^2 . If $p_{AN}^1 < p_{AN}^2$, then the ad network is better off offering to all advertisers contract C_i^2 . Because advertisers have identical revenue functions, if contract C_i^2 is such that one advertiser accepts it, then all advertisers accept it. Due to diminishing returns from advertising, the ad network spreads publishers' advertising levels equally across advertisers. We now show that the equilibrium contract is such that $m_i = q_i$. First, by assumption q_i represents the capacity constraint per consumer of publisher i , hence it cannot be $m_i > q_i$. Moreover, it cannot be that $m_i < q_i$ because $r_i(m_i)$ and $r_{12}^\sigma(m_i, m_j)$ are increasing in m_i . Hence, the ad network can always increase its revenue by offering more ads for $p_{AN} + \varepsilon$, with $\varepsilon > 0$, up to $m_i = q_i$. Finally, the price p_{AN} is equal to the difference between the revenues of the advertiser when it accepts the contract, $r_i(q_i) D_i(q_i, q_j) + r_j(q_j) D_j(q_i, q_j) + r_{12}^\sigma(q_1, q_2, \beta) D_{12}(q_1, q_2)$, and its outside option, which is equal to zero because AN manages ad inventory of both publishers. Hence, we get (9).

Following the previous proof, one can prove that each publisher offers one contract to all advertisers, such that $m_i = q_i$. p_i is equal to the difference between the revenues of the advertiser when it advertises on both publishers, $r_i(q_i) D_i(q_i, q_j) + r_j(q_j) D_j(q_i, q_j) + r_{12}^\sigma(q_1, q_2) D_{12}(q_1, q_2)$, and its outside option, that is, the revenues of the advertiser if it acquires only ad spaces on publisher j $r_j(q_j) D_j(q_i, q_j) + \hat{r}_j(q_j) D_{12}(q_1, q_2)$. As a result, we get (8). ■

A.2 Proof of Proposition 1

We denote by FOD_i the first order derivatives of the maximization problem of p_{AN} with respect to q_i . By the Implicit Function Theorem, we have that:

$$\frac{\partial q_i}{\partial \beta} = \frac{\frac{\partial FOD_i}{\partial \beta} \frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_j}{\partial \beta} \frac{\partial FOD_i}{\partial q_j}}{\frac{\partial FOD_i}{\partial q_j} \frac{\partial FOD_j}{\partial q_i} - \frac{\partial FOD_i}{\partial q_i} \frac{\partial FOD_j}{\partial q_j}}. \quad (27)$$

The denominator is always negative, because $\frac{\partial FOD_i}{\partial q_j} \frac{\partial FOD_j}{\partial q_i} - \frac{\partial FOD_i}{\partial q_i} \frac{\partial FOD_j}{\partial q_j} > 0$ by the second order conditions (SOC) of the problem. To simplify, we assume symmetry, implying that the numerator

can be rewritten as $\frac{\partial FOD_i}{\partial \beta} \left(\frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_i}{\partial q_j} \right)$. Again, by the SOC, we have $\frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_i}{\partial q_j} < 0$. This implies that if $\frac{\partial FOD_i}{\partial \beta} = D_{12} \frac{\partial^2 r_{12}^T}{\partial q_i \partial \beta} + \frac{\partial D_{12}}{\partial q_i} \frac{\partial r_{12}^T}{\partial \beta} > 0$, then $\frac{\partial q_i}{\partial \beta} > 0$. ■

A.3 Proof of Lemma 3

Let (q_1^C, q_2^C) and $(\bar{q}_1^*, \bar{q}_2^*)$ be the couples of quantities that satisfy, respectively, $\frac{\partial \pi_i^C}{\partial q_i} = 0$ and $\frac{\partial p_{AN}}{\partial q_i} = 0$, with $\beta = 0$, for $i = 1, 2$. Profit functions are strictly concave with respect to q_i , $i = 1, 2$. By concavity, it can be shown that these couples are unique. We compute (12) in (q_1^C, q_2^C) and with $\beta = 0$. Using (23), we can rewrite this expression as

$$\frac{\partial p_{AN}}{\partial q_i} = \frac{\partial D_{12}(q_1, q_2)}{\partial q_i} \hat{r}_j(q_j) + \frac{\partial D_j(q_1, q_2)}{\partial q_i} r_j(q_j), \quad i = 1, 2. \quad (28)$$

where all functions are evaluated in (q_1^C, q_2^C) . Using the fact that $-\frac{\partial D_{12}(q_1, q_2)}{\partial q_i} = \frac{\partial D_j(q_1, q_2)}{\partial q_i}$, then (28) can be rewritten as $\frac{\partial D_{12}(q_1, q_2)}{\partial q_i} [\hat{r}_j(q_j) - r_j(q_j)]$. By concavity, if $\frac{\partial p_{AN}}{\partial q_i}$ computed in $\beta = 0$ is strictly positive when evaluated in (q_1^C, q_2^C) , then $(\bar{q}_1^*, \bar{q}_2^*)$ must be such that $q_i^{ANT} > q_i^C$, $i = 1, 2$. Therefore $\bar{q}_i^* > q_i^C$ if and only if $\hat{r}_j(q_j) < r_j(q_j)$. ■

A.4 Proof of Lemma 4

Let $(\bar{q}_1^*, \bar{q}_2^*)$ and (q_1^*, q_2^*) be the couples of quantities that satisfy, respectively, $\frac{\partial p_{AN}}{\partial q_i} = 0$ computed in $\beta = 0$, and $\frac{\partial p_{AN}}{\partial q_i} = 0$ computed in $\beta > 0$. By concavity, it can be shown that these couples are unique. Evaluating (12) in (q_1^{ANT}, q_2^{ANT}) , and using the fact that

$$\left. \frac{\partial D_i}{\partial q_i} r_i + D_i \frac{\partial r_i}{\partial q_i} + \frac{\partial D_j}{\partial q_j} r_j \right|_{q_i=q_1^*, q_2=q_2^*} = -D_{12} \frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial D_{12}}{\partial q_i} r_{12}^{nT} \Big|_{q_i=\bar{q}_1^*, q_2=\bar{q}_2^*}$$

we have, for $i = 1, 2$:

$$\begin{aligned} \frac{\partial p_{AN}}{\partial q_i} &= D_{12}(q_1, q_2) \left[\frac{\partial r_{12}^T(q_1, q_2, \beta)}{\partial q_i} - \frac{\partial r_{12}^{nT}(q_1, q_2)}{\partial q_i} \right] \Big|_{q_i=\bar{q}_1^*, q_2=\bar{q}_2^*} \\ &\quad + \frac{\partial D_{12}(q_1, q_2)}{\partial q_i} [r_{12}^T(q_1, q_2, \beta) - r_{12}^{nT}(q_1, q_2)] \Big|_{q_i=\bar{q}_1^*, q_2=\bar{q}_2^*} \end{aligned}$$

By concavity, if $\frac{\partial p_{AN}}{\partial q_i}$ is strictly positive when evaluated in $(\bar{q}_1^*, \bar{q}_2^*)$, then (q_1^*, q_2^*) must be such that $q_i^* > \bar{q}_i^*$, $i = 1, 2$. ■

A.5 Proof of Proposition 5

Consider the FOD (12). Let (q_1^C, q_2^C) and (q_1^*, q_2^*) be the couples of quantities that satisfy, respectively, $\frac{\partial \pi_i^C}{\partial q_i} = 0$ and $\frac{\partial p_{AN}}{\partial q_i} = 0$, $i = 1, 2$. By assumption, profit functions are strictly concave with respect to q_i , $i = 1, 2$. Hence, it can be shown that these couples are unique. Evaluating (12) in (q_1^C, q_2^C) and using (23), we have, for $i = 1, 2$,

$$D_{12}(q_1, q_2) \left[\frac{\partial r_{12}^T(q_1, q_2, \beta)}{\partial q_i} - \frac{\partial r_{12}^{nT}(q_1, q_2)}{\partial q_i} \right] + \frac{\partial D_{12}(q_1, q_2)}{\partial q_i} [r_{12}^T(q_1, q_2) - r_{12}^{nT}(q_1, q_2) + \hat{r}_j(q_j)] + \frac{\partial D_j(q_1, q_2)}{\partial q_i} r_j(q_j) \Big|_{q_i=q_1^C, q_2=q_2^C} \quad (29)$$

By concavity, if $\frac{\partial p_{AN}}{\partial q_i} > 0$ when evaluated in (q_1^C, q_2^C) , then (q_1^*, q_2^*) is such that $q_i^* > q_i^C$, $i = 1, 2$. Using the equality $-\frac{\partial D_{12}(q_1, q_2)}{\partial q_i} = \frac{\partial D_j(q_1, q_2)}{\partial q_i}$, we find that $q_i^* > q_i^C$ if and only if

$$D_{12} \left[\frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i} \right] + \frac{\partial D_{12}}{\partial q_i} [r_{12}^T - r_{12}^{nT} + \hat{r}_j - r_j] \Big|_{q_i=q_1^C, q_2=q_2^C} > 0. \quad \blacksquare$$

B Online Appendix

B.1 Microfoundation for the advertising revenues

We provide a micro-foundation for the advertising technology presented in the text. Assume there are $N \geq 1$ advertisers, indexed by $a \in \{1, \dots, N\}$. Each advertiser has two messages to send. We normalize the return produced by the first message (e.g., an ad about a mainstream product) when informing a consumer to one, whereas the second message's return is $y < 1$ (e.g., an ad about a niche product). We assume each consumer has a probability x_i^a of registering an ad by a on publisher $i = 1, 2$, where $x_i^a \sim U[0, 1]$, $\forall i, a$, and x_i^a is i.i.d. for all consumers, advertisers, and publishers. In this framework, advertising is behavioral and contextual: x_i^a is determined by the consumer's behavior on website i . For example, a consumer is more likely to register a political ad when visiting a politics-related webpage, and an ad on sporting goods when visiting a sports webpage. For simplicity, we assume that the behavior of a consumer on a website is independent of her behavior on the other website. Furthermore, x_i^a is orthogonal to the consumers' preferences for content u_1 and u_2 , as well to the preference for privacy θ . For simplicity, we also assume that each

advertiser wants no more than one impression on a consumer on a given publisher.³³ The publisher can meet this requirement because it is able to perfectly track the quantity of impressions on a consumer on its own website (i.e., perfect internal tracking). These assumptions imply that there are no repeated impressions of the same ad on a publisher.³⁴ In this setting the total capacity of the publishers is $q_i (D_i + D_{12})$, and the mass of advertisers has to be smaller than the mass of the consumers.

Publisher Competition. Suppose advertiser a buys m_i^a impressions on publisher i . Intuitively, the publisher places these impressions on the consumers with the largest x_i^a , maximizing the campaign's reach (i.e., the expected quantity of consumers who become informed about a product), but does not send more than one impression per advertiser on a consumer. Let \bar{x}_i^a denote the value of x_i^a such that all consumers with higher x_i^a receive an impression on i from advertiser a . The following holds

$$m_i^a = (1 - \bar{x}_i^a) (D_i + D_{12}) \Rightarrow \bar{x}_i^a = 1 - \frac{m_i^a}{D_i + D_{12}}. \quad (30)$$

Intuitively, the more impressions an advertiser buys, the higher is the share of consumers that are impressed on a given publisher. However, the probability the impression is registered by the marginal consumer decreases with m_i^a .

Assume that advertiser a single-homes on publisher i . Because only consumers such that $x_i^a \geq \bar{x}_i^a$ receive an impression, a 's total return is

$$R_i^a = D_i \alpha_i \int_{\bar{x}_i^a}^1 x dx + D_{12} \gamma_i \int_{\bar{x}_i^a}^1 x dx, \quad (31)$$

where α_i (respectively, γ_i) is the probability that a single-homing (multi-homing) consumer watches advertisements on publisher i (which may depend on the time a consumer spends on the publisher). Thus, the integral $\alpha_i \int_{\bar{x}_i^a}^1 x dx$ (resp. $\gamma_i \int_{\bar{x}_i^a}^1 x dx$) is the expected probability that a single-homing (resp. multi-homing) consumer registers an ad by a , that is the expected share of single- (multi)-homers

³³Consumers might lose interest in ads from the same advertiser on the same publisher, or even be annoyed by them.

³⁴In the main text we consider a more general model, where an advertiser may want more than one impression on the same publisher.

that are informed on publisher i . Using (30), we can write these shares as

$$\alpha_i \int_{\bar{x}_i^a}^1 x dx = \alpha_i \left(\frac{1}{2} - \frac{(\bar{x}_i^a)^2}{2} \right) = \alpha_i \left(\frac{m_i^a}{D_i + D_{12}} - \frac{1}{2} \left(\frac{m_i^a}{D_i + D_{12}} \right)^2 \right),$$

$$\gamma_i \int_{\bar{x}_i^a}^1 x dx = \gamma_i \left(\frac{m_i^a}{D_i + D_{12}} - \frac{1}{2} \left(\frac{m_i^a}{D_i + D_{12}} \right)^2 \right).$$

These functions are increasing and concave in $\frac{m_i^a}{D_i + D_{12}}$ (and, hence, in m_i^a). Observe that, under our assumptions, if the advertiser single-homes it only sends impressions of the first message (with value one).

Consider now an advertiser that multi-homes. We assume that the advertiser prefers to send the same type of message on both publishers.³⁵ Given this assumption, we have

$$R_{MH}^a = D_1 \alpha_1 \int_{\bar{x}_1^a}^1 x dx + D_2 \alpha_2 \int_{\bar{x}_2^a}^1 x dx + D_{12} \left(\gamma_1 \int_{\bar{x}_1^a}^1 x dx + \gamma_2 \int_{\bar{x}_2^a}^1 x dx - \gamma_2 \int_{\bar{x}_2^a}^1 x dx \gamma_1 \int_{\bar{x}_1^a}^1 x dx \right), \quad (32)$$

where the last term is the expected share of multi-homing consumers informed about a 's first message. This is equal to the probability that the consumer is informed while on publisher $i = 1, 2$, minus the probability that she is informed on both. The latter is equal to $\gamma_2 \int_{\bar{x}_2^a}^1 x dx \gamma_1 \int_{\bar{x}_1^a}^1 x dx$ because x_2^a and x_1^a are i.i.d..³⁶ Following the same steps as above, we have

$$R_{MH}^a = D_1 \alpha_1 M_1^a + D_2 \alpha_2 M_2^a + D_{12} (\gamma_1 M_1^a (1 - \gamma_2 M_2^a)) + D_{12} \gamma_2 M_2^a.$$

where we have denoted $M_i^a \equiv \frac{m_i^a}{D_i + D_{12}} - \frac{1}{2} \left(\frac{m_i^a}{D_i + D_{12}} \right)^2$ to save notation. The expression for R_{MH}^a tells us that the revenue of an advertiser that multi-homes is increasing and concave in $\frac{m_1^a}{D_1 + D_{12}}$ and $\frac{m_2^a}{D_2 + D_{12}}$ (and, hence, in m_1^a and m_2^a). Defining $r_i \left(\frac{m_i^a}{D_i + D_{12}} \right) \equiv \alpha_i M_i^a$ and $\hat{r}_i \left(\frac{m_i^a}{D_i + D_{12}} \right) \equiv \gamma_i M_i^a$, we can rewrite

$$R_i^a = D_i r_i \left(\frac{m_i^a}{D_i + D_{12}} \right) + D_{12} \hat{r}_i \left(\frac{m_i^a}{D_i + D_{12}} \right)$$

³⁵ A sufficient condition for this assumption to be satisfied is that y is not too large. Specifically, the loss from not informing a consumer with the first message is bigger than the gain from sending a different message to a multi-homer that is already informed by the first message.

³⁶ To avoid repetition, the publisher may decide to target a group of consumers with lower probability x_i^a . However, given the i.i.d. assumption, the share of consumers on i that is already informed on publisher j is the same regardless of the set of consumers on i that is chosen.

and

$$R_{MH}^a = D_1 r_1 \left(\frac{m_1^a}{D_1 + D_{12}} \right) + D_2 r_2 \left(\frac{m_2^a}{D_2 + D_{12}} \right) + D_{12} \left(\hat{r}_1 \left(\frac{m_1^a}{D_1 + D_{12}} \right) \left(1 - \hat{r}_2 \left(\frac{m_2^a}{D_2 + D_{12}} \right) \right) + \hat{r}_2 \left(\frac{m_2^a}{D_2 + D_{12}} \right) \right).$$

Because of diminishing advertising returns, it is optimal for both publishers to allocate the same number of impressions to each advertiser. Thus, $m_i^a = \frac{(D_i + D_{12})q_i}{N}$, where q_i represents the advertising capacity of publisher i . Hence, we can rewrite $r_i \left(\frac{q_i}{N} \right)$ and $\hat{r}_i \left(\frac{q_i}{N} \right)$. Each publisher extracts from each advertiser the incremental value of placing $m_i^a = \frac{(D_i + D_{12})q_i}{N}$ impressions on its platform, that is

$$R_{MH}^a - R_i^a = \left[\sum_{i=1,2} D_i r_i \left(\frac{q_i}{N} \right) + D_{12} \left(\hat{r}_1 \left(\frac{q_1}{N} \right) \left(1 - \hat{r}_2 \left(\frac{q_2}{N} \right) \right) + \hat{r}_2 \left(\frac{q_2}{N} \right) \right) \right] - \left[D_i r_i \left(\frac{q_i}{N} \right) + D_{12} \hat{r}_i \left(\frac{q_i}{N} \right) \right].$$

Ad Network. Consider now the case where AN sells impressions on behalf of both publishers. We assume that AN can track consumers across publishers. Specifically, for each advertiser a , AN observes x_i^a for each consumer on the respective publisher, and chooses which consumers receive an impression. Following this choice, AN observes whether a consumer selected to receive an impression on publisher i has already been informed by a 's messages on j . This event takes place with probability β (i.e. the probability the a consumer allows tracking). Based on this information, AN decides whether to impress the consumer with either the first or the second type of message from a . Given these assumptions, because there are no internal repetitions, and because x_i^a are i.i.d., AN cannot do better than send the m_i^a impressions a buys on publisher i to the consumers with the largest x_i^a . Therefore, we have

$$m_i^a = (1 - \bar{x}_i^a) (D_i + D_{12}) \Rightarrow \bar{x}_i^a = 1 - \frac{m_i^a}{D_i + D_{12}}.$$

However, these impressions may not all be of the same type: although informing a consumer with the first message is more profitable, the second message avoids repetition on multi-homers. Specifically, suppose that a consumer has been selected to receive an impression on publisher i (because x_i^a is above \bar{x}_i^a). With probability $(1 - \beta)$, AN does not observe whether the consumer has already been informed on the other publisher. In that case, AN will send the first type of message (as long as y is not too large, see footnote 35). With probability β , AN observes instead whether the consumer has been informed on publisher j (for instance, it learns whether a consumer has clicked on the ad). If the consumer has not already been informed, AN sends her an impression containing

the first message (because this impression generates the largest expected value to the advertiser). However, if the consumer has already been informed, AN sends her the second message, because the value of repeating the first impression is zero. Hence, the total revenue of an advertiser buying m_1^a and m_2^a impressions is equal to (denoting again $M_i^a \equiv \frac{m_i^a}{D_i + D_{12}} - \frac{1}{2} \left(\frac{m_i^a}{D_i + D_{12}} \right)^2$)

$$R_{MH}^a = D_1 \alpha_1 M_1^a + D_2 \alpha_2 M_2^a + D_{12} (1 - \beta) (\gamma_1 M_1^a (1 - \gamma_2 M_2^a)) + \\ + D_{12} (1 - \beta) \gamma_2 M_2^a + D_{12} \beta \gamma_2 M_2^a + D_{12} \beta (\gamma_1 M_1^a (1 - \gamma_2 M_2^a)) + D_{12} \beta \gamma_1 M_1^a \gamma_2 M_2^a y,$$

which simplifies to

$$R_{MH}^a = D_1 \alpha_1 M_1^a + D_2 \alpha_2 M_2^a + D_{12} (\gamma_1 M_1^a (1 - \gamma_2 M_2^a)) + D_{12} \gamma_2 M_2^a + D_{12} \beta \gamma_1 M_1^a \gamma_2 M_2^a y.$$

In words, conditional on being able to track consumers across publishers, AN raises the value of impressions on multi-homers: it can replace ads containing a message from an advertiser with ads from the same advertiser, but containing another message. By using the definition of $r_1(\cdot)$ and $\hat{r}_1(\cdot)$ above, we get

$$R_{MH}^a = \sum_{i=1,2} D_i r_i \left(\frac{m_i^a}{D_i + D_{12}} \right) + D_{12} \hat{r}_1 \left(\frac{m_1^a}{D_1 + D_{12}} \right) \left(1 - \hat{r}_2 \left(\frac{m_2^a}{D_2 + D_{12}} \right) \right) + \\ + D_{12} \hat{r}_2 \left(\frac{m_2^a}{D_2 + D_{12}} \right) + D_{12} \beta \hat{r}_1 \left(\frac{m_1^a}{D_1 + D_{12}} \right) \hat{r}_2 \left(\frac{m_2^a}{D_2 + D_{12}} \right) y.$$

B.2 Privacy as intermediate good

Given our assumptions, the utility of a multi-homing consumer that allows tracking is

$$u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta. \quad (33)$$

Except for this change, the structure of the model is the same as described in Section 3.

As a first step, it is useful to describe a consumer's choice in terms of tracking. Using (33), we

have that a consumer with preference parameter θ behaves as follows

$$\begin{aligned}
&\text{if } \theta \leq c && \text{the consumer allows tracking, } \forall q_1, q_2, \\
&\text{if } c < \theta \leq c + z(q_1 + q_2) && \text{the consumer allows tracking only when multi-homing,} \\
&\text{if } \theta > c + z(q_1 + q_2) && \text{the consumer never allows tracking.}
\end{aligned} \tag{34}$$

The consumer's choice concerning tracking depends on the advertising capacity q_i on the two publishers. Below, we prove that the demand system in this version of the model is as follows (the subscripts B and NB stand for “block tracking” and “not block tracking”, respectively):

- if $\theta > c + z(q_1 + q_2)$, then

$$\begin{aligned}
D_{12,B} &= Pr \{u_1 - \delta q_1 - c \geq 0, u_2 - \delta q_2 - c \geq 0, \theta > c + z(q_1 + q_2)\}, \\
D_{i,B} &= Pr \{u_i - \delta q_i - c \geq 0, u_j - \delta q_j - c < 0, \theta > c + z(q_1 + q_2)\}, \quad i, j = 1, 2 \text{ and } i \neq j.
\end{aligned} \tag{35}$$

- if $c < \theta \leq c + z(q_1 + q_2)$, then

$$\begin{aligned}
D_{12,NB} &= Pr \{u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta \geq \max(u_1 - \delta q_1 - c; u_2 - \delta q_2 - c; 0), \\
&\quad c < \theta \leq c + z(q_1 + q_2)\}, \\
D_{i,B} &= Pr \{u_i - \delta q_i - c \geq 0, u_i - \delta q_i - c > u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta, \\
&\quad c < \theta \leq c + z(q_1 + q_2)\}, \quad i, j = 1, 2 \text{ and } i \neq j.
\end{aligned} \tag{36}$$

- if $\theta \leq c$, then

$$\begin{aligned}
D_{12,NB} &= Pr \{u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta \geq \max(u_1 - \delta q_1 - \theta; u_2 - \delta q_2 - \theta; 0), \theta \leq c\}, \\
D_{i,NB} &= Pr \{u_i - \delta q_i - \theta \geq 0, u_i - \delta q_i - \theta > u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta, \\
&\quad \theta \leq c\}, \quad i, j = 1, 2 \text{ and } i \neq j.
\end{aligned} \tag{37}$$

with $D_0 = 1 - D_{12,NB} - D_{12,B} - \sum D_{i,B} - \sum D_{i,NB}$. We define $D_{12} = D_{12,B} + D_{12,NB}$.

Consumer surplus is

$$CS = CS_B + CS_{NB}, \tag{38}$$

where CS_B (resp. CS_{NB}) is the surplus of consumers who block (resp. allow) tracking. These terms are rather involved, but not essential for the results the follow, so relegate the expressions below.

We can now proceed to the comparison of the competitive equilibrium with the first- and second-best allocations, as we did in Section 5. The first best allocation is such that $W = AS + CS$ is maximized with respect to the share of multi-homing consumers that allow tracking, i.e. $\beta \equiv \frac{D_{12,NB}}{D_{12}}$, and to the advertising capacities on both publishers (q_1, q_2) . Following the same steps as in Section 5.1, we can show that the equilibrium quantity of consumers that allow tracking is strictly smaller than the first-best one, whereas the comparison between advertising capacities is ambiguous. The intuition is the same as in Section 5.1.

Consider now the case where a regulator can control the quantity of multi-homers who allow tracking β , but not the advertising capacities. Again, the analysis follows the same steps as in Section 5.2. We obtain that the comparison among the second best level of tracking and the equilibrium one depends on the sign of $\frac{\partial q_i^*}{\partial \beta}$ (see (14)). Again, the intuition is the same as in the baseline model, so we refer to Section 5.2 for further discussion.

Proofs. To derive (36), the full set of conditions characterizing consumer demands is

$$\begin{aligned} D_{i,B} = & \Pr\{u_i - \delta q_i - c \geq 0, u_j - \delta q_j - c < 0, u_i - \delta q_i - c > u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta, \\ & c < \theta \leq c + z(q_1 + q_2)\}, \quad i, j = 1, 2 \text{ and } i \neq j. \end{aligned}$$

However, condition $u_j - \delta q_j - c < 0$ is implied by $u_i - \delta q_i - c > u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta$. To see this, consider that the last expression can be rewritten as $u_j - \delta q_j < \theta - c - z(q_1 + q_2)$. The right hand side of this expression is smaller than c , because $\theta - c - z(q_1 + q_2) < c \Leftrightarrow \theta < 2c + z(q_1 + q_2)$. The latter condition holds because we are assuming that $\theta \leq c + z(q_1 + q_2)$. Therefore, $u_j - \delta q_j < c$. Hence, we can write $D_{i,B}$ as in (36).

To derive (37), the full set of conditions characterizing consumer demands is

$$\begin{aligned} D_{i,NB} = & \Pr\{u_i - \delta q_i - \theta \geq 0, u_j - \delta q_j - \theta < 0, u_i - \delta q_i - \theta > u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta, \\ & \theta \leq c\}, \quad i, j = 1, 2 \text{ and } i \neq j. \end{aligned}$$

However, condition $u_j - \delta q_j - \theta < 0$ can be ignored because it is implied by $u_i - \delta q_i - \theta > u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta$. To see this, consider that the last expression can be rewritten as $u_j - \delta q_j < -z(q_1 + q_2)$. The right hand side of this expression is smaller than θ , because $\theta > 0$. Therefore, $u_j - \delta q_j < \theta$. Hence, we can write $D_{i,NB}$ as in (37).

Using the demand system in (35)-(37), we can write the consumer surpluses, see (38). The first term, CS_B , refers to individuals who block tracking. We have

$$\begin{aligned}
CS_B = & \int_{c+z(q_1+q_2)}^{\bar{\theta}} \int_{\delta q_1+c}^{\infty} \int_0^{\delta q_2} (u_1 - \delta q_1 - c) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_{c+z(q_1+q_2)}^{\bar{\theta}} \int_0^{\delta q_1} \int_{\delta q_2+c}^{\infty} (u_2 - \delta q_2 - c) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_{c+z(q_1+q_2)}^{\bar{\theta}} \int_{\delta q_1}^{\delta q_1+c} \int_{\delta q_2+c}^{\infty} (u_1 + u_2 - \delta(q_1 + q_2) - c) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_{c+z(q_1+q_2)}^{\bar{\theta}} \int_{\delta q_1+c}^{\infty} \int_{\delta q_2}^{\delta q_2+c} (u_1 + u_2 - \delta(q_1 + q_2) - c) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_{c+z(q_1+q_2)}^{\bar{\theta}} \int_{\delta q_1+c}^{\infty} \int_{\delta q_2+c}^{\infty} (u_1 + u_2 - \delta(q_1 + q_2) - c) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_c^{c+z(q_1+q_2)} \int_{\delta q_1+c}^{\infty} \int_0^{\theta-c-z(q_1+q_2)+\delta q_2} (u_1 - \delta q_1 - c) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_c^{c+z(q_1+q_2)} \int_0^{\theta-c-z(q_1+q_2)+\delta q_1} \int_{\delta q_2+c}^{\infty} (u_2 - \delta q_2 - c) h(u_1, u_2) du_2 du_1 g(\theta) d\theta.
\end{aligned}$$

The first five terms capture the surplus of those whose preference for privacy is so strong that they block tracking regardless of which content they browse. Specifically, the first two terms capture the surplus of these consumers who single home, whereas the remaining terms capture the surplus of multi-homers who block tracking. The second term in (38), CS_{NB} , represents the surplus of

consumers who allow tracking:

$$\begin{aligned}
CS_{NB} = & \int_c^{c+z(q_1+q_2)} \int_{\theta-c-z(q_1+q_2)+\delta q_1}^{\delta q_1+c} \int_{\delta q_2+c}^{\infty} (u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_c^{c+z(q_1+q_2)} \int_{\delta q_1+c}^{\infty} \int_{\theta-c-z(q_1+q_2)+\delta q_2}^{\delta q_2+c} (u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_c^{c+z(q_1+q_2)} \int_{\delta q_1+c}^{\infty} \int_{\delta q_2+c}^{\infty} (u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_0^c \int_{\delta q_1+\theta}^{\infty} \int_0^{\delta q_2-z(q_1+q_2)} (u_1 - \delta q_1 - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_0^c \int_{\delta q_1+\theta}^{\infty} \int_0^{\delta q_2-z(q_1+q_2)} (u_2 - \delta q_2 - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_0^c \int_{\delta q_1+c}^{\infty} \int_{\delta q_2-z(q_1+q_2)}^{\delta q_2+\theta} (u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_0^c \int_{\delta q_1-z(q_1+q_2)}^{\delta q_1+\theta} \int_{\delta q_2+c}^{\infty} (u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta + \\
& \int_0^c \int_{\delta q_1+c}^{\infty} \int_{\delta q_2+c}^{\infty} (u_1 + u_2 - (\delta - z)(q_1 + q_2) - \theta) h(u_1, u_2) du_2 du_1 g(\theta) d\theta.
\end{aligned}$$

The first three terms in this expression capture the surplus of consumers who allow tracking only when multi-homing. The remaining terms capture the surplus of consumers whose preference for privacy is so low that they allow tracking regardless of which content they browse.