Ad networks, consumer tracking, and privacy*

Anna D'Annunzio[†] Antonio Russo[‡]

This version: February 15, 2017

Abstract

We study the role of ad networks and consumer tracking in the online advertising market, in presence of multi-homing consumers and advertisers. We consider two ad-supported publishers, selling ads to advertisers either directly or via an ad network. Differently from publishers, the ad network can track consumers across websites using third-party cookies. Tracking increases revenues from multi-homers. We show that, when the sale of ads is outsourced to the ad network, the ad level on a publisher may decrease compared to the case where publishers compete. In particular, ad level can decrease with the effectiveness of tracking. The ability of the ad network to track depends on the (endogenous) choice of consumers to block tracking. We find that blocking by consumers may be too high or too low in equilibrium. In fact, when consumers decide to block tracking, they do not internalize the externalities they impose on other consumers (by affecting publishers' incentives to advertise) and on publishers (by affecting advertising revenues). If ad levels decrease with the effectiveness of tracking, there is always too much blocking at equilibrium. This analysis implies that a privacy policy that reduces tracking by reducing the cost of blocking cookies may have adverse effects on welfare and consumer surplus.

Keywords: two-sided markets, advertising, ad network, Internet, tracking, multi-homing, privacy.

JEL Classification: D43, D62, L82, M37

*An updated version of the paper is available at https://sites.google.com/site/annadannunzioresearch/research. We thank Stefan Behringer, Jan Cauwer, Alessia Russo, Marianne Verdier, and the participants at Telenor Research seminar (Oslo), NHH seminar (Bergen), CSEF seminar (Naples), the Economics of IP, Software and Internet Conference 2016 (Toulouse), EARIE 2016 (Lisbon), the Media Economics Workshop 2016 (Zurich), and the 6th EIEF-UNIBO-IGIER Bocconi workshop on IO for useful comments and discussions. All errors are ours alone. The views expressed in this paper are those of the authors only and not those of Telenor.

†Telenor Research

‡ETH Zurich and CESifo

1 Introduction

An increasingly large share of advertising revenue is generated online: the International Advertising Board estimates that the US online advertising market was worth about 90% of the TV market in 2015 (IAB, 2015). Online advertising has many distinguishing features from traditional markets. In this paper, we focus on two aspects which, in our view, have received too little attention in previous work. The first is that several publishers outsource their inventory to advertising networks, such as Google's AdSense. The second is that the effectiveness of digital advertising hinges on publishers' ability to monitor the behavior of individual consumers and, in turn, on consumers' privacy preferences. Our objective is to provide a theory of the online advertising market that incorporates the role of ad networks and of consumers' preferences for privacy.[1]

One of the key factors behind the growing importance of ad networks is that consumers commonly consult several online contents in a short time frame. This fact makes it difficult for advertisers to plan a campaign, because an ad placed on multiple publishers can hit the same consumer too many times.[2] Ad networks address this problem in at least two ways. First, they centralize the sale of advertising space on different publishers. Furthermore, they use data-driven tracking technologies to improve the reach of campaigns.[3] For example, by using third-party cookies, ad networks are able to keep track of which ads a consumer has already

[1] When evaluating the Google-DoubleClick merger in 2008, the European Commission stated that "[t]oday, roughly 60-75% of all display ad space is sold directly and 25-40% is sold through intermediation. However, the market shows a trend towards an increase of intermediation, partly resulting from improved targeting technology leading to better monetization of the inventory". Roesner, Kohno, and Wetherall (2012) find that users across approximately 89% of the (Alexa) top 500 sites include at least one cross site tracker and 40% are tracked by Google ad network.

[2] Although each publisher can monitor a consumer on its own web-pages, and thus potentially avoid internal repetition, it can hardly do so when the consumer visits other content. The consequences of this aspect have been studied by several recent papers, including Calvano and Jullien, 2012; Ambrus, Calvano and Reisinger, 2016; Anderson, Foros and Kind, forthcoming; Athey, Calvano and Gans, forthcoming.

[3] As noted by Yuan, Wang, and Zhao (2013), "advertisers traditionally set a frequency cap on their campaigns (the maximum times to display the ad to the same user)". Also, advertisers may want to show ads to consumers 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.

been exposed to, avoiding wasteful repetition.[4] However, privacy-concerned consumers can take measures to avoid being tracked, at least partially. They can, for instance, adjust their browser settings to block third-party cookies.

The issues outlined above raise several interesting questions. Do ad networks expose consumers to more advertising than competing publishers would? How does the ad network's tracking ability affect the advertising level on the Internet? What is the impact of the ability of consumers to avoid tracking on consumer surplus and welfare? Which privacy policy is desirable when the effects on the advertising market are considered? Our paper provides an analytical framework to tackle these questions. We show that, despite its centralizing function, an ad network may bring to higher advertising levels than when publishers compete directly. Also, improved tracking effectiveness does not necessarily lead to higher advertising levels. We also show that consumers may have excessive incentives to avoid being tracked, henceforth reducing consumer surplus and welfare.[5]

We consider a market with two ad-financed online publishers, an ad network, and homogeneous advertisers who intend to reach consumers. Advertisers and consumers possibly multi-home. In the first part of the analysis, we assume the ad network's tracking capability is exogenous. We compare the case where publishers compete directly to that where they outsource advertising to the ad network, and identify two main differences. First, because it centralizes the sale of ads, the ad network internalizes the effects of advertising on audience composition. Namely, if a single-homer is more valuable than a multi-homer on a publisher, then the ad network increases the level of advertising so as to lose multi-homers and gain single-homers. By contrast, each publisher only takes into account that multi-homers become exclusive consumers of the other platform. Second, cross-publisher tracking reduces the quantity of wasted impressions. Interestingly, the effect of tracking on advertising levels

[4] Cookies are pieces of code placed on a device or browser while a consumer visits a website, that can be used to identify her afterward. First-party cookies are issued by the website, whereas third-party cookies are generally issued by an ad network with the objective of monitoring the consumer across several websites.

[5] In reality, of course, ad networks also perform targeting functions, by matching advertisements with consumers based on information regarding their interests. However, in order to sharpen the focus of our investigation, we concentrate on tracking technologies in this paper.

is ambiguous. The intuition is that tracking increases the marginal revenue from ads hitting multi-homers, but also the revenue from each multi-homer. Because the quantity of multi-homers shrinks when advertising levels go up, the two effects are countervailing. When the latter dominates, more effective tracking reduces the equilibrium quantity of ads. To the best of our knowledge, this finding uncovers a novel effect of data-driven advertising technologies on advertising levels..

There exists an important policy debate on how to regulate collection and storage of consumers information. To inform this debate, we study the link between privacy regulation and the advertising market. In the second part of the paper we extend the model to account for the fact that, in reality, the ability of ad networks to track consumers depends also on the extent to which consumers choose to protect their privacy.[6] As mentioned, consumers may take steps to avoid being tracked while browsing the Internet, for instance by blocking third-party cookies. By so doing, consumers reduce the ability of the ad network to avoid cross-outlet repetition of impressions. In our model, users block third-party cookies if the disutility from being tracked is bigger than the cost of blocking (capturing, for instance, the effort of installing a plug-in or the cost of buying an effective software). The ad network is able to track only consumers that choose not to block. Therefore, the quantity of consumers that block cookies affects the advertising revenue and, in turn, the equilibrium quantity of ads. We begin by comparing the laissez-faire equilibrium to the first-best allocation, whereby the quantity of ads and of consumers blocking are chosen to maximize social welfare. We find that the equilibrium quantity of consumers that block is too high. The intuition is straightforward: consumers block without taking into account the effect on advertising revenue. However, the equilibrium quantity of advertising may be higher or smaller than the first-best one. Next, we turn to the comparison between laissez-faire and a second-best allocation where the quantity of consumers that block is chosen to maximize welfare, but

[6] Many market analysis find that consumers are privacy concerned, as, e.g., Tucker (2014). Also, Turow 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.

advertising levels are unregulated. We find that, from a second-best point of view, the level of tracking may be too low. This occurs because there are two externalities that consumers produce by blocking. First, blocking affects other consumers by affecting the equilibrium quantity of ads. Furthermore, as mentioned, blocking reduces advertising revenues.

These results have interesting policy implications. To begin, they suggest that a policy intervention promoting blocking of third-party cookies is not always beneficial for social welfare. In reality, regulators can choose to reduce the cost of blocking, for instance imposing that each browser does it by default or creating a "do not track mechanism" for online advertising.[7] Furthermore, regulators can reduce the disutility from intrusion, for instance banning the most intrusive types of cookies.[8] According to our model, these two policies have comparable effects on consumers' disutility from being tracked, but a different effect on the advertising market (provided that less intrusive cookies are informative enough for the ad network). In particular, if the advertising level decreases with the tracking ability of the ad network, then the interests of consumers and publishers are aligned, and both are better off when the tracking effectiveness increases. In this case, a policy intervention that reduces the cost of blocking third party cookies decreases social welfare and consumer surplus. By contrast, if the advertising level highly increases with the tracking effectiveness, than the interests of consumers and publishers are misaligned, and it may be desirable to decrease the cost of blocking cookies.

The remainder of this paper is organized as follows. Section 2 presents a brief literature review. In Section 3 we describe the baseline model and we solve it in Section 4. Section 5 studies welfare effects. Section 6 contains the extension where we endogenize the consumers' decision to block third-party cookies. Section 7 concludes. Unless otherwise stated, Proofs of Propositions and Lemmas are in Appendix A.

[7] In the US, the FTC discussed the introduction of a "do not track mechanism", with the objective of making inexpensive for consumers to avoid intrusive advertising.

[8] For instance, some cookies can store more personal information on consumers than others, and may be more difficult to remove from a computer.

2 Literature

Many papers have studied the provision of advertising in the media industry. Anderson and Coate (2005) and Peitz and Valletti (2008) study the provision of advertising in a two-sided market where consumers single-home. A more recent strand of literature takes into account multi-homing by consumers. Anderson, Foros and Kind (forthcoming) and Ambrus, Calvano and Reisinger (2016) study how multi-homing by consumers affects platforms' choices, entry and mergers in the sector. Anderson, Foros and Kind (forthcoming) find that a merger of two platforms increases ad prices. Ambrus, Calvano and Reisinger (2016) show that the equilibrium advertising level in duopoly is lower than under joint ownership when the per-viewer benefit from a single-homer is bigger than that from a multi-homer. In the present paper, we also consider multi-homing by consumers and we account for the fact that repeated impressions are partially wasted. Differently from the previous papers, we analyze how the presence of an ad network to which platforms can outsource the sale of ads and that can track consumers affects the provision of ads in the market.

The literature on online advertising has studied targeted advertising (see, e.g., Bergemann and Bonatti, 2011). However, the analysis of other data-driven technologies used in online advertising markets is rather limited. In this paper we concentrate on the analysis of tracking technologies. Both tracking and targeting allow more efficient advertising, but while targeting allows to improve the match between ads and consumers, tracking allows to manage more efficiently the reach and the frequency of exposure of an advertising campaign. Athey, Calvano and Gans (forthcoming) consider publishers which can track consumers internally, but they are not able to track consumers when they switch. With heterogeneous advertisers, as more consumers multi-home, more advertisers single-home. In general, having a larger set of single-homers allows a publisher to gain competitive advantage in the advertising market. Differently from them, we consider the effects of cross-publishers tracking by an ad network on the advertising level. Our analysis is complementary to that in those paper:

across publishers tracking increases the value of multi-homers for advertisers.

Some papers have analyzed the presence of a common distributors on one side of the market. Kind, Nilssen and Sørgard (2016) consider a common distributor of TV channels to which platforms can delegate the choice over the subscription price for viewers. George and Hogendorn (2012) study news aggregators. Differently from us, they consider an aggregator on the consumer market, implying that the role of the aggregator is different in those models. Also, de Cornière and Taylor (2014) study a model where a search engine allows consumers to access content provided by two publishers. When the search engine controls one publisher, ad volumes fall compared to the case when all firms are independent. Differently from us, they consider only a partial coordination in the advertising market.

In this paper we also study consumers' choice to block third-party cookies and the effects this choice has in the advertising market. A relevant strand of literature has studied targeting and advertising avoidance technologies. Johnson (2013) shows that there may be too little use of advertising-avoidance technologies from a second-best social standpoint. This is because consumers do not take into account that by blocking advertising they discourage advertising, which may benefit other consumers. Differently from Johnson (2013), we introduce intermediaries among consumers and advertisers. Moreover, we consider blocking of third-party cookies and not of ads by consumers: consumers decision only indirectly affect the level of advertising they are exposed to. Ad avoidance and third-party cookies blocking have different effects, for instance on the level of advertising displayed to consumers.

Our paper studies the link between privacy policies and the advertising market (see Tucker, 2012). There are few papers on this topic. Goldfarb and Tucker (2011) study how the Privacy Directive (2002/58/EC), which restrict the ability of advertisers to collect and use data for targeting purposes, decreased the effectiveness of advertising in Europe relative to other geographical markets. The effect was more pronounced for websites showing general content. Campbell, Goldfarb, and Tucker (2015) study the effect of privacy regulation on the market structure, finding that forced consent-gathering from consumers will have nega-

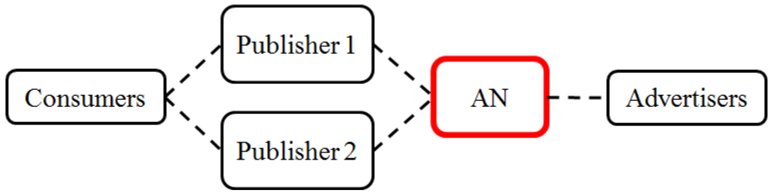

Figure 1: Market structure.

tive effects on smaller firms. We analyze a different issues, concentrating on the link among advertising and privacy, when cookies are used for tracking.

3 Model setup

We consider a market with two publishers, indexed by $i = 1, 2$, and one ad network, denoted by AN. Publishers provide free content to consumers and sell ads to advertisers. Figure 1 provides an illustration. We study two scenarios: in the first one, publishers compete directly in the advertising market; in the second one, publishers outsource the sale of ads to the ad network. Our baseline setup closely follows that of Ambrus, Calvano and Reisinger (2016), which we henceforth refer to as ACR for short. ACR do not consider the ad network.

Consumers. There is a unit mass of consumers. Each consumer can 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. These demands are formally described as follows

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

where u_i is the reservation utility from browsing content on publisher i. We assume that $(u_1, u_2) \sim h\left(u_1, u_2\right)$, a bivariate joint distribution with smooth density. The parameter δ

captures the disutility generated by an ad and q_i is the quantity of ads a consumer is exposed to when visiting publisher i. We assume that demands satisfy the necessary conditions to ensure existence and uniqueness of an equilibrium with interior solutions.[9] The demand functions in (1) have the following properties. When q_i increases, publisher i loses multi-homers, who become single-homers on publisher j, i.e. $\frac{\partial D_{12}}{\partial q_i} < 0$ and $\frac{\partial D_j}{\partial q_i} > 0$. Hence, when q_i changes, the composition of j's audience changes. However, j's total demand does not change, i.e. $\frac{\partial (D_{12}+D_i)}{\partial q_j} = 0$. Total consumer surplus is given by

$$\begin{aligned} CS_c = & \int_{\delta q_1}^{\infty} \int_0^{\delta q_2} \left(u_1 - \delta q_1\right) h\left(u_1, u_2\right) du_2 du_1 + \int_0^{\delta q_1} \int_{\delta q_2}^{\infty} \left(u_2 - \delta q_2\right) h\left(u_1, u_2\right) du_2 du_1 \\ & + \int_{\delta q_1}^{\infty} \int_{\delta q_2}^{\infty} \left(u_1 - \delta q_1 + u_2 - \delta q_2\right) h\left(u_1, u_2\right) du_2 du_1. \end{aligned} \tag{2}$$

Advertisers. There is a unit mass of homogeneous advertisers. In our framework advertising is informative, and advertisers fully extract the information benefit from advertising as, e.g., in Anderson and Coate (2005) and in Crampes, Haritchabalet, and Jullien (2009). Following ACR, we decompose the advertisers' revenue in several components, each related to a subset of viewers. Consider, to begin, the case where the advertiser buys m_i impressions only from publisher i. We denote by $r_i\left(m_i\right)$ (resp. $\hat{r}_i\left(m_i\right)$) the expected revenue the advertiser obtains from a single-homing (resp. a multi-homing) consumer. These revenues are given by the probability that the consumer is informed about the advertiser's product on publisher i, times the return from informing her. Thus, a single-homing advertiser's total revenue is

$$r_i\left(m_i\right) D_i + \hat{r}_i\left(m_i\right) D_{12},\ i = 1, 2. \tag{3}$$

To capture the fact that there are diminishing returns to advertising, we assume that $\frac{\partial r_i}{\partial m_i} > 0$, $\frac{\partial^2 r_i}{\partial m_i^2} < 0$, $\frac{\partial \hat{r}_i}{\partial m_i} > 0$, and $\frac{\partial^2 \hat{r}_i}{\partial m_i^2} < 0$.[10]

[9] See Vives (2000) and ACR (see footnote 15) for a discussion of these conditions.

[10] Diminishing 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 the product than inframarginal ones. See, e.g., Bagwell (2007) and Renault (2015).

Consider now the case where the advertiser buys m_i impressions on both publishers $i = 1, 2$. The expected revenue from single-homers is again equal to $r_i(m_i)$, because they are exposed only to impressions on publisher i. The expected revenue from a multi-homer is $r_{12}^h(m_1, m_2, \beta)$, given by the probability of informing her on some outlet, times the subsequent return. Thus, a multi-homing advertiser's total revenue is

$$r_1(m_1) D_1 + r_2(m_2) D_2 + r_{12}^h(m_1, m_2, \beta) D_{12}. \tag{4}$$

Again, we assume that $r_{12}^h(.)$ is increasing and concave in m_i, $i = 1, 2$. We also assume that $\frac{\partial^2 r_{12}^h}{\partial m_i \partial m_j} \leq 0$, for $i \neq j$ and $i, j = 1, 2$. This 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 (ACR, and de Cornière and Taylor, 2014). Hence, the relation $\hat{r}_i \leq r_{12}^h \leq \hat{r}_1 + \hat{r}_2$ holds.

The revenues from a multi-homing consumer when advertisers multi-home are affected by tracking. Parameter β represents the degree of cross-outlet tracking. Upon identifying a consumer already informed on another outlet thanks to tracking technologies (for instance, because it has observed whether she has bought the advertised product or clicked on the ad), a publisher may decide to expose her to an impression from a different advertiser she has not already seen, or to a different message from the same advertiser (e.g., to an ad containing a specific offer on a certain product the consumers is interested to, as in the case of retargeting, or information about a different product).[11] Formally, we assume that tracking has three effects. First, it increases the value of inframarginal impressions, i.e. $\frac{\partial r_{12}^T}{\partial \beta} > 0$. Thus, it increases the revenue from each multi-homer. Second, tracking 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

[11] Evidence suggests that ad campaigns that send ad messages in a sequence to tell a brand story before asking to buy a product are more effective than ad campains focused only on driving a purchase (see, e.g. https://scontent-arn2-1.xx.fbcdn.net/t39.2365-6/10333119_1457635661160496_168768318_n.pdf). Also, Chandra and Kaiser (2014) find that, as the share of multi-homing consumers increases, advertisers may see offline and online companion webistes as complements, because they may allow cross-media advertising campains.

value of a marginal impression on a multi-homer on publisher i as the number of impressions on j increases, i.e. $\frac{\partial^3 r_{12}^T}{\partial m_i \partial m_j \partial \beta} > 0$. Thus, the two publishers become less substitutable as the precision of tracking increases.

In reality, publishers are relatively effective at monitoring the behavior of consumers on their own websites, but have limited ability to do so on other outlets. However, by operating on multiple websites, ad networks perform the latter function with much greater effectiveness, e.g. by exploiting third-party cookies. To capture these differences in a simple way, we assume that there is no cross-outlet tracking unless both publishers outsource to the ad network. We use superscript $h = T$ when the ad network is able to track consumers (i.e., $\beta > 0$), while $h = nT$ when there is no cross-publisher tracking (i.e. $\beta = 0$). 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). We assume that revenues from single-homers are unaffected by cross-outlet tracking (i.e., the ad network cannot track single-homers more effectively than the publishers).

Finally, note that we treat the degree of cross-outlet tracking β as exogenous in the first part of the analysis. We endogenize this parameter in section 5.

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 *total* advertising level on publisher i. We refer to q_i as publisher i's "advertising capacity", and assume it is chosen at an initial stage (see below): it cannot be exceeded by the ad spaces sold to advertisers in the following stages, in the spirit of Kreps and Scheinkman (1983). For simplicity, we assume that i exposes each consumer to the same quantity of impressions.[12] We assume publishers sustain no costs.

In the case where they sell impressions directly, we follow ACR in assuming that they offer advertisers a menu of contracts that specify an advertising intensity m_i, i.e. the average

[12] More explicitly, we assume that all consumers visit the same number of web-pages on a publisher, and that 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 (although the type of impressions may not be).

quantity of impressions on a consumer visiting i's pages, in exchange for a payment p_i. We denote publisher i's profit function as π_i^C.

In the case where a publisher outsources advertising to AN, it relinquishes control of its ad inventory and of all the ensuing revenue, but receives a transfer x_i from AN. Publisher i's profits are denoted by π_i^A. AN offers a a menu of contracts to advertisers specifying advertising intensities (m_1, m_2) in exchange for a payment p_{AN}. We denote its profit by π_{AN}. We characterize the contract between AN and the publishers below.

Social welfare. We define welfare as the sum of consumer surplus (2) and advertiser surplus, plus publishers' and ad network's profits. Because all payments collected by the latter are transfers from other parties, welfare boils down to the sum of consumer surplus and gross advertiser surplus $AS \equiv \sum_i r_i D_i + r_{12}^h D_{12}$. Hence,

$$W = CS_c + AS. \tag{5}$$

Timing. The timing of the game is as follows. In the scenario where publishers compete directly, at stage 1 they simultaneously choose the advertising capacity q_i, $i = 1, 2$. At stage 2, consumers observe (q_1, q_2) and decide which publishers to join, if any. At stage 3, publishers simultaneously offer a menu of contracts (p_i, m_i) to advertisers. The advertisers decide which contract to accept, if any. At stage 4, consumers get exposed to ads and all payoffs are realized. In the scenario where publishers outsource the sale of ad spaces to the ad network, the timing is identical except for the following differences. There is a preliminary stage 0 where AN offers transfers to each publisher i, in exchange for their advertising inventory. Furthermore, at stage 3, AN (instead of the publishers) offers a menu of contracts (p_{AN}, m_1, m_2) to advertisers.

We solve the model by backward induction, and adopt SPNE as the solution concept.

4 Solving the model

We now solve the game described before. Note that, because stage 2 is fully described by the specification of consumer demands in (1), in the following we focus on the other stages of the game. We use superscripts C, AT, and ANT to denote equilibrium variables in, respectively, the scenario where publishers compete in the advertising market, the scenario where publishers outsource to AN that tracks consumers across outlets and the scenario where publishers outsource to AN but the latter does not track consumers across outlets.

The objective of the analysis is to understand how an AN and the intensity of tracking affects the supply of advertising.

4.1 Publishers compete directly

To set the stage, we consider here the scenario where AN is inactive and the publishers compete directly on the advertising market.

4.1.1 Stage 3

At stage 3 publishers offer a menu of contracts to each advertiser, specifying an advertising intensity m_i in exchange for a payment p_i. Because advertising revenues are increasing in m_i, the publishers fill all the available advertising capacity. Furthermore, due to diminishing returns, publishers divide their capacity equally across all advertisers, offering a single contract in equilibrium. This contract specifies $m_i = q_i$, and p_i equal to an advertisers' incremental revenue from q_i impressions on i's consumers. This revenue is given by the total advertiser surplus when multi-homing minus the advertiser surplus when single-homing on j. Hence, p_i is equal to the return from impressions on i's single-homers, plus the extra return from sending q_i impressions on multi-homers, given that the advertiser already send q_j impressions (per consumer) on platform j. We summarize in the following Lemma.

Lemma 1. *In the scenario where publishers compete directly, any SPNE of the game is such that publisher $i = 1, 2$ offers a single contract to all advertisers, whereby $m_i = q_i$ and*

$$p_i^C = r_i (q_i) D_i (q_i, q_j) + \left(r_{12}^{nT} (q_1, q_2) - \hat{r}_j (q_j)\right) D_{12} (q_1, q_2) . \tag{6}$$

All advertisers accept this contract. Thus, the publisher's profit is $\pi_i^C = p_i^C$.

Observe from (6) that multi-homers create an interdependence between the advertising revenues of the two publishers: the higher the revenue from reaching a multi-homer exclusively on platform j, the less advertisers are willing to pay for impressions on platform i.

4.1.2 Stage 1

Consider now stage 1. The publishers simultaneously decide their advertising capacity, maximizing π_i^C. The system of FOCs is

$$\frac{\partial \pi_i^C}{\partial q_i} = \left[\frac{\partial r_i}{\partial q_i} D_i + \frac{\partial r_{12}^{nT}}{\partial q_i} D_{12}\right] + \left[r_i \frac{\partial D_i}{\partial q_i} + \left(r_{12}^{nT} - \hat{r}_j\right) \frac{\partial D_{12}}{\partial q_i}\right] = 0,\ i, j = 1, 2,\ i \neq j. \tag{7}$$

Each publisher thus chooses q_i such that the marginal revenue from an additional impression (per consumer) is equal to zero. The marginal revenue depends on how a change in advertising levels affects revenues and consumer demands. As q_i increases, the marginal advertising revenue from consumers increases (the first term in brackets of (7) is positive), but the consumer demand for the platform decreases (the second term in brackets of (7) is negative).

4.2 Advertising outsourced to the ad network

We now consider the case where AN manages the ad inventory of both publishers. As will become clear below, in equilibrium both publishers prefer to outsource.

4.2.1 Stages 4 and 3

The ad network offers a menu of contracts to each advertiser, specifying a pair of advertising intensities (m_1, m_2) and a total payment p_{AN}. As in the previous scenario, AN offers a single contract to all advertisers in equilibrium, such that $m_i = q_i, i = 1, 2$. Now, p_{AN} is equal to the *total* revenue the advertiser gains from these impressions. The reason is the ad network is the unique gateway to the audience of both publishers, and is thus able to extract the entire advertiser surplus.

Lemma 2. *When publishers outsource their advertising inventory to AN, any SPNE of the game is such that the latter offers a single contract to all advertisers, whereby $m_i = q_i$ and*

$$p_{AN} = r_i(q_i) D_i(q_i, q_j) + r_j(q_j) D_j(q_i, q_j) + r_{12}^h(q_1, q_2, \beta) D_{12}(q_1, q_2). \tag{8}$$

All advertisers accept this contract. AN's profit is $\pi_{AN} = p_{AN}$.

Now, AN captures all advertising revenues, as discussed above. This implies that the ad network internalizes the effects that additional impressions on a publisher produce on advertisers' willingness-to-pay for impressions on the other. Also, p_{AN} is influenced by the AN's ability to track consumers across outlets. This ability allows it to identify with some accuracy the multi-homers who have already been informed by an ad on one publisher. Because $r_{12}^T > r_{12}^{nT}$, tracking raises the willingness-to-pay for ads by increasing the revenue that advertisers obtain from multi-homers. The size of this effect depends on the tracking ability of the ad network, captured by β. The tracking ability β may depend on the technology available to the ad network and/or the extent to which consumers protect their privacy (e.g. by blocking third party cookies).[13]

[13] We explore the implications of this last issue in Section 5 below.

4.2.2 Stages 1 and 0

We now describe the contracting stage between AN and the publishers, and how the latter subsequently choose the respective advertising capacity q_i. At stage 0, the ad network simultaneously makes a take-it-or-leave-it offer to each publisher, in exchange for the revenues that AN generates on the publisher's website. This offer specifies a schedule of prices $x_i(q_i)$ contingent on the q_i ad spaces per consumer outsourced. In equilibrium, both publishers accept the offer if they receive a transfer higher or equal than their outside option.[14] At stage 1 of the game, they simultaneously choose the respective advertising capacity. Conditional on outsourcing to AN, each publisher chooses q_i maximizing $x_i(q_i)$. We denote the resulting advertising capacities as (q_1^e, q_2^e). Consider now stage 0. Because q_i^e depends only on the schedule $x_i(.)$, the ad network can implement any value of q_i^e by appropriately choosing this price schedule. However, $x_i(.)$ must be such that, when $q_i = q_i^e$, T_i is at least equal to i's outside option. The latter consists in the revenue obtained when rejecting AN's offer (given that the other publisher accepts it), which we denote by π_i^o. Therefore, at stage 0 the ad network solves the following problem

$$\max_{q_i^e, x_i(.)} \quad p_{AN} - T_i, \quad \text{s.t. } T_i \geq \pi_i^o. \tag{9}$$

The solution is such that q_i^e satisfies $\frac{\partial p_{AN}}{\partial q_i} = 0$, and $x_i(q_i^e)$ is such that $T_i = \pi_i^o$.[15]

Thus, the equilibrium advertising capacities, which we denote $\left(q_1^{AT}, q_2^{AT}\right)$, maximize total revenues collected by the ad network, p_{AN}. Using the property that $\frac{\partial D_j}{\partial q_i} = -\frac{\partial D_{12}}{\partial q_i}$, these capacities satisfy the following FOCs

$$\frac{\partial p_{AN}}{\partial q_i} = \left[D_i \frac{\partial r_i}{\partial q_i} + D_{12} \frac{\partial r_{12}^T}{\partial q_i}\right] + \left[\frac{\partial D_i}{\partial q_i} r_i + \frac{\partial D_{12}}{\partial q_i}\left(r_{12}^T - r_j\right)\right] = 0, \; i = 1, 2. \tag{10}$$

[14] Outsourcing by both publishers is pareto-efficient, because AN internalizes the external effects of repeated impressions on different publishers, and because it is able to track consumers across outlets. Hence, in equilibrium AN is able to compensate each publisher for the revenue it loses by relinquishing control of its ad inventory.

[15] Results are equivalent if AN's offer specifies a schedule of per-impression prices $x_i(q_i)$ contingent on the q_i.

A first key difference with respect to the competitive scenario is that AN fully internalizes the impact of a change in the advertising level on market composition, hence on advertiser surplus. In particular, as q_i increases, some multi-homers disconnect from publisher i, becoming single-homers on j. Because the ad network controls the sale of impressions on all consumers, it takes this effect into account, as testified by the last term in the second bracket in (10). We analyze the implications of this internalization for the advertising market below.

In addition, the advertising network is able to track consumers across publishers. Specifically, tracking has two effects. On the one hand, it raises the marginal revenue from multi-homers $\frac{\partial r_{12}^T}{\partial q_i}$, thereby inflating the first term in brackets of (10). On the other hand, tracking increases the loss produced by losing multi-homers as q_i increases, thereby making the second term in brackets smaller. Hence, intuition suggests that the net impact of tracking on the advertising market is ambiguous. Using the Implicit Function theorem, we look at how a marginal change in the level of tracking affects the advertsing level in equilibrium. Determining the sign of this effect is generally rather complex, because it depends on how impressions on one publisher affect the marginal revenues generated on the other. Nevertheless, by focusing on the special case where publishers are symmetric, it is possible to establish an intuitive result.

Proposition 1. *Assume that publishers are symmetric. In the scenario where they outsource to the ad network, the equilibrium advertising level on each publisher q_i^{AT} increases with the tracking effectiveness β if and only if*

$$D_{12}\frac{\partial^2 r_{12}^T}{\partial q_i \partial \beta} + \frac{\partial D_{12}}{\partial q_i}\frac{\partial r_{12}^T}{\partial \beta} > 0. \tag{11}$$

Proposition 1 suggests, interestingly, that advertising levels q_i^{AT} do not necessarily increase with the ad network's tracking capability. As anticipated, tracking raises the profitability of a marginal impression on a multi-homer, i.e. $\frac{\partial^2 r_{12}^T}{\partial q_i \partial \beta} > 0$, but it also increases the value of the infra-marginal impressions, i.e. $\frac{\partial r_{12}^T}{\partial \beta} > 0$. In words, tracking increases the value of advertising on shared consumers, but the size of this group decreases with the level of advertising. Therefore, tracking raises the cost of losing multi-homers when the quantity of impressions increases (recall that $\frac{\partial D_{12}}{\partial q_i} < 0$). Intuitively, when the second effect dominates, the ad network prefers to reduce the volume of advertising on both websites at the margin. Thus, q_i^{AT} decreases with β.

To the best of our knowledge, our paper insulate a new effect of data-driven advertising technologies, and of cross-outlet tracking in particular. Specifically, we assess that the impact of tracking effectiveness on the quantity of advertising depends on how tracking increases the value of advertising on shared consumers and on how multi-homers respond to advertising. Previous papers have studied the effects of targeting technologies on the supply 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. These results are distinctive of targeting technologies, that increase the probability of a good match among consumers and advertisers. Furthermore, they do not account for competition between different publishers, because they either model consumer demands as exogenous (e.g., Athey and Gans, 2010) or ignore publishers altogether (e.g., Johnson, 2013). Our results suggest these effects should be taken into account in order to understand the net effect of data-driven technology on the provision of advertising.

4.3 The ad network's impact on the level of advertising

The objective of this section is to compare the equilibrium advertising levels when the publishers compete directly and when they outsource to the ad network. For illustrative purposes, we proceed in steps. To begin, we assume the ad network is unable to track consumers cross-outlets (i.e., $\beta = 0$). This assumption allows us to concentrate on one of the roles played by the ad network, namely that of centralizing the sale of ads on multiple outlets. The following result emerges.

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

$$r_j(q_j) > \hat{r}_j(q_j). \tag{12}$$

The result in Lemma 3 is identical to that of ACR (see Proposition 2). The reason is that, in our framework, the ad network is able to extract the whole advertising surplus. Hence, outsourcing to an ad network generates similar effects as a merger. The comparison of (7) and (10) suggests that, when $\beta = 0$, the incentives of the ad network and of publisher i differ only due to how a change in q_i affects revenues from the marginal multi-homer. As q_i increases, some multi-homers switch and become single-homers on j. Because the ad network sells impressions on both publishers, this switch produces a net loss equal to $r_{12}^{nT} - r_j$. That is, the difference between the revenue generated from a multi-homer (without tracking) and a single-homer on j. By contrast, when publishers compete head-to-head, each multi-homer switching imposes a loss equal to $r_{12}^{nT} - \hat{r}_j$ to publisher i. In fact, each publisher is only able to extract the incremental revenue generated by multi-homers (see (6)). In sum, the ad network expands the supply of ads if and only if the switch imposes a higher loss in case of competing publishers, that is when $r_{12}^{nT} - r_j < r_{12}^{nT} - \hat{r}_j$, that simplifies to (12). In the following, we refer to this effect on the equilibrium advertising levels, captured by (12), as the "joint control" effect.

Let us now turn to cross-outlet tracking. To isolate its effect, we make use of (10), and compare the case where the ad network is able to track consumers ($\beta > 0$) to that where it is not ($\beta = 0$). We get the following result.

Lemma 4. *At equilibrium, the supply of ads when the ad network is able to track consumers exceeds that in the case when it is not able to do so if and only if*

$$D_{12}\left(\frac{\partial r_{12}^{T}}{\partial q_i}-\frac{\partial r_{12}^{nT}}{\partial q_i}\right)>-\frac{\partial D_{12}}{\partial q_i}\left(r_{12}^{T}-r_{12}^{nT}\right), \tag{13}$$

where all functions are evaluated in $\left(q_1^{ANT}, q_2^{ANT}\right)$.

Because tracking reduces the waste from repeated impressions, it produces two countervailing effects. First, the revenue from additional ads on infra-marginal multi-homers increases (see the left hand side of (13)). Second, tracking increases the total revenue from impressions on multi-homers. Given that the latter react to an increase in the level of advertising by becoming single-homers, this effect discourages the publishers from raising q_i (see the right hand side of (13)). In line with our previous findings (see Proposition 1), the equilibrium advertising levels may be lower than in the absence of tracking, as long as the second effect dominates. In the following, we refer to this combined effect on the equilibrium advertising levels, captured by (13), as the "tracking effect".

Summing up, the fully-fledged comparison of the equilibrium levels of advertising when publishers compete and with outsourcing depends on the magnitudes of the joint control and the tracking effects. Using (10), we conclude that:

Proposition 2. *At equilibrium, the supply of ads when the ad network is active exceeds that in the case where publishers compete directly if and only if*

$$D_{12}\left[\frac{\partial r_{12}^{T}}{\partial q_i}-\frac{\partial r_{12}^{nT}}{\partial q_i}\right]>-\frac{\partial D_{12}}{\partial q_i}\left[r_{12}^{T}-r_{12}^{nT}+\hat{r}_j-r_j\right], \tag{14}$$

where all functions are evaluated in $\left(q_i^{C}, q_j^{C}\right)$.

Expression (14) puts together the two effects identified previously. The left hand-side captures the extra marginal revenue from infra-marginal multi-homers that comes with tracking.

The right hand side captures the difference in forgone revenues from losing multi-homers, which are due to both tracking, i.e. $r_{12}^{T} - r_{12}^{nT}$, and joint control, i.e. $\hat{r}_j - r_j$. Overall, Proposition 2 suggests that the effect of the ad network on the level of advertising is ambiguous, because the forces we just described may push in opposite directions. The sign of this effect may also depend ultimately on the strength of the ad network's tracking ability, captured by the parameter β.

4.4 Welfare analysis

To conclude the first part of the analysis, we compare welfare levels in the different scenarios considered above. As a first step, we characterize the socially optimal advertising levels, which satisfy the following condition:

$$\frac{\partial W}{\partial q_i} = \frac{\partial CS_c}{\partial q_i} + \frac{\partial AS_c}{\partial q_i} = \frac{\partial CS_c}{\partial q_i} + \left[D_i \frac{\partial r_i}{\partial q_i} + D_{12} \frac{\partial r_{12}^h}{\partial q_i} + \frac{\partial D_i}{\partial q_i} r_i + \frac{\partial D_j}{\partial q_i} r_j + \frac{\partial D_{12}}{\partial q_i} r_{12}^h \right] = 0,\ i, j = 1, 2, i \neq j, \tag{15}$$

where $h = T$ when $\beta > 0$, and $h = nT$ otherwise. As discussed previously, the ad network fully internalizes the effects of changes in advertising levels on total advertiser surplus. Indeed, the terms in square brackets in (15) are the same as (10) (taking into account that $\frac{\partial D_j}{\partial q_i} = -\frac{\partial D_{12}}{\partial q_i}$). But the ad network does not account for consumer surplus, that is decreasing in the advertising level. Hence, although gross advertiser surplus is maximized, advertising is overprovided when it is outsourced to AN. However, it is harder to establish whether advertising is over- or underprovided when publishers compete head-to-head. As indicated by Proposition 2, publishers may adopt advertising levels either lower or higher than when the ad network is involved. Hence, even if gross advertising surplus is not maximized, consumers may be either worse or better off, because their surplus decreases with the quantity of impressions. Specifically, they are better off when publishers compete if and only if $q_i^{AT} \geq q_i^{C}$ holds. Summarizing, we conclude the following:

Proposition 3. *Advertising is overprovided when outsourced to the ad network, but may be under- or overprovided when publishers compete directly. Furthermore, if* $q_i^{AT} < q_i^C$ *consumer, advertiser and social welfare are higher when advertising is outsourced than when it is not. However, if* $q_i^{AT} \geq q_i^C$ *, the comparison is ambiguous: consumer (resp. advertiser) surplus is higher when publishers compete directly (resp. outsource to the ad network).*

Also, we analyse the welfare effects of tracking. Proposition 1 establishes that the level of advertising may increase or decrease with the effectiveness of tracking. This implies that consumers may be better off when the effectiveness of tracking increases, because they are exposed to less ads while browsing. Hence, we conclude the following (we omit the proof, that simply follows from Proposition 1):

Proposition 4. *Consumers surplus may increase or decrease with* β*, depending on whether* q_i^{AT} *decreases or increases with* β*.*

5 Privacy concerned consumers

Till now, we have treated the ability β of the ad network to track consumers as exogenous. In general, the tracking ability is due both to the precision of the tracking technology available to the ad network and to consumers' preferences. When publishers outsource the sale of their ads to an ad network, the latter is authorized to follow the consumers on the publishers. Ad networks use third-party cookies to track consumers' activity online. A cookie allows the ad network to know which publishers a consumer visits and they can be used to deliver ads to the consumers. However, consumers can decide to block those cookies using tracker-blocking browser plug-ins.

In this section, we only consider the scenario with an ad network, and we allow consumers to block cookies. The blocking decision is determined by the privacy preferences of the consumers.[16] Consumers have an intrinsic disutility from third-party cookies, that has an economic effect on the advertising market: by blocking third-party cookies, consumers prevent the ad network from following them across publishers. We assume that the ad network has a perfect tracking technology, entailing that its ability to track is entirely determined by consumers' choice to block third-party cookies.[17] We make this assumption to insulate preference-driven effects.

Formally, consumers have a disutility θ from being tracked by the ad network through third-party cookies. We assume that θ is distributed according to the c.d.f. $F(.)$ on $[0,1]$, with density $f(.)$. For simplicity, we assume that θ is independent of (u_i, u_j). Also, θ is independent of the probability of being informed by an ad on a publisher. Consumers can block third-party cookies at a cost $c > 0$. The decision of blocking cookies affects the revenues from multi-homers because tracked and non tracked multi-homers are valued in a different way in the advertising market: on the former, but not on the latter the ad network can avoid wasteful impressions. In Section 3 we have stated that the revenues from multi-homers depend on a parameter β, that should be interpreted in this section as the share of consumers that at equilibrium does not block third-party cookies and are therefore tracked. The total disutility for consumers from third-party cookies is $CS_p = \int_0^\beta \theta f(\theta)\, d\theta + (1-\beta)\, c$. Hence, now total welfare is $W_p = CS_c - CS_p + AS$.

The timing of the game is as described in Section 3, with the only difference that at stage 1, simultaneously with the publishers' choice on the advertising levels, consumers choose whether to block third-party cookies.

[16] Farrell (2012) argues that privacy is both a final good ("Consumers care about privacy in part for its own sake: many of us at least sometimes feel it's just icky to be watched and tracked. [...] 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."). On the economic impact of privacy concerns, see Lohr (2010).

[17] Publishers also use first-party cookies, that allow them to recognize a consumer when she visits again a website. We assume that first-party cookies cannot be blocked, implying that internal tracking is not affected by consumers' decisions. In fact, blocking first-party cookies often compromises the basic functionality of a website, making the visit impossible. Instead, a consumer has many ways of blocking third-party cookies without affecting her experience of the Internet.

By backward induction, the solution of the game is as described in Section 4.2. Now, at stage 1 we also determine the marginal consumer β^{AT} who is indifferent between blocking cookies or not. A consumer chooses to block if the disutility from being tracked is bigger than the cost of blocking, i.e. if $\theta \geq c$. Hence, the indifferent consumer is $\beta^{AT} = c$. This implies that the share $\left(1 - \beta^{AT}\right)$ of consumers who at equilibrium block third-party cookies is equal to $(1 - c)$.

We are interested in understanding how the equilibrium levels of advertising and tracking compare to the optimum. There are three relevant externalities that drive this comparison. First, as we have already seen in Section 5, publishers do not take into account the effect on consumers when they decide the advertising levels. Second, when consumers decide whether to block third-party cookies, they do not consider the effect of their decision on all consumers: by blocking cookies consumers affect the incentives of publishers to advertise, hence the level of ads each consumer is exposed to when visiting a publisher. Third, when consumers decide whether to block third-party cookies, they do not take into account the direct effect on advertisers, because tracking increases advertising effectiveness on multi-homers, hence revenues from advertising. The second and the third externalities come into play when we endogenize the decision of the consumers of blocking cookies. This decision, that is motivated by privacy preferences, has effect in the advertising market, hence on all the players in those market. Depending on how advertising incentives are affected by consumers' blocking decision, the equilibrium may deviate from the optimum in different directions.

5.1 First-best

Now, we compute the first-best. At stage 1 the regulator chooses β and (q_1, q_2) to maximize total welfare $W_p = CS_c - CS_p + AS$. Denoting the first-best as (β^o, q_1^o, q_2^o), we conclude that

Proposition 5. *Compared to the first-best, at equilibrium tracking is too low.*

The first-best is such that

$$\begin{aligned} \frac{\partial W_p}{\partial q_i} &= \frac{\partial CS_c}{\partial q_i} + \frac{\partial AS}{\partial q_i}, \quad i = 1,2 \\ \frac{\partial W_p}{\partial \beta} &= -\frac{\partial CS_p}{\partial \beta} + \frac{\partial AS}{\partial \beta}. \end{aligned} \tag{16}$$

The intensity of tracking in this model has two direct effects: on consumer disutility from being tracked and on advertisers' revenues (CS_c does not depend on β). When deciding whether to block third-party cookies, consumers only minimize CS_p (hence maximize $-CS_p$), while they ignore the effect that blocking has on $r_{12}^T(q_1, q_2)$. Because $\frac{\partial r_{12}^T(q_1,q_2)}{\partial \beta} D_{12} > 0$, there is too much blocking of cookies at equilibrium, implying that there is too low tracking. Because privacy tastes affect economic outcomes in the advertising market, it is important to take them into consideration in a model that aims at having relevant policy implications. When considering the effects in the advertising market of consumers' privacy decisions, we find that, compared to the first best, it would be preferable to induce less cookies blocking to allow more effective advertising.

Then, the comparison among (q_1^o, q_2^o) and $\left(q_1^{AT}, q_2^{AT}\right)$ is less straightforward. First, when deciding the advertising levels, publishers do not consider the direct effect of advertising on consumer surplus, while the regulator endogenizes the fact that consumers dislike ads, i.e. that $\frac{\partial CS_c}{\partial q_i} < 0$. This effect pushes the optimum down. There is a second effect: Proposition 1 states that the equilibrium advertising level can increase or decrease with β, and tracking is higher at optimum than at equilibrium, i.e. $\beta^o > \beta^{AT}$. Hence, $argmax\left\{\frac{\partial AS}{\partial q_i}\right\}$ calculated in β^{AT} can be higher or lower than $argmax\left\{\frac{\partial AS}{\partial q_i}\right\}$ calculated in β^o. In particular, if $argmax\left\{\frac{\partial AS}{\partial q_i}\right\}$ decreases with β, then the advertising levels solving $\frac{\partial AS}{\partial q_i} = 0$ when calculated in β^o are lower than the equilibrium ones. In this case, both effects bring to $(q_1^o, q_2^o) < \left(q_1^{AT}, q_2^{AT}\right)$. Otherwise, the two effects push in different directions, and the optimal level of advertising can be higher or lower than the equilibrium one.

5.2 Second-best level of tracking

Now, we assume that the regulator cannot control the level of advertising directly, but it can intervene on the level of tracking. This case is highly relevant: we observe that there are no advertising caps on online publishers, while there is a discussion on how to regulate third-party cookies. For instance, the FTC has discussed the possibility of introducing a Do Not Track mechanism that would register persistently the consumers' choice not to be tracked in their online activity (see https://www.ftc.gov/news-events/media-resources/protecting-consumer-privacy/do-not-track). Also, European Data Protection Authorities encouraged the development of "do-not-track" technical tools allowing consumers to easily manage online cookies (see http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2016/wp240_en.pdf). Moreover, because default can influence consumers' behaviours, there is a policy discussion on opt-in and opt-out strategies for obtaining consumers' consent to tracking. In general, authorities are uncertain on how to deal with the collection of consent online, because the interests of many different groups have to be considered.

We study the second-best where the regulator maximizes welfare with respect to β, anticipating how publishers will set their advertising level afterward. In the second best, the regulator controls perfectly the mass of consumers who do not block cookies, that determines the effectiveness of tracking. We compare this second best to the equilibrium level of tracking, determined by the choice of the consumers to block cookies. We conclude the following:

Proposition 6. *The second-best tracking level is higher than the equilibrium one if* $\frac{\partial q_i^{AT}}{\partial \beta} < 0$ *for* $i = 1, 2$.

To obtain the result in Proposition 5, we maximize the total welfare W_p with respect to β and compute it in $\left(q_1^{AT}, q_2^{AT}\right)$. Taking into account that, by the first order conditions of publishers, $\frac{\partial AS}{\partial q_i} = 0$ for $i = 1, 2$ in $\left(q_1^{AT}, q_2^{AT}\right)$, we obtain that:

$$-\frac{\partial CS_p}{\partial \beta} + \frac{\partial CS_c}{\partial q_i}\frac{\partial q_i^{AT}}{\partial \beta} + \frac{\partial CS_c}{\partial q_j}\frac{\partial q_j^{AT}}{\partial \beta} + \frac{\partial AS}{\partial \beta}. \tag{17}$$

When consumers decide whether to block third-party cookies at equilibrium, they internalize the effect on CS_p. Then, there are two externalities. First, the effect on advertiser surplus AS. Because $\frac{\partial AS}{\partial \beta} > 0$ (i.e. $r_{12}^T(q_1, q_2, \beta)$ increases in β), consumers impose a negative externality on publishers. Second, there is an externality imposed on consumers in the advertising market, that may be a positive or a negative externality depending on the sign of $\frac{\partial q_i^{AT}}{\partial \beta}$ for $i = 1, 2$. We already know that $\frac{\partial CS_c}{\partial q_i} < 0$ for $i = 1, 2$. If $\frac{\partial q_i^{AT}}{\partial \beta} < 0$, then the sign of the second and third terms of (17)is positive: if consumers block less cookies and the effectiveness of tracking increases, consumers are exposed to less ads, hence they experience lower disutility from advertising. Instead, if $\frac{\partial q_i^{AT}}{\partial \beta} > 0$, consumers will be exposed to more ads as β increases. Hence, the comparison among the second best level of tracking and the equilibrium one depends on the sign of $\frac{\partial q_i^{AT}}{\partial \beta}$. If $\frac{\partial q_i^{AT}}{\partial \beta} < 0$, we find that $\beta^{SB} > \beta^{AT}$, meaning that there is too low tracking in equilibrium because both externalities are negative. Otherwise, if $\frac{\partial q_i^{AT}}{\partial \beta} > 0$, the comparison is ambiguous. Indeed, while $\frac{\partial AS}{\partial \beta} > 0$, we find that $\frac{\partial CS_c}{\partial q_i}\frac{\partial q_i^{AT}}{\partial \beta} + \frac{\partial CS_c}{\partial q_j}\frac{\partial q_j^{AT}}{\partial \beta} < 0$. If the effect of β on AS is stronger than that on CS_c, then $\beta^{SB} > \beta^{AT}$, otherwise it can be that $\beta^{SB} < \beta^{AT}$. Proposition 1 gives the conditions determining the sign of $\frac{\partial q_i^{AT}}{\partial \beta}$ under the assumption of symmetric publishers.

These results are important for their implications for privacy regulation. Our analysis implies that an intervention in the market promoting blocking of third-party cookies is not always beneficial for social welfare. The regulatory intervention should depend on the effect that tracking has on the equilibrium level of advertising. Tracking may indeed be too low from a second-best point of view. When increasing tracking is socially desirable, the regulator should intervene in the market to decrease the disutility of consumers from being tracked θ. It may do so, for instance, by forbidding some kind of intrusive third-party cookies.[18] It may even be desirable to increase the perceived cost of blocking cookies: the negative effect that

[18] There are many studies in computer science trying to design privacy-preserving tracking tools.

this intervention has on CS_p may be overcome by the positive effects for both consumers and advertisers on the advertising markets. For instance, this objective may be reached by making opt-out strategies for collecting consumers' consent the default option. Instead, in the case where the regulator wants to induce less tracking, it might intervene in the market by reducing the cost for the consumers of blocking third-party cookies. This can be done, for instance, by making blocking of third-party cookies the default option in every browser or by creating a Do Not Track mechanism, on the lines of that discussed by the FTC.

It is interesting to consider the case where the regulator wants to maximize consumer surplus, defined as $-CS_p + CS_c$. The level of tracking may be higher or lower than that maximizing consumer surplus, depending on the sign of $\frac{\partial q_i^{AT}}{\partial \beta}$. Hence, promoting blocking of third-party cookies might not be beneficial even if the objective of the regulator is to increase total consumer surplus. In fact, when deciding whether to block cookies, consumers only consider their privacy taste $-CS_p$, but they neglect the effect that their decision has on the incentives of the publishers to advertise, hence on consumer surplus from visiting content CS_c. Hence, if $\frac{\partial q_i^{AT}}{\partial \beta} < 0$, inducing more blocking of cookies may have adverse effects on consumers.

5.3 Second-best level of advertising

Now, we assume that the regulator controls the level of advertising directly, but it leaves to consumers the choice of blocking third-party cookies. The regulator maximizes total welfare with respect to (q_1, q_2), anticipating consumers' choice over β. Comparing with the equilibrium levels of cookies $\left(q_1^{AT}, q_2^{AT}\right)$, we conclude that:

Proposition 7. *The second-best advertising level is lower than the equilibrium one.*

Indeed, by the maximization of total welfare W_p with respect to (q_1, q_2), we obtain

$$\frac{\partial CS_c}{\partial q_i} + \frac{\partial AS}{\partial q_i} = 0, \quad i = 1, 2. \tag{18}$$

At equilibrium, publishers do not internalize the effect of advertising on consumers. Because $\frac{\partial CS_c}{\partial q_i} < 0$ for $i = 1, 2$, then Proposition 7 holds. The discussion of this result goes along the lines of the discussion in Section 5.

6 Concluding remarks

We have investigated how the presence of an ad network affects the online advertising market, focusing, in particular, on the ad network's ability to track multi-homing consumers across publishers. In the first part of the analysis, we have considered tracking intensity as exogenous. We have shown that the intensity of advertising depends on how the quantity of multi-homers is affected by the advertising level and how advertising revenues change due to tracking. Interestingly, we have found that ad levels can decrease with the tracking effectiveness. By comparing the scenarios where publishers compete and that where they outsource the sale of ads to an ad network, we have identified two effects: the joint control effect, due to the fact that the ad network centralizes the sale of ads, and the tracking effect, due to the fact that tracking technologies increase the value of multi-homers on the advertising market. The combination of these effects is such that the level of advertising may increase or decrease when ads are outsourced to an ad network, compared to the case of competition.

In the second part of the analysis, we have endogenized tracking. Specifically, we assumed that consumers can block third-party cookies, thereby affecting the ad network's ability to track. We have shown that tracking at equilibrium may be too low, compared to both the first-best and second-best allocations. This is due to the fact that consumers do not consider how advertisers' willingness to pay for impressions and publishers' incentives to set ad levels are affected by their decision to block cookies. The results have interesting policy implications: there is an intense political debate on how to protect consumers' privacy in the online advertising market, where data-driven advertising technologies are in use. According to our model, when advertising intensity decreases with the intensity of tracking, promoting tracking is beneficial for consumers and society. Hence, promoting third-party cookies blocking may not always be socially desirable, both from a welfare and consumer surplus point of view, because of the implications of this policy on the advertising market. As we discuss above, these findings suggest that policymakers should carefully consider which conditions apply before choosing the type of privacy-related regulatory intervention.

This analysis can be extended in many directions. First, it would be interesting to consider the effect of privacy as an intermediate good (depending on their privacy choices, different consumers may be exposed to different ads). Second, one could endogenize the choice of the ad network over its tracking technology. Third, it would be interesting to introduce targeting, possibly micro-founding the model to make it tractable. Fourth, we 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 kind of inventory are outsourced. Finally, one may consider how results are affected if advertisers are different.

Appendix A

Proof of Lemma 1

We begin by proving that each publisher offers a single contract. Assume the publisher offers two contracts, $C_i^1 \equiv (p_i^1, m_i^1)$ and $C_i^2 \equiv (p_i^2, m_i^2)$, and that some advertisers accept C_i^1, while others accept C_i^2. If $p_i^1 < p_i^2$, then the publisher 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, publishers spread their advertising level 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 $r_{12}^{nT}(m_i, m_j)$ are increasing in m_i. Hence, the publisher can always increase its revenue by offering more ads for $p_i + \varepsilon$, with $\varepsilon > 0$, up to $m_i = q_i$. Finally, p_i is equal to the difference between the revenues of the advertiser when it advertises on both publishers, $r_i(q_i) D_i(q_i, q_j) + r_j(q_j) D_j(q_i, q_j) + r_{12}^{nT}(q_1, q_2) D_{12}(q_1, q_2)$, 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 (6) ■

Proof of Lemma 2

Following the Proof of Lemma 1, one can prove that AN offers one contract to all advertisers, such that $m_i = q_i$. The price p_{AN} is equal to the difference between the revenues of the advertiser when it accepts the contract, $r_i(q_i) D_i(q_i, q_j) + r_j(q_j) D_j(q_i, q_j) + r_{12}^h(q_1, q_2, \beta) D_{12}(q_1, q_2)$, and its outside option, which is equal to zero because AN manages ad inventory of both publishers. Hence, we get (8) ■

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:

$$\frac{\partial q_i}{\partial \beta} = \frac{\frac{\partial FOD_i}{\partial \beta}\frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_j}{\partial \beta}\frac{\partial FOD_i}{\partial q_j}}{\frac{\partial FOD_i}{\partial q_j}\frac{\partial FOD_j}{\partial q_i} - \frac{\partial FOD_i}{\partial q_i}\frac{\partial FOD_j}{\partial q_j}}. \tag{19}$$

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 $\frac{\partial q_i}{\partial \beta} > 0$. ■

Proof of Lemma 3

Let $\left(q_1^C, q_2^C\right)$ and $\left(q_1^{ANT}, q_2^{ANT}\right)$ 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 (10) in $\left(q_1^C, q_2^C\right)$ and with $\beta = 0$. Using (7), we can rewrite this expression as

$$\frac{\partial p_{AN}}{\partial q_i} = \frac{\partial D_{12}\left(q_1, q_2\right)}{\partial q_i}\hat{r}_j\left(q_j\right) + \frac{\partial D_j\left(q_1, q_2\right)}{\partial q_i} r_j\left(q_j\right), \ i = 1, 2. \tag{20}$$

where all functions are evaluated in $\left(q_1^C, q_2^C\right)$. 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 (20) can be rewritten as $\frac{\partial D_{12}(q_1,q_2)}{\partial q_i}\left[\hat{r}_j\left(q_j\right) - r_j\left(q_j\right)\right]$. By concavity, if $\frac{\partial p_{AN}}{\partial q_i}$ computed in $\beta = 0$ is strictly positive when evaluated in $\left(q_1^C, q_2^C\right)$, then $\left(q_1^{ANT}, q_2^{ANT}\right)$ must be such that $q_i^{ANT} > q_i^C$, $i = 1, 2$. Therefore $q_i^{ANT} > q_i^C$ if and only if $\hat{r}_j\left(q_j\right) < r_j\left(q_j\right)$. ■

Proof of Lemma 4

Let $\left(q_1^{ANT}, q_2^{ANT}\right)$ and $\left(q_1^{AT}, q_2^{AT}\right)$ 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 (10) in $\left(q_1^{ANT}, q_2^{ANT}\right)$, and using the fact that $\left.\frac{\partial D_i}{\partial q_i} r_i + D_i \frac{\partial r_i}{\partial q_i} + \frac{\partial D_j}{\partial q_j} r_j\right|_{q_i=q_1^{NTR}, q_2=q_2^{NTR}} = \left.-D_{12}\frac{\partial r_{12}^{nT}}{\partial q_i} - \frac{\partial D_{12}}{\partial q_i} r_{12}^{nT}\right|_{q_i=q_1^{NTR}, q_2=q_2^{NTR}}$ we have, for $i = 1, 2$:

$$\frac{\partial p_{AN}}{\partial q_i} = D_{12}\left(q_1, q_2\right)\left[\frac{\partial r_{12}^{T}\left(q_1, q_2, \beta\right)}{\partial q_i} - \frac{\partial r_{12}^{nT}\left(q_1, q_2\right)}{\partial q_i}\right]\Bigg|_{q_i=q_1^{NTR}, q_2=q_2^{NTR}}$$
$$+\frac{\partial D_{12}\left(q_1, q_2\right)}{\partial q_i}\left[r_{12}^{T}\left(q_1, q_2, \beta\right) - r_{12}^{nT}\left(q_1, q_2\right)\right] \|_{q_i=q_1^{NTR}, q_2=q_2^{NTR}}$$

By concavity, if $\frac{\partial p_{AN}}{\partial q_i}$ is strictly positive when evaluated in $\left(q_1^{ANT}, q_2^{ANT}\right)$, then $\left(q_1^{AT}, q_2^{AT}\right)$ must be such that $q_i^{AT} > q_i^{ANT}$, $i = 1, 2$. ■

Proof of Proposition 2

Consider the FOD (10). Let $\left(q_1^{C}, q_2^{C}\right)$ and $\left(q_1^{AT}, q_2^{AT}\right)$ 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 (10) in $\left(q_1^{C}, q_2^{C}\right)$ and using (7), we have, for $i = 1, 2$,

$$\begin{aligned} & D_{12}\left(q_1, q_2\right)\left[\tfrac{\partial r_{12}^{T}(q_1,q_2,\beta)}{\partial q_i} - \tfrac{\partial r_{12}^{nT}(q_1,q_2)}{\partial q_i}\right] + \\ & \left. + \tfrac{\partial D_{12}(q_1,q_2)}{\partial q_i}\left[r_{12}^{T}\left(q_1, q_2\right) - r_{12}^{nT}\left(q_1, q_2\right) + \hat{r}_j\left(q_j\right)\right] + \tfrac{\partial D_j(q_1,q_2)}{\partial q_i} r_j\left(q_j\right)\right|_{q_i=q_1^C, q_2=q_2^C} \end{aligned} \tag{21}$$

By concavity, if $\frac{\partial p_{AN}}{\partial q_i} > 0$ when evaluated in $\left(q_1^{C}, q_2^{C}\right)$, then $\left(q_1^{AT}, q_2^{AT}\right)$ is such that $q_i^{AT} > q_i^{C}$, $i = 1, 2$. 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 $q_i^{AT} > q_i^{C}$ if and only if $\left.D_{12}\left[\frac{\partial r_{12}^{T}}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i}\right] + \frac{\partial D_{12}}{\partial q_i}\left[r_{12}^{T} - r_{12}^{nT} + \hat{r}_j - r_j\right]\right|_{q_i=q_1^C, q_2=q_2^C} > 0.$

References

[1] 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.

[2] Anderson, S.P. and Coate, S., 2005. Market Provision of Broadcasting: A Welfare Analysis. Review of Economic Studies, 72, 947-972.

[3] Anderson, S.P., Foros, Ø. and Kind, H. J., forthcoming. Competition for Advertisers in Media Markets. Economic Journal.

[4] Athey, S., E. Calvano and J.S. Gans, forthcoming. The Impact of Consumer Multi-Homing on Advertising Markets and Media Competition, Management Science.

[5] Athey, S., and J.S. Gans, 2010. The Impact of Targeting Technology on Advertising Markets and Media Competition. The American Economic Review, 100, 608-613.

[6] Bagwell, K., 2007. The Economic Analysis of Advertising, in Mark Armstrong and Rob Porter (eds.). Handbook of Industrial Organization, Vol. 3, Amsterdam: North-Holland.

[7] Bergemann, D. and A. Bonatti, 2011.Targeting in Advertising Markets: Implications for Offline versus Online Media. RAND Journal of Economics, 42, 417-443.

[8] 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.

[9] Campbell, J., A. Goldfarb, and C. Tucker, 2015. Privacy Regulation and Market Structure. Journal of Economics & Management Strategy, 24, 47–73.

[10] Crampes, C., C. Haritchabalet, and B. Jullien, 2009. Advertising, Competition, and Entry in Media Industries. Journal of Industrial Economics, 57, 7-31.

[11] de Cornière, A. and G. Taylor, 2014. Integration and search engine bias. RAND Journal of Economics, 45, 576-597.

[12] Farrell, J., 2012. Can privacy be just another good? Journal on Telecommunications and High Technology Law, 10, 251-264.

[13] Foros, Ø., H. J. Kind, and G. Shaffer, 2014. Apple's Agency Model and the Role of Resale Price Maintenance. NHH Dept. of Business and Management Science Discussion Paper No. 2015/32.

[14] George, L.M. and C. Hogendorn, 2012. Aggregators, search and the economics of new media institutions. Information Economics and Policy, 24, 40–51.

[15] Goldfarb, A. and C. E. Tucker 2011. Privacy regulation and online advertising. Management Science 57, 57–71.

[16] Johnson, J.P., 2013. Targeted advertising and advertising avoidance. RAND Journal of Economics, 44, 128–144.

[17] Kind, H. J., Nilssen, T., and Sørgard, L., 2016. Inter-Firm Price Coordination in a Two-Sided Market. International Journal of Industrial Organization, 44, 101-112.

[18] Kreps, D.M. and Scheinkman, J.A.,1983. Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes. The Bell Journal of Economics, 14, 326-337.

[19] International Advertising Board, 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

[20] Lohr, S., 2010. Privacy Concerns Limit Online Ads, Study Says, New York Times.

[21] Peitz, M. and Valletti, T.M., 2008. Content and advertising in the media: Pay-TV versus free-to-air. International Journal of Industrial Organization, 26, 949-965.

[22] Renault, R., 2015. Advertising in markets. In Anderson, S., Stromberg, D., and J. Waldfogel (eds.). Handbook of Media Economics. Amsterdam: Elsevier.

[23] Roesner, F., Kohno, T. and Wetherall D., 2012. Detecting and Defending Against Third-Party Tracking on the Web. NSDI 2012.

[24] Tucker, C. E., 2012. The economics of advertising and privacy. International Journal of Industrial Organization 30, 326–329.

[25] Tucker, C. E., 2014. Social Networks, Personalized Advertising, and Privacy Controls. Journal of Marketing Research, 51, 546-562.

[26] Turow, J., King, J., Hoofnagle, C.J., Bleakley, A., Hennessy, M., 2009. Americans Reject Tailored Advertising and Three Activities that Enable It. Mimeo, Berkeley

[27] Vives, X., 2000. Oligopoly Pricing: Old Ideas and New Tools. MIT Press: Cambridge, Massachusetts.

[28] Yuan, S., Wang J., and Zhao, X., 2013. Real Time Bidding for Online Advertising: Measurement and Analysis, Proceedings of the Seventh International Workshop on Data Mining for Online Advertising.