This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Ad networks and consumer tracking

PLEASE CITE THE PUBLISHED VERSION

<https://doi.org/10.1287/mnsc.2019.3481>

PUBLISHER

INFORMS

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

This paper was accepted for publication in the journal Management Science and the definitive published version is available at <https://doi.org/10.1287/mnsc.2019.3481>.

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

D'Annunzio, Anna, and Antonio Russo. 2019. "Ad Networks and Consumer Tracking". Loughborough University. <https://hdl.handle.net/2134/9938015.v1>.

Ad Networks and Consumer Tracking

Abstract

We study the role of ad networks in the online advertising market. Our baseline model considers two publishers that can outsource the sale of their ad inventories to an ad network, in a market where consumers and advertisers multi-home. The ad network increases total advertising revenue by tracking consumers across outlets and reduces competition between publishers by centralizing the sale of ads. Consequently, outsourcing to the ad network is beneficial to the publishers, but may penalize the advertisers. We show that the ad network's ability to track consumers may either expand or reduce the provision of ads, depending on consumers' preferences for the publishers and how advertisers use tracking information. Specifically, tracking is more likely to expand (resp. reduce) the provision of ads when consumers' preferences for the publishers are positively (resp. negatively) correlated. Tracking is also more likely to expand (resp. reduce) the provision of ads when advertisers use tracking information to cap the frequency of impressions (resp. target specific consumers). Furthermore, we study the implications of consumers' choice to block tracking. Generally, blocking generates a negative impact on the advertising industry, by making the allocation of ads less effective. Blocking also entails an externality on consumers, which is negative when tracking reduces the provision of ads. Given these conditions, regulatory restrictions on tracking may reduce consumer surplus as well as advertising revenue. These findings contrast with the presumption that regulation should make it easier for consumers to avoid tracking. We propose further extensions, including competing ad networks, more than two publishers and networks that do not sell ads, but only tracking information to the advertisers.

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

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

Anna D'Annunzio
Toulouse Business School / France
dannunzio.anna@gmail.com

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

We thank Stefan Behringer, Alberto Bennardo, Emilio Calvano, Grazia Cecere, Jan Cauwer, Lapo Filistrucchi, Øystein Foros, Bruno Jullien, Hans Jarle Kind, Tore Nilssen, Marco Pagnozzi, Jerome Pouyet, Alessia Russo, Marianne Verdier, Ying Lei Tho, Liad Wagman, and participants to presentations at Toulouse Business School, CSEF, NHH, Rotterdam School of Management, CREST-ENSAI, ETH Zurich, Telenor Research, Pescara Summer Workshop 2018, the 8th Searle Center Conference on Internet Commerce and Innovation, CRESSE 2017, ICT Telecom-ParisTech conference 2017, EUI Media Conference 2017, CESifo Applied Microeconomics Conference 2017, the 6th EIEF-UNIBO-IGIER Bocconi workshop on IO, the 2016 Economics of IP, Software and Internet Conference, EARIE 2016 and the 2016 Media Economics Workshop for useful comments and discussions. We are grateful to the Editors and to three anonymous referees for insightful comments that substantially improved the paper. All errors are ours. Part of this research was carried out while Anna D'Annunzio was at the CSEF (University Federico II, Naples). The online appendix to this paper is available at <https://sites.google.com/site/antoniorussoweb/>.

1 Introduction

A large share of advertising takes place online: the Interactive Advertising Bureau estimates that the US online advertising market was worth more than the TV advertising market in 2017 (IAB, 2017). Several features of this market distinguish it from traditional ones. In this paper, we focus on two aspects that have received little attention in previous research. First, most digital publishers outsource the sale of their ad inventories to ad networks, e.g. Google AdSense.¹ Furthermore, ad networks use advanced technologies to provide advertising, whose effectiveness hinges on monitoring consumers’ behavior online.

One of the key factors behind the prominence of ad networks is that consumers commonly visit several websites in a short time frame. Previous literature (Ambrus et al., 2016; Athey et al., 2018) has pointed out that a publisher’s ad inventory loses value as the share of multi-homers increases. One reason is that, whereas each publisher is the unique gatekeeper to the attention of exclusive users, it must compete with other publishers for advertisers that intend to reach multi-homers. Secondly, the same ad can hit a multi-homer too many times on different publishers, being wasted with the consumer’s attention. Although each publisher can monitor a consumer on its own web pages, potentially avoiding internal repetitions, it can hardly do so when the consumer visits other websites. The impact of multi-homing on advertising revenues is clearly relevant for digital publishers, many of which are entirely financed through advertising.

Ad networks address the above issues in two ways. First, they centralize the sale of advertising, reducing competition among advertising outlets. In addition, ad networks use tracking technologies to improve the effectiveness of ad campaigns. *Across-outlet* tracking technologies, such as third-party cookies, allow to follow consumers on multiple websites. By using these technologies, ad networks can, for instance, keep track of which ads a consumer has already been exposed to on each publisher, avoiding wasteful repetition.²

However, tracking has non-trivial implications for consumers. For example, tracking affects the mix of ads they are exposed to and, thus, the utility they get from browsing the publishers. Moreover, tracking consumers entails storing information about their behavior and preferences, potentially invading their privacy.³ A growing number of consumers take measures to avoid tracking, such as adjusting their browser settings to block third-party cookies or using anonymizing software (Goldfarb and Tucker, 2012).⁴ By affecting the ad network’s ability to allocate ads efficiently,

¹Roesner et al. (2012) report that 89% of the (Alexa) top 500 sites includes at least one cross-site tracker and 40% is tracked by Google’s ad network. The latter reaches 94% of total Internet audience (Comscore, 2016).

²Furthermore, tracking consumers on multiple publishers allows the ad network to profile them more accurately and improve the matching with advertisers (targeting). Tracking is also useful when advertisers intend to show ads in a given sequence (https://support.google.com/dfp_premium/answer/1665531?hl=en) or reach a consumer who has shown interest in a generic ad with a personalized offer. In addition, tracking can be used to show to a consumer a given ad on multiple publishers based on previous interactions with their website (re-targeting).

³Turow 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 Gallup survey shows that 61% of respondents argued that online behavioral tracking is unjustifiable (Morales, 2010).

⁴According to a 2011 USA Today.com article (see <https://usatoday30.usatoday.com/tech/news/story/2011-12-29/internet-privacy/52274608/1>), Ghostery, one of the most common block-tracking tools, was “being downloaded by

consumers' blocking of tracking may have far-reaching implications for the advertising market.

The issues outlined above raise several interesting questions. What role do ad networks play in the online advertising market? How do they generate value for digital publishers, advertisers and consumers? How do they affect key market outcomes, such as the provision of ads on digital publishers? What is the effect of tracking technologies and the implications of consumers avoiding tracking on market outcomes? Our paper provides an analytical framework to tackle these questions. Essentially, our contribution is to explore the role of ad networks in the digital ecosystem. We develop a model of an ad network that centralizes the sale of ads and tracks consumers on multiple publishers. In a nutshell, our analysis shows that ad networks make the allocation of ads more efficient, but also change the distribution of surplus among publishers and advertisers by increasing concentration on the supply side of the market. Furthermore, ad networks can have unexpected effects on key market outcomes, such as the provision of ads on digital publishers, due to the two-sided nature of digital advertising markets, multi-homing and consumers' response to tracking.

In the baseline model, we consider two ad-financed publishers and an ad network. Consumers and advertisers possibly multi-home and ads are annoying to consumers. We analyze the impact of the ad network on publishers and advertisers, looking at its effect on the size and the distribution of advertising revenues. The network can exploit tracking technologies to improve the allocation of ads. Furthermore, it can coordinate the provision of ads on both publishers. Therefore, the ad network increases total advertising revenue. The network also extracts a larger share of this revenue from the advertisers, because it weakens competition among publishers by centralizing the sale of ads. As a result, the ad network can guarantee to each publisher a share of the advertising revenue large enough that they are better off outsourcing their ad inventories. On the other hand, given the network's market power, the advertisers may lose when the publishers outsource. This is the case, in particular, if the advertisers have little bargaining power, suggesting that small advertisers are more likely to be worse off with ad networks than larger ones.

We then analyze the effect of the ad network on the provision of ads on digital publishers. This is an important market outcome, notably because it affects the impact of ad networks on consumers. We disentangle the ad network's effect on the quantity of ads which is due to tracking from that which is due to centralizing their sale. The former depends on two forces stemming from the two-sided nature of digital publishers. On the one hand, tracking increases the revenue from additional ads that hit multi-homers. On the other hand, tracking also increases the total revenue generated from each multi-homer and, hence, the cost of losing these consumers when the publishers carry more ads. The first effect tends to expand the provision of ads on digital publishers, whereas the second tends to restrict it.

We identify some observable determinants of the magnitude of the effects of tracking described above. We argue that tracking is more likely to expand the provision of ads when consumers' preferences for the two publishers are positively correlated (which makes demand by multi-homing consumers less elastic) and advertisers use tracking information to cap the frequency of impressions

140000 new users each month, with total downloads doubling to 4.5 million in the past 12 months".

or implement retargeting across outlets (implying that tracking affects primarily the revenue from additional ads that hit multi-homers). By contrast, we expect tracking to have a restrictive effect on the provision of ads when consumers’ preferences for the two publishers are negatively correlated and when tracking information is used for targeting ads to specific consumers, based on their preferences for products.

The ad network also affects the provision of ads by centralizing their sale. When the quantity of ads carried by a publisher increases, some multi-homers stop visiting it, so the revenue these consumers generate on the publisher is lost. Differently from the case where the publishers sell their own ad inventories directly, the ad network internalizes that these consumers still generate revenue as single-homers on the other publisher. On the other hand, the ad network generates more revenue from a multi-homer than the publishers, due to its market power. If the net cost of increasing the quantity of ads is smaller for the ad network, centralization expands the provision of ads. We argue that this outcome is more likely when multi-homers devote significantly less attention to each publisher (and, hence, generate less revenue on each publisher) than single-homers do.

Next, we study the implications of consumers avoiding tracking, taking into account the related changes in the quantity and mix of ads they are exposed to. In our model, consumers may dislike tracking per se, e.g. because it reduces their privacy. In addition, tracking can make ads more (e.g., in the case of retargeting) or less (e.g., in the case of frequency capping) annoying. Therefore, tracking can change consumers’ incentives to multi-home and, thus, the composition of the publishers’ audiences. Due to this effect, tracking can reduce total advertising revenue in specific circumstances, particularly when it increases the nuisance from ads, discouraging consumers from browsing the publishers, without increasing advertising revenues from multi-homers substantially. Nevertheless, in practice, it is likely that the ad network avoids using tracking in ways that reduce total ad revenue. Hence, also in light of tracking’s potential to increase the value of impressions hitting multi-homers, we expect consumers avoiding tracking to generally exert a negative externality on the advertising industry. More interestingly, we find that blocking tracking also entails an externality on all consumers. The reason is that tracking affects the provision of ads and, thus, the surplus all users get from browsing content. Consequently, consumers may be better off and social welfare may increase with more tracking when the latter restricts the provision of ads.

The above findings have direct implications for the debate on privacy-protection rules designed to either reduce the costs of blocking tracking or induce the advertising industry to adopt less intrusive practices.⁵ These measures might penalize the advertising industry as well as consumers, if tracking reduces the provision of ads. By contrast, if the provision of ads increases, restrictions on tracking benefit consumers. We also evaluate the implications of advertising caps, that might mitigate the unintended consequences of privacy regulation on the supply of ads.

We propose a series of extensions in the last part of the analysis. We first consider a tracking network, that sells only tracking information to the advertisers but lets the publishers sell their

⁵For example, the Federal Trade Commission proposed in 2012 a “do not track” mechanism. The EU General Data Protection Regulation (GDPR, 2016) establishes opt-in policies for collecting consumers’ consent.

own ad inventories. The publishers do not benefit from letting this network track their audience if tracking reduces the total advertising revenue. By contrast, an ad network that centralizes the sale of ads is attractive to the publishers even if tracking reduces ad revenues, because the ad network softens competition for advertisers. On the other hand, the advertisers benefit from competition between the publishers and, hence, may be better off with a tracking network.

In a further extension, we consider competition between ad networks. Due to the gains from centralization and tracking, the ad networks compete fiercely to attract the publishers and the equilibrium is such that both publishers outsource to the same one. However, the allocation of impressions and total advertising revenues are the same as with a monopolist network. Thus, competition benefits primarily the publishers. We also show that our results are robust to introducing more than two publishers and to various modifications of the structure of contracts between publishers, the ad network and advertisers.

The remainder of this paper is organized as follows. Section 2 presents a brief literature review. Section 3 describes the model. Section 4 analyzes the impact of the ad network on the size and distribution of advertising revenues, as well as on the provision of advertising. Section 5 endogenizes the extent of tracking, focusing on consumers' choice to block. Section 6 presents the extensions. Section 7 discusses the policy implications of our analysis, and Section 8 concludes. Unless otherwise stated, proofs are in Appendix 1. Appendix 2, 3 and 4 are available online.

2 Literature

Our paper belongs to a recent literature that studies the provision of advertising in two-sided markets with multi-homing consumers and advertisers (Anderson et al., 2018; Ambrus et al., 2016; Athey et al. 2018). This literature acknowledges that multi-homing reduces the value of ad inventories by increasing platforms' competition for advertisers and making the allocation of ads less efficient. Furthermore, previous work emphasizes the role of demand composition (single- and multi-homers) in determining market outcomes. Our model incorporates these aspects and contributes by considering ad networks, that can address multi-homing by centralizing the sale of ads and exploiting tracking technologies.

There is a growing literature studying the impact of technologies that use consumer data to improve the effectiveness of advertising. This literature shows that firms do not necessarily benefit from collecting and exploiting consumer data. Athey and Gans (2010) show that targeting has value for general publishers only when they are constrained in advertising space. Bergemann and Bonatti (2011) analyze how a change in the targeting ability of media firms affects competition for advertising space. Bergemann and Bonatti (2015) consider a data-broker that collects and sells consumer information to advertisers. They show that more precise information may not benefit the data-broker. An important difference with our paper is that these papers focus on ad targeting, disregarding the implications of multi-homing. We focus instead on technologies that address the publishers' inability to monitor multi-homers outside their own domain. Furthermore, by consid-

ering publishers that bring consumers and advertisers together, our model differs from Bergemann and Bonatti (2011 and 2015), and is, in this regard, closer to Athey and Gans (2010). However, we differ from the latter by endogenizing participation on the consumer side. Given the response of consumers as well as advertisers, we find that the supply of ads can increase or decrease with tracking intensity.

Although we do not focus on the privacy implications of tracking for consumers, we contribute to the literature on the economics of privacy by studying the link between consumers avoidance of tracking and the provision of ads on digital publishers. Most papers in this literature (see Acquisti et al., 2016, for a survey) focus on firms' incentives to collect consumer data (de Cornière and de Nijs, 2016; Levin and Milgrom, 2010), but not on consumers' reaction. Our approach is closer to Casadesus-Masanell and Hervas-Drane (2015), who study the gathering of data by firms as well as the response by consumers. Differently from us, the authors concentrate on how firms set prices for consumers and for selling data to third parties, showing that firms collecting more data tend to set lower prices. Conitzer et al. (2012) and Montes et al. (2019) study settings where firms track consumers to implement personalized prices and consumers can choose whether to stay anonymous. Conitzer et al. (2012) find that increasing the cost of protecting their privacy can benefit consumers. Montes et al. (2019) show that, in duopoly, a larger cost of protecting privacy decreases firms' profits because competition for each consumer becomes more intense. Differently from these papers, we focus on the relation between consumers avoiding tracking and advertising quantities. Aziz and Telang (2019) consider an exogenous share of consumers that blocks tracking, showing that blocking can affect the accuracy of ad targeting. In our paper the decision whether to block is endogenous. Moreover, we differ by analyzing the behavior of digital platforms (publishers as well as the ad network) and studying the effect of tracking on advertising quantities.

A consumer's decision to avoid tracking imposes an externality on other consumers, by affecting the platforms' provision of ads. This aspect connects our paper to a recent literature that studies the implications of consumers choice to avoid ads (e.g. with ad-blockers), although there is no ad avoidance in our model. Johnson (2013) shows that there may be too little avoidance because consumers ignore that blocking ads discourages marginal advertisers, benefiting other consumers. Hann et al. (2008) study consumer efforts to deflect marketing, finding that ad blocking may affect other consumers by changing firms' incentive to advertise. In our model, differently from these papers, the interplay between consumer choices and market outcomes is driven by the strategies of two-sided platforms. Anderson and Gans (2011) consider ad avoidance in a two-sided model, but assume consumers single-home. In their model, platforms always respond to ad avoidance by raising ad quantities, because consumers who do not block are less sensitive to ads. By contrast, we find opting out from tracking can also bring to less advertising, depending on conditions that are intimately related to multi-homing and the two-sidedness of the market.

Our study contributes to the literature on common distributors in media markets. Kind et al. (2016) consider a distributor to which TV channels can delegate the choice over subscription prices. They show that using the distributor is not always profitable because it allows inter-firm price

Figure 1: Illustration of the baseline market structure.

coordination on the consumer side of the market, but prevents intra-firm coordination. George and Hogendorn (2012) study news aggregators, showing that they facilitate multi-homing and can thus reduce a publisher’s advertising revenue if many of its original viewers are exclusive. In contrast to these papers, we consider an aggregator on the advertising side of the market. In our model, by centralizing the supply of ads, the aggregator benefits ad-financed publishers.

Finally, our paper also relates to a recent literature on advertising attribution, i.e. the ex-post allocation of advertising revenue to publishers when advertisers multi-home (Kireyev et al., 2016). Consistently with our findings, Berman (2018) shows that increasing the efficiency in the allocation of ads may increase profits for publishers, but decreases those of advertisers. Our approach is complementary to this literature, because we study the ad networks’ role of improving the (ex-ante) allocation of ads on multiple publishers.

3 Baseline setup

We consider a market with two publishers, indexed by $i = 1, 2$, and an ad network, denoted by AN (see Figure 1).⁶ The basic structure of our model borrows from the setting of Ambrus et al. (2016).

Publishers and the ad network. Publishers provide free content to consumers and generate revenue from advertising. We denote by q_i the advertising level on publisher i , i.e. the quantity of impressions that the publisher can expose each consumer to. We shall at times refer to q_i as the publisher’s “advertising capacity”: in our model, the publishers decide this capacity at an early stage and later sell the available impressions to the advertisers.⁷

The publishers can either sell the impressions directly or outsource this function to the ad network.⁸ When a publisher outsources, the ad network sells the publisher’s advertising inventory and retains the ensuing revenues, in exchange for a fee. In the baseline model, we allow this fee to depend on the ad capacity chosen by the publisher and on the outcome of the negotiation with the other publisher. We characterize the contract between the publishers and the ad network in Section 4.2. We denote publisher i ’s profit as π_i , and the ad network’s profit as π_{AN} .

⁶We introduce competing ad networks and more than two publishers in Sections 6.2 and 6.3, respectively.

⁷We have in mind a situation where the publishers first design the layout of their web pages, including the allocation of space between content and ads. Once consumers visit their webpages, the publishers generate impression opportunities for the advertisers. Note that q_i is the same for all consumers visiting i . We discuss this assumption in Section 4.2.

⁸We consider ad networks that do not sell advertising, but only tracking information, in Section 6.1.

Unlike the publishers, the ad network can track consumers across outlets (Athey et al., 2018). As explained below, this feature has relevant implications for advertising revenues.

Consumers. There is a unit mass of consumers. A consumer visits publisher i if and only if she obtains a non-negative utility from doing so. We denote by D_{12} the quantity of multi-homers, D_i the quantity of single-homers on publisher i , and D_0 the quantity of consumers visiting no publisher (zero-homers). We specify consumer demands as follows⁹

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

where u_i is the reservation utility for content on publisher i . We assume that u_1 and u_2 are distributed according to a joint distribution with smooth density, $g(u_1, u_2)$. For simplicity, we assume the net utilities from browsing the publishers are additively separable. The term δq_i captures the disutility generated by ads when visiting publisher i , where $\delta > 0$.¹⁰

The demand system in (1) has 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, the composition of j 's audience changes with q_i . However, j 's total demand does not, because $\frac{\partial(D_{12}+D_j)}{\partial q_i} = 0$.

The total surplus consumers get from browsing the publishers, CS , is

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

Advertisers. There is a unit mass of advertisers. Advertising is informative: the revenue generated from ads depends on the probability that consumers are informed about the advertiser's product as well as on the return from informing a consumer. We assume the advertisers capture the entire surplus generated when a consumer is informed.¹¹ We assume the advertisers are symmetric and decompose their expected revenue in several parts, each related to a subset of consumers. If an advertiser does not place ads on any publisher, its revenue is zero. Suppose an advertiser buys m_i impressions (per consumer) only from publisher i (i.e. the advertiser single-homes). 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

⁹We assume these demands satisfy the necessary conditions for existence and uniqueness of an equilibrium with interior solutions. See Ambrus et al. (2016) for a discussion of these conditions.

¹⁰The assumption that consumers dislike ads is consistent with the increasing diffusion of ad-blockers and the fact that several publishers offer ad-free versions of their websites for a price, suggesting that internet ads tend to be annoying. We recognize that digital technologies employed by ad networks may make some ads more valuable to users (e.g. in the case of personalized offers) or more annoying (e.g. in the case of retargeting). Accordingly, in Section 5 we let the disutility δ depend on whether a consumer is tracked by the ad network.

¹¹Whether consumers obtain a real benefit from the information conveyed by ads is uncertain. Accordingly, we follow the literature assuming advertisers capture the entire surplus from informing consumers (see, e.g., Anderson and Coate, 2005). Our analysis is unchanged if consumers' retain part of this surplus, as long as the net effect of exposure to ads is negative (given δ).

multi-homing) consumer.¹² A single-homing advertiser's revenue is

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

We assume that r_i and $\hat{r}_i$ are increasing and concave functions of m_i , to capture diminishing returns to advertising.¹³ Assume now the advertiser multi-homes, buying m_i impressions (per consumer) on publisher i . The advertiser's total revenue is

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

The expected revenue from a single-homing consumer is still $r_i(m_i)$, because she is exposed only to impressions on i . The expected revenue from a multi-homer is $r_{12}(m_1, m_2, \beta)$, which is increasing and concave in m_1 and m_2 . This revenue also depends on whether consumers are tracked across outlets. Let $\beta \in [0, 1]$ denote the expected share of tracked consumers. Also, let $r_{12}^T(m_1, m_2)$ be the revenue from ads that hit a tracked multi-homer, and $r_{12}^{nT}(m_1, m_2)$ the revenue from ads on a non-tracked multi-homer. Hence, the expected revenue from a multi-homer is

$$r_{12}(m_1, m_2, \beta) = (1 - \beta) r_{12}^{nT}(m_1, m_2) + \beta r_{12}^T(m_1, m_2). \quad (5)$$

We assume that $\hat{r}_i \leq r_{12}^{nT} \leq \hat{r}_1 + \hat{r}_2$. This assumption captures the fact that ads shown to a consumer on different publishers are typically imperfect substitutes, for example because a multi-homer may receive an impression that she has already registered while visiting another publisher (Ambrus et al., 2016; de Cornière and Taylor, 2014). The ad network can improve the effectiveness of ads on multi-homers by tracking them across outlets (e.g. with third-party cookies). We assume tracking increases the expected revenue from the ads that hit each multi-homer, i.e. $r_{12}^T > r_{12}^{nT}$, as well as the value of additional impressions on such consumer, i.e. $\frac{\partial r_{12}^T}{\partial m_i} > \frac{\partial r_{12}^{nT}}{\partial m_i}$. That is, $\frac{\partial r_{12}}{\partial \beta} > 0$ and $\frac{\partial^2 r_{12}}{\partial m_i \partial \beta} > 0$. Note that we model the effect of tracking on ad revenues in a simple, reduced-form way. This formulation enables us to study the impact of tracking technologies in a setting where the behavior of platforms, advertisers and consumers is endogenous, while keeping the analysis tractable.

The above assumptions are consistent with several ways of exploiting tracking information. For instance, tracking can avoid wastefully repeated impressions: upon identifying a consumer informed on a publisher (e.g., by observing that she has already clicked on a given ad), the ad network can expose her to a different message on the other publisher (e.g., an ad about a related product, possibly from the same advertiser). Tracking may also be useful for re-targeting, i.e. proposing an ad containing a specific offer on a certain product to a consumer who has previously shown interest in it.¹⁴ Finally, tracking can improve the ad network's ability to identify consumers' preferences for

¹²The difference between $r_i(m_i)$ and $\hat{r}_i(m_i)$ may stem from the different amount of attention or time that single- and multi-homers devote to a given publisher.

¹³Diminishing returns may arise because ads are more likely to reach already informed consumers as the size of an advertising campaign increases. See, e.g., Bagwell (2007) and Renault (2015).

¹⁴Another example is the case of ad campaigns sending messages in a sequence to tell a brand story (see, e.g.,

products and match them with the relevant advertisers (targeting). Observe that these examples are consistent with the assumption that tracking has a positive effect on the revenue generated from each multi-homer as well as from each additional impression that hits these consumers. For instance, by reducing wasteful across-outlet repetition, tracking increases the share of multi-homers that are informed about a product for a given number of impressions ($r_{12}^T > r_{12}^{nT}$), but it also increases the likelihood that each additional impression from the advertiser hits consumers who are not already informed ($\frac{\partial r_{12}^T}{\partial m_i} > \frac{\partial r_{12}^{nT}}{\partial m_i}$).¹⁵ However, the magnitude of these two effects may vary, depending on the way advertisers exploit tracking information. We return to this point below.

The share of tracked consumers, β , is exogenous (we relax this assumption in Section 5). However, the ad network can track consumers across outlets only if both publishers allow it on their websites. Therefore, $\beta = 0$ whenever either publisher does not outsource. Given our focus on across-outlet tracking, we also assume ad revenues from single-homers are unaffected by tracking.¹⁶

Social welfare. After netting out transfers between players (including firms' profits), welfare is the sum of consumer surplus (see (2)) and gross advertiser surplus, AS :

$$W = CS + AS. \quad (6)$$

Indexing advertisers by a and letting M (resp., S_i) be the set of advertisers that multi-homes (resp., single-homes on i), we can write AS as follows:

$$AS = \sum_{a \in M} \left(\sum_i r_i D_i + r_{12} D_{12} \right) + \sum_i \sum_{a \in S_i} (r_i D_i + \hat{r}_i D_{12}). \quad (7)$$

Timing and equilibrium concept. At stage 1, AN offers a transfer to each publisher in exchange for its advertising inventory. The publishers accept or refuse. At stage 2, the publishers choose their advertising capacities, q_i . At stage 3, consumers observe these capacities (q_1, q_2) on each website and decide which to visit, if any. At stage 4, if ad inventories are outsourced, the ad network sells the impressions to advertisers. Otherwise, the publishers sell the impressions. Consumers then get exposed to ads and all payoffs are realized. We adopt Subgame-Perfect Nash Equilibrium (SPNE) as the solution concept and solve the model by backward induction.

https://scontent-arn2-1.xx.fbcdn.net/t39.2365-6/10333119_1457635661160496_168768318_n.pdf).

¹⁵Similarly, if tracking improves the targeting of ads, it makes each impression better matched to the consumers' preferences ($r_{12}^T > r_{12}^{nT}$), but also increases the likelihood that any additional impression hits the right consumers ($\frac{\partial r_{12}^T}{\partial m_i} > \frac{\partial r_{12}^{nT}}{\partial m_i}$). Finally, a re-targeted additional impression is more likely to be valuable to the advertiser than one which is not, so that $\frac{\partial r_{12}^T}{\partial m_i} > \frac{\partial r_{12}^{nT}}{\partial m_i}$. We provide a foundation for the revenue functions for multi-homers in Appendix D.

¹⁶Digital publishers are typically effective at monitoring consumers and allocating ads efficiently on their own websites, but unable to monitor consumers when they visit other outlets (Athey et al., 2018). Moreover, publishers are often able to share with the ad network the information collected using first-party cookies (see, e.g., <https://clearcode.cc/blog/first-party-cookie-approach>). Hence, the ad network's ability to track affects primarily the revenue generated by ads hitting multi-homers.

4 Solving the model

Our first objective is to study the impact of ad networks on digital publishers, advertisers and consumers. We consider a benchmark model where the ad network is unavailable. Next, we allow the publishers to outsource their inventories to the network. Because consumer demands (stage 3) are fully described in (1), in the following we focus on the other stages of the game.

4.1 No ad network

Assume the ad network is inactive. At stage 4, each publisher offers a contract to each advertiser, specifying an advertising quantity m_i in exchange for a payment p_i .¹⁷ The choice of m_i is characterized as follows. First, because advertising revenues are increasing in m_i (see (3) and (4)), each publisher fills the ad capacity on each consumer. Furthermore, due to diminishing returns, each publisher splits this capacity equally across all advertisers. Thus, in equilibrium each publisher offers the same contract to all advertisers, such that $m_i = q_i$ (recall that there is a unit mass of advertisers).

To determine the payment p_i , we evaluate the advertiser's payoff when rejecting i 's offer. If the advertiser places ads only on the other publisher, j , its revenue is given by (3), net of the payment p_j . By contrast, if the advertiser also accepts i 's offer, it gets (4) net of p_i and p_j . Note that $r_{12} = r_{12}^{nT}$, because $\beta = 0$ without the ad network. Hence, each publisher sets

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

All advertisers accept the above contract, thus i obtains $\pi_i = p_i$ in equilibrium (see Appendix B.1 for a proof). This payment equals the advertisers' incremental revenue from placing q_i impressions per consumer on i . The publishers are therefore unable to extract the whole surplus generated by ads on their platform because there are non-exclusive consumers. Such consumers create an interdependence between the advertising revenues on the two publishers: the higher the revenue from reaching a multi-homer only on j , the less advertisers are willing to pay for impressions on i .

At stage 2, each publisher chooses its advertising capacity, q_i , maximizing (8) given the q_j chosen by the other publisher. Let (q_1^C, q_2^C) be the resulting equilibrium capacities, where C is the superscript denoting the equilibrium variables without the ad network. This pair satisfies the following system of first-order conditions:

$$\frac{\partial \pi_i}{\partial q_i} = \left[\frac{\partial r_i}{\partial q_i} D_i + \frac{\partial r_{12}^{nT}}{\partial q_i} D_{12} \right] + \left[r_i \frac{\partial D_i}{\partial q_i} + (r_{12}^{nT} - \hat{r}_j) \frac{\partial D_{12}}{\partial q_i} \right] = 0, \quad i, j = 1, 2, \quad i \neq j. \quad (9)$$

As q_i increases, the revenue from impressions on publisher i goes up (the first bracket of (9) is positive), but fewer consumers join the publisher's platform (the second bracket is negative). The optimal q_i for publisher i optimizes the trade-off between these two effects.

¹⁷We analyze alternative formulations of contracts between publishers, the ad network and advertisers in Section 6.4.

Lemma 1. *Assume the publishers do not outsource to the ad network. The SPNE is such that ad capacities (q_1^C, q_2^C) satisfy (9). Each advertiser acquires q_i^C ads per consumer on publisher i , with the payment p_i^C characterized in (8) computed at (q_1^C, q_2^C) . Publisher i 's profit, π_i^C , equals p_i^C .*

4.2 Introducing the ad network

Assume now the ad network is active. Consider stage 4 in the subgame where both publishers outsourced. AN offers a contract to each advertiser, specifying a pair (m_1, m_2) and a payment p_{AN} . Following the same logic as above, AN fills the available ad capacity and divides it equally across all advertisers. Hence, the equilibrium contracts are such that $m_i = q_i$. Furthermore, AN extracts the entire advertiser surplus, because it is the unique gateway to the audience of both publishers. In equilibrium, all advertisers accept the above contract. Thus, the ad network's revenue, p_{AN} , coincides with the advertiser surplus, AS (see (7)):

$$p_{AN} = AS = r_1(q_1) D_1(q_1, q_2) + r_2(q_2) D_2(q_1, q_2) + ((1 - \beta) r_{12}^{nT}(q_1, q_2) + \beta r_{12}^T(q_1, q_2)) D_{12}(q_1, q_2). \quad (10)$$

The comparison of (8) and (10) provides a basic insight: all else given, the ad network can expand total advertising revenues because tracking increases the effectiveness of ads on multi-homers. In addition, the network centralizes the sale of ads, thereby capturing a larger share of this revenue than competing publishers can. In fact, in our baseline model AN captures the entire advertiser surplus, given it has all the bargaining power when dealing with the advertisers (we discuss the implications of relaxing this assumption below).

Recall that, by assumption, a publisher's ad capacity, q_i , is the same for all consumers. Hence, in equilibrium a tracked multi-homer is exposed to the same quantity of ads on each publisher as a non-tracked one (see Appendix B.2 for a proof).¹⁸ However, the mix of ads these consumers are exposed to may be different, because tracking enables the ad network to adjust the allocation of ads on multi-homers. Hence, the ad network charges more for ads hitting a tracked multi-homer than for those hitting a non-tracked one (recall that $r_{12}^T \geq r_{12}^{nT}$).

We now describe the outcome of the earlier stages of the game. At stage 1, AN makes simultaneous take-it-or-leave-it offers to each publisher, specifying a schedule of transfers $x_i(q_i)$ contingent on the publisher's advertising capacity, in exchange for the publisher's ad inventory and the ensuing revenue.¹⁹ Furthermore, we assume the contract with publisher i is void if the other publisher rejects AN 's offer. This feature captures the fact that the ad network can make an advantageous offer to each publisher only if it manages their inventories jointly. If only one publisher outsources,

¹⁸We are not aware of evidence that digital publishers significantly change the quantity of ads on their web pages according to whether a consumer is tracked by third parties. Theoretically, one could allow the publishers to set different advertising capacities for tracked and non-tracked consumers. However, doing so would complicate the analysis by doubling the size of the publishers' choice set and prevent a clean comparison of advertising quantities between the case where the ad network is available and that where it is not.

¹⁹This assumption is consistent with the idea that the ad network can provide incentives to the publishers to optimize their ad capacities (see, e.g., Google Ad Sense optimization tips to publishers at <https://urly.it/3192c>). We relax the assumption in Section 6.5, where we explore the implications of alternative contract formulations (including per-impression transfers and revenue sharing contracts).

the ad network cannot generate more revenue than the publisher itself, given that it cannot track consumers across outlets and has to compete with the other publisher for advertisers. Hence, by including this multilateral dimension in the contract we take into account that the ad network and a publisher may want to renegotiate their contract if the other publisher does not outsource.²⁰ The contractual structure we assume allows the ad network to coordinate the provision of advertising on both publishers. The assumption simplifies the exposition without affecting our results substantially (see Section 6.5).

We now argue that the unique SPNE of this game is such that both publishers outsource.²¹ Let q_i^* be the equilibrium ad capacity chosen by publisher i (to be characterized below). Publisher i accepts AN 's offer if and only if the transfer it obtains in equilibrium, $x_i(q_i^*)$, is greater or equal than its outside option, π_i^o . To determine the latter, suppose the publisher rejects the offer. Because the ad network's contracts are void, the continuation game is identical to that described in Section 4.1. It follows that $\pi_i^o = \pi_i^C$, as characterized in Lemma 1. However, AN can generate more advertising surplus and extract a larger share of it than the publishers can (see (8) and (10)). Therefore, we have $p_{AN} > \sum_{i=1,2} \pi_i^o$. In equilibrium, AN pays to each publisher a transfer equal to π_i^C and earns

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

The ad network designs the schedules $x_i(q_i)$ to induce the publishers to choose ad capacities that maximize π_{AN} . Noting that π_i^C is independent of the capacities chosen on the equilibrium path, (q_1^*, q_2^*) maximize p_{AN} at stage 2. Specifically, (q_1^*, q_2^*) satisfy the following first-order conditions:

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

Note that we have used the property that $\frac{\partial D_j}{\partial q_i} = -\frac{\partial D_{12}}{\partial q_i}$ and that $r_{12} = (1 - \beta) r_{12}^n + \beta r_{12}^T$. The equilibrium values of q_i^* are the result of the following trade-off: increasing the provision of ads on one publisher generates more revenue from its consumers (first bracket), but also reduces the quantity of consumers that visit it (second bracket) as some switch from multi-homing to single-homing on the other publisher.

Lemma 2. *When the ad network is available, the SPNE is such that both publishers outsource their inventories, earning π_i^C in equilibrium, whereas the ad network earns $p_{AN} - \pi_1^C - \pi_2^C$. The advertising quantities, (q_1^*, q_2^*) , maximize total ad revenues, p_{AN} , as defined in (10).*

There is a similarity between the effect of the ad network on the advertising market and that

²⁰Excluding this multilateral dimension from the contract between the ad network and each publisher does not imply major changes in our findings (see Section 6.5). See Bernheim and Whinston (1985) on multilateral contracting in multi-agent relationships.

²¹In our model each publisher fully outsources its ad inventories. This is profitable because of the gains from centralization and tracking. Observe that selling directly the inventory that comes from single-homers (for which publishers face no competition and there are no inefficiencies due to across-outlet repetition) is impossible because the publishers cannot know whether a consumer multi-homes without the ad network's tracking technology.

of joint ownership of the two publishers. This similarity follows from the ad network’s ability to centralize the sale of ads and, given the contract structure we assume, to coordinate the publishers’ ad capacities. In Section 6.5, we check the robustness of our results to alternative contracts between the ad network and the publishers, showing that the main results do not change. Furthermore, we provide below additional extensions that further characterize the difference with joint ownership. These extensions include tracking networks, that only sell tracking information but do not centralize the sale of ads (Section 6.1) and competing ad networks (Section 6.2). We also consider more than two publishers (Section 6.3). Although many publishers may join an ad network, it is unlikely that they would merge under the same owner in reality. Also, mergers between large digital publishers are likely to face constraints imposed by competition authorities. Therefore, in practice, competing publishers cannot replicate the benefits delivered by ad networks simply by merging.

4.3 The effect of the ad network on publishers and advertisers

In our model, the ad network increases gross advertiser surplus, by tracking consumers and coordinating the provision of ads on multiple publishers. However, the ad network also changes the distribution of this surplus, by reducing competition on the seller side. We now consider the implications of these changes for advertisers and digital publishers.

By centralizing the sale of ads, the ad network can extract more revenue from the advertisers than the publishers can. In our baseline model, the ad network can in fact extract all such revenue. Hence, the advertisers are worse off when the publishers outsource, even though the ad network makes their ads more effective with tracking. This result can be seen by comparing the payments charged by the ad network, (10), to those charged by the publishers, (8). However, this outcome depends on our assumption that the ad network offers take-it-or-leave-it contracts to the advertisers. We relax this assumption in Section 6.4, showing that when an advertiser has sufficient bargaining power, it can retain a large enough share of the additional revenues generated by tracking and, thus, also benefit when the publishers’ outsource.

In equilibrium, the publishers receive at least as much revenue from the ad network as they earn when they do not outsource. The reason is that each publisher has the alternative of selling its inventory directly and the ad network must compensate it accordingly. Therefore, the publishers are naturally in a stronger position than the advertisers when dealing with the ad network: when the publishers outsource, the advertisers have no alternative but going through the network to reach consumers. As we argue below (Section 6), these conclusions hold also if one considers more than two publishers, competing ad networks, and alternative assumptions on the structure of contracts.

4.4 The effect of the ad network on the provision of advertising

We now turn to the effect of the ad network on the provision of advertising on digital publishers. This is an important market outcome and is key to determine the impact of ad networks on consumers. The ad network’s effect on the supply of ads is not straightforward, as it depends on its ability to

track consumers and to centralize the sale of ads. For the sake of exposition, we shall first focus on the effect of tracking and analyze the effect of centralization later.

4.4.1 The effect of tracking on the provision of advertising

We now investigate how a change in the extent of tracking by the ad network, β , affects the equilibrium advertising quantities, q_i^* . Assuming the publishers are symmetric and using (12), we establish the following result:²²

Proposition 1. *Assume that publishers are symmetric. The advertising quantity on each publisher increases with the extent of tracking (i.e., $\frac{\partial q_i^*}{\partial \beta} > 0$) if and only if*

$$\underbrace{D_{12} \left(\frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i} \right)}_{(+)} + \underbrace{\frac{\partial D_{12}}{\partial q_i} (r_{12}^T - r_{12}^{nT})}_{(-)} > 0, \quad i = 1, 2. \quad (13)$$

where all functions are evaluated at (q_1^*, q_2^*) .

The equilibrium ad quantities may increase or decrease with tracking. Tracking raises the revenue from marginal ads hitting the D_{12} multi-homers ($\frac{\partial^2 r_{12}}{\partial q_i \partial \beta} = \frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i} > 0$), implying a stronger incentive to expand the provision of ads. On the other hand, tracking also raises the expected revenue from all ads that hit multi-homers ($\frac{\partial r_{12}}{\partial \beta} = r_{12}^T - r_{12}^{nT} > 0$). As a result, the cost of carrying additional ads increases, given that the quantity of multi-homers decreases with q_i ($\frac{\partial D_{12}}{\partial q_i} < 0$).

The sign of (13) depends on the elasticity of the quantity of multi-homing consumers to the quantity of ads and on how tracking affects the advertising technology. It is useful to evaluate these two aspects separately. The elasticity of D_{12} with respect to q_i relates to the distribution of preferences for the two publishers. Suppose these preferences are negatively correlated (e.g., in the case of polarized news websites such as MSNBC.com and FOXnews.com). Because consumers tend to visit either one publisher or the other, multi-homers are a relatively small share of the audiences, but a relatively large share of the change in demand triggered by more ads. Formally, D_{12} is small relative to $\frac{\partial D_{12}}{\partial q_i}$ in absolute value. Hence, all else given, it is more likely that tracking restricts the provision of ads. By contrast, many consumers that visit one publisher tend to visit also the other if preferences are positively correlated (e.g. with two websites providing complementary sport coverage, such as ESPN.com and Yahoo!sports.com). All else given, tracking is more likely to increase the provision of ads in this case.

Tracking's effect on the provision of ads also depends on how it influences the value from ads on a multi-homer ($r_{12}^T - r_{12}^{nT}$), relatively to the value of the marginal impressions on such consumer ($\frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i}$). The magnitude of these effects can be related to the objective of advertising campaigns. For instance, firms advertising mainstream products typically aim to maximize the reach of their campaigns. These advertisers value controlling the frequency of impressions on each consumer

²²Without symmetry, the effect of β on q_i^* gets more involved, because it depends on how changes in the quantity of impressions on one publisher affect the marginal revenues generated on the other.

(as too few or too many ads on the same consumer are wasteful), with relatively little regard for hitting specific individuals. In this case, it is reasonable to expect tracking to mainly increase the value of marginal impressions, because the likelihood of repetition increases with the total volume of impressions acquired. By contrast, firms advertising niche products tend to target specific consumers. In this case, we expect tracking to increase the value of any impression (and not just marginal ones) on multi-homers, because each single impression is worthless if it does not hit the right consumer. Finally, when advertisers care for retargeting the same consumer or sending ads in sequence to tell a brand story, we expect tracking to increase the value of marginal impressions primarily.

To summarize, expression (13) suggests that tracking is more likely to expand the provision of advertising on digital publishers when consumer preferences for the two publishers are positively correlated and advertisers use tracking information to cap the frequency of impressions or implement retargeting across outlets. By contrast, we expect a restrictive effect on the provision of ads when consumers' preferences for the publishers are negatively correlated and tracking information is used for targeting ads to consumers with specific preferences.

Proposition 1 identifies a novel effect of tracking technologies on the provision of ads. Athey and Gans (2010) find that it is optimal for publishers to cut the supply of ads when ads can be targeted. Johnson (2013) finds that the level of advertising increases with targeting if the value of the marginal ad is higher than the value of an ad on a random consumer. Our result differs primarily because it stems from the two-sided nature of digital publishers, that must account for the response of consumers as well as advertisers (consumer demands are exogenous in Athey and Gans, 2010, and publishers are not modeled in Johnson, 2013).

4.4.2 The overall effect of the ad network on the provision of advertising

We now study the overall effect on ad quantities of the publishers' decision to outsource to the ad network. Formally, we compare the equilibrium ad quantities (q_1^*, q_2^*) when the publishers outsource (satisfying (12)) to the quantities (q_1^C, q_2^C) when they do not (satisfying (9)). In so doing, we shall isolate the effect of centralizing the sale of ads from that of tracking. Note that we do not require symmetry to derive the result that follows.

Proposition 2. *When the publishers outsource to the ad network, the provision of ads exceeds that in the case where they do not, i.e. $q_i^* > q_i^C$, if and only if*

$$\beta \left[\underbrace{D_{12} \left(\frac{\partial r_{12}^T}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i} \right) + \frac{\partial D_{12}}{\partial q_i} (r_{12}^T - r_{12}^{nT})}_{\text{Tracking}} \right] > \underbrace{\frac{\partial D_{12}}{\partial q_i} (r_j - \hat{r}_j)}_{\text{Centralization}}, \quad i, j = 1, 2, \quad i \neq j, \quad (14)$$

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

The effect of tracking on ad quantities is captured by the left hand side of (14). This part of the expression contains the same terms as Proposition 1 (see the commentary thereafter). To focus on the effect of centralization, suppose there is no tracking, i.e. $\beta = 0$. As (14) suggests, AN 's incentives regarding the choice of q_i differ from the publishers' only due to their different ability to generate revenue from multi-homers. As the quantity of advertising on publisher i , q_i , increases, some multi-homers become single-homers on j ($\frac{\partial D_{12}}{\partial q_i} < 0$). Because the network centralizes the sale of ad inventories, this increase in q_i produces a net loss of $r_{12}^{nT} - r_j$ (see (12)). By contrast, higher q_i imposes on publisher i a loss equal to the incremental value of a multi-homer, $r_{12}^{nT} - \hat{r}_j$ (see (9)). Note that, unlike the ad network, each publisher ignores the impact of this consumer becoming a single-homer on the other publisher. The ad network perceives a smaller cost of increasing the ad quantity than the publishers if and only if the former loss is smaller, i.e. $r_j > \hat{r}_j$. Hence, under this condition, the provision of ads increases when the publishers outsource.²³

We expect the advertising revenue generated by single- and multi-homers within a given website to differ. For example, these two types of consumers may devote a different amount of attention or time to each website: if multi-homers devote the same total amount of attention to the publishers as single-homers, but spread it on several websites, we expect the probability that they register an ad on a given website to be smaller than the probability that single-homers do (i.e. $r_j > \hat{r}_j$). By contrast, if multi-homers devote substantially more time in total to browsing the publishers than single-homers do, the inequality may be reversed.

4.5 The effect of the ad network on consumers and welfare

We briefly summarize the implications of the above findings for consumers and social welfare. In our baseline model, the impact of the ad network for consumers depends on how it influences the quantity of ads they are exposed to. Therefore, consumers benefit when the publishers outsource to the ad network if and only if $q_i^* < q_i^C$. This condition is sufficient (though not necessary) for social welfare (see (6)) to increase as well, because the ad network increases the gross advertiser surplus.

Note that our baseline model ignores that consumers may dislike third-party tracking per se (e.g. for privacy reasons) and that tracking might make ads more or less annoying. We introduce these aspects in the following section, and briefly discuss their impact on consumer surplus.

5 Endogenous consumer tracking

We now extend the model allowing consumers to avoid tracking. For instance, consumers can block tracking by using browser extensions to reject third-party cookies or anonymize their online activity. We are interested in the implications of blocking for the advertising industry as well as for consumers.

²³When $\beta = 0$, (14) is identical to Proposition 2 in Ambrus et al. (2016), who analyze the effects of joint ownership of media outlets.

5.1 Modified setup

We extend the baseline model as follows. At stage 1, we let consumers decide whether to block across-outlet tracking.²⁴ Let $\beta_{cons} \in [0, 1]$ be the share of consumers that does not block. At the same stage, the ad network chooses the expected share of consumers it tracks, denoted by $\beta_A \in [0, 1]$. Therefore, the share of tracked consumers is $\beta = \beta_A \cdot \beta_{cons}$. However, to focus on the implications of consumers' choice whether to avoid tracking, throughout most of this section we assume the ad network tracks everyone who does not block (consumers are aware of this condition). That is, $\beta_A = 1$, so $\beta = \beta_{cons}$. We also take the publishers' choice to outsource to the ad network as given. We relax these two assumptions at the end of the section.

By assuming consumers choose whether to block tracking in stage 1, i.e. before deciding which content to browse, we capture the fact that most privacy-related choices (e.g., 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. In a similar vein, we model blocking as a binary choice.

We consider two channels whereby tracking affects consumer utility.²⁵ First, by tracking consumers, the ad network collects possibly sensitive information, which some consumers may perceive as invading their privacy (Goldfarb and Tucker, 2012). Accordingly, we assume consumers suffer a disutility θ when tracked, distributed according to the c.d.f. $F(\theta)$, with support $[0, \bar{\theta}]$ and density $f(\theta)$. This distribution is independent of consumers' preferences for content. Consumers who dislike tracking can block it at a cost $b \in (0, \bar{\theta})$. The way we model the privacy implications of tracking for consumers is deliberately simple, because our focus is not on such implications per se. Rather, we are interested in the consequences of consumers' avoidance of tracking for the advertising market.

The second channel whereby tracking can affect consumers is by changing the nuisance cost from ads for multi-homers. This effect may arise because tracking modifies the mix of impressions consumers are exposed to across the two publishers. To isolate the implications of each channel, we shall proceed in steps. We first consider the benchmark case where the disutility from ads is unaffected by tracking (Case *I*), i.e. $\delta^I = \delta$. Next, we consider case *R*, whereby tracked multi-homers suffer a greater disutility from ads, $\delta^R > \delta$. This case can capture, for example, situations where tracking is used for retargeting, which consumers tend to find annoying (Goldfarb and Tucker, 2011). Finally, in case *FC* tracking reduces the disutility from ads, $\delta^{FC} < \delta$. This case can capture situations where tracking is used mainly for frequency capping, making ads less repetitive. Thus, a tracked multi-homer's surplus from browsing the publishers is

$$u_1 + u_2 - \delta^x (q_1 + q_2), \quad x = I, R, FC. \quad (15)$$

Single-homers' utility from browsing any publisher remains the same as in our baseline model,

²⁴In principle, consumers could also block within-outlet tracking (e.g., first-party cookies). However, it is much more difficult for a consumer to block within-website tracking in practice, because it is often intertwined with the site's basic functionalities. Moreover, most anti-tracking tools only block third-party tracking.

²⁵Farrell (2012) argues that privacy is both a final good (i.e. consumers care about it for its own sake) and an intermediate one (i.e. consumers care about it in an instrumental way, depending on how the data collected are used).

because cross-outlet tracking does not change the type of ads they are exposed to.

Conceivably, when consumers set their browser's privacy settings, or decide whether to install privacy-protection software, they may not be fully informed about which websites they end up visiting and on how these websites match to their preferences. Accordingly, we solve the model under the assumption that when consumers choose whether to block, they are unaware of the utility from browsing publisher i , u_i , that will be realized only at stage 3.²⁶ However, consumers have rational expectations and can correctly calculate the expected utilities, u_i^e .²⁷ Furthermore, they can anticipate the equilibrium advertising levels. Given these assumptions, at stage 1 the expected utility of a consumer blocking and allowing tracking, respectively, are²⁸

$$u_1^e + u_2^e - \delta (q_1 + q_2) - b, \quad (16)$$

$$u_1^e + u_2^e - \delta^x (q_1 + q_2) - \theta, \quad x = I, FC, R. \quad (17)$$

Given the information available at stage 1, the ex ante consumer surplus, CS^e , is defined as:²⁹

$$CS^e = (u_1^e + u_2^e) - \left(\int_0^{\theta^x} \theta + \delta^x (q_1 + q_2) dF(\theta) \right) - (b + \delta (q_1 + q_2)) (1 - F(\theta^x)), \quad x = I, R, FC, \quad (18)$$

where θ^x is a threshold such that consumers allow tracking if and only if $\theta \leq \theta^x$ (we characterize this threshold below). The first term of this expression capture the expected gross utility from content. The second term captures the expected disutility of ads and from being tracked for those who do not block. The last term captures the expected disutility of ads and the cost of blocking for those who avoid tracking. The effect of tracking on consumer surplus (given the advertising quantities) is rather straightforward: a higher (lower) disutility from ads has a negative (positive) impact. There is also a negative impact due to the disutility from being tracked per se as well as the cost of blocking.

5.2 Case I: tracking does not affect the disutility from ads

Consider the benchmark case where the disutility from ads, δ , does not depend on tracking. The latter has no direct effect on consumers' choice of publishers, hence consumer demands are as in (1). The model is therefore identical to that in Section 4, except that β is endogenously determined

²⁶For completeness, we solve the model assuming that consumers have full information when choosing whether to block in Appendix C. The choice of blocking gets significantly more involved, but the main results do not change. Assuming incomplete information is in line with previous literature on internet media (see, e.g. de Cornière and Taylor, 2014, and Anderson et al., 2018).

²⁷Formally, this is the mean of the marginal distribution of u_i . That is, $u_i^e = \int_0^\infty E(u_i|v)h_j(v)dv$, $i = 1, 2; j \neq i$, where $E(u_i|v)$ is the mean of u_i conditional on $u_j = v$ and $h_j(\cdot)$ is the marginal density of u_j .

²⁸These expressions assume that consumers expect to browse both publishers ex ante. However, once the u_i are realized at stage 3, some consumers end up visiting only one, or none. The analysis is mostly unchanged if consumers expect to visit only one content. Details are available from the authors.

²⁹When evaluating regulatory interventions on tracking (which we consider below), CS^e is the relevant measure of consumer surplus. A regulator does not have more information than the consumers concerning their own preferences: at stage 1, the regulator does not observe the utility from content, u_i , $i = 1, 2$, which is realized only at stage 3.

at stage 1. Specifically, given (16) and (17), a consumer blocks tracking if and only if $\theta > b$, so that $\theta^I = b$ and $\beta = F(b)$.

We now analyze the implications of consumers avoiding tracking. Focus first on the effect of β on gross advertiser surplus, AS (characterized in (10)). Noting that $\frac{\partial AS}{\partial q_i} = 0$ in equilibrium, we have:³⁰

$$\frac{dAS}{d\beta} = (r_{12}^T - r_{12}^{nT}) D_{12}^I > 0. \quad (19)$$

In case I , increasing β only raises the returns from ads hitting multi-homers. Thus, consumers avoiding tracking impose a negative externality on the advertising industry. Note that any change in AS accrues to the ad network, due to its monopolistic position and bargaining power. In a more general setting, e.g. competing ad networks and advertisers with bargaining power (see Section 6), part of the change in gross advertiser surplus would accrue to advertisers and publishers as well.

The effect of avoiding tracking on consumer surplus is less straightforward. Differentiating (18) with respect to β and evaluating the resulting expression at equilibrium, we obtain

$$\frac{dCS^e}{d\beta} = \underbrace{\frac{\partial CS^e}{\partial \beta}}_{=0} + \sum_{i=1,2} \underbrace{\frac{\partial CS^e}{\partial q_i}}_{(-)} \underbrace{\frac{\partial q_i}{\partial \beta}}_{(+)/(-)}. \quad (20)$$

The first term on the right hand side, $\frac{\partial CS^e}{\partial \beta}$, captures the direct effect of changing the share of tracked consumers. This term equals zero in equilibrium, because only consumers who are indifferent between blocking and not modify their choice after a marginal change in β . The last term in (20) captures the fact that, as the extent of tracking changes, so does the quantity of ads and, thus, the surplus consumers get from the publishers. Given that $\frac{\partial CS^e}{\partial q_i} < 0$, this effect is positive if and only if tracking reduces the provision of ads, i.e. $\frac{\partial q_i}{\partial \beta} < 0$ (see Proposition 1 for the conditions driving the sign of this derivative). Each consumer ignores the effect of her own choice regarding tracking on the quantity of ads carried by digital publishers. Therefore, this decision entails an externality on all consumers. When $\frac{\partial q_i^*}{\partial \beta} > 0$ holds, blocking leads to less advertising and, hence, the externality is positive. Consequently, the equilibrium is such that consumers as a whole gain with less extensive tracking. However, if $\frac{\partial q_i^*}{\partial \beta} < 0$, blocking produces a negative externality on consumers, that are better off with *more extensive* tracking. This externality has so far gone unnoticed in the literature.

Proposition 3. *If $\frac{\partial q_i^*}{\partial \beta} < 0$, $i = 1, 2$ (respectively, $\frac{\partial q_i^*}{\partial \beta} > 0$, $i = 1, 2$), the equilibrium is such that a marginal increase in the extent of tracking leaves consumers better (resp. worse) off.*

Finally, let us consider the effect of avoiding tracking on social welfare, W . After netting out

³⁰All the following expressions are evaluated at the equilibrium ad capacities (q_1^*, q_2^*) , that maximize (10).

transfers between players, $W = CS^e + AS$. Evaluating at the equilibrium allocation, we have

$$\frac{\partial W}{\partial \beta} = \sum_{i=1,2} \underbrace{\frac{\partial CS^e}{\partial q_i}}_{(-)} \underbrace{\frac{\partial q_i}{\partial \beta}}_{(+)/(-)} + \underbrace{\frac{\partial AS}{\partial \beta}}_{(+)}.$$
 (21)

The first term is positive if and only if $\frac{\partial q_i^*}{\partial \beta} < 0$, as discussed above. Furthermore, we have $\frac{\partial AS}{\partial \beta} > 0$. Consequently, if $\frac{\partial q_i^*}{\partial \beta} < 0$, welfare increases with *fewer* consumers blocking. We discuss the policy implications of these findings in Section 7.

5.3 Case R and FC

We now present the main changes in the analysis when tracking affects consumers' disutility from ads, considering cases R and FC .

Each publisher's audience is composed of tracked and non-tracked consumers. Demands by non-tracked consumers remain as in (1). Instead, demands by tracked consumers change, because the disutility from ads depends on tracking. Specifically, a tracked consumer's utility from multi-homing is given by (15), whereas the net utility from multi-homing if non-tracked and from single-homing are, respectively, $u_1 + u_2 - \delta(q_1 + q_2)$ and $u_i - \delta q_i$, as in our baseline model. Utility maximization brings to the following demands by tracked consumers in case R :

$$\begin{aligned} D_i^R &= Pr \{u_i - \delta q_i \geq \max(0; u_j - \delta q_j); u_j < (\delta^R - \delta) q_i + \delta^R q_j\}, \quad i, j = 1, 2 \text{ and } i \neq j, \\ D_{12}^R &= Pr \{u_1 > (\delta^R - \delta) q_2 + \delta^R q_1; u_2 > (\delta^R - \delta) q_1 + \delta^R q_2\}. \end{aligned}$$
 (22)

The main difference with respect to the demand system in (1) is that, because tracking increases the disutility from ads when multi-homing, a tracked consumer that visits i is less likely to visit j as well, all else given. Indeed, visiting another publisher comes with an additional nuisance from ads equal to $(\delta^R - \delta) q_i + \delta^R q_j$. This extra nuisance exceeds that for a non-tracked consumer, δq_j .

In case FC , demands are given by

$$\begin{aligned} D_i^{FC} &= Pr \{u_i - \delta q_i > 0; u_j < \delta^{FC} q_j - (\delta - \delta^{FC}) q_i\}, \quad i, j = 1, 2 \text{ and } i \neq j, \\ D_{12}^{FC} &= Pr \{u_1 \geq \delta^{FC} q_1 - (\delta - \delta^{FC}) q_2; u_2 \geq \delta^{FC} q_2 - (\delta - \delta^{FC}) q_1; u_1 + u_2 \geq \delta^{FC} (q_1 + q_2)\}. \end{aligned}$$
 (23)

Contrary to case R , tracking reduces the disutility from ads for multi-homers. Hence, comparing (23) to (1), a tracked consumer that visits one publisher is more likely to visit the other as well, because the additional nuisance from ads, $\delta^{FC} q_j - (\delta - \delta^{FC}) q_i$, is smaller than that for a non-tracked consumer, δq_j .

To summarize, using superscript nT to denote demands by non-tracked consumers, the following inequalities hold:

$$D_i^R > D_i^{nT}, \quad i = 1, 2; \quad D_{12}^R < D_{12}^{nT};$$
 (24)

$$D_i^{FC} < D_i^{nT}, \quad i = 1, 2; \quad D_{12}^{FC} > D_{12}^{nT}.$$
 (25)

Note also that, in case FC , the consumer market might expand, i.e. $D_0^{nT} > D_0^{FC}$, because tracking reduces the disutility of ads when multi-homing.

Consider now the decision whether to block tracking. Given (16) and (17), we find that a consumer blocks tracking if and only if $\theta > \theta^x$, where θ^x is such that

$$\theta^x = \begin{cases} b - (\delta^R - \delta)(q_1 + q_2) & \text{Case } R, \\ b + (\delta - \delta^{FC})(q_1 + q_2) & \text{Case } FC. \end{cases} \quad (26)$$

The share of consumers that allows tracking is $\beta = F(\theta^x)$. Consumers are therefore more (respectively, less) likely to avoid tracking when the latter makes ads more (less) annoying.

Publisher i 's total audience is given by the weighted sum of demands from tracked and non-tracked consumers, i.e. $\beta D_i^x + (1 - \beta) D_i^{nT} + \beta D_{12}^x + (1 - \beta) D_{12}^{nT}$. Hence, tracking affects the composition of the publishers' audiences and, therefore, the gross advertiser surplus. This surplus, captured by the ad network in equilibrium, is given by:

$$p_{AN}^x = AS = \sum_{i=1,2} (\beta D_i^x + (1 - \beta) D_i^{nT}) r_i + \beta D_{12}^x r_{12}^T + (1 - \beta) D_{12}^{nT} r_{12}^{nT}, \quad x = R, FC. \quad (27)$$

Recalling that $\frac{\partial AS}{\partial q_i} = 0$ in equilibrium, we find that

$$\text{Case } R: \frac{dAS}{d\beta} > 0 \Leftrightarrow (r_{12}^T - r_{12}^{nT}) D_{12}^R > (r_{12}^{nT} - r_i) (D_{12}^{nT} - D_{12}^R), \quad (28)$$

$$\text{Case } FC: \frac{dAS}{d\beta} > 0 \Leftrightarrow (r_{12}^T - r_{12}^{nT}) D_{12}^{FC} > (r_{12}^{nT} - r_i) (D_{12}^{nT} - D_{12}^{FC}) - r_i (D_0^{nT} - D_0^{FC}). \quad (29)$$

These expressions suggest that, in specific circumstances, the gross advertiser surplus might decrease when the extent of tracking increases. To illustrate, let us focus on case R (see (28)). In this case, some multi-homers become single-homers when tracked ($D_{12}^R < D_{12}^{nT}$). Suppose the returns from advertising to non-tracked multi-homers exceed those from single-homers (i.e., $r_{12}^{nT} > r_i$), so the right hand side of (28) is positive. Hence, tracking increases gross advertiser surplus only if the increase in the value of ads that hit multi-homers (left hand side of (28)) is large enough. The interpretation of (29) for case FC is similar, the main difference being that tracking increases the share of multi-homers ($D_{12}^{FC} > D_{12}^{nT}$). Furthermore, given that $D_0^{nT} > D_0^{FC}$, the right hand side decreases, making a negative effect of tracking less likely.

Without tracking, the returns from advertising to single- and multi-homers may differ for several reasons. If multi-homers spread their attention across multiple websites, ads hitting them on either outlet may lose effectiveness compared to ads hitting single-homers. This can be the case, for instance, when the content provided by the two publishers overlaps, so that after visiting one publisher, the consumer spends relatively little time on the other. In addition, cross-outlet repetition makes ad campaigns less efficient (Athey et al., 2018). Thus, impressions hitting multi-homers may substantially lose value compared to those on single-homers. To summarize, we expect that $r_{12}^{nT} > r_i$

hold if multi-homers devote enough attention to each publisher compared to single-homers and if repetition across outlets is of relatively little concern to advertisers. However, the inequalities may reverse, e.g., when ads repeatedly hitting consumers on different outlets are highly wasteful and multi-homers are not very attentive.

Given tracking's potential to increase the value of impressions on multi-homers and how pervasive multi-homing is on the internet, we expect conditions such that $\frac{\partial AS}{\partial \beta} < 0$ hold rather rarely, and mostly when tracking is associated to practices that consumers strongly dislike (e.g., aggressive re-targeting). Furthermore, although our model is too stylized to analyze the ad network's optimal tracking strategy, we expect the network would tend to avoid using tracking in ways that reduce total advertising revenue.³¹

Finally, consider the effect of avoiding tracking on consumer surplus and welfare. Given that $\frac{\partial AS}{\partial q_i} = 0$ and $\frac{\partial CS^e}{\partial \beta} = 0$ in equilibrium (see Appendix A.3 for the proof), expressions (20) and (21) are unchanged and so is their interpretation. However, the effect of tracking on advertising quantities is more involved. As we show in Appendix B.3, we have that $\frac{\partial q_i^*}{\partial \beta} > 0$ if and only if

$$\begin{aligned} & D_{12}^x \left(\frac{\partial r_{12}^x}{\partial q_i} - \frac{\partial r_{12}^{nT}}{\partial q_i} \right) + \frac{\partial D_{12}^x}{\partial q_i} (r_{12}^x - r_{12}^{nT}) + \\ & + (D_{12}^x - D_{12}^{nT}) \left(\frac{\partial r_{12}^{nT}}{\partial q_i} - \frac{1}{2} \frac{\partial r_i}{\partial q_i} \right) + \left(\frac{\partial D_{12}^x}{\partial q_i} - \frac{\partial D_{12}^{nT}}{\partial q_i} \right) (r_{12}^{nT} - r_i) + \\ & + (D_0^{nT} - D_0^x) \frac{1}{2} \frac{\partial r_i}{\partial q_i} + \left(\frac{\partial D_0^{nT}}{\partial q_i} - \frac{\partial D_0^x}{\partial q_i} \right) r_i > 0 \quad x = R, FC. \end{aligned} \quad (30)$$

The terms in the first row of expression (30) replicate those in (13) (see the commentary to Proposition 1). The terms in the second and third rows of (30) indicate that, when tracking affects the disutility of ads, there is an additional effect due to the direct impact of β on the composition of audiences (see (27)). To illustrate, we focus on case R . Given $\delta < \delta^R$, some multi-homers stop browsing one publisher when tracked (i.e. $D_{12}^R < D_{12}^{nT}$). In addition, tracked multi-homers are more sensitive to an increase in advertising quantities, making them more likely to switch to single-homing ($\frac{\partial D_{12}^R}{\partial q_i} < \frac{\partial D_{12}^{nT}}{\partial q_i} < 0$). Suppose ads hitting multi-homers without tracking are sufficiently effective, compared to ads hitting single-homers (i.e. $\frac{\partial r_{12}^{nT}}{\partial q_i} > \frac{1}{2} \frac{\partial r_i}{\partial q_i}$ and $r_{12}^{nT} > r_i$). Then, by discouraging multi-homing, tracking is more likely to reduce the platforms' supply of ads. See our discussion above for why r_{12}^{nT} and r_i (and their derivatives) may differ. The last row of (30) is irrelevant in case R , because no consumers drop out of the market when tracked (i.e. $D_0^{nT} = D_0^R$).

In case FC , the interpretation for the terms in the second row of (30) is the same as above, but their sign is reversed because $\delta > \delta^{FC}$ and, therefore, $D_{12}^{FC} > D_{12}^{nT}$ and $\frac{\partial D_{12}^{FC}}{\partial q_i} > \frac{\partial D_{12}^{nT}}{\partial q_i}$. Moreover, because of the potential market expansion when tracking makes advertising less annoying (i.e. $D_0^{nT} > D_0^{FC}$), the last two terms in (30) are positive, raising the likelihood that the quantity of ads increases with tracking.

³¹The choice of how to use tracking information also depends on the objectives of the advertisers' campaigns. Modeling these aspects is beyond the scope of our analysis.

5.3.1 Should the publishers outsource to the ad network?

We have assumed the publishers outsource their ad inventories to the ad network. However, if the conditions such that $\frac{\partial AS}{\partial \beta} < 0$ hold, one may suspect that the publishers may be better off not outsourcing. Notice, though, that a negative impact of tracking on ad revenues does not imply that the publishers avoid outsourcing. As we show in Appendix B.4, there always exists a value of β_A the ad network can choose that makes outsourcing profitable. In the extreme, the ad network may avoid tracking altogether (i.e. set $\beta_A = 0$), while still making outsourcing worthwhile to the publishers given the centralization of the sale of ads.³² This outcome can be seen by comparing (8) to (27): the advertising revenue collected by the ad network exceeds the publishers' joint revenue when competing if $\beta = 0$.

6 Extensions

We now present further extensions to the baseline model of Section 3.

6.1 Tracking networks

In our baseline setting, the ad network sells the publishers' ad inventories. However, there are several third-party platforms that only sell tracking information to advertisers (e.g. Adobe Analytics). We now consider this kind of platform, referring to it as a *tracking network* (TN). We modify the baseline model of Section 3 as follows. At stage 1, the publishers may allow TN to track their consumers, in exchange for a payment.³³ At stage 4, the advertisers can acquire tracking information, paying p_{TN} to TN . At the same stage, they acquire ad impressions from the publishers.

We summarize the main results from this model (see Appendix B.5 for proofs). All advertisers acquire ads on both publishers and, if available, tracking information, generating the ad revenue described by (10). The equilibrium payments each advertiser makes to the publishers and TN , respectively, are:

$$p_i = r_i D_i + (r_{12}^{nT} - \hat{r}_j) D_{12}, \quad i = 1, 2, \quad p_{TN} = \beta D_{12} (r_{12}^T - r_{12}^{nT}). \quad (31)$$

Due to its monopolist position, TN captures all the value that tracking brings to the ads hitting the multi-homers, $r_{12}^T - r_{12}^{nT}$. Consequently, competing publishers earn as much as they do without the tracking network: p_i in (31) is as in (8). Therefore, TN can induce each publisher to grant access to its consumers paying a very small transfer. Furthermore, the advertisers can retain a larger share of the advertising revenue than with an ad network (equal to $(2\hat{r}_j - r_{12}^{nT}) D_{12}$ in equilibrium).

At stage 1, the publishers choose the same advertising capacities, (q_1^C, q_2^C) , as when there is no network. Hence, the tracking network does not affect the provision of ads, unlike an ad network

³²This result holds for any share of consumers blocking tracking, β_{cons} . Hence, it is not necessary to make assumptions about whether each consumer knows to be tracked or not.

³³This payment is not contingent on ad capacities, because TN does not manage the publishers' ad inventories.

that centralizes their sale (see Section 4.4). It follows that consumers are better off with a tracking network than with an ad network if and only if $q_i^C > q_i^*$ (see Proposition 2 for the conditions driving the sign of this inequality). We discuss the policy implications of this finding in Section 7.

6.1.1 Should the publishers allow the tracking network on their websites?

We have shown that the tracking network can obtain access to the publishers' consumers at virtually no cost if tracking does not affect the disutility from ads. However, if tracking is used in ways that change such disutility, the publishers may be better off avoiding the tracking network. To grasp the intuition, consider the model described in Section 5.1 and focus on case R , whereby tracking makes ads more annoying to consumers. The tracking network's revenue and that of publisher i are, respectively (see Appendix B.5 for the proof):

$$\beta D_{12}^R (r_{12}^T - r_{12}^{nT}), \quad (32)$$

$$r_i (\beta D_i^R + (1 - \beta) D_i^{nT}) + (r_{12}^{nT} - \hat{r}_j) (\beta D_{12}^R + (1 - \beta) D_{12}^{nT}), i = 1, 2. \quad (33)$$

Expression (33) suggests the publishers' revenue might decrease with tracking. In case R , tracking reduces the share of multi-homers and increases that of single-homers (see (24)). If a multi-homer generates substantially more revenue to a publisher than a single-homer (i.e. $r_{12}^{nT} - \hat{r}_j$ is large enough compared to r_i), tracking hurts the publisher. Hence, TN must pay a compensation to gain access to i 's consumers, which is too expensive if the extra revenue from the ads that hit multi-homers (see (32)) is not large enough. In this case, i is better off not accepting TN on its platform.

As in the discussion in Section 5.3, we expect that the revenue generated from multi-homers on each publisher (without tracking) exceeds that from single-homers if the former devote enough attention to each publisher compared to single-homers and if repetition across outlets is of relatively little concern to advertisers. However, the publishers may generate more revenues from single-homers when ads repeatedly hitting consumers on different outlets are highly wasteful and multi-homers are not very attentive.

Although the conditions such that tracking reduces advertising revenues might be quite specific, the finding that the publishers may prefer to avoid a tracking network differs from the case where the publishers' outsource their ad inventories to an ad network. In Section 5.3.1, we argued that outsourcing is profitable to the publishers even if tracking reduces total ad revenues. The reason is that the ad network collects more revenue from the advertisers by centralizing the sale of ads, even without tracking. Tracking networks cannot offer the same advantage because they do not reduce competition between the publishers (even if they control the extent of tracking).

6.2 Competing ad networks

We now introduce two (identical) ad networks in the baseline model (see Appendix B.6 for the proofs). An ad network increases gross advertiser surplus, by tracking multi-homers, and raises the share of this surplus captured by the supply side of the market. Intuitively, these functions

can be performed only if both publishers outsource to the same network. Otherwise, the networks compete to attract the advertisers and are unable to track consumers across outlets. As a result, the publishers outsource to the same network in equilibrium. This network deals with the advertisers as a monopolist. Hence, it implements the same ad quantities and collects the same revenue as in the baseline model. However, competition to attract the publishers forces the ad networks to transfer all this revenue to them.³⁴ In sum, competition between ad networks benefits the publishers more than advertisers and consumers.

Proposition 4. *If two ad networks are available, in equilibrium both publishers outsource to the same one. This network implements the same ad quantities, (q_1^*, q_2^*) , and collects the same revenue, p_{AN} , as a monopolist network. However, the publishers capture all this revenue.*

6.3 More than two publishers

Our results are robust to the inclusion of more than two publishers (see Appendix B.7). To convey the main ideas, consider the baseline model and assume there are $N = 3$ publishers. There are three sets of consumers on each publisher: single-homers, multi-homers shared between two publishers, and multi-homers shared among all publishers. If one publisher does not outsource to the ad network, the other two can still benefit from the network’s cross-outlet tracking technology and from centralizing the sale of ads. However, competition for advertisers gets more intense when some publishers do not outsource. As a result, we find that all the publishers outsource in equilibrium, as in our baseline setting. In addition, the revenues generated by the ad network and the forces that drive the choice of advertising quantities do not change. Hence, Propositions 1 and 2 are qualitatively unchanged.

We also study the case where the number of publishers is very large, that is $N \rightarrow \infty$. In this case, each publisher represents a tiny fraction of the market. Hence, its decision whether to outsource to the ad network has practically no effect on the profit and advertising capacity of the other publishers, as well as on the allocation of ads and revenue.³⁵

6.4 Contracts and negotiation between advertisers and platforms

In the baseline model, we assume the publishers and the ad network make take-it-or-leave-it offers to the advertisers, specifying a transfer in exchange for a certain quantity of impressions on their audience. In Appendix B.8, we show that the analysis is robust to other specifications of these contracts. In particular, we consider the case where the publishers or the ad network do not have all the bargaining power when they negotiate with advertisers. The advertisers thus retain a positive

³⁴As in our baseline model, each publisher outsources its entire advertising space to this network. Even if a publisher could outsource to different ad networks, it is better off dealing with a single one (see footnote 21).

³⁵To evaluate the case where the ad network is an essential link between publishers and advertisers, we also consider a scenario where a publisher cannot generate any revenue when it does not outsource to the ad network. The only aspect of the model that is affected is that the ad network has a stronger position when contracting with the publishers and thus retains a larger share of the gross advertiser surplus.

share of the gross advertiser surplus. However, because this share does not depend on advertising capacities, the first-order conditions regarding the choice of such capacities are as in the baseline model (see (12) and (9)). The equilibrium allocation of ads is also identical, and the main results of the analysis do not change. In Appendix B.8, we also provide an informal discussion of how our model could accommodate impressions being sold via second-price auctions.

6.5 Contracts between publishers and the ad network

We have assumed that the ad network offers to the publishers payments contingent on their ad capacity. We now discuss the robustness of our baseline results to other formulations of contracts between the ad network and the publishers (see Appendix B.9 for the proofs).

Suppose the ad network does not offer payments contingent on capacities, but a revenue sharing contract whereby the publishers receive a fixed share of the ad network’s total advertising revenue (see (10)). The publishers’ first-order conditions when choosing the ad capacity are as in (12). Thus, equilibrium capacities are identical to those in the baseline model and so are the main results.

Assume now the ad network shares with each publisher a percentage of the revenue earned from impressions on its domain. Proposition 1 remains the same. The intuition is that the publishers internalize the effect of their ad capacities on the total advertising revenue from multi-homers (and how this effect changes with tracking), because the ad network transfers to them a fixed share of this revenue. However, Proposition 2 is slightly modified. The main reason is that each publisher ignores the effect of its own ad capacity on the revenue from single-homers on the other outlet.

Finally, we consider the case where the transfers from the ad network to a publisher are not contingent on the outcome of negotiation with the other publisher (that is, we rule out the multilateral dimension from the contract). Again, Proposition 1 is unaffected. Even if the ad network takes into account that changes in the advertising capacity on one publisher affect the other publisher’s outside option, the latter is unaffected by tracking, because it refers to a scenario where the ad network is not present on both publishers. This effect changes Proposition 2 slightly, but not the main forces driving the effect of the ad network on ad quantities.

7 Policy implications

Privacy regulation and consumer consent. There is an ongoing debate regarding the effect of tracking technologies on consumers. Several regulators have considered mandating the inclusion of tools that facilitate the deletion of third-party cookies in Internet browsers.³⁶ A related intervention consists in implementing opt-in policies that decrease the perceived cost for consumers to avoid tracking, as opposed to opt-out policies.³⁷ In our model, these policies would reduce the cost of

³⁶The FTC recently considered a Do Not Track mechanism that would persistently record on a consumer’s device her choice not to be tracked. 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).

³⁷For example, the EU’s GDPR (2016) prescribes that “Consent should be given by a clear affirmative act”.

blocking, b , raising the share of consumers who block.³⁸

Assuming the online advertising market is unregulated (as it is currently the case), our findings of Section 5 suggest that interventions aiming to facilitate opting out from tracking may hurt consumers if they induce advertising quantities to go up. According to our model (see Proposition 1), this should be the case when consumers' preferences for the publishers are negatively correlated and tracking information is used mainly for targeting ads to specific consumers. The unintended consequences of privacy regulation may nonetheless be mitigated by regulating advertising quantities as well (see below). By contrast, if advertising quantities tend to go down when tracking is restricted, consumers would get a double dividend from privacy regulation.

Regulators can also increase the incentives to accept tracking while minimizing consumers' privacy losses (captured by θ in our model). Regulators may try to reduce the disutility that users perceive when tracked, e.g. by fostering the provision of more transparent information to consumers, limiting the timespan of data retention or requiring data to be anonymized and/or aggregated. Presumably, these measures can be implemented without drastically reducing the effectiveness of tracking.

Ad caps. It is straightforward to show that, in our model, the quantity of ads is excessive from a welfare standpoint, given platforms ignore the disutility suffered by consumers when deciding on ad quantity. This result lends support to the introduction of advertising caps. These caps are common in traditional media markets, but not in the commercial internet.³⁹ Interestingly, if tracking is likely to expand the quantity of ads (see Proposition 1), regulation facilitating opt out should move this quantity in the right direction. By contrast, if tracking tends to reduce the provision of ads, ad caps become more desirable with stricter privacy rules.

Configuring the supply-side of the advertising market. Setting privacy considerations aside, our analysis also has implications for the configuration of the supply-side of the advertising market. Should a regulator discourage the publishers from outsourcing the sale of ads to ad networks? In our baseline model, total advertising revenue is higher with the ad network. Furthermore, consumers benefit if the ad network restricts the provision of ads. Hence, when this condition applies (see Proposition 1 and 2), regulators should not impede outsourcing. Otherwise, regulators might favor the use of tracking networks and encourage the publishers to sell their own ad inventories directly.

8 Summary and Conclusions

We have studied the publishers' decision to make use of third-party platforms, distinguishing between ad networks and tracking networks. We found that outsourcing the sale of ads to an ad network is profitable for digital publishers. In contrast, the publishers do not necessarily gain by

³⁸The effect of reducing b is equivalent to exogenously reducing β , except for the mechanical effect on the costs sustained by those who block tracking. We formally characterize this effect in Appendix B.10.1.

³⁹See Anderson (2007) for an analysis of the effects of ad caps on the TV market.

allowing a tracking network on their websites. This is the case when tracking reduces the value of the publishers' ad inventory, because, unlike an ad network, the tracking network does not soften competition between advertising outlets. We found this adverse effect is more likely when advertisers use tracking information in ways that make ads more annoying for consumers (e.g. retargeting) and when a multi-homer generates substantially more revenue to the publishers than a single-homer (e.g., because cross-outlet repetition is limited and multi-homers devote a similar amount of attention to each publisher as single-homers).

The implications of third-party platforms for advertisers are rather different. Due to the higher concentration on the supply side, advertisers may not benefit when the publishers outsource to ad networks, particularly if their bargaining power is limited. This logic suggests that small advertisers are more likely to be worse off than large ones when the publishers use ad networks. To increase their bargaining power, small advertisers may therefore join demand aggregators, such as demand-side platforms. By contrast, tracking networks enable the advertisers to take advantage of tracking technologies without increasing concentration on the supply side of the market.

We have shown that an ad network's ability to allocate ads efficiently improves with the number of publishers the network attracts, provided their audiences overlap. Hence, connecting a large number of publishers results in a competitive advantage versus other networks. However, competition between networks makes attracting these publishers significantly more expensive (Proposition 4). These findings help to interpret the current configuration of the ad network market, characterized by one dominant ad network (Google Ad Sense) and several smaller networks that tend to avoid competition with other networks by focusing on specific audiences (e.g. Jumpstart for automotive advertising and RGM for luxury advertising).

Finally, the policy implications of this paper are relevant for the whole advertising industry, as it debates with regulators the design of policies intended to restrain online tracking. Regulation that limits tracking may harm the advertising industry as well as consumers if it induces digital platforms to carry more ads. We uncover a novel mechanism working through the advertising market and provide conditions such that tracking either expands or reduces online advertising levels. The latter outcome is more likely when consumers' preferences for the publishers are negatively correlated and when advertisers exploit tracking primarily for targeting purposes (see Proposition 1). These findings provide an argument in favor of a cautious approach in the design of regulation aiming to protect consumer privacy, contrasting with the general presumption that regulators should necessarily make it easier for consumers to avoid tracking.

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: Microeco-*

- nomics, 8, 189-222.
- [3] Anderson, S. P., 2007. Regulation of television advertising. In P. Seabright and J. von Hagen (Eds.), *The Economic Regulation of Broadcasting Markets*, 189–224. Cambridge University Press, Cambridge.
 - [4] Anderson, S.P. and S. Coate, 2005. Market Provision of Broadcasting: A Welfare Analysis. *Review of Economic Studies*, 72, 947-972.
 - [5] Anderson, S.P., Ø. Foros and Kind, H. J., 2018. Competition for Advertisers in Media Markets. *Economic Journal*, 128, 34-54.
 - [6] Anderson, S.P. and J. S. Gans, 2011. Platform Siphoning: Ad-Avoidance and Media Content. *American Economic Journal: Microeconomics*, 3, 1-34.
 - [7] Athey, S., E. Calvano and J.S. Gans, 2018. The Impact of Consumer Multi-Homing on Advertising Markets and Media Competition. *Management Science*, 64, 1477-1973.
 - [8] 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.
 - [9] Aziz, A. and R. Telang, 2019. Irrelevant Advertising: How More Consumer Tracking Can Make Ads Less Effective. Mimeo.
 - [10] Bagwell, K., 2007. The Economic Analysis of Advertising, in Mark Armstrong and Rob Porter (eds.). *Handbook of Industrial Organization*, Vol. 3, Amsterdam: North-Holland.
 - [11] Bergemann, D. and A. Bonatti, 2011. Targeting in Advertising Markets: Implications for Offline versus Online Media. *RAND Journal of Economics*, 42, 417-443.
 - [12] Bergemann, D. and A. Bonatti, 2015. Selling Cookies. *American Economic Journal: Microeconomics*, 7, 259-294.
 - [13] Berman, R., 2018. Beyond the Last Touch: Attribution in Online Advertising. *Marketing Science*, 37, 685-853.
 - [14] Bernheim D. and M.D. Whinston, 1985. Common Marketing Agency as a Device for Facilitating Collusion. *RAND Journal of Economics*, 16, 269-281.
 - [15] Casadesus-Masanell, R., A. Hervas-Drane, 2015. Competing with Privacy. *Management Science*, 61, 229-246.
 - [16] 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>
 - [17] 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.

- [18] de Cornière, A. and R. de Nijs, 2016. Online advertising and privacy. *RAND Journal of Economics*, 47, 48-72.
- [19] de Cornière, A. and G. Taylor, 2014. Integration and search engine bias. *RAND Journal of Economics*, 45, 576-597.
- [20] Farrell, J., 2012. Can privacy be just another good? *Journal on Telecommunications and High Technology Law*, 10, 251-264.
- [21] General Data Protection Regulation (GDPR), 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. Available at <http://data.europa.eu/eli/reg/2016/679/oj>.
- [22] George, L.M. and C. Hogendorn, 2012. Aggregators, search and the economics of new media institutions. *Information Economics and Policy*, 24, 40-51.
- [23] Goldfarb, A. and C. Tucker, 2011. Online display advertising: Targeting and obtrusiveness. *Marketing Science* 30, 389-404.
- [24] Goldfarb, A., and C. E. Tucker, 2012. Shifts in Privacy Concerns. *American Economic Review*, 102, 349-53.
- [25] 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.
- [26] Interactive Advertising Bureau (IAB), 2017. Internet advertising revenue report. Available at: https://www.iab.it/wp-content/uploads/2018/07/IAB-2017-Full-Year-Internet-Advertising-Revenue-Report.REV2_.pdf
- [27] Johnson, J.P., 2013. Targeted advertising and advertising avoidance. *RAND Journal of Economics*, 44, 128-144.
- [28] Kind, H. J., T. Nilssen, and L. Sørsgard, 2016. Inter-Firm Price Coordination in a Two-Sided Market. *International Journal of Industrial Organization*, 44, 101-112.
- [29] Kireyev, P., K. Pauwels, S. Gupta, 2016. Do display ads influence search? Attribution and dynamics in online advertising. *International Journal of Research in Marketing*, 33, 475-490.
- [30] Levin, J., and P. Milgrom, 2010. Online Advertising: Heterogeneity and Conflation in Market Design. *American Economic Review*, 100, 603-07.
- [31] Montes, R., W. Sand-Zantman and T.M. Valletti, 2019. The Value of Personal Information in Online Markets with Endogenous Privacy. *Management Science*, 65, 955-1453.
- [32] 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/>

- [33] Renault, R., 2015. Advertising in markets. In Anderson, S., Stromberg, D., and J. Waldfogel (eds.). Handbook of Media Economics. Amsterdam: Elsevier.
- [34] Roesner, F., T. Kohno and D. Wetherall, 2012. Detecting and Defending Against Third-Party Tracking on the Web. NSDI 2012.
- [35] Tucker, C. E., 2014. Social Networks, Personalized Advertising, and Privacy Controls. Journal of Marketing Research, 51, 546-562.
- [36] 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.

Appendix - Proofs of Propositions

A.1 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:

$$\frac{\partial q_i}{\partial \beta} = \frac{\frac{\partial FOD_i}{\partial \beta} \frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_j}{\partial \beta} \frac{\partial FOD_i}{\partial q_j}}{\frac{\partial FOD_i}{\partial q_j} \frac{\partial FOD_j}{\partial q_i} - \frac{\partial FOD_i}{\partial q_i} \frac{\partial FOD_j}{\partial q_j}}. \quad (34)$$

The denominator is always negative, because $\frac{\partial FOD_i}{\partial q_i} \frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_i}{\partial q_j} \frac{\partial FOD_j}{\partial q_i} > 0$ by the second order conditions (SOC) of the problem. Given symmetry, the numerator can be rewritten as $\frac{\partial FOD_i}{\partial \beta} \left(\frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_i}{\partial q_j} \right)$. Using the SOC, one can show that $\frac{\partial FOD_j}{\partial q_j} - \frac{\partial FOD_i}{\partial q_j} < 0$. Hence, $\frac{\partial q_i}{\partial \beta} > 0$ if and only if $\frac{\partial FOD_i}{\partial \beta} = D_{12} \frac{\partial^2 r_{12}^T}{\partial q_i \partial \beta} + \frac{\partial D_{12}}{\partial q_i} \frac{\partial r_{12}^T}{\partial \beta} > 0$. ■

A.2 Proof of Proposition 2

Let (q_1^C, q_2^C) be the couples of quantities that satisfy $\frac{\partial \pi_i^C}{\partial q_i} = 0$, for $i = 1, 2$, where π_i is given by (8). Also, let (q_1^*, q_2^*) be the couples of quantities that satisfy $\frac{\partial \pi_{AN}}{\partial q_i} = 0$, for $i = 1, 2$, with π_{AN} is defined in (11). As mentioned in the text (see footnote 9), we assume conditions such that the equilibrium couples (q_1^C, q_2^C) and (q_1^*, q_2^*) exist and are unique hold. By assumption, therefore, both π_i and π_{AN} are strictly concave with respect to q_i . Evaluating (12) in (q_1^C, q_2^C) and using (9), we have, for $i = 1, 2$,

$$\begin{aligned} \frac{\partial p_{AN}}{\partial q_i} = & D_{12}(q_1, q_2) \beta \left[\frac{\partial r_{12}^T(q_1, q_2)}{\partial q_i} - \frac{\partial r_{12}^{nT}(q_1, q_2)}{\partial q_i} \right] + \\ & + \frac{\partial D_{12}(q_1, q_2)}{\partial q_i} \left[\beta (r_{12}^T(q_1, q_2) - r_{12}^{nT}(q_1, q_2)) + \hat{r}_j(q_j) \right] + \frac{\partial D_j(q_1, q_2)}{\partial q_i} r_j(q_j) \Big|_{q_1=q_1^C, q_2=q_2^C}. \end{aligned} \quad (35)$$

By concavity of π_{AN} , $q_i^* > q_i^C$ if and only if $\frac{\partial p_{AN}}{\partial q_i} > 0$ when evaluated at (q_1^C, q_2^C) , for $i = 1, 2$. Therefore, using the equality $-\frac{\partial D_{12}(q_1, q_2)}{\partial q_i} = \frac{\partial D_j(q_1, q_2)}{\partial q_i}$, $q_i^* > q_i^C$ if and only if (14) is satisfied. ■

A.3 Proof of Proposition 3

Ex ante consumer surplus, (18), is such that $\frac{\partial CS^e}{\partial q_i} < 0$. By Proposition 1, the equilibrium quantity of ads on publisher i , q_i^* , can increase or decrease in β . We here show that $\frac{\partial CS^e}{\partial \beta} = 0$ when evaluated at equilibrium. Consider case $x = \{I, R, FC\}$, where $\beta = F(\theta^x)$. Because $\theta^x = F^{-1}(\beta)$ and $(F^{-1}(\beta))' = 1/f(\theta^x)$, we can write the derivative of consumer surplus with respect to β as:

$$\frac{\partial CS^e}{\partial \beta} = -(\theta^x + \delta^x(q_1 + q_2)) \frac{f(\theta^x)}{f(\theta^x)} + (b + \delta(q_1 + q_2)).$$

In case I , we have $\theta^I = b$, and $\delta^I = \delta$, hence we find $\frac{\partial CS^e}{\partial \beta} = 0$. The same can be shown for case R , where $\theta^R = b - (\delta^R - \delta)(q_1 + q_2)$, and case FC , where $\theta^{FC} = b + (\delta - \delta^{FC})(q_1 + q_2)$. ■