Intermediaries in the Online Advertising Market*

Anna D'Annunzio

Toulouse Business School, CSEF (University Federico II) and Toulouse School of Economics

Antonio Russo

Loughborough University and CESifo

March 30, 2022

Abstract

A large share of the ads displayed by digital publishers (e.g., newspapers and blogs) are sold via intermediaries (e.g., Google), that have large market power and reportedly allocate the ads in an opaque way. We study the incentives of an intermediary to disclose consumer information to advertisers when auctioning ad impressions. In turn, we study how disclosure affects the incentives of publishers to outsource the sale of their ads to an intermediary, and relate these incentives to the extent of consumer multi-homing, the competitiveness of advertising markets and the ability of platforms to profile consumers. We show that disclosing information that enables advertisers to optimize the allocation of ads on multi-homing consumers is profitable to the intermediary only if advertising markets are sufficiently thick. Even when most consumers multi-home, the publishers may be worse off by outsourcing to the intermediary, in particular if they operate in thin advertising markets. Finally, we study how the intermediary responds to policies designed to enhance transparency or consumer privacy, and the implications of these policies for the online advertising market.

Keywords: online advertising, intermediary, multi-homing, privacy, transparency

JEL Classification: D43, D62, L82, M37

*D'Annunzio: dannunzio.anna@gmail.com, Russo: a.russo@lboro.ac.uk. We gratefully acknowledge the contribution by Christian Peukert to the first stages of this work. We also thank Doh-Shin Jeon, Bruno Jullien, Leonardo Madio and participants to presentations at Nottingham University Business School, the Louvain and Paris Online Seminar, the OsloMet seminar, TSE digital seminar, the Scientific Seminar on the Economics, Law and Policy of Communications and Media at EUI in Florence, the Workshop on the economics of platforms (ECOP) in Bologna, for valuable comments. Part of this research was carried out while Anna D'Annunzio was at the University Federico II, Naples.

1 Introduction

The functioning of the online advertising market is a fascinating and important topic.[1] Advertising is key to many digital content providers and publishers, such as newspapers, blogs and review websites, with significant spillovers on the rest of society.[2] In addition, advertising affects the prices that consumers pay for goods and services (Bagwell, 2007).

Digital publishers typically provide "display" ads.[3] A striking feature of the market for this kind of advertising is the major role played by intermediaries. More than 60 percent of display ad spend is traded via intermediaries (IAB, 2017). Furthermore, although there is a complex chain of intermediaries connecting advertisers to publishers, a major firm (Google) has a dominant position in every link of this chain (see Figure 1 for an illustration).[4] In this paper, we study the behavior of a large advertising intermediary, focusing on the transparency of auctions for ad impressions. In addition, we study whether and when the publishers gain by using the intermediary, and consider the implications for the advertising market. Finally, we examine how transparency and privacy regulation can affect the choice of information disclosure by the intermediary.

Conceivably, a large intermediary like Google presents several attractive features to digital publishers, particularly when consumers visit multiple online outlets in a short time frame (multi-home). The intermediary centralizes the sale of ads from multiple publishers and can typically achieve a more precise profiling of consumers. Furthermore, it can coordinate consumers' exposure to ads on each publisher, e.g., by limiting the cross-outlet ad repetition that hinders the reach of advertising campaigns (Athey et al., 2018),[5] or allowing more effective re-targeting. Hence, the intermediary can increase total surplus in the advertising market, as well as the share captured by the platforms on the supply side. In practice, however, the functioning of intermediaries is quite complex and obscure. Regulators and market operators have pointed out the lack of transparency in how intermediaries run the auctions for ad impressions and, more generally, in how they allocate these impressions. One of the main concerns is that these platforms strategically retain valuable information from advertisers,

[1] Global digital advertising spending amounted to about USD 280 billion in 2018, and about USD 330 billion in 2019. https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/.

[2] For example, ad revenue drives investment decisions in the quality of content by digital publishers in crucial domains such as journalism. On the relation between competition in the online advertising market - with particular regard to the role of large platforms - and the viability of high quality journalism, see chapter 4 of the Cairncross Review (Cairncross, 2019).

[3] Display ads represent one of the three main segments of the digital advertising market (the other two being "search" ads and "classified" ads). In the UK, the display advertising market was worth GBP 5.5 billion in 2019, compared to GBP 7.3 billion for the search advertising market (CMA, 2020).

[4] The chain includes supply-side platforms (SSPs) that collect ad inventories from publishers and run ad auctions; demand-side platforms (DSPs) that allow advertisers to buy ad inventories; publisher ad servers, that manage publishers' inventory and decide which ad to serve, based on the bids received from SSPs and direct deals between the publisher and advertisers. Google has virtually a monopoly in the ad server market, and also dominates the SSP and DSP segments (CMA, 2020).

[5] See the emphasis on *unique* users and impression repetition in Google's own campaign evaluation tools. See https://support.google.com/google-ads/answer/2472714?hl=en&ref_topic=3123050.

Figure 1: Google is the dominant intermediary in online display advertising. Source: CMA (2020)

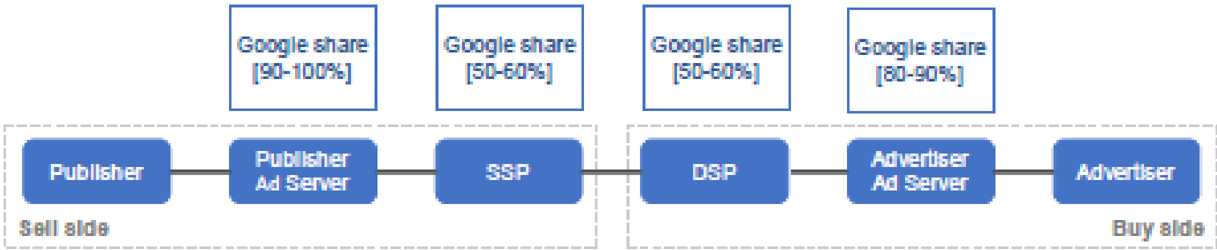

for example by making it difficult to assess the effectiveness of ad campaigns and their reach (CMA, 2020).

The observations above raise several interesting questions. What drives the incentives of an intermediary to retain information from advertisers when auctioning ad impressions? How does transparency (or lack thereof) by the intermediary affect the efficiency of the market and the distribution of profits? Do digital publishers necessarily gain by selling their advertising space via an intermediary? To address these questions, we consider a simple setting with two publishers and an intermediary. In our setting, consumers either visit one publisher or both, being exposed to one ad impression per visit. Each consumer is characterized by a type, that corresponds to an advertising market, i.e., a set of advertisers that intend to reach (only) her type. We characterize the advertising markets by their thickness (i.e., the number of advertisers) and by the marginal returns to advertising. Ad impressions are allocated to advertisers via auctions, possibly by the intermediary if the publishers outsource their ad inventories. If they do, the intermediary can gather more accurate information about consumers than the publishers, which allows to sell more targeted ads, and observes whether a consumer receives ad impressions on different publishers. This information enables the advertisers to manage consumers' frequency of exposure to their messages across different outlets.

In the baseline model, we assume there are diminishing returns to advertising to the same consumer (e.g., because the consumer may have be already informed by the ad). Consequently, disclosing information about a consumer's browsing history entails a key trade-off for the intermediary: bids increase from advertisers that are not reaching the consumer on another outlet, but those that already reach her elsewhere lose interest, so the market effectively "thins out". If the consumer belongs to a thin market and multi-homes, the market-thinning effect of information disclosure dominates, resulting in a sharp drop in the price of impressions. Therefore, if many consumers belong to thin markets, the intermediary is better off reducing the amount of information disclosed to the advertisers for impressions on multi-homers.

We explore two dimensions along which the intermediary can dilute the information available to advertisers. First, by retaining information about the consumers' browsing history, preventing the advertisers from capping the frequency of exposure to their ads. Alternatively, the intermediary can "thicken" the set of advertisers interested in a consumer by revealing less granular information about her preferences. Despite the key role of the potential repetition of ads in determining the market-thinning effects of disclosure, we find that, conditional on disclosing only part of the information available, the intermediary should opt for less granluar preference targeting, but let the advertisers control ad frequency. The intuition is that, due to the risk of repetition, impressions on the two publishers are substitutes for the advertisers if they cannot control frequency. Hence, in equilibrium the advertisers single-home on different publishers. Thus, each advertiser is unwilling to bid a high price for the impressions sold by the other publisher that fall on multi-homers, because the advertiser already reaches these consumers on one publisher. If the market is thin, the result is a drop in the revenue from impressions similar to the case of full disclosure.

It is natural to wonder whether the publishers gain by using the intermediary to sell their ads, since the superior information at its disposal make backfire. We find that outsourcing is more likely to be profitable to the publishers if their audience belongs to thick advertising markets. In such markets, the intermediary can not only sell a higher volume of targeted impressions (given its superior ability to profile consumers), but also at a higher price, by disclosing information that allows the advertisers to avoid repetition. However, if a substantial share of the publishers' audience belongs to thin advertising markets, outsourcing to the intermediary may reduce the equilibrium price of targeted impressions, particularly when many consumers multi-home.

The intermediary affects the size and distribution of surplus in the advertising market. The publishers outsource if and only if the intermediary increases the revenue generated from their ad inventory. However, the advertisers are not necessarily better off, because competition among advertising outlets weakens. Moreover, despite its superior ability to target consumers and allocate ads, the intermediary may reduce the total surplus in the advertising market. The reason is that some impressions are wrongfully targeted when the intermediary dose not fully disclose information on multi-homing consumers.

Regulation imposing full transparency to the intermediary (see, e.g., the proposed remedies in CMA, 2020, p. 395), may either increase or reduce total advertising surplus. Given that outsourcing occurs, full disclosure raises the efficiency of targeted ads on multi-homers. However, by reducing the revenue the intermediary is able to generate, may induce the publishers not to outsource, reducing the extent to which the industry as a whole benefits from the intermediary's superior tracking capabilities.

We then extend the model allowing consumers to take steps to protect their privacy by blocking tracking by the intermediary (e.g., rejecting third-party cookies). We let the intermediary's disclosure policy affect the consumers' incentives to block, e.g., by determining the relevance of the ads that tracked consumers see. Therefore the intermediary must adopt a two-sided logic when choosing its degree of transparency, considering not only the effect on the revenue from ads on tracked consumers, but also how consumers react. Within this setting, we analyze the consequences of regulation designed to protect consumer privacy, such as the European Union's GDPR, on the choice of transparency by the intermediary. We show that the regulation may result in either more or less transparency, depending on whether consumers prefer non-targeted ads or target ads, respectively.

Finally, we present several extensions to the baseline model. These include reserve prices in advertising auctions, increasing returns to advertising, competition among advertisers in the product market and heterogeneous advertising returns within each advertising market.

The remainder of the paper is organized as follows. Section 2 discusses previous related literature. Section 3 presents the model, that we solve in Section 4. We provide a welfare analysis and an evaluation of the effects of different regulatory policies in Section 5. Section 6 provides an overview of the extensions. Section 7 concludes, and discusses the policy and managerial implications of our analysis. Proofs of lemmas and propositions not given in the text are in the Appendix.

2 Literature

This paper contributes to several strands of literature. A recent literature studies the incentives of platforms to disclose information about consumers to advertisers. Bergemann and Bonatti (2015) consider an online data provider that sells information to advertisers targeting specific consumers. The authors study the demand for this information by advertisers and the optimal pricing policy of the data provider. Xiang and Sarvary (2013) consider competing data providers and analyze how competition among advertisers affects the price of data. Gu et al. (2019) show that the nature of competition between data brokers depends on the value of merged versus separate data-sets. Unlike these papers, we focus on platforms that sell advertising impressions, as well as information, to advertisers. de Cornière and de Nijs (2016) consider an advertising-financed publisher that has detailed information about consumers. When disclosed to an advertiser, this information puts the advertiser in a monopolist position when selling its products. The publisher thus discloses the information if and only if there is sufficient competition among advertisers, which enables the publisher to capture the ensuing monopoly rents.[6] Although information disclosure by platforms is central in our paper as well, we focus on an intermediary that connects publishers to advertisers and study how the managing of information by the intermediary affects the structure of the market, i.e. whether the publishers choose to use an intermediary to sell their ad inventories.

Few previous papers have considered the disclosure of information by third-party platforms that connect advertisers to publishers and how this disclosure affects the market for ad impressions. Levin and Milgrom (2010) discuss informally the incentives by platforms to conflate ad impressions, referring to the case where "similar but distinct products are treated as identical in order to make markets thick or reduce cherry-picking". We model the relation between the attributes of advertising markets (e.g.,

[6] There is also a literature studying the effects of targeted advertising on consumers and product markets without considering platforms. See, e.g., Johnson (2013) and Karle and Reisinger (2019).

thickness and multi-homing) and the platform's choice of information disclosure. Furthermore, we embed this relation in a model that studies the interplay between digital publishers, the intermediary platform and advertisers. Ghosh et al. (2015) consider a setting with digital publishers and a third-party platform that collects information about consumers in the publishers' audience. They show that the platform can use the information gathered about consumers visiting a publisher to enable the advertisers to target the same consumer on other publishers at a lower price (information leakage). This result is somewhat related to our findings, although the mechanism we focus on is different. Furthermore, we explore the conditions such that the publishers rely on the third-party platform to sell their ad inventories and the ensuing effects on the volume and distribution of surplus in the advertising market.

Only a limited number of papers study how large intermediaries affect key outcomes in the online advertising market. Marotta et al. (2021) study the welfare implications of consumer data sharing in targeted advertising, evaluating the role of an advertising intermediary as well the effects of restricting data sharing. The authors model the effects of information on competition among the advertisers on the product market, but do not consider digital publishers. Sharma et al. (2019) consider an asymmetric duopoly of ad networks. The authors focus on the analysis the contractual arrangements between publishers and ad networks, and on how the publishers sort across the different networks. Furthermore, they study the implications of data frictions (such as data sharing regulations or digital ad taxes). However, they do not consider the allocation of ads and the disclosure of information by the ad networks. D'Annunzio and Russo (2020) consider an ad network that centralizes the sale of ads in presence of multi-homing by consumers and advertisers, focusing on how the ad network affects the equilibrium quantity of ads and on the implications of consumers' avoidance of third-party tracking. Unlike the above papers, we explore the type and the extent of information disclosed by a dominant intermediary to advertisers, and how this decision relates to market fundamentals. We can therefore study the interplay between transparency and the choice of digital publishers to rely on intermediary for the sale of ads. Furthermore, we focus on how the intermediary affects the allocation and price, rather than the quantity of ads displayed by digital publishers. In a recent paper, Peitz and Reisinger (2020) show that an ad network that tracks consumers across outlets can have a negative effect on the equilibrium price of impressions on multi-homers, because advertisers do not want to reach the same consumer too many times. We take a different approach to modeling advertising auctions, and find a similar outcome only if advertising markets are thin, in particular when the intermediary fully discloses information to advertisers. However, we let the intermediary choose the degree of information disclosure, and show that full disclosure is suboptimal in thin advertising markets.

Our analysis also relates to the literature on information disclosure in auctions. Ganuza and Penalva (2010) show that disclosing more information about the object for sale increases some buyers' willingness-to-pay, but may also increase the buyers' information rents. As a result, they find that disclosing information is not profitable when there are few bidders.[7] Bourreau et al. (2017) find similar results in a setting with competing auctioneers. Hummel and McAfee (2016) analyze the incentives to disclose information for different auction formats. Rafieian and Yoganarasimhan (2021) show theoretically and empirically an ad network's revenues may decrease when user information is used to target advertising. We confirm these fundamental insights in a context where platforms choose whether to disclose not only information on consumers' characteristics, but also their browsing history and exposure to ads on multiple outlets. Moreover, we analyze how the intermediary's choice to disclose information affects publishers' choice to outsource the sale of ads.

Finally, in a broader perspective, our analysis relates to the literature that analyzes the incentives of platforms to deter advertising click fraud and to disclose related information to advertisers when auctioning ad impressions. Wilbur and Zhu (2009) show that the incentives to disclose click fraud by a platform that sells search ads increase with the competitiveness of the advertising market, which is consistent with our findings.

3 The Model

3.1 Setup

We consider a setting with three platforms: two digital publishers, $i = 1, 2$, and an intermediary, IN. The publishers provide free content to consumers and sell ad impressions to the advertisers, either directly or via IN.

Consumers. There is a unit mass of consumers. Let m be the (exogenous) quantity of multi-homers and $\frac{1-m}{2}$ the quantity of single-homers on each publisher. Each consumer is characterized by a type, θ, which summarizes a set of characteristics such as interests (culture, sports, news, etc.), demographics and geographic location. These characteristics determine the consumer's relevance to the advertisers. We let θ be distributed among consumers according to the uniform distribution $G(\theta)$ with support $[0, 1]$ (so that $g(\theta) = 1$, $\forall\theta$) and assume this distribution is independent of the allocation of consumers on the publishers. The quantity of impressions that each publisher exposes a consumer to is given and set to one. Therefore, single-homing consumers receive one impression and multi-homers receive two impressions in total.

Advertisers. Ads inform consumers about products. Let $k(\theta)$ be the set of advertisers that intend to reach type-θ consumers. An ad generates a positive return to an advertiser in $k(\theta)$ only if it informs a type-θ consumer, and zero otherwise. We refer to each type θ as a separate *advertising market*, because only advertisers in $k(\theta)$ intend to reach consumers of that type. We assume each consumer

[7] Chen and Stallaert (2014) obtain a similar result in a setting where a publisher decides whether to use behavioral information when selling ads.

belongs to one and only one market (for instance, the consumer can only be in one location at a time). We say that an ad impression is *targeted* if the platform selling the impression reveals to the advertisers that the consumer's type belongs to a finite set of values. Given a continuum of types, the expected return from informing a consumer without this information is zero. Hence, non-targeted impressions are worthless to the advertisers.

Each advertising market is characterized by two parameters. First, the number of advertisers, n, that we refer to as the market's *thickness*. We refer to markets such that $n = 2$ as "thin", to markets such that $n = 3$ as "intermediate" and to markets such that $n \geq 4$ as "thick".[8] Let x, y and $1 - x - y$ be the share of thin, intermediate and thick markets, respectively. At times, we will refer to n also as the market's competitiveness. The thickness of an advertising markets may depend, for example, on the number of sellers of a given product operating in a certain geographical area. To ease the exposition, we solve the baseline model assuming $y = 0$, then we consider the full model including intermediate markets in Appendix B.

The second parameter characterizing advertising markets is the advertisers' marginal return from informing a relevant consumer. We denote this return by v and assume it is distributed according to a distribution $F(v)$ with smooth density, $f(v)$, support $[0, v_H]$, and mean $\overline{v} \equiv \int_0^{v_H} v\, dF(v)$. The advertising return may depend, e.g., on product margins and/or on the probability that informing a consumer converts into a sale. In the baseline model, v is homogeneous among advertisers within a market (see Section 6.5 for a rubustness check).

The distributions of n and v across markets are independent and common knowledge. However, the realization of these parameters is private information of the advertisers in each market, and thus unobservable to the platforms.

Each advertiser maximizes its expected return from ads, net of the prices paid to the platforms. We assume impressing a consumer with one ad is enough to inform her. Sending the same ad twice to the consumer is thus wasteful. Therefore, as will become clear, an advertiser's willingness-to-pay for an ad impression depends on whether the consumer (i) belongs to the relevant market and (ii) may receive the same ad while visiting another publisher. In the baseline model, an advertiser's marginal return from informing a consumer does not depend on the consumer being exposed to ads from other advertisers in the same market. In other words, the advertisers do not compete in the product market. This assumption is relaxed in an extension, summarized in Section 6.2.

Publishers. The publishers earn revenue only from the sale of ads and incur no costs. If a publisher does not outsource to IN, it sells each impression in a first-price auction. All auctions take place simultaneously. If there is more than one winning bid for an impression, the publisher allocates the impression randomly to one of the top bidders.

[8] The definition of market thickness is relative to the maximum quantity of impressions available on each consumer, which is two in our model.

When selling its impressions directly, each publisher generates a signal σ for each consumer (i.i.d. across consumers), that conveys information about the consumer's type based on the information collected on the website (e.g. using first-party cookies). This signal is perfectly informative (i.e., $\sigma = \theta$) with probability q and is pure noise otherwise. When $\sigma = \theta$, we say that the consumer is *profiled*. We assume σ is i.i.d. across consumers and publishers. However, each publisher does not observe whether a consumer visits -and is thus exposed to ads on- the other publisher.

Intermediary. At the beginning of the game, IN makes to each publisher $i = 1, 2$ a simultaneous take-it-or-leave-it offer specifying a transfer T_i for the publisher's ad inventory. If a publisher outsources to the intermediary, the latter sells each available impression in simultaneous first-price auctions. Similarly to the publishers, the intermediary generates a signal about the consumer's θ, σ^{IN}, that is perfectly informative with probability $\tilde{q}$ and uninformative otherwise. If only one publisher outsources, IN obtains the same information about consumers as the publisher does, and profiles each consumer with their same probability, $\tilde{q} = q$. If both publishers outsource, however, the intermediary can track consumers on both outlets, which allows it to profile each consumer with higher probability, i.e. $\tilde{q} > q$, and to observe which publishers a consumer visits and control which ads she is exposed to during each visit.

For each profiled consumer (i.e., for which $\sigma^{IN} = \theta$), the intermediary chooses the degree of transparency to advertisers. We let the intermediary decide whether to fully reveal information on (i) the consumer's type and (ii) her browsing history. The first piece of information allows the advertisers to target relevant conusmers. The latter piece of information is valuable due to diminishing returns of repeated impressions across publishers, and allows to control the frequency of the ad campaign. Specifically, the intermediary decides among the following information disclosure regimes for each impression on a profiled consumer:

1. *Full Disclosure* (F): IN discloses θ and whether the consumer is a single- or a multi-homer.

2. *Partial Disclosure on Type* (PT): IN discloses whether the consumer is a single- or a multi-homer and that σ^{IN} takes one among a finite set Θ of $t \geq 2$ values, including θ.

3. *Partial Disclosure on History* (PH): IN discloses the consumer's θ, but not whether the consumer is a single-or a multi-homer.

The F regime implies maximum transparency to advertisers. If the intermediary chooses some form of partial disclosure, impressions are conflated, meaning that impressions on different consumers are sold as if they were the same object (see Levin and Milgrom, 2010). In the PT regime, the intermediary limits the granularity of ad targeting, by conflating different advertising markets. An example is the conflation of geographically differentiated markets. For instance, instead of revealing the consumer's exact location, IN could only reveal her postcode. Finally, in the PH regime, IN

conflates impressions along a different dimension, by not distinguishing between single- and multi-homing consumers. This prevents the advertisers from managing the frequency of ad exposure on consumers that switch across multiple publishers.

Timing. We summarize the model by describing the timing of moves:

1. IN offers T_i to publisher $i = 1, 2$ in exchange for the publisher's ad inventory. Each publisher accepts or refuses.
2. Consumers visit the publishers and all impressions opportunities are generated.
   (a) If one or no publisher outsourced, the platforms profile each consumer with probability q.
   (b) If both publishers outsourced, the intermediary profiles each consumer with probability $\tilde{q}$.
3. If both publishers outsourced, IN chooses the information regime, r, for each impression. The platforms (publishers or intermediary) sell each available impression in simultaneous first-price auctions. The advertisers bid for each impression simultaneously.
4. Impressions take place and all payoffs are realized.

We focus on pure strategy Subgame-Perfect Nash Equilibria.

3.2 Discussion of the setup

We briefly discuss some of our assumptions. Consistently with stylized facts of the market for "open display" digital ads, we assume the intermediary is a monopolist.[9]
We focus on first-price auctions as the mechanism for allocating ad impressions because most digital intermediaries in the digital display market run auctions based on this format. Google's ad echange moved to first-price auctions in 2019 (CMA, 2020). Our results would be identical if we considered second-price auctions, since the revenue equivalence principle applies in our setting. We assume the platforms do not impose reserve prices in the baseline model, but allow for such prices in an extension (see Section 6.1) and show that results are qualitatively unchanged.

In keeping with the literature on advertising-financed platforms (e.g., Anderson and Coate, 2005; Ambrus et al., 2016; Athey et al., 2018), we assume there are diminishing returns to advertising. In particular, our model is geared towards analyzing the implications of diminishing returns across multiple outlets. By normalizing the number of impressions received by a consumer on each publisher to one, we effectively rule out the possibility that a consumer receives the same ad multiple times on the same publisher. This possibility is typically not a major concern in realty, since digital publishers have the means to manage the frequency of ads within their own domains, e.g. using

[9] Google is by far the largest intermediary. It has among 90-100% market share in the publisher ad server market, 80-90% in the advertiser ad server market and 50-60% as SSP and DSP (CMA, 2020).

first-party cookies. We focus on cross-outlet repetition, which is a more relevant challenge in the management of ad campaigns (Athey et al., 2018).

To economize on notation, we set to zero the marginal return to sending an ad to a consumer more than once. Letting this return be positive would have no fundamental impact on our results as far as there are diminishing returns. In an extension (see Section 6.4), we assume increasing returns to advertising on the same consumer (as in, e.g., re-targeting). For the sake of simplicity, the baseline model also assumes that the return from informing consumers is homogeneous across advertisers within the same market. We relax this assumption in Section 6.5, showing that are main conclusions qualitatively hold.

To streamline the exposition, we assume the signal σ about the consumer's type, is either perfectly informative or completely uninformative. Allowing for some noise in the signal (e.g., allowing the signal to show an incorrect value of θ) would reduce the advertisers' willingness-to-pay for a targeted impression, because the impression could end up on the "wrong" consumer. However, the structure of the equilibrium bids, and the allocation of ads would not be affected.

We model the contractual arrangements between the publishers and the intermediary in a stylized way, but this assumption is not crucial for our results. For instance, we would obtain the same results if we assumed that the intermediary transfers to the publishers a given share of the revenue from the sale of impressions, as we argue in Section 4.3.

Finally, we assume the intermediary can apply its superior targeting technology ($\tilde{q} > q$) only if it can gather data about consumers from both publishers. Alternatively, one could assume that the probability the intermediary profiles consumers is $\tilde{q}$ regardless of how many publishers outsource, even if the intermediary would not observe whether a consumer is a single- or a multi-homer if only one publisher outsources. This assumption would make the analysis slightly more lengthy (we should also solve the subgame where only one publisher outsources) without yielding different results.[10]

4 Solving the model

We begin by considering the equilibrium of the subgame where the publishers outsource to the intermediary and characterizing its profit-maximizing disclosure regime. Next, we consider the equilibrium in the subgame where the publishers do not outsource. Finally, we consider the first stage of the game, where the publishers decide whether to outsource or not.

Before proceeding, it is useful to state two simple results that apply throughout the analysis. First, in a first-price auction the advertiser with the highest willingness-to-pay wins the impression (see Appendix A.1 for the proof) and the equilibrium price of each targeted impression equals the second-highest willingness-to-pay of advertisers. The equilibrium outcome is therefore identical to that of a second-price auction.

[10]Details are available upon request.

Secondly, all non-targeted impressions are sold at zero price, regardless of the disclosure regime and the outsourcing decision, because no advertiser is willing to place a positive bid on them. Therefore, in the following we are going to focus on targeted impressions.

The subgames we consider can admit multiple equilibria. To deal with this multiplicity, we restrict attention to equilibria such that no advertiser can deviate by acquiring a larger volume of impressions while making the same profit. As we shall argue below, this refinement comes at virtually no loss of generality.

4.1 Equilibria when the publishers outsource to the intermediary

In the following, we will first analyse each disclosure regime, then derive the profit-maximizing choice of the intermediary. Let $w^r_{i,j}$ be advertiser j's willingness-to-pay for an impression sold by the intermediary and delivered on publisher $i = 1, 2$, under the disclosure regime r, where $r \in \{F, PT, PH\}$. We characterize $w^r_{i,j}$ in each regime and the resulting market equilibria.

4.1.1 Full Disclosure ($r = F$)

We start from the full disclosure regime. The advertisers in a given market bid for each targeted impression knowing θ and whether the consumer is a single- or a multi-homer. Only the advertisers in $k(\theta)$ are thus willing to place positive bids and advertisers bid differently depending on whether the consumer is a single- or a multi-homer. If the consumer is a single-homer, their willingness-to-pay for a targeted impression is $w^F_{i,j} = v$, for any i, since the impression cannot be repeated. If the consumer is a multi-homer, however, an advertiser is willing to pay v for one of the two impressions available if and only if it is not already acquiring the other impression. Consequently, the equilibrium price of each targeted impression under full disclosure depends crucially on the thickness of the advertising market. If the market is thick, the equilibrium price of each impression is v, since there are at least two advertisers willing to pay such price. By contrast, if the market is thin, the price of the impressions on multi-homers drops considerably, since for each impression only one of the advertisers is willing to pay v (see Table 1). Letting p denote the equilibrium price of an impression, we have

$$p^F_{n=2,MH} = 0, \; p^F_{SH} = p^F_{n\geq 4,MH} = v. \tag{1}$$

Recalling that only targeted impressions sell at a positive price, we find that the revenues earned on single-homers consumers, both in thin and thick markets, are equal to $v\tilde{q}(1-m)$, while revenues from multi-homers consumers are 0 in thin markets and $2mv\tilde{q}$ in thick markets. Hence, letting R denote the expected revenue earned by the platform in a given market (conditional on v and n), we have

$$R^F_{SH} = v\tilde{q}(1-m), \quad R^F_{MH} = (1-x)\, 2mv\tilde{q}. \tag{2}$$

Table 1: Equilibrium willingness-to-pay and price for impression on a multi-homer in a thin market under F.

	Publisher 1	Publisher 2
$w_{i,a}^F$	v	0
$w_{i,b}^F$	0	v
$p_{n=2,MH}^F$	0	0

With full disclosure, the intermediary enables the advertisers to control the frequency of ads on each consumer. This regime thus maximizes the advertisers' willingness-to-pay for impressions that do not fall on consumers already exposed. However, it also reduces the number of advertisers that compete for impressions on multi-homers. If the market is thin, the result is a sharp reduction in the equilibrium price of impressions.

4.1.2 Partial Disclosure on Types ($r = PT$)

In this regime, the advertisers bid for each targeted impression knowing whether another one is available on the same consumer. Therefore, as with full disclosure, the advertisers can control the frequency of exposure to their ads on different outlets. However, targeting is less granular: rather than knowing the consumer's θ with certainty, advertisers only know that this parameter belongs to a finite set Θ, comprising $t \geq 2$ values. To avoid inessential complexities, in the baseline model we assume that the advertisers in the conflated markets have the same return from informing consumers (we relax this assumption in Appendix XXX, showing that there is no significant change in the results).[11]

Therefore, because the impression falls on a consumer that belongs to the "right" market with probability $1/t$, each advertiser is willing to bid $w_{i,j}^{PT} = v/t$, conditional on not acquiring another impression on the same consumer, and $w_{i,j}^{PT} = 0$ otherwise.

By conflating the consumer's type, θ, with other $t-1$ "wrong" types, the intermediary ensures that there are at least two advertisers with positive willingness-to-pay for *any* targeted impression, on single- and multi-homers. Hence, even if the consumer belongs to a thin market, the equilibrium price of a targeted impression never drops to zero, as it could under full disclosure because the intermediary makes the market thicker. The drawback is that the advertisers are not willing to bid as much as they would for a non-repeated impression targeted more granularly. Hence, the highest willingness-to-pay drops, reducing the revenues that can be extracted in thick markets. At equilibrium, the intermediary sets $t = 2$ because conflating more than two markets would reduce profits. The equilibrium price,

[11] A fitting example is the conflation of geographically differentiated markets. Suppose the consumer is interested in a specific type of restaurant (e.g., Thai) within a certain travel time (e.g., ten minutes) from her home. Suppose also that, instead of disclosing the consumer's precise location, the intermediary only discloses her postcode. Conceivably, the Thai restaurants within this enlarged geographical area sell similar products and have therefore similar returns from informing the consumer, provided she is not too far.

regardless of which publishers the consumer visits and of the thickness of the market, is

$$p^{PT} = \frac{v}{2}. \tag{3}$$

Intuitively, the intermediary dilutes the quality of the information it delivers to the advertisers just enough to "thicken" the set of advertisers that bid for impressions on multi-homers in thin markets, but not more. Hence, equilibrium revenues from targeted impressions are equal to $\frac{v\tilde{q}}{2}(1-m)$ and $mv\tilde{q}$ for single-homers and multi-homers respectively, both in thin and thick markets. Hence, revenues are independent of the thickness of the market and equal to as follows

$$R_{SH}^{PT} = (1-m)\,\tfrac{v\tilde{q}}{2}, \quad R_{MH}^{PT} = mv\tilde{q}. \tag{4}$$

4.1.3 Partial Disclosure on History ($r = PH$)

In this regime, the advertisers in a given market bid for each targeted impression on a given publisher knowing the consumer's θ, but not whether the consumer receives another impression on a different publisher. The targeting of impressions is granular, but the advertisers cannot control the frequency of exposure to ads across outlets. The willingness-to-pay does not depend on whether the consumer is a single- and multi-homer, as this information is not available to the advertisers under PH. Therefore, $w_{i,j}^{PH}$ depends on the probability the consumer is already informed while visiting the other publisher, $i' \neq i$. This probability is zero if the consumer visits i only, in which case the impression is worth v. However, a multi-homer may be exposed to the same ad on i'. Indeed, when the intermediary profiles a consumer, it serves a targeted impression on such consumer on both publishers. Therefore, the probability of repetition depends on two factors. First, the likelihood that an impression falls on a multi-homer. Given each publisher sells $\frac{1+m}{2}$ impressions in total, this probability is $\frac{m}{\frac{1+m}{2}}$ (as opposed to $\frac{\frac{1-m}{2}}{\frac{1+m}{2}}$ for a single-homer). Second, the probability that the targeted impression the consumer receives on the other publisher is from the same advertiser, j. Letting $S_{i'j}$ be the share of targeted impressions advertiser j acquires on i', this probability is $S_{i'j}$. We have

$$w_{i,j}^{PH} = v\left(\frac{\frac{1-m}{2} + m\left(1 - S_{i'j}\right)}{\frac{1+m}{2}}\right) = v\left(1 - \frac{2mS_{i,j}}{1+m}\right), \quad i, i' = 1, 2; i' \neq i. \tag{5}$$

Note that $w_{i,j}^{PH}$ is independent on the volume of impressions acquired on publisher i, because there is no repetition within a given outlet. However, $w_{i,j}^{PH}$ decreases in $S_{i'j}$: since repeated impressions are wasteful, the impressions on the two publishers are substitutes (Ambrus et al., 2016; Athey et al., 2018).

This substitutability is key to characterize the equilibrium bidding strategies for targeted impressions under this regime. Consider two advertisers, a and b, in the same market. Advertiser

a outbids b for every targeted impression on publisher i if and only if b acquires a larger share of targeted impressions on publisher i' than a. Consequently, there can only be two sets of bidding strategies in equilibrium. The first is such that all advertisers in a market single-home, i.e. each places winning bids on all impressions on one publisher only. The second possible set of equilibrium strategies is such that all advertisers in a market place equal bids (given by (11)) for each targeted impression, independently of the publisher where it takes place. Therefore, advertisers multi-home, acquiring identical shares of such impressions from each publisher (i.e. $S_{ij} = 1/n, \forall i, j$). However, in this latter equilibrium candidate, each advertiser would earn zero net payoff, because it pays exactly w_{ij} for each impression on both publishers. Any advertiser can thus profitably deviate by raising its bid for all impressions on one publisher while bidding zero on the other (and thus single-homing). We summarize these findings in the following (see Appendix A.2 for a formal argument):

Lemma 1. *If the intermediary adopts partial disclosure on history, all advertisers in each market single-home on different publishers.*

This finding is in line with Athey et al. (2018), who show that, due to diminishing returns from repeated impressions on consumers that switch across publishers, advertisers in a given market single-home unless their return from informing consumers is sufficiently larger than that of their competitors.

We are now in a position to characterize the equilibrium bids, the allocation of impressions and the revenue earned by the intermediary with PH. In the following, we distinguish advertising markets only according to their thickness because, as will become clear, the allocation of impressions does not depend on the realization of v. Furthermore, to ease exposition, we present only one of the symmetric equilibria in each market, since the equilibrium prices and profits for all parties are identical.

Thin markets ($n = 2$). Given Lemma 1 and expression (5), advertiser a (resp. b)'s willingness-to-pay in equilibrium is v for each targeted impression on publisher 1 (resp. 2). Furthermore, advertiser a (resp. b)'s willingness-to-pay is $v\left(1 - \frac{2m}{1+m}\right) = v\frac{1-m}{1+m}$ for each targeted impression on publisher 2 (resp. 1). To understand this expression, consider that, given $S_{1a} = 1$, advertiser a informs all the profiled consumers in this market that visit publisher 1. Hence, any targeted impression on publisher 2 is worthless to the advertiser if it hits a multi-homer. The equilibrium price of impressions is therefore

$$p^{PH}_{n=2} = v\frac{1-m}{1+m}. \tag{6}$$

The intermediary therefore earns the following expected revenue in a given market

$$R^{PH}_{n=2} = v\tilde{q}\,(1-m)\,. \tag{7}$$

Table 2: Advertiser's willingness-to-pay, equilibrium price and intermediary's revenues under PH.

Thin ($n=2$)	Publisher 1	Publisher 2	**Thick** ($n=4$)	Publisher 1	Publisher 2
w_{ia}	v	$v\left(1-\frac{2m}{1+m}\right)$	w_{ia}	v	$v\left(1-\frac{m}{1+m}\right)$
w_{ib}	$v\left(1-\frac{2m}{1+m}\right)$	v	w_{ib}	v	$v\left(1-\frac{m}{1+m}\right)$
$p_{n=2}$	$v\left(1-\frac{2m}{1+m}\right)$	$v\left(1-\frac{2m}{1+m}\right)$	w_{ic}	$v\left(1-\frac{m}{1+m}\right)$	v
			w_{id}	$v\left(1-\frac{m}{1+m}\right)$	v
			$p_{n\geq 4}$	v	v
$R^{PH}_{n=2}$	$v\tilde{q}\,(1-m)$		$R^{PH}_{n\geq 4}$	$v\tilde{q}\,(1+m)$	

When the advertising market is thin, the intermediary can only capture the surplus generated by impressions on single-homers, just like under full disclosure. Each advertiser reaches *all* the profiled multi-homers by placing ads on a single publisher. When determining how much to bid for the impressions on the other publisher, the advertiser heavily discounts those expected to fall on multi-homers. Hence, the equilibrium price is equal to the incremental value of ads on the other publisher. We summarize the equilibrium bids and revenue in Table 2.

Thick markets ($n \geq 4$). Suppose now that $n = 4$. Given Lemma 3, we focus on the equilibrium such that two advertisers (say, a and b) single-home on publisher 1 while the other two (say, c and d) single-home on 2.[12] The willingness-to-pay by a and b (resp. c and d) equals v for each targeted impression on publisher 1 (resp. 2), since they do not acquire any impression on the other publisher. Given (5) and $S_{ij} = 1/2$, $\forall i, j$, the willingness-to-pay by these advertisers for each impression on publisher 2 equals $v\left(1-\frac{m}{1+m}\right) = \frac{v}{1+m}$. The equilibrium price of impressions is therefore

$$p^{PH}_{n=4} = v, \tag{8}$$

and each advertiser receives half the targeted impressions supplied by the respective publisher. The intermediary therefore earns the following expected revenue in a given market

$$R^{PH}_{n\geq 4} = v\tilde{q}\,(1+m)\,. \tag{9}$$

We obtain the same equilibrium prices and revenues in markets with $n > 4$. The reason is that there are at least two single-homing advertisers willing to pay v for each targeted impression per publisher.

In a thick market, therefore, the intermediary can extract the full value of the targeted impressions

[12] Equilibria with more than two advertisers on the same publisher would not satisfy our requirement that no advertiser can acquire a larger volume of impressions by deviating and make at least as much profit. Suppose three advertisers single-home on publisher 1 and one single-homes on 2. All advertisers bid v for each relevant impression on the respective publisher. Hence, if one of the advertisers on publisher 1 deviates and single-homes on 2, it gets the same profit (zero) but half the available impressions, rather than one third.

under PH. Again, this is the same outcome as under full disclosure. We summarize the bids and profits of the publishers in Table 2.

4.1.4 Revenue comparision and choice among disclosure regimes

We are now going to consider the choice of disclosure regime, r, conditional on the information available to the intermediary when selling each impression. Recall that the intermediary observes the type θ of profiled consumers, and the consumer's browsing history (single- or multi-homing). However, it does not observe the values of n and v characterizing the consumer's market, and thus cannot condition the choice of the disclosure regime on this information.

First, we compare F and PH. From (2) and Table 2, we observe that the F and PH regimes yield the same expected revenue to the intermediary in every market, that is, $\overline{v}q\left(1+m\left(1-2x\right)\right)$. Indeed, under both regimes, the full value of each impression can be extracted from the advertisers if the market is thick, while only the impressions on single-homers generate value to the intermediary if the market is thin. Hence, the intermediary has nothing to gain from conflating impressions on history. Therefore, in the following we are going to assume that, when indifferent among full disclosure and a partial disclosure regime, the intermediary chooses the most transparent regime F.[13]

Compare now F and PT. Considering (2) and (4), the intermediary earns more revenue with F when selling impressions on single-homers. However, the intermediary earns more revenue from multi-homers with PT if and only if the consumer belongs to a thin market. Conflating different markets counteracts the market-thinning effect that takes place under full disclosure when $n=2$. However, it reduces the price of impressions when market are thick enough. The expected revenue from adopting F for impressions on multi-homers is $2\bar{v}m\tilde{q}\left(1-x\right)$, whereas the revenue from adopting PT is $\bar{v}m\tilde{q}$. Intuitively, therefore, F maximizes the expected revenue from multi-homers if and only if the share of thin market is relatively small. We conclude the following:

Proposition 1. *When both publishers outsource, the intermediary sells all targeted impressions on single-homing consumers under full disclosure. Instead, it sells the impressions on multi-homing consumers under partial disclosure on type if and only if $x>\frac{1}{2}$, and full disclosure otherwise.*

This result establishes an important link between consumer multi-homing, the thickness of the advertising market and the intermediary's choice to conceal consumer information from the advertisers. By tracking consumers, the intermediary allows advertisers to avoid wasteful impressions

[13] We find there that the indifference among F and PH breaks down in favor of F when we consider the complete version of the model including intermediate markets (i.e., markets with $n=3$ advertisers). In Appendix B we show that under F,the intermediary is able to extract all revenues from multi-homers in intermediate markets, because at least two advertisers are willing to pay v for each impression. Instead, under PH, the value of impressions is fully extracted only on one publisher. Indeed, two advertisers single-home on publisher i and one single-homes on i' if the market is intermediate. Hence, the second highest bid on publisher i' is not equal to v, but only to $v\left(1-\frac{m}{1+m}\right)$, meaning that impressions on i' are only sold at a price equal to their incremental value.

on those who switch across publishers. However, the ensuing market-thinning effect prevents the intermediary from monetizing the value of these impressions in thin markets, and creates an incentive to reduce the quality of the information disclosed to the advertisers.[14] Interestingly, the intermediary reduces the granularity of targeting, but not the quality of information about the frequency of exposure to ads.

We can now compute the intermediary's total revenue, aggregating all markets. Given that the distributions of v and n are independent, we get the following

Lemma 2. *When both publishers outsource, the intermediary's total revenue is*

$$R_{IN} = \begin{cases} \overline{v}\tilde{q}\left[x\left(1-m\right)+\left(1-x\right)\left(1+m\right)\right] = \overline{v}\tilde{q}\left[1+m\left(1-2x\right)\right], & \text{if } x \leq \frac{1}{2}, \\ \overline{v}\tilde{q}\left[\left(1-m\right)+\frac{2m}{2}\right] = \overline{v}\tilde{q}, & \text{if } x > \frac{1}{2}. \end{cases} \tag{10}$$

4.2 No intermediary

We now consider the equilibrium that emerges if no publisher outsources to IN. Note that if only one publisher outsourced, the equilibrium would be identical, because the intermediary would be exactly in the same position as the competing publisher. That is, the intermediary would have the same probability of profiling consumers, q, and be unable to track consumers across publishers. Hence, the ensuing subgame is identical to that where the two publishers compete directly.

4.2.1 Advertisers' willingness-to-pay for impressions

Each publisher discloses to advertisers the consumer's type, θ, but they do not know whether a consumer is a single- or a multi-homer. Hence, the information available to the advertisers when bidding on a targeted impression is essentially the same as in the case where the intermediary adopts PH: advertisers know the consumer's type, θ, and on which publisher the impression takes place, but not whether the consumer visits multiple outlets.

Consider an advertiser $j \in k\left(\theta\right)$ and let w_{ij} be its willingness-to-pay for a targeted impression delivered by publisher i. Each targeted impression on publisher i is worth v to advertiser j if the consumer is a single-homer (probability $\left(1-m\right)/\left(1+m\right)$). However, if the consumer is a multi-homer (probability $2m/\left(1+m\right)$), she receives the same impression on i' with probability $qS_{i'j}$, i.e. the probability that the same consumer is profiled by the other publisher i' and impressed by advertiser j on i' as well. Hence, we have

$$w_{ij} = v\left(\frac{\frac{1-m}{2}+m\left(1-qS_{i'j}\right)}{\frac{1+m}{2}}\right) = v\left(1-q\frac{2mS_{i'j}}{1+m}\right), \quad i,i'=1,2; i'\neq i. \tag{11}$$

[14] This finding is in line with previous literature on the theory of auctions showing that the seller has an incentive to disclose less information about the object for sale when the set of buyers shrinks (Ganuza, 2004; Ganuza and Penalva, 2010; Bourreau et al., 2017).

This expression is similar to (5), the willingness-to-pay under the PH regime. Just like in that scenario, w_{ij} decreases in the share of targeted impressions acquired by j on the other publisher, $S_{i'j}$, implying that impressions on the two publishers are substitutes. However, there is an important difference: for a given share of targeted impressions acquired on a platform, an advertiser faces a *higher* probability of repetition when the intermediary adopts PH than when the publishers do not outsource. In the latter case, a multi-homer can receive the same ad twice only if she is identified by *both* publishers independently. By contrast, when the intermediary profiles a multi-homer, both impressions are identified. As we shall see, this implies that the revenue the publishers can generate on multi-homers may exceed the revenue generated by the intermediary.

4.2.2 Market equilibrium without the intermediary

Given the substitutability of targeted impressions on the two publishers, we can use similar arguments as in Section 4.1.3 to claim that

Lemma 3. *If neither or only one of the publishers outsources to the intermediary, all advertisers in each market single-home on different publishers.*

Then, we can follow the same steps as in Section 4.1.3 to characterize the equilibrium bids, the allocation of impressions and the profits earned by the publishers when neither outsources to the intermediary. Table 3 summarizes willingness-to-pay, equilibrium prices and revenues for different level of competition in the advertising market. Specifically, in a given thin market, the equilibrium price of a targeted impression on each publisher is

$$p_{n=2} = v\left(1 - q\frac{2m}{1+m}\right). \tag{12}$$

Each publisher therefore earns

$$R_{i,n=2} = \frac{vq}{2}\left(1 + m\left(1 - 2q\right)\right), \quad i = 1, 2. \tag{13}$$

If the market is thick (that is, there are at least four advertisers), the equilibrium price of impressions on both publishers is v, and each publisher earns

$$R_{i,n=4} = \frac{vq}{2}\left(1 + m\right), \quad i = 1, 2. \tag{14}$$

We can now compute the aggregate profits earned by the publishers. Recalling that v and n are independently distributed, we get the following:

Table 3: Advertiser's willingness-to-pay, equilibrium price and publishers's revenues.

Thin $(n=2)$	Publisher 1	Publisher 2	**Thick** $(n=4)$	Publisher 1	Publisher 2
w_{ia}	v	$v\left(1-\frac{2mq}{1+m}\right)$	w_{ia}	v	$v\left(1-\frac{mq}{1+m}\right)$
w_{ib}	$v\left(1-\frac{2mq}{1+m}\right)$	v	w_{ib}	v	$v\left(1-\frac{mq}{1+m}\right)$
$p_{n=2}$	$v\left(1-\frac{2mq}{1+m}\right)$	$v\left(1-\frac{2mq}{1+m}\right)$	w_{ic}	$v\left(1-\frac{mq}{1+m}\right)$	v
			w_{id}	$v\left(1-\frac{mq}{1+m}\right)$	v
			$p_{n\geq 4}$	v	v
$R^{PH}_{n=2}$	$\frac{vq}{2}\left(1+m\left(1-2q\right)\right)$	$\frac{vq}{2}\left(1+m\left(1-2q\right)\right)$	$R^{PH}_{n\geq 4}$	$\frac{vq}{2}\left(1+m\right)$	$\frac{vq}{2}\left(1+m\right)$

Lemma 4. *If neither publisher joins the intermediary, each publisher collects the following total revenue in equilibrium*

$$\begin{aligned} R_i = x\left[\frac{\overline{v}q}{2}\left(1+m\left(1-2q\right)\right)\right]+\left(1-x\right)\left[\frac{\overline{v}q}{2}\left(1+m\right)\right] = \\ = \frac{\overline{v}q}{2}\left[1+m\left(1-2xq\right)\right],\ i=1,2. \end{aligned} \tag{15}$$

4.3 The publishers' decision to join the intermediary

We now consider the first stage of the game and investigate the following question: should the publishers outsource the sale of their ad inventories to the intermediary?

At the beginning of the game, the intermediary makes a take-it-or-leave-it offer T_i to each publisher i. In our setting, each publisher outsources if and only if IN offers a payment at least as large as the revenue the publisher could earn by selling its ad inventory directly, i.e. $T_i \geq R_i, i=1,2$, where R_i is characterized in expression (15). As explained in Section 4.2, this revenue is the same whether the other publisher outsources or not.[15] The intermediary chooses T_i to maximize its profit $\pi_{IN} = R_{IN} - \sum_{i=1,2} T_i$,subject to $T_i \geq R_i, i=1,2$. The solution to this problem is such that $T_i = R_i$ for $i=1,2$. Therefore, the intermediary can profitably induce each publisher to outsource if and only if

$$R_{IN} \geq \sum_{i=1,2} R_i. \tag{16}$$

Using (15) and (10), we find that this condition holds if and only if IN's ability to profile consumers is above a threshold $\tilde{q}_T$, which is characterized as follows

$$\tilde{q} \geq \tilde{q}_T \equiv \begin{cases} q\left(1+\frac{2mx(1-q)}{1+m(1-2x)}\right), & \text{if } x \leq \frac{1}{2}, \\ q\left(1+m\left(1-2qx\right)\right), & \text{if } x > \frac{1}{2}. \end{cases} \tag{17}$$

[15]This observation also implies that the situation where only one publisher outsources cannot arise on the equilibrium path, since the intermediary always prefers either getting both publishers on board or none.

Figure 2: Threshold $\tilde{q}_T$, variation with respect to q.

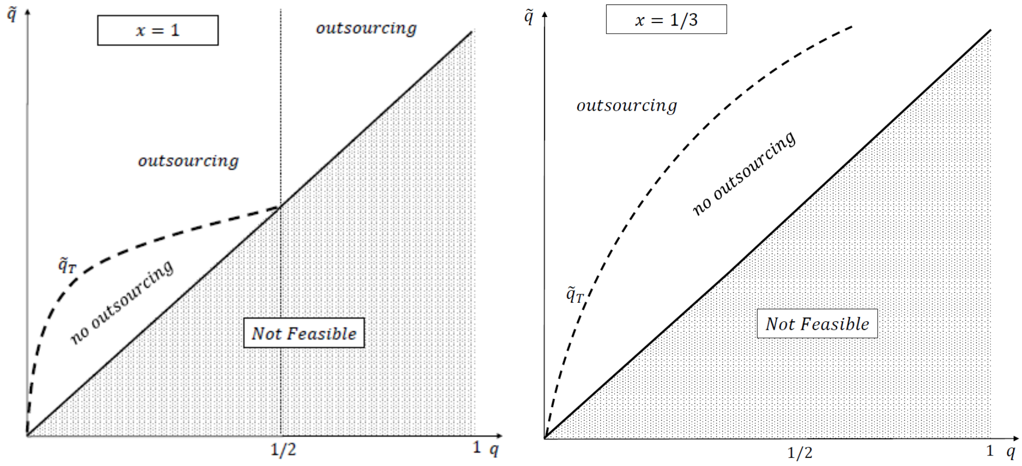

Note: figures obtained setting $m = 1/2$ and $y = 0$.

The following proposition summarizes these findings and describes how $\tilde{q}_T$ varies with the parameters of the model (see Figures 2 and 3 for an illustration).

Proposition 2. *Assume the share of consumers belonging to thin markets is small enough, i.e. $x \leq \frac{1}{2}$. Then, there exists a threshold $\tilde{q}_T \equiv q\left(1 + \frac{2mx(1-q)}{1+m(1-2x)}\right)$ such that the publishers outsource if and only if $\tilde{q} \geq \tilde{q}_T$ holds. This threshold increases with x and the share of multi-homers, m. Assume, instead, the share of consumers belonging to thin markets is high enough, i.e. $x > \frac{1}{2}$. Then, the publishers always outsource if their ability to identify consumers is large enough, i.e. $q > \frac{1}{2x}$. Otherwise, there exists a threshold $\tilde{q}_T \equiv q\left(1 + m\left(1 - 2qx\right)\right)$, such that the publishers outsource if and only if $\tilde{q} \geq \tilde{q}_T$ holds. This threshold increases with m and decreases with x.*

Consider the following polar cases to get the intuition. First, suppose there are only thick markets ($x = 0$). In this case, the intermediary chooses F and $\tilde{q}_T = q$, meaning that the publishers always outsource. Both the publishers and the intermediary can extract the full value of the impressions on profiled consumers. However, given $\tilde{q} > q$, the intermediary can profile consumers with a higher probability. Thus, outsourcing is necessarily advantageous to the supply side of the market (publishers and intermediary) if advertising markets are thick. As x increases (up to $x \leq \frac{1}{2}$, where IN keeps choosing F), some thin markets are included in the analysis. However, in each thin market the intermediary is able to extract zero revenues from multi-homers under F, while the independent publishers get some positive revenues: the higher m is, the higher the average price of each target impression for independent publishers. Hence, the higher m and x are, the lower is the incentives to outsource to an intermediary.

Focus now on the opposite polar case where there are only thin markets, i.e. $x = 1$. In this case, the intermediary chooses PT and $\tilde{q}_T$ boils down to $q\left(1 + m\left(1 - 2q\right)\right)$. Rearranging the expression, we find that the publishers choose to outsource if and only if $\tilde{q} - q > mq\left(1 - 2q\right)$. The term on

Figure 3: Threshold $\tilde{q}_T$, variation with respect to m.

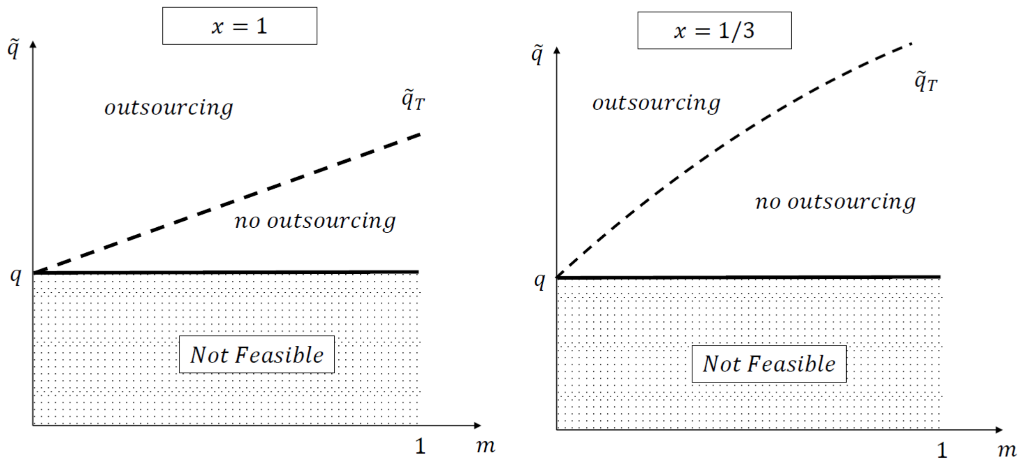

Note: figures obtained setting $q = 1/3$ and $y = 0$.

the left-hand side captures the fact that the intermediary increases the probability of profiling a consumer, i.e. $\tilde{q} > q$. Then, the term on the right-hand side is positive if and only if $q < 1/2$. Although the intermediary can sell a larger volume of targeted impressions than the publishers (given $\tilde{q} > q$), it may have to sell each impression at a lower price, as it can be seen by comparing (12) to (3). When markets are thin, the intermediary conflates impressions from different markets, and therefore can only extract half of their value. Instead, when q is low enough, the probability of repetition across publishers is so small that the expected value of a targeted impression to the advertisers is higher than it would be with the intermediary. Indeed, the price the publishers are able to extract for each targeted impression decreases with q. The higher m is, the bigger the advantage of independent publishers. The intermediary can thus collect more revenue than the publishers on aggregate only by profiling a sufficiently larger number of consumers, i.e. if and only if $\tilde{q} \geq \tilde{q}_T$ holds. On the other hand, if $q \geq 1/2$, not only the volume, but also the equilibrium price of targeted impressions sold by the publishers fall short of the level the intermediary can achieve. Now, let x decrease below one, with $x > \frac{1}{2}$, where IN keeps choosing PT. In this case, we introduce some thick markets in the analysis. Outsourcing becomes less profitable because each impression is sold at $\frac{v}{2}$ under PT and at v by independent publishers (see (3) and (12)). Hence, as x decreases, the probability to outsource decreases as well.

From the above discussion, it is easy to undestand how x and m affect the outsourcing decision of the publishers. Let us summarize the discussion concerning how the share of multi-homers, m, affects the threshold $\tilde{q}_T$. In thick markets, both the publishers and the intermediary can extract the full value of impressions on either single- and multi-homers. However, in thin markets, the publishers sell each targeted impression at a higher price than the intermediary when it adopts full disclosure. Hence, as the share of multi-homers increases, outsourcing becomes less profitable. When the share of thin markets increases, the intermediary adopts partial disclosure on type. Again, unless$q > \frac{1}{2x}$, the publishers sell each impression at a higher average price, implying that more multi-homing makes outsourcing less profitable. Hence, we conclude that $\tilde{q}_T$ increases with m. The implication of this result is that more multi-homing by consumers should not be associated by publishers with the necessity to rely on an intermediary to sell ads, even if the latter can follow consumers across publishers and repeated impressions are wasteful.

5 Welfare analysis and the effects of regulating the intermediary

In this section, we evaluate how the intermediary affects the distribution of surplus among the supply side (publishers and intermediary) and the demand side (advertisers) of the market, and the total surplus in the advertising market. Furthermore, we evaluate the implications of two regulatory policies: one that mandates full transparency to the intermediary and one that restricts the amount of information that the intermediary can collect from consumers, to protect their privacy.

5.1 How does the intermediary affect advertising surplus and its distribution?

We consider how outsourcing to an intermediary affects the total surplus generated on the advertising market (i.e., the total profits of publishers, advertisers and the intermediary) and the distribution of this surplus.

By costruction, the publishers outsource to the intermediary if and only if the total revenue on the supply side of the advertising market increases. Hence, the supply side of the market is better off when outsoucing occurs at equilibrium.

Now, consider the demand side of the market. Advertiser surplus can increase or decrease under outsourcing, depending on the disclosure regime.

Given (12), each advertiser in thin markets gets the following net surplus per each targeted impression when the publishers do not outsource $v - v\left(1 - \frac{2mq}{1+m}\right) = v\frac{2mq}{1+m}$. Instead, in thick markets all surplus is extracted by the intermediary, and advertisers are let with zero surplus. Hence, advertiser surplus under no outsourcing is

$$AS = \int_0^{v_H} \left(2xmvq^2\right) dv \quad . \tag{18}$$

When there is outsourcing, all surplus is extracted from advertisers for impressions sold on single-homers. Instead, some surplus from impressions on multi-homers might be left to advertisers. Specifically, in thin markets under full disclosure the advertisers retain all the surplus from each targeted impression on multi-homers, since the price drops to zero (see (1)). By contrast, under PT the advertisers get no surplus, since all the impressions sell at a price equal to the advertisers' willingness-to-pay (see (3)). The advertisers also get zero net surplus in thicker markets. Hence, advertiser surplus under outsourcing is

$$\begin{cases} AS^F = \int_0^{v_H} (2xmv\widetilde{q})\, dv & \text{if } x \leq \frac{1}{2}, \\ (1-m)\, AS^F + 2mAS^{PD} = 0 & \text{if } x > \frac{1}{2}. \end{cases} \tag{19}$$

Advertiser surplus increases with outsourcing if and only if $x \leq \frac{1}{2}$, otherwise it decreases because all surplus is extracted by the intermediary from single-homers and under PT, while advertisers retain some surplus in thin markets when publishers sell impressions independently. Hence, despite selling a larger volume of targeted impressions ($q \leq \tilde{q}$), the intermediary does not necessarily make the advertisers better off, in particular if advertising markets are thin.

Finally, total surplus in the advertising market may either increase or decrease when the publishers outsource to the intermediary. If the publishers operate independently, each targeted impression produces the maximum total surplus, since the impressions are granularly targeted and not repeated across publishers (recall that advertisers would single-home in that scenario). The same occurs if the intermediary adopts full disclosure. However, the volume of targeted impressions is higher with outsourcing, implying that total surplus is higher when outsourcing and full disclosure occurs. However, if the intermediary adopts partial disclosure by type, some impressions are wasted being sold to the "wrong" advertisers. Hence, total surplus may decrease with outsourcing the PT occurs. We thus obtain the following result (see Appendix A.3 for the proof):

Proposition 3. *Total surplus from advertising increases when the publishers outsource to the intermediary if and only if either (i) $x \leq \frac{1}{2}$, or (ii) $x > \frac{1}{2}$ and $\tilde{q} > \tilde{q}_W \equiv q\,(1+m)$.*

5.2 The effects of transparency regulation

Regulators have considered mandating greater transparency by intermediaries towards advertisers. See, for example, the proposed draft of the Digital Market Act (European Commission, 2020) and the remedies discussed by the UK Competition and Markets Authority (CMA, 2020, pag. 395).[16] Formally, we study the possible implications of these proposals, assuming that regulation imposes full disclosure to the intermediary. Proposition 1 shows that the intermediary adopts full disclosure if $x \leq 1/2$, but may blur the information shared with advertisers when auctioning impressions on multi-homers if the share of thin markets is high enough ($x > 1/2$). Hence, we focus on the effect of regulation when $x > 1/2$.

Our model suggests that mandating greater transparency reduces the profits earned by the

[16] In the proposal for the Digital Market Act, the European Commission states that gatekeepers should "provide advertisers and publishers, upon their request and free of charge, with access to the performance measuring tools of the gatekeeper and the information necessary for advertisers and publishers to carry out their own independent verification of the ad inventory".

intermediary and the publishers, to the benefit of advertisers. Obviously, the regulation tightens the region of parameters where the publishers outsource, because the profits earned by the supply side of the market decreases. Indeed, under outsourcing and full disclosure the price of impressions on multi-homers drops to zero in thin markets. If the regulation deters outsourcing, the advertisers are better off because they earn a positive surplus when the publishers sell impressions independently (see (18)), whereas they would earn zero under partial disclosure with the intermediary (see second line in (19)). Also if publishers outsource despite the regulation, advertisers are better off because they retain more surplus under full disclosure than under partial disclosure by type.

Consider now the effects of transparency regulation on total surplus from advertising. If the publishers outsource even when the regulation is imposed, full disclosure avoids the waste of inaccurately targeted impressions on multi-homers. Hence, total surplus increases. However, if the regulation induces the publishers not to use the intermediary, the effect on welfare can be either positive or negativebecase there are fewer targeted ads without the intermediary, but more of those ads are targeted in a precise way. Hence, not imposing transparency may increase total surplus of the advertising market, because it expands the use of intermediaries that increase the efficiency of the market by being able to identify more consumers. We summarize the results in the following proposition (see Appendix A.4 for the proof).

Proposition 4. *Imposing full disclosure is detrimental to the supply side of the advertising market (publishers and intermediary), but beneficial to the advertisers in thin markets. Total surplus increases if the publishers use the intermediary despite the regulation. Otherwise, total surplus decreases with the regulation if and only if* $\tilde{q}$ *exceeds the threshold* $\tilde{q}_W$ *characterized in Proposition 3.*

5.3 The effects of privacy regulation

Consumer can block third-party tracking, hence affecting the ability of an intermediary to collect information about them. [17] We now explore the implications of a privacy regulation increasing consumer control over personal data. One of the main provisions of many recent privacy regulations (as the European Union's GDPR, and similar laws adopted in Chile, Japan, Brazil, South Korea, and some US States, as California) is to increase consumer control over personal data and to facilitate avoiding tracking by third-parties.[18]

We introduce privacy considerations in the model, allowing consumers to block third-party tracking, e.g., by installing browser extensions that block cookies or anonymize online activity.

[17] Advertising firm Flashtalking estimates that as many as 64 percent of third-party cookies get blocked or rejected daily (https://www.mediapost.com/publications/article/316757/64-of-tracking-cookies-are-blocked-deleted-by-we.html). In principle, consumers could also block first-party tracking done by publishers. However, this tracking is often intertwined with the publisher site's basic functionalities. Moreover, most antitracking tools only block third-party tracking.

[18] For instance, the GDPR (European Parliament, 2016) mandates opt-in policies for collecting consumers' consent regarding third-party cookies, as opposed to opt-out ones.

Depending on the type of disclosure regime, consumers mat be reached by a different type of advertising that may affect their perceived benefit from blocking tracking. The objective of the analysis is to understand how privacy policy may affect the degree of transparency chosen by the intermediary.

We augment the baseline model as follows. At stage 1, we let consumers decide whether to block cross-outlet tracking. We assume that, if a consumer blocks, the intermediary cannot observe whether she visits multiple publishers, and can profile her only with probability q. If a consumer does not block, the intermediary collects the same information as in the baseline model. To streamline the exposition, we focus (without loss) only on the two disclosure regimes that emerge in equilibrium in our baseline model, i.e. F and PT. Furthermore, to concentrate on the choice of disclosure regime (rather than on whether the publishers outsource), we assume that outsourcing occurs, that is, that the revenue earned by the intermediary under each regime is at least equal to the revenue earned by the publishers independently.[19]

Let c be the idiosyncratic cost (or effort) of blocking tracking, distributed uniformly among consumers on the $[0, \overline{c}]$ interval. Let b_r,with $r = F, PT$, be the private benefit from opting out as a function of the disclosure regime chosen by the intermediary.[20] This benefit captures consumers' enhanced privacy, but also the perceived relevance/intrusiveness of ads. For instance, under full disclosure consumers are more likely to see more granularly targeted ads than under PT. Consumers may appreciate such ads because they may perceive them as more relevant, or they may dislike target ads perceiving them as exceedingly intrusive (Goldfarb and Tucker, 2012; Turow et al., 2009). Accordingly, b_r depends on the disclosure regime chosen by the intermediary, and b_F may be bigger or lower than b_{PT} depending on how consumer perceive ads relevance/intrusiveness.

We model privacy policy as a parameter, t, that reduces the cost of blocking for every consumer. A consumer thus blocks tracking if and only if her cost of blocking tracking net of the effect of policy,$c - t$, is smaller than b_r. Hence, the fraction of consumers blocking is $\frac{b_r+t}{\overline{c}}$ (assumed to be smaller than one for consistency). Consider now IN's choice between F and PT. The latter results in the higher amount of revenue for the intermediary if and only if the following holds

$$\begin{gathered} \left(1 - \tfrac{b_P+t}{\overline{c}}\right) \tilde{q} + \left(\tfrac{b_P+t}{\overline{c}}\right) q \left(1 + m \left(1 - 2xq\right)\right) \geq \\ \left(1 - \tfrac{b_F+t}{\overline{c}}\right) \tilde{q} \left(1 + m \left(1 - 2x\right)\right) + \left(\tfrac{b_F+t}{\overline{c}}\right) q \left(1 + m \left(1 - 2xq\right)\right). \end{gathered} \tag{20}$$

Rearranging the above inequality, we obtain that IN chooses PT if and only if

$$x > x_T \equiv \frac{1}{2} - \frac{b_F - b_P}{2m\tilde{q}\left(\overline{c} - (b_F + t)\right)} \left(\tilde{q} - q \left(1 + m \left(1 - 2xq\right)\right)\right). \tag{21}$$

[19] Formally, we assume that $\tilde{q} > q\left(1 + m\left(1 - 2xq\right)\right)$ and $\tilde{q}\left(1 + m(1 - 2x)\right) > q\left(1 + m\left(1 - 2xq\right)\right)$.

[20] We assume b_r to be the same for all consumers for simplicity. Relaxing this assumption would not qualitatively affect our results.

Given the last term in brackets is positive by assumption (see footnote 19), we find

$$x_T < \frac{1}{2} \Longleftrightarrow b_F > b_P.$$

Hence, the intermediary is less likely to choose F than in baseline model if $b_F > b_P$, and viceversa. The upshot is that the intermediary now adopts a two-sided logic when choosing its disclosure regime, considering not just the revenue from targeted ads, but also how consumers' reaction affects the total volume of such ads. Quite intuitively, whenever more consumers block tracking under F (i.e., $b_F > b_P$ holds), the intermediary is less likely to choose such regime than in our baseline model. By contrast, whenever $b_F < b_P$, the intermediary is more likely to choose F.

Consider now the effect of the privacy policy, than decreases the cost of blocking cookies by t. Expression (21) shows that

$$\frac{\partial x_T}{\partial t} < 0 \Longleftrightarrow b_F > b_P.$$

Hence, we find that a tighter privacy policy results in a lower probability that IN chooses F whenever $b_F > b_P$, and viceversa. We conclude that a privacy policy that facilitates blocking moves information disclosure decisions by the intermediary toward consumers' preferences. To grasp the intuition, focus on (20). This inequality shows that whenever t increases, the intermediary suffers a net loss of income, since more consumers block tracking. If $b_F > b_P$ (that implies $x = x_T < \frac{1}{2}$), the loss is greater under F than under P. Therefore, the set of values of x such that IN adopts F shrinks. Hence, if consumers dislike targeted ads, a tighter privacy policy reduces the likelihood that the intermediary adopts full disclosure toward advertisers, and decreases the level of targeting also on those who do not block.

Proposition 5. *Privacy policy that facilitates blocking of third-party tracking moves the information disclosure decision toward consumer preferences.*

To sum up, this result shows that privacy policy can either reinforce or diminish the intermediary's incentives to disclose consumer information to the advertisers. Furthermore, consumer's reaction to targeted ads is an important driver of the intermediary's response to privacy policy.

6 Extensions and robustness checks

6.1 Reserve prices in advertising auctions

In a first extension, we allow the intermediary to introduce reserve prices in advertising auctions.[21] We have shown in Section 4.1 that the intermediary adopts partial disclosure on type for impressions

[21] Reserve prices are sometimes adopted in advertising auctions ran by intermediaries, but some evidence suggests that the platforms do not necessarily set them at revenue-maximizing levels (see, e.g., Ostrovsky and Schwarz, 2016). Google lets each publisher decide the reserve price (if any) for the impressions on its own webpages (https://support.google.com/admanager/answer/9298008?hl=en).

on multi-homers when there are enough thin advertising markets (Proposition 1). Reserve prices are reduntant in auctions on multi-homers under PT because the price paid is equal to the highest willingness-to-pay. Also, given full disclosure, a reserve price is redundant for impressions on single-homers and on multi-homers in thick markets (where the intermediary can extract the whole advertising surplus regardless), but can limit the price drop for impressions on multi-homers in thin markets. However, because the intermediary do not observe v, it chooses the same reserve price for all impressions sold. Hence, the drawback is that the higher the reserve price, the higher the share of markets where advertisers drop out from the auction because their return from informing consumers, v, is beneath the reserve price. Due to this trade-off, we show in Appendix C.1 that, even if the intermediary sets the revenue-maximizing reserve price (given the available information), it is not obvious that this revenue exceeds the revenue with partial disclosure for impressions on multi-homers in thin markets. For instance, full disclosure with a reserve price is weakly dominated by partial disclosure given a continuous Bernoulli distributions on v. We conclude that the possibility to adopt reserve prices does not necessarily induce the intermediary to be more transparent when auctioning ad impressions.

6.2 Advertiser competition on the product market

We assumed that the advertisers' marginal return from informing a relevant consumer is independent from the consumer being exposed to ads from competitors. However, if competing advertisers provide substitute products, advertising may intensify price competition on the product market (see de Cornière and de Nijs, 2016). Also, it may affect the probability of selling a product, if consumers only choose one of the advertised product.

To account for these possible effects, we modify the model assuming that, if a consumer is exposed to two ads from different advertisers in the same market, each advertiser gets a return of αv, with $\alpha \in [0, 1]$. The model is solved in Appendix C.2. On the one hand, if competition on the product market dissipates advertising returns ($\alpha < 1$), it reduces the willingness-to-pay to acquire an impression on a consumer who is already exposed to an ad by a rival. On the other hand, this effect also creates an incentive for advertisers who inform a consumer to prevent competitors from doing the same.

We find that if the effect of competition on advertising returns is relatively small, i.e. $\alpha \geq \frac{1}{2}$, there is no qualitative change in the analysis. However, if the effect of competition is strong, i.e. $\alpha < \frac{1}{2}$, some of our results change. Both when the publishers operate independently and when the intermediary adopts F, we find that in each market one advertiser acquires all the targeted impressions on multi-homers, because the effect of competition makes such impressions complementary across publishers, rather than substitutes: each advertiser has a strong incentive to exclude its rivals from informing consumers. This finding, however, also implies that the price of impressions on such consumers is low, because competition strongly reduces the willingness-to-pay of all the other bidders in the same market.

Quite interestingly, there is an additional incentive to adopt PT for the intermediary: if markets are conflated, each advertiser anticipates that, with positive probability, a consumer that already receives an impression is not informed by a competitior, but from an advertiser in a different market. Hence, we find that $\alpha < \frac{1}{2}$ is sufficient for the intermediary to sell all the impressions on multi-homers under PT, even if advertising markets are thick.

6.3 Revenue sharing arrangements between publishers and intermediary

In the baseline model, we assumed that the publishers receive a lump-sum transfer from the intermediary in exchange for the right to sell their ad inventories. In Appendix C.3, we modify the setting by introducing revenue sharing agreements, whereby the intermediary transfers to each publisher a given share of the revenue earned by selling the impressions.[22] Let this share be $\alpha \in (0, 1)$. When choosing the disclosure regime for each impression, the intermediary would have the same incentives as in our baseline model, since it would earn a fraction, $1 - \alpha$, of the revenue and because the revenue sharing agreement does not affect the willingness-to-pay by advertisers. Therefore, Proposition 1 does not change. Then, the result stated in Proposition 2 also applies: the intermediary is able to induce the publishers to outsource (by setting α sufficiently large) and still make a profit if and only if it earns at least as much total revenue larger as the publishers can earn independently, i.e. condition (17) holds. Another condition is needed for outsourcing to occur: the publishers will outsource if and only if the intermediary leaves them a high enough share of their profits, that is, if and only if α is not too low.

6.4 Increasing returns to advertising on the same consumer and re-targeting

Throughout the analysis, we assumed diminishing returns to advertising on the same consumer. However, some advertisers may want to propose an ad containing a specific offer to a consumer previously exposed to their product. Also, advertisers might want to send ads in sequence to tell a brand story. In both cases, advertisers may put a premium on hitting the same consumer more than once (re-targeting). To capture this possibility, in Appendix C.4 we assume each advertiser gets a higher return from the second impression that hits a consumer than from the first one. Under this assumption, we find the intermediary always benefits from adopting full disclosure, which enables it to extract the maximum advertising surplus, even in thin markets. The reason is that each advertiser now has an incentive to acquire both impressions available on the same consumer (indeed, advertisers multi-home also when the publishers sell their impressions independently). For the same reason, the

[22] Google's AdSense, for instance, pays to publishers a fixed percentage (68%) of the revenue from each ad they display. See `https://support.google.com/adsense/answer/180195`.

publishers always benefit from outsourcing to the intermediary.

6.5 Heterogeneous advertising returns within markets

In Appendix C.5, we relax the assumption that advertisers within the same market get the same return from informing consumers. We allow for a subset of markets where one of the advertisers (say, a) gets a larger return from informing a profiled consumer, v^+, than the others. More precisely, we assume there is a share x' of thin markets where one of the advertisers has a return from informing consumers large enough that it acquires all the impressions on all platforms in equilibrium, and a share x of thin markets with homogeneous advertisers. Similarly, we assume there is a share z of thick markets with a dominant advertiser.

With homogeneous returns, all advertisers in a market single-home when the publishers do not outsource their ad inventories (see Lemma 3). However, if v^+ is large enough compared to v, advertiser a outbids the others for each targeted impression on *both* publishers. The resulting effect on ad prices is positive, because even in a thin market, the advertisers who do not acquire any impression are willing to bid the full value, v. We find the same positive effect when the intermediary adopts PH. There is instead no increase in revenue with full disclosure. Therefore, in this setting the intermediary prefers PH to F, and chooses PT if and only if $x \geq 1/2$. However, there is no change in Proposition 2.

7 Concluding remarks

We have studied the incentives of a large online advertising intermediary to disclose consumer information to advertisers when auctioning ad impressions. We have shown that fully disclosing information that enables advertisers to optimize the allocation of ads on multi-homing consumers is profitable to the intermediary only if thera are enough thick advertising markets. Retaining information on the type of the consumer may be profitable, while a lack of transparency on the browsing history does not allow to raise profits. We have also considered how disclosure affects the incentives of publishers to outsource the sale of their ads to an intermediary, and related these incentives to the extent of consumer multi-homing, the competitiveness of advertising markets and the ability of platforms to profile consumers. We have shown that, even when most consumers multi-home, the publishers do not necessarily benefit from outsourcing to the intermediary.

Policy implications. First, the paper contributes to the debate on regulating transparency in the digital advertising market. Transparency requirements to advertising gatekeepers are considered, for instance, in the Digital Market Acts by the European Commission (European Commission, 2020). We have found that the intermediary prefers not to disclose all the available information on consumers when auctioning impressions in thin advertising markets. Imposing full transparency moves surplus from the supply side (publishers and intermediary) to the demand side (advertisers) of the advertising market. However, total surplus from advertising may decrease because transparency makes using an intermediary less profitable. Indeed, an intermediary may increase the efficiency of the advertising market because of their higher ability to profile consumers.

Another important ongoing debate concerns the possibility of restricting third-party tracking to protect consumer privacy. For example, the European Union's GDPR of 2018 and the California Consumer Privacy Act of 2020 gives to consumers more control over the information collected by third-parties for the purpose of targeted advertising. If consumers can block tracking, an intermediary considers their taste when choosing which information to disclose to advertisers. Any privacy policy making blocking easier moves information disclosure decision closer to consumer preferences. However, it is difficult to get clear cur indications on the welfare effects of stricter privacy policy: consumer surplus should be compared to the surplus from advertising.

Managerial implications. The analysis has interesting managerial implications for several players in the digital advertising market. Our results suggest that the pervasiveness of multi-homing by consumers should not necessarily induce digital publishers to rely on intermediaries for selling their advertising space. We have found that the publishers could end up selling the impressions on multi-homers in thin markets at a higher price when operating independently. The intermediary can use the data collected on one publisher to profile a consumer and sell targeted impressions on such consumer when she visits another publisher. This sort of "information leakage" can result in a higher risk of repetition and, hence, lower ad prices. More generally, if there are diminishing returns to advertising, e.g. if advertisers prefer to avoid excessive repetition on the same consumer (frequency capping), the thickness of the advertising market should be a key parameter driving the decision whether to use an intermediary (Proposition 2). On the other hand, if the advertisers value impressing the same consumer multiple times (re-targeting), these conclusions change: the publishers unambiguously benefit from outsourcing to the intermediary, because its ability to track consumers across outlets results not only in a higher volume of targeted impressions, but also in higher prices of such impressions.

The fact that digital publishers sell their impressions via an intermediary has mixed implications for the advertisers. On the one hand, the intermediary has a technological and informational advantage compared to the publishers in terms of profiling consumers, targeting the impressions to the right consumer and managing the frequency of exposure to the same ad. On the other hand, competition on the supply side is significantly weakened when the publishers outsource to the intermediary. In our model, the net effect of these forces on the surplus of advertisers may be negative or positive. However, the advertisers should benefit from (and should thus be supportive of) regulation that encourages greater transparency. The advertisers should also favor regulation enhancing privacy, as long as the latter discourages digital publishers from relying on the intermediary.

Our model also provides some insights concerning the optimal disclosure of information about the behavior and preferences of consumers in advertising auctions. We have shown that the thickness of advertising markets is an important determinant of this disclosure (Proposition 1). Disclosing information to advertisers is not necessarily beneficial to the intermediary and, more generally, to the supply side of the advertising market. If there are many thin markets, disclosing granular information on the type of the consumers can increase the equilibrium price of ad impressions. By contrast, disclosure is advantageous to all players when there are many thick advertising markets.

References

Ambrus, A., Calvano, E., and Reisinger, M. (2016). Either or Both Competition: a Two-Sided Theory of Advertising With Overlapping Viewerships. *American Economic Journal: Microeconomics*, 8:189?–222. 3.2, 4.1.3

Anderson, S. P. and Coate, S. (2005). Market provision of broadcasting: A welfare analysis. *Review of Economic Studies*, 72(4):947–972. 3.2

Athey, S., Calvano, E., and Gans, J. (2018). The Impact of Consumer Multi-homing on Advertising Markets and Media Competition. *Management Science*, 64(4):1574–1590. 1, 3.2, 4.1.3, 4.1.3

Bagwell, K. (2007). The economic analysis of advertising. volume 3 of *Handbook of Industrial Organization*, pages 1701–1844. Elsevier. 1

Bergemann, D. and Bonatti, A. (2015). Selling Cookies. *American Economic Journal: Microeconomics*, 7(3):259?–294. 2

Bourreau, M., Caillaud, B., and de Nijs, R. (2017). The Value of Consumer Data in Online Advertising. *Review of Network Economics*, 16(3):269–289. 2, 14

Cairncross, F. (2019). The Cairncross Review: a sustainable future for journalism. UK Department for Digital, Culture, Media and Sport. 2

Chen, J. and Stallaert, J. (2014). An Economic Analysis of Online Advertising Using Behavioral Targeting. *MIS Quarterly*, 38(2):429–450. 7

CMA (2020). Online platforms and digital advertising market study. Market study final report, Competition and Markets Authority. 3, 4, 1, 1, 3.2, 9, 5.2

D'Annunzio, A. and Russo, A. (2020). Ad Networks and Consumer Tracking. *Management Science*, 66(11):5040–5058. 2

de Cornière, A. and de Nijs, R. (2016). Online Advertising and Privacy. *RAND Journal of Economics*, 15(3):311–327. 2, 6.2

European Commission (2020). Proposal for the Digital Markets Act. Communication 842-2020. 5.2, 7

European Parliament (2016). General Data Protection Regulation. Regulation 2016/679. 18

Ganuza, J.-J. (2004). Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations. *RAND Journal of Economics*, 35(3):583–598. 14

Ganuza, J.-J. and Penalva, J. S. (2010). Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions. *Econometrica*, 78(3):1007–1030. 2, 14

Ghosh, A., Mahdian, M., McAfee, R. P., and Vassilvitskii, S. (2015). To Match or Not to Match: Economics of Cookie Matching in Online Advertising. *ACM Transactions on Economics and Computation*, 3(2):article 12. 2

Goldfarb, A. and Tucker, C. (2012). Shifts in privacy concerns. *American Economic Review*, 102(3):349–53. 5.3

Gu, Y., Madio, L., and Reggiani, C. (2019). Data brokers co-opetition. Available at SSRN: `https://ssrn.com/abstract=3308384`. 2

Hummel, P. and McAfee, R. P. (2016). When Does Improved Targeting Increase Revenue? *ACM Transactions on Economics and Computation*, 5(1). 2

IAB (2017). IAB Europe Report: European Programmatic Market Sizing 2017. available at `https://iabeurope.eu/research-thought-leadership/iab-europe-report-european-programmatic-market-sizing-2017/`. 1

Johnson, J. P. (2013). Targeted advertising and advertising avoidance. *The RAND Journal of Economics*, 44(1):128–144. 6

Karle, H. and Reisinger, M. (2019). Excessive Targeting. Working Paper. 6

Levin, J. and Milgrom, P. (2010). Online Advertising: Heterogeneity and Conflation in Market Design. *American Economic Review*, 100:603–07. 2, 3.1

Marotta, V., Wu, Y., Zhang, K., and Acquisti, A. (2021). The welfare impact of targeted advertising technologies. Available at SSRN: `https://ssrn.com/abstract=2951322`. 2

Ostrovsky, M. and Schwarz, M. (2016). Reserve prices in internet advertising auctions: A field experiment. Working Paper, Stanford University. 21

Peitz, M. and Reisinger, M. (2020). Advertising networks and consumer tracking. Mimeo. 2

Rafieian, O. and Yoganarasimhan, H. (2021). Targeting and Privacy in Mobile Advertising. *Marketing Science*, 4(2):193–218. 2

Sharma, P., Sun, Y., and Wagman, L. (2019). The Differential Effects of New Privacy Protections onPublisher and Advertiser Profitability. Working Paper, Illinois Institute of Technology. Available at SSRN: `https://ssrn.com/abstract=3503065`. 2

Turow, J., King, J., Hoofnagle, C. J., Bleakley, A., and Hennessy, M. (2009). Americans reject tailored advertising and three activities that enable it. Available at: `http://repository.upenn.edu/cgi/viewcontent.cgi?article=1138&context=asc_papers`. 5.3

Wilbur, K. C. and Zhu, Y. (2009). Click Fraud. *Marketing Science*, 28(2):293–308. 2

Xiang, Y. and Sarvary, M. (2013). Buying and Selling Information Under Competition. *Quantitative Marketing and Economics*, 11:321–351. 2

Appendix

A Proofs of Lemmas and Propositions

A.1 Proof of the claim "the equilibrium price equals the second-highest willingness-to-pay"

Consider Let w_{II} be the second-highest value of the willingness-to-pay (WTP) for an impression delivered on publisher i, w_{ij}, among all advertisers. We first establish the existence of an equilibrium such that the price of the impression is w_{II} and that this bid is placed by the advertiser(s) with the highest value of w_{ij}, that we denote by w_I. Suppose first that $w_I > w_{II}$. In a first price auction, the advertiser with WTP w_I wins that impression and earns $w_I - w_{II} > 0$. Deviating by bidding more than w_{II} does not increase the probability of winning for such advertiser, but raises the equilibrium price and is therefore not profitable. Bidding less than w_{II} is also not profitable because the advertiser would not win the impression and thus earn zero. There is also no profitable deviation by the other advertisers. If any of them bids more than w_{II}, it wins the impression but earns a negative profit, whereas if it bids less than w_{II} the deviation has no effect on the outcome of the auction. Suppose now that $w_I = w_{II}$, i.e. there are two or more advertisers with the highest willingness-to-pay. Each of these advertisers wins the impression with positive probability, but earns zero profit since the price equals their willingness-to-pay. Hence, none can gain by either raising their bid, which would result in negative profit, or reducing it, which would result in not winning. The remaining advertisers, whose WTP is smaller than w_{II}, cannot gain by raising their bid above that level either.

We now establish that the equilibrium we just described is unique. Any candidate such that the price exceeds w_{II} cannot be an equilibrium. Suppose now that $w_I > w_{II}$. If the advertiser with WTP w_I is placing the highest bid, and $w_I > w_{II}$, the advertiser can profitably deviate by bidding $\varepsilon > 0$ (arbitrarily small) less than its current bid. If another advertiser is placing the highest bid, and the bid is below w_I, then the advertiser with WTP w_I can deviate by bidding w_I minus $\varepsilon > 0$ (arbitrarily small) and win the impression. Suppose now that $w_I = w_{II}$. If the price exceeds w_{II}, the winning advertiser must be bidding more than its WTP. Hence, the advertiser can deviate profitably by bidding zero. In a similar way, it can be established that there cannot be an equilibrium such that the price of the impression is smaller than w_{II}.

A.2 Proof of Lemma 1

Consider advertiser j's bidding strategy for each relevant impression delivered on publisher i. The advertiser gets an expected return of (11) from each such impression. As established above, the equilibrium bidding strategies of the advertisers must be such that the equilibrium price of the impression equals the second-highest willingness-to-pay. However, because the willingness-to-pay for ads on one publisher are conditional on $S_{i'j}$, there are potentially multiple equilibrium strategies for each advertiser. To characterize them, we have to establish which values of $S_{i'j}$ can emerge in any equilibrium of the subgame.

Thin markets. Consider a market such that $n = 2$. Let a, b be the set of advertisers in this market. Focus, without loss of generality, on the relation between the share of relevant impressions acquired by an advertiser in this market on publisher 2 and the value of (11) for relevant impressions on publisher 1, and consider the bidding strategy of the advertisers on such publisher. There are two possible cases:

- A: if $S_{2a} > S_{2b}$, the advertisers' willingness-to-pay each relevant impression on 1 are such that $w_{1a} = v\left(1 - \frac{2mS_{2a}}{1+m}\right) < w_{1b} = v\left(1 - \frac{2mS_{2b}}{1+m}\right)$. Hence, b outbids a for all such impressions, so $S_{1a} = 0 < S_{1b} = 1$.

- B: if $S_{2a} = S_{2b}$, the advertisers' bid for each relevant impression on 1 are such that $w_{1a} = v\left(1 - \frac{2mS_{2a}}{1+m}\right) = w_{1b} = v\left(1 - \frac{2qS_{2b}}{1+m}\right)$. Hence, a and b place equal bids for all such impressions, so $S_{1a} = S_{1b} = 1/2$.

In equilibrium, these bidding strategies must be consistent with the bidding strategies (and the ensuing shares S_{2j}) on publisher 2. If case A applies, since $S_{1a} < S_{1b}$, by the same reasoning as above we must have $S_{2a} = 1 > S_{2b} = 0$. This case constitutes an equilibrium candidate such that the advertisers single-home, i.e. place winning bids on a single publisher. Given (11), in a first-price auction the equilibrium winning bid is (12). Therefore, each advertiser earns $v\frac{2m}{1+m}$ per impression acquired. Given there are $\tilde{q}\frac{1+m}{2}$ relevant impressions per publisher in this market, each advertiser earns $vm\tilde{q}$ in total.

If Case B applies, each advertiser bids $v\left(1-\frac{m}{1+m}\right)$ for all relevant impressions on each publisher. The latter is also the price of relevant impressions on both publishers. Therefore, the advertisers make zero profit in this equilibrium candidate.

We have thus identified two equilibrium candidates and must now establish whether these are indeed equilibria. The candidate associated to case B cannot be an equilibrium because, given the bids placed by the rival, an advertiser can deviate by bidding $v\left(1-\frac{m}{1+m}\right)+\varepsilon$, where $\varepsilon>0$ and arbitrarily small, for all relevant impressions on one publisher (and thus win them all) and zero for all impressions on the other. The advertiser would earn $v\frac{m}{1+m}\left(\tilde{q}\frac{1+m}{2}\right)=\frac{vm\tilde{q}^2}{2}$ by deviating, so the deviation is profitable.

As for the candidate associated to case A, there is no profitable deviation: each advertiser cannot profitably outbid the other on the publisher where it is not acquiring any impression, because the winning bid is (12) on that publisher. Therefore, to win those impressions the advertiser would have to pay more than it is already paying, and would thus not earn more. Nor would the advertiser gain by reducing its bid for the impressions it is already acquring. There exists also the symmetric equilibrium obtained from swapping a and b.

Intermediate markets. Consider a market such that $n=3$. Let a,b,c be the set of advertisers in this market. Focus, without loss of generality, on the relation between the share of relevant impressions acquired by an advertiser in this market on publisher 2 and the value of (11) for relevant impressions on publisher 1, and consider the bidding strategy of the advertisers on such publisher. There following cases are possible

- A: if $S_{2c}>S_{2b}>S_{2a}$, advertisers' willingness-to-pay each relevant impression on 1 are such that $w_{1a}=v\left(1-\frac{2mS_{2c}}{1+m}\right)<w_{1b}=v\left(1-\frac{2mS_{2b}}{1+m}\right)<w_{1c}=v\left(1-\frac{2mS_{2a}}{1+m}\right)$. Hence, a outbids the other advertisers for all such impressions, i.e. $S_{1c}=S_{1b}=0<S_{1a}=1$.

- B: if $S_{2c}>S_{2a}=S_{2b}$, advertisers' willingness-to-pay each relevant impression on 1 are such that $w_{1a}=v\left(1-\frac{2mS_{2a}}{1+m}\right)=w_{1b}=v\left(1-\frac{2mS_{2b}}{1+m}\right)>w_{1c}=v\left(1-\frac{2mS_{2c}}{1+m}\right)$. Hence, a and b bid equally and outbid c for for all such impressions, i.e. $S_{1a}=S_{1b}=1/2>S_{1c}=0$.

- C: if $S_{2a}=S_{2b}=S_{2c}$, advertisers' willingness-to-pay each relevant impression on 1 are all equal to $v\left(1-\frac{2mS_{2j}}{1+m}\right)$. So they all bid equally and we obtain $S_{1a}=S_{1b}=S_{1c}=1/3$.

In equilibrium, these bidding strategies (and the ensuing shares S_{2j}) must be consistent with the bidding strategies on publisher 2. If case A applies, since $S_{1c}=S_{1b}=0<S_{1a}=1$, by the same

reasoning as above we must have $S_{2c} = S_{2b} = 1/2 > S_{2a} = 0$. Since this outcome is inconsistent with the assumption that $S_{2c} > S_{2b} > S_{2a}$, we disregard this case.

If Case B applies, we have $S_{1a} = S_{1b} = 1/2 > S_{1c} = 0$ and thus $S_{2c} = 1 > S_{2a} = S_{2b} = 0$. In this equilibrium candidate, all advertisers single-home. Given (11), the price of relevant impressions on publishers 1 (second highest willingness-to-pay) is v. As for publisher 2, the second-highest willingness-to-pay (and the price of impressions), is $v\left(1 - \frac{m}{1+m}\right)$, but advertiser c's willingness-to-pay is v. In this equilibrium candidate, advertisers a and b earn zero while advertiser c earns $v\frac{m}{1+m}$ for each impression acquired. Given there are $\tilde{q}\frac{1+m}{2}$ relevant impressions per publisher in this market, advertiser c earns $vm\tilde{q}/2$ in total.

If Case C applies, all advertisers place identical bids, equal to $v\left(1 - \frac{2m}{3(1+m)}\right)$, for each relevant impression on each publisher. The latter is also the price of relevant impressions on both publishers. This price equals the expected return from each impression for all advertisers. Therefore, the advertisers make zero profit in this equilibrium candidate.

We have thus identified two equilibrium candidates (case B and case C) and must now establish whether these are indeed equilibria. The candidate associated to case C cannot be an equilibrium because, given the bids placed by the rivals, each advertiser can deviate by bidding $v\left(1 - \frac{2m}{3(1+m)}\right) + \varepsilon$, where $\varepsilon > 0$ and arbitrarily small, for all relevant impressions on one publisher (and thus win them all) and zero for all impressions on the other. The advertiser would earn a strictly positive profit by deviating becuase each such impression would be worth v, so the deviation is profitable.

As for the candidate associated to case B, there is no profitable deviation: no advertiser can profitably outbid the others on the publisher where it is not acquiring any impression, because the winning bid equals v. Nor would the advertiser gain by reducing its bid for the impressions it is winning. Hence, the candidate associated to case B is indeed an equilibrium. Note that also all the other candidates associated to case B, obtained by permutations of a, b, c and $1, 2$, are equilibria.

Thick markets. Consider a market such that $n = 4$. Let a, b, c, d denote the set of advertisers in the market. Focus again on the relation between S_{2j} and the value of (11) for relevant impressions on publisher 1, and consider the bidding strategy of the advertisers on this publisher. The following cases are possible:

- A: if $min\,(S_{2d}, S_{2c}) > S_{2b} > S_{2a}$, advertisers' willingness-to-pay each relevant impression on 1 are such that $max(v(1 - \frac{2mS_{2d}}{1+m}), v\left(1 - \frac{2mS_{2c}}{1+m}\right)) < w_{1b} = v\left(1 - \frac{2mqS_{2b}}{1+m}\right) < w_{1a} = v\left(1 - \frac{2mS_{2a}}{1+m}\right)$. Hence, a outbids the other advertisers for all such impressions, i.e. $S_{1d} = S_{1c} = S_{1b} = 0 < S_{1a} = 1$.

- B: if $min\,(S_{2d}, S_{2c}) > S_{2b} = S_{2a}$, advertisers' willingness-to-pay each relevant impression on 1 are such that $max\left(v(1 - \frac{2mS_{2d}}{1+m}), v\left(1 - \frac{2mS_{2c}}{1+m}\right)\right) < w_{1b} = v\left(1 - \frac{2mS_{2b}}{1+m}\right) = w_{1a} =$

$v\left(1-\frac{2mS_{2a}}{1+m}\right)$. Hence, a and b bid equally and outbid c and d for for all such impressions, i.e. $S_{1a}=S_{1b}=1/2>S_{1c}=S_{1d}=0$.

- C: if $S_{2d}>S_{2c}=S_{2b}=S_{2a}$, advertisers' willingness-to-pay each relevant impression on 1 are such that $w_{1d}=v(1-\frac{2mS_{2d}}{1+m})<w_{1c}=v\left(1-\frac{2mS_{2c}}{1+m}\right))=w_{1b}=v\left(1-\frac{2mS_{2b}}{1+m}\right)=w_{1a}=v\left(1-\frac{2mS_{2a}}{1+m}\right)$. Hence, a,b and c bid equally and outbid d for for all such impressions, i.e. $S_{1a}=S_{1b}=S_{1c}=1/3>S_{1d}=0$.

- D: if $S_{2a}=S_{2b}=S_{2c}=S_{2d}=1/4$, advertisers' willingness-to-pay each relevant impression on 1 are all identical, and equal to $v\left(1-\frac{m}{2(1+m)}\right)$. So the advertisers also place identical bids, such that $S_{1a}=S_{1b}=S_{1c}=S_{1d}=1/4$.

If case A applies, since $S_{1d}=S_{1c}=S_{1b}=0<S_{1a}=1$, by the same reasoning as above we must have $S_{2d}=S_{2c}=S_{2b}=1/3>S_{2a}=0$. As this outcome is inconsistent with the assumption that $min\,(S_{2d},S_{2c})>S_{2b}>S_{2a}$, we can disregard this case.

If Case B applies, we have $S_{1a}=S_{1b}=1/2>S_{1c}=S_{1d}=0$ and thus $S_{2c}=S_{2d}=1/2>S_{2a}=S_{2b}=0$. This case constitutes an equilibrium candidate such that all advertisers single-home, i.e. place winning bids on a single publisher. Given (11), advertisers a and b's willingness-to-pay is v on publisher 1 while c and d's is $v\left(1-\frac{m}{1+m}\right)$. The price of relevant impressions on publisher 1 is thus v. By the same token, the price of relevant impressions on publisher 2 is v. Therefore, the advertisers make zero profit in this equilibrium candidate.

If Case C applies, we have $S_{1a}=S_{1b}=S_{1c}=1/3>S_{1d}=0$ and thus $S_{2d}=1>S_{2a}=S_{2b}=S_{2c}=0$. In this equilibrium candidate, all advertisers single-home, i.e. place winning bids on a single publisher. Given (11), advertisers a, b and c's willingness-to-pay is v on publisher 1 while d's is equals $v\left(1-\frac{2m}{1+m}\right)$. The price of relevant impressions on publisher 1 is thus v. On publisher 2, d's willingness-to-pay is v while a, b and c's equals $v(1-\frac{2m}{3(1+m)})$. The latter is the price of relevant impressions on publisher 2. Therefore, in this equilibrium candidate advertisers a, b and c make zero profit, while d makes $v\frac{2m}{3(1+m)}$ per each of the $\tilde{q}\frac{1+m}{2}$ impressions acquired on publisher 2.

If Case D applies, all advertisers place identical bids on both publishers. Given (11) each advertiser's willingness-to-pay is $v(1-\frac{m}{2(1+m)})$ for all relevant impressions on each publisher. The latter is also the equilibrium price of such impressions. Therefore, the advertisers make zero profit in this equilibrium candidate.

We have thus identified possible equilibrium candidates in cases B, C and D, and must now establish whether these are indeed equilibria. The candidate associated to case D cannot be an equilibrium because, given the bids placed by the rivals, each advertiser can deviate by bidding $v\left(1-\frac{m}{2(1+m)}\right)+\varepsilon$, where $\varepsilon>0$ and arbitrarily small, for all relevant impressions on one publisher (and thus win them all) and zero for all impressions on the other. The advertiser would single-home in this deviation, and earn a strictly positive profit, so the deviation is profitable.

Similarly, the candidate associated to case C cannot be an equilibrium because each advertiser among a, b and c can deviate by bidding $v\left(1-\frac{2m}{3(1+m)}\right)+\varepsilon$, where $\varepsilon>0$ and arbitrarily small, for all relevant impressions on publisher 2 (and thus win the impressions on this publisher) and zero for impressions on the other. The deviating advertiser would earn strictly positive profit since each impression on publisher 2 would now be worth v.

As for the candidate associated to case B, there is no profitable deviation: no advertiser can profitably outbid the others on the publisher where it is not acquiring any impression, for the winning bid there equals v already. Nor would the advertiser gain by reducing its bid for the impressions it is winning. Hence, the candidate associated to case B is indeed an equilibrium. Note that also all the other candidates associated to case B, obtained by permutations of a, b, c, d and $1, 2$, are equilibria.

Finally, following a similar reasoning as above one can show that, if $n>4$, the equilibria are again such that advertisers single-home and there are at least two advertisers winning impressions and bidding v on each publisher. Hence, the equilibrium price of all impressions is still equal to v, as in the case where $n=4$.

A.3 Proof of Proposition 3

Consider first the total surplus on the advertising market when the publishers do not outsource. Given Lemma 3, no impression is wastefully repeated and thus generates a value v in equilibrium. Therefore, total surplus from advertising when the publishers do not outsource is

$$W=\int_0^{v_H} vq\,(1+m)\,dv\ . \tag{22}$$

Consider now the total surplus generated when the publishers outsource. Given Proposition 1, the total surplus is

$$\begin{cases} W^F=\int_0^{v_H} v\widetilde{q}\,(1+m)\,dv & \text{if } x\leq\frac{1}{2}, \\ (1-m)\,W^F+2mW^{PDT}=\int_0^{v_H} v\widetilde{q}dv & \text{if } x>\frac{1}{2}. \end{cases} \tag{23}$$

Therefore, if $x\leq\frac{1}{2}$, given $\tilde{q}>q$, total surplus increases. Suppose now that $x>\frac{1}{2}$, the total surplus generated on multi-homers given PT is equal to half the total returns from ads on such consumers, because some impressions are sold to the "wrong" advertiser, hence lost. Hence, $\int_0^{v_H} vq\,(1+m)\,dv<\int_0^{v_H} v\widetilde{q}dv$ if and only if $\tilde{q}>\tilde{q}_w\equiv q\,(1+m)$.

A.4 Proof of Proposition 4

When restricted to F, the IN earns the revenue $\overline{v}\tilde{q}\,[1+m\,(1-2x)]$ when publishers outsource. This revenue is smaller than in the baseline model for $x\geq 1/2$. Hence, the region of parameters such that the publishers outsource gets smaller when the regulation is adopted.

Let us concentrate in the region such that $x \geq 1/2$, where there is PT without the regulation. Given that outsoucing occurs, the share of identified impressions is the same, however total surplus on each impression on multi-homers is v under full disclosure and $\frac{v}{2}$ under partial disclusure on type. Given that outsourcing does not occur, total surplus is $W = \int_0^{v_H} vq(1+m)\,dv$ when there is no outsourcing and $(1-m)W^F + 2mW^{PDT} = \int_0^{v_H} v\widetilde{q}dv$ without transparency regulation. Hence, total surplus decreases with the regulation $\int_0^{v_H} vq(1+m)\,dv < \int_0^{v_H} v\widetilde{q}dv$ if and only if $\widetilde{q} \equiv q(1+m) < \widetilde{q}$.

B Intermediate markets

We consider now the case where there are $n = 3$ advertisers in a market, and we solve the game considering all possible values of n.

Full Disclosure ($r = F$). Following the same steps as in 4.1.1, we find that if the consumer is a single-homer, advertisers' willingness-to-pay for a targeted impression is $w_{i,j}^F = v$, for any i, since the impression cannot be repeated. If the consumer is a multi-homer, in an intermediate market there are two advertisers willing to pay v. Hence, the equilibrium price of an impression in intermediate markets is

$$p_{SH}^F = p_{n=4,MH}^F = v.$$

Hence, intermediate markets behave as thick markets. Considering all markets, the expected revenue earned by the platform are

$$R_{SH}^F = v\tilde{q}(1-m), \quad R_{MH}^F = (1-x)\,2mv\tilde{q}. \tag{24}$$

Partial Disclosure on Types ($r = PT$). In this case, there are at least two advertisers willing to bid $w_{i,j}^{PT} = v/2$ (again, $t = 2$ in equilibrium) on each impression sold by the intermediary. Hence, also in intermediate markets, the equilibrium price is

$$p^{PT} = \frac{v}{2}. \tag{25}$$

Considering all markets, we find that revenues are independent of the thickness of the market and equal to

$$R_{SH}^{PT} = (1-m)\,\tfrac{v\tilde{q}}{2}, \quad R_{MH}^{PT} = mv\tilde{q}. \tag{26}$$

Partial Disclosure on History ($r = PH$). In intermediate markets, two advertisers (say, a and b) single-home on publisher 1, whereas advertiser c single-homes on 2. The equilibrium willingness-to-pay for impressions on 1 by a and b is v, whereas their willingness-to-pay for these impressions on publisher 2 is $v\left(1 - \frac{m}{1+m}\right)$. This expression is derived from (11), noting that a and b each win half

Table 4: Advertiser's willingness-to-pay, equilibrium price and intermediary's revenues under PH.

Intermediate ($n = 3$)	Publisher 1	Publisher 2
w_{ia}	v	$v\left(1-\frac{m}{1+m}\right)$
w_{ib}	v	$v\left(1-\frac{m}{1+m}\right)$
w_{ic}	$v\left(1-\frac{2m}{1+m}\right)$	v
$p_{n=3}$	v	$v\left(1-\frac{m}{1+m}\right)$
$R_{n=3}$	$v\tilde{q}\frac{(1+m)}{2}+v\left(1-\frac{m}{1+m}\right)\tilde{q}\frac{(1+m)}{2}=v\tilde{q}\left(1+\frac{m}{2}\right)$	

the available targeted impressions on publisher 1, so $S_{1a} = S_{1b} = 1/2$. Then, the willingness-to-pay by advertiser c for each targeted impression on publisher 2 is v, whereas on publisher 1 it equals (6), since $S_{2c} = 1$. Consequently, the equilibrium price of targeted impressions on publisher 1 is v, whereas it equals $v\left(1-\frac{m}{1+m}\right)$ on publisher 2. Thus, revenues of the intermediary are $v\tilde{q}\frac{(1+m)}{2}$ for impressions sold on publisher 1, but only $v\left(1-\frac{m}{1+m}\right)\tilde{q}\frac{(1+m)}{2}$ for impressions on publisher 2. Table

Hence, considering all markets, total revenues are

$$R^{PH} = x\overline{v}\tilde{q}\,(1-m) + yv\tilde{q}\left(1+\frac{m}{2}\right) + (1-x-y)\,\overline{v}\tilde{q}\,(1+m)\,.$$

The publisher that ends up serving two advertisers extracts the full advertising surplus, but the other publisher does not. We summarize the bids and profits of the publishers in the middle panel of Table 4.

Choice among disclosure regimes. First, we compare F and PH. Revenues in intermediate markets are higher under F than under PH, because the intermediary are able to extract all surplus in the former but not in the latter regime. Hence, because in thin and thick markets (see (2) and Table 2) F and PH yield the same expected revenue to the intermediary, then when considering also intermediate markets, F turns out to be the most profitable regime.

Compare now F and PT. Comparing revenues that include intermediate markets (see (24) and (26)) with revenues in the baseline mode (see (2) and (4)), we see that revenues are not affected. Indeed, revenues in intermediate markets are as in thick markets. Hence, Proposition 1 and Lemma 2 are not affected.

Market equilibrium without the intermediary. Following the same steps as in 4.2 and above for the regime PH with intermediate markets, we find that only one of the two publishers (say 1) is able to extract all surplus, while the other (say 2) is not (see Table 5 for the equilibrium price of impressions.). Hence, we find that $R_{1,n=3} = \frac{vq}{2}(1+m)$ and $R_{2,n=3} = \frac{vq}{2}(1+m(1-q))$.

We can now compute the aggregate profits earned by the publishers in all markets, we find

Table 5: Advertiser's willingness-to-pay, equilibrium price and publishers's revenues for intermediate markets.

Intermediate $(n=3)$	Publisher 1	Publisher 2
w_{ia}	v	$v\left(1-\frac{mq}{1+m}\right)$
w_{ib}	v	$v\left(1-\frac{mq}{1+m}\right)$
w_{ic}	$v\left(1-\frac{2mq}{1+m}\right)$	v
$p_{n=3}$	v	$v\left(1-\frac{mq}{1+m}\right)$
$R_{i,n=3}$	$\frac{vq}{2}(1+m)$	$\frac{vq}{2}(1+m(1-q))$

$$
\begin{aligned}
R_1+R_2 = 2x\left[\frac{\overline{v}q}{2}(1+m(1-2q))\right]+y\left[\frac{\overline{v}q}{2}(1+m)+\frac{\overline{v}q}{2}(1+m(1-q))\right]+ \\
+2(1-x-y)\left[\frac{\overline{v}q}{2}(1+m)\right]=\overline{v}q\left[1+m\left(1-q\left(2x+\frac{y}{2}\right)\right)\right].
\end{aligned}
\tag{27}
$$

Hence, total surplus are affected by the presence of intermediate markets.

Outsourcing decision. Comparing the revenues with the intermediary in 10 to the revenues of the publishers in 27. We find that the publishers outsource if and only if $\tilde{q}$ is high enough, that is

$$
\tilde{q} \geq \tilde{q}_T \equiv \begin{cases} q\left(1+\frac{2mx(1-q)-mq\frac{y}{2}}{1+m(1-2x)}\right), & \text{if } x \leq \frac{1}{2}, \\ q\left(1+m\left(1-q\left(2x+\frac{y}{2}\right)\right)\right), & \text{if } x > \frac{1}{2}. \end{cases}
$$

In the baseline model, the share of thin markets is x, and all the others are thick. This is like pooling intermediate and thick markets. Because profits in intermediate markets are lower than under thick markets for publishers, while they are the same for the intermediary, the region where outsourcing occurs gets bigger when intermediate markets are considered in the analysis.

C Proofs of robustness checks

C.1 Reserve price

We establish that adopting F with a revenue-maximizing reserve price does not necessarily increase IN's revenue. We proceed as follows. We characterize the equilibrium bidding strategies of advertisers conditional on the reserve price p. Next, we compute the equilibrium value of p chosen by IN. Note that p can only be made conditional on whether the consumer is a single- or multi-homer, since the platforms do not observe v and n. Finally, we compare the revenue earned with the revenue-maximizing reservation price under F to the revenue earned with PT. For concreteness, we focus on impressions on multi-homers in thin advertising markets, where, as demonstrated in the baseline model, there is a strong drop in ad prices on multi-homers under F that gives the intermediary an incentive to choose PT. Also, for the sake of brevity, we do not consider PH.

Consider an impression on a profiled consumer and assume IN adopts F.If the consumer is a multi-homer, the equilibrium willingness-to-pay are as characterized in Table 1 for the couple of impressions on this consumer. Hence, the equilibrium price of each such impression is zero if there is no reserve price. With a reserve price p, the equilibrium price is p, but the impressions are sold only if the market is such that $v \geq p$. Remark that, if the consumer is a single-homer, all the advertisers in the given market bid v. Hence, the equilibrium price of the impression is v and IN cannot do better by imposing a reserve price.

Let us now compute IN's expected revenue from multi-homers in each (thin) market when adopting F and setting a reserve price p. This revenue is $\tilde{q}\,(2mp)$ if $v \geq p$, and 0 if $v < p$, since in the latter case IN sells impressions on multi-homers at zero. Because impressions on multi-homers in thin markets are sold at zero in the baseline model, it follows that, conditional on adopting F, the intermediary is better off imposing a reserve price such that $v_H \geq p > 0$, for all impressions on multi-homers.

Let us now calculate the revenue-maximizing p. The revenue on multi-homers in thin advertising markets is

$$R^{F}_{n=2,MH}(p) = x\tilde{q}\left(2mp\int_{p}^{v_H} dF(v)\right) = x\tilde{q}\,(2pm\,(1-F(p)))\,. \tag{28}$$

The revenue-maximizing reserve price, p^*, is such that

$$\frac{d\pi^{F}_{n=2}(p)}{dp} = 1 - F(p) - pf(p) = 0 \tag{29}$$

and therefore

$$p^* = \frac{1-F(p^*)}{f(p^*)}. \tag{30}$$

Suppose now IN adopts PT. There is no scope for a reserve price to increase revenues under PT for multi-homers, because the equilibrium price is equal to the highest bid. . Hence, conditional on choosing PT, IN does not impose any reserve price. Given the price of an impression under PT (3) in the main text, the total revenues from multi-homers in thin markets under PT is $R^{PT}_{n=2,MH} = x\tilde{q}\overline{v}m$.

Given the above findings, IN adopts F with the reserve price p^* in (31) for each impression on multi-homers if and only if $\pi^{F}_{n=2}(p^*) > \pi^{PT}_{n=2}$. Otherwise, IN adopts PT without a reserve price for each such impression. We find that

$$R^{F}_{n=2,MH}(p^*) \geq R^{PT}_{n=2,MH} \iff x\tilde{q}m\left(2\frac{(1-F(p^*))^2}{f(p^*)} - \overline{v}\right) \geq 0. \tag{31}$$

At this level of generality, it is difficult to establish whether the above condition holds, as this ultimately depends on the distribution of advertising returns, $F(v)$. However, the above inequality does not hold for a well-known distribution:

- Suppose that $F(v)$ is a continuous Bernoulli distribution with $\lambda = \frac{1}{2}$. Hence, $E(v) = \frac{1}{2}$, $F(p) = p$ and $f(p) = 2$. Replacing these values in (30), we obtain $p^* = \frac{1}{3}$. Replacing these values in (31), we obtain that $\pi^F_{n=2}(p^*) < \pi^{PD}_{n=2}$.

C.2 Competition between advertisers

No intermediary. Let w_{ij} be the willingness-to-pay of advertiser $j \in k(\theta)$ for an impression on a profiled consumer of type θ on publisher i. w_{ij} depends on whether the consumer is profiled on the other publisher as well, and thus exposed to ads from advertisers in the same market:

$$w_{ij} = v\left(\frac{\frac{1-m}{2} + m\left[(1-q) + q\left((1-\alpha)S_{i'j} + \alpha\left(1 - S_{i'j}\right)\right)\right]}{\frac{1+m}{2}}\right) \; i, i' = 1, 2; i' \neq i \tag{32}$$

To understand this expression, consider that the return from informing multi-homers not profiled by the other publisher is v. The same return characterizes single-homers, who receive just one impresion by assumption. However, the return is αv if the consumer is informed by another advertiser in the same market. In addition, if the advertiser already informs the consumer with an impression on the other publisher, the return is $(1-\alpha)v$, i.e. the value of avoiding that the consumer receives an impression from a competitor. Note that w_{ij} decreases in $S_{i'j}$ if and only if $\alpha > \frac{1}{2}$. Therefore, as in our baseline model, impressions on the two publishers are substitutes for the advertisers if $\alpha \geq \frac{1}{2}$. By contrast, if $\alpha < \frac{1}{2}$, the impressions are complementary. Hence, unlike in the baseline model, the equilibrium bidding strategies must be such that a single advertiser in each market acquires all the targeted impressions on both publishers. We summarize in the following

Lemma 5. *If neither or only one of the publishers outsources to the intermediary, in each market the equilibrium price of a targeted impression equals the second-highest willingness-to-pay among the advertisers in $k(\theta)$. Furthermore, if $\alpha \geq \frac{1}{2}$, all advertisers in each market single-home on different outlets. By contrast, if $\alpha < \frac{1}{2}$, in each market a single advertiser acquires all the targeted impressions.*

Given the above claims, we can follow similar steps as in Section 4.2.2 to establish the following. Suppose that $\alpha \geq \frac{1}{2}$. In thin markets, each advertiser bids v on one publisher and $v\left(\frac{(1-m)+2m(1-\alpha q)}{1+m}\right)$ on the other publisher. Hence, at equilibrium, the price of impressions is $p_{n=2} = v\left(\frac{(1-m)+2m(1-\alpha q)}{1+m}\right)$ and each publisher earns a revenue equal to $R_{i,n=2} = \overline{v}\left(\frac{1-m}{2} + m(1-\alpha q)\right)$. In thick markets, the equilibrium price of impressions is $p_{n\geq 4} = v\left(\frac{(1-m)+2m(\alpha q+1-q)}{1+m}\right)$ and revenues are $R_{i,n\geq 4} = \overline{v}\left(\frac{1-m}{2} + m(\alpha q + 1 - q)\right)$.

Suppose now that $\alpha < \frac{1}{2}$. A single advertiser multi-homes and wins all impressions. The second-highest willingness-to-pay comes from an advertiser that acquires zero ads in equilibrium. By replacing $S_{i'j} = 0$ in (32), the equilibrium price and revenues from impressions are $p = v\left(\frac{(1-m)+2m(\alpha q+1-q)}{1+m}\right)$ and $R = \overline{v}\left(\frac{1-m}{2} + m(\alpha q + 1 - q)\right)$ respectively, in either thin or thick markets.

Intermediary and Full Disclosure. Assume the intermediary chooses F and consider the return that advertisers in a given market get from a targeted impression. This return is v if the consumer is a single-homer (who does not receive any other impression). If the consumer is a multi-homer, however, the advertiser gets a total return of v if it is the only one to inform the consumer (i.e., buys both impressions), and αv if a competitor informs the consumer too. Therefore, the willingness-to-pay for an impression on a on multi-homer is αv if the advertiser is not already acquiring the other impression on the given consumer (which implies that a rival is), and $v(1-\alpha)$ otherwise (this is the value of preventing a competitor from informing the consumer as well). Hence, impressions on single-homers are sold at v in equilibrium, while the price of impressions on multi-homers depends on α. If $\alpha \geq \frac{1}{2}$, each impression is bought by a different advertiser. If the market is thin, the price for impressions on multi-homers is $p = v(1-\alpha)$ and revenues are $R^F = \overline{v}(1 - m + 2m(1-\alpha))$. In thick markets, the price is $p = v\alpha$ and $R^F_{n=2} = \overline{v}(1 - m + 2m\alpha)$. If $\alpha < \frac{1}{2}$, in thin and thick markets, both impressions on multi-homers are sold to same advertiser at $p = v\alpha$ and revenues are $R^F = \overline{v}(1 - m + 2m\alpha)$.

Intermediary and Partial Disclosure on History. Suppose now that the intermediary adopts PH. Following the same reasoning as when characterizing the willingness-to-pay with the publishers operating independently, the willingness-to-pay by advertiser $j \in k(\theta)$ is

$$w^{PH}_{i,j} = v\frac{\left(\frac{1-m}{2}\right) + m\left((1-\alpha)S^{PH}_{i'j} + \alpha\left(1 - S^{PH}_{i'j}\right)\right)}{\frac{1+m}{2}}, \, i' \neq i.$$

Observe that $w^{PH}_{i,j}$ decreases in $S^{PH}_{i'j}$ if and only if $\alpha \geq \frac{1}{2}$. Therefore, if $\alpha \geq \frac{1}{2}$, advertisers in each market single-home on different publishers. In thin markets, the equilibrium price and revenues are $p_{n=2} = v\left(\frac{(1-m)+2m(1-\alpha)}{1+m}\right)$ and $R^{PH} = \overline{v}(1 - m + 2m(1-\alpha))$ respectively. In thick markets, the equilibrium price and revenues are $p_{n\geq 4} = v\left(\frac{(1-m)+2m\alpha}{1+m}\right)$ and $R^{PH} = \overline{v}(1 - m + 2m\alpha)$ respectively. Suppose now that $\alpha < \frac{1}{2}$. One advertiser multi-homes and wins all the impressions. In both thin and thick markets, the equilibrium price and revenues are $p = v\left(\frac{(1-m)+2m\alpha}{1+m}\right)$ and $R^{PH} = \overline{v}(1 - m + 2m\alpha)$ respectively.

Intermediary and Partial Disclosure on Types. Suppose now that the intermediary adopts PT for targeted impressions on multi-homers. Recall that impressions on single-homers, sold under F, will all be sold at a price equal to v. At equilibrium, each market gets conflated with one more other market. The return from informing the consumer is $v/2$ if the consumer does not receive any other impression from competitiors. Note, however, that if the consumer receives an impression from an advertiser in the conflated market, who is not a competitor, the return is still $v/2$. The return is instead $\alpha v/2$ if the consumer receives another impression from a competitior. At equilibrium, therefore, each advertiser anticipates correctly that if it buys only one impression on a multi-homer, the other will be bought by an advertiser from the other conflated market. Indeed, in equilibrium each impression is acquired by an advertiser from a different conflated market, and the equilibrium price of each impression is $p^{PT} = \frac{v}{2}$. Independently of market thickness, total revenues are $R^{PT} = (1-m)\overline{v} + \frac{\overline{v}}{2}\tilde{q}2m \ = \tilde{q}\overline{v}$.

Choice of disclosure regime. As in the baseline model, Full Disclosure dominates Partial Disclosure on History. Then, impressions on single-homers are sold under full disclosure. Then, for impressions on multi-homers, IN chooses F over PT if and only if

$$\begin{cases} \overline{v}\tilde{q}\left(x\left(2m\left(1-\alpha\right)\right)+\left(1-x\right)\left(2m\alpha\right)\right)>\overline{v}\tilde{q}m, & \text{if } \alpha \geq \frac{1}{2}, \\ \overline{v}\tilde{q}\left(2m\alpha\right)>\overline{v}\tilde{q}m, & \text{if } \alpha < \frac{1}{2}. \end{cases}$$

Hence, if $\alpha \geq \frac{1}{2}$, there is F if and only if $x+\alpha \leq \frac{1}{2}$, similarly to our baseline model. Istead, if $\alpha < \frac{1}{2}$, PT dominates F.

C.3 Proof of the claim in Section 6.3

Let α be the share of the revenue from each impression that the intermediary transfers to each publisher. It is straightforward to show that Proposition 1 does not change, since, given α, the intermediary's objective, $(1-2\alpha)\,R_{IN}$ for any disclosure regule. Hence, the choice of disclosure regime is independent of α. Let us now establish how α should be set in order to induce the publishers to outsource. Each publisher outsources if and only if $\alpha R_{IN} \geq R_i$. Since the intermediary's objective is decreasing in α, in equilibrium we will have $\alpha = {}^{R_1}/_{R_{IN}} = {}^{R_2}/_{R_{IN}}$, R_1 and R_2 being identical in our setting. Replacing in the intermediary's profit, $(1-2\alpha)\,R_{IN}$, we obtain that this profit is positive if and only if $R_{IN} \geq \sum_{i=1,2} R_i$, which is the same as (16).

C.4 Increasing returns to advertising on the same consumer and re-targeting

We assume that the first impression by an advertiser on a consumer is worth v and the second impression is worth $v_s > v$.

No outsourcing by the publishers. We now characterize the equilibrium conditional on the publishers not outsourcing to the intermediary. The willingness-to-pay for a targeted impression on publisher i by advertiser j is

$$w_{ij} = v\frac{\frac{1-m}{2}+m\left[\left(1-q\right)+q\left(\frac{v_s}{v}S_{i'j}+1-S_{i'j}\right)\right]}{\frac{1+m}{2}},\ i,i'=1,2;i'\neq i.$$

To understand this expression, consider that the return from informing single-homers is v. For impressions that hit multi-homers, the return is v if the consumer is not already receiving an impression from the same advertiser. However, if the impression is repeated, the return from the second impression is $v_s > v$. This implies that impressions on the two publishers are complements: w_{ij} increases in $S_{i'j}$. It follows that the only possible equilibrium bidding strategies are such that

	Publisher 1	Publisher 2
$w_{i,a}$	$v\frac{\frac{1-m}{2}+m\left[(1-q)+q\frac{v_s}{v}\right]}{\frac{1+m}{2}}$	$v\frac{\frac{1-m}{2}+m\left[(1-q)+q\frac{v_s}{v}\right]}{\frac{1+m}{2}}$
$w_{i,b}$	v	v
$p_{n=2}$	v	v
$R_i\,(n=2)$	$\frac{vq}{2}\,(1+m)$	$\frac{vq}{2}\,(1+m)$

Table 6: Equilibrium willingness-to-pay without intermediary and with $n = 2$.

	Publisher 1	Publisher 2
$w_{i,a}$	$v\frac{\frac{1-m}{2}+m\left[(1-q)+q\frac{v_s}{v}\right]}{\frac{1+m}{2}}$	$v\frac{\frac{1-m}{2}+m\left[(1-q)+q\frac{v_s}{v}\right]}{\frac{1+m}{2}}$
$w_{i,b}$	v	v
$w_{i,c}$	v	v
$w_{i,d}$	v	v
$p_{n\geq 4}$	v	v
$R_i\,(n\geq 4)$	$\frac{vq}{2}\,(1+m)$	$\frac{vq}{2}\,(1+m)$

Table 7: Equilibrium willingness-to-pay without intermediary and with $n \geq 4$.

one advertiser outbids all the others in a given market for all the targeted impressions, on both publishers. Since $S_{i'j} = 0$ for all the other advertisers on both publishers, in a first-price auction the equilibrium price of impressions is

$$p = v\frac{\frac{1-m}{2}+m}{\frac{1+m}{2}} = v.$$

Tables 6 and 7 summarize the willingness-to-pay, equilibrium bids and profits of the publishers in a given thin and thick market, respectively. In these tables, we focus (without loss) on the equilibrium where advertiser a wins all the targeted impressions. The total revenue earned by each publisher in this equilibrium is $\frac{vq}{2}\,(1+m)$.

Intermediary and publishers' choice of outsourcing. We now characterize the equilibrium conditional on the publishers outsourcing to the intermediary. Suppose the intermediary adopts F. In this case, the equilibrium price for impressions on a single-homer is v. If the consumer is a multi-homer, an advertiser already placing an impression on the consumer has willingness-to-pay v_s for the second impression, while the willingness-to-pay of any other advertiser is v. It follows that the equilibrium price is v for all targeted impressions on multi-homers. Hence, the intermediary can earn the following revenue in any given market:

$$\tilde{q}v\,(1+m)\,. \tag{33}$$

It is straightforward to show that this revenue is weakly larger than the revenue the platform could earn under either PT or PH. Hence, F is the disclosure regime chosen by the platform in equilibrium.

This revenue also clearly exceeds the revenue the publishers can earn independently, given $\tilde{q} > q$. Hence, the publishers outsource to the intermediary in this scenario.

C.5 Heterogeneous returns to advertising within each advertising market

We modify the baseline setting by allowing for markets where one advertiser (that we take to be a without loss of generality) has a higher valuation, v^+, than the remaining $n-1$ advertisers, whose valuation is $v < v^+$. Specifically, we assume there is a share x of thin markets such that advertisers are homogeneous and a share x' of thin markets where advertiser a (the "dominant" one) has a return v^+ from informing consumers, large enough that it acquires all the impressions on all platforms. Similarly, there is a share z of thick markets where advertisers are homogeneous, and a share $1-x-x'-z$ of thick markets with a dominant advertiser. For simplicity, we assume advertisers within each market are aware of the presence of a dominant advertiser (if any) and of the value of v^+. We ignore intermediate markets for ease of exposition.

C.5.1 Intermediary

We now revisit the equilibrium prices and revenues earned by the intermediary in the three disclosure regimes considered in the baseline model.

Full disclosure. The presence of a dominant advertiser does not change the equilibrium prices and revenue under F. If the market is thin, since only one advertiser is willing to place a positive bid per each targeted impression on a multi-homer, the equilibrium price for that impression is zero. In thicker markets, with a single dominant advertiser, the second highest willingness-to-pay for each such impression is v. It follows that the revenue under F is the same as in expressions (2).

Partial Disclosure on Types. The presence of a dominant advertiser does not change the equilibrium prices and revenue under PT. When the intermediary conflates different markets, even with a dominant advertiser, the second-higest willingness-to-pay for each targeted impression on a multi-homer or a single-hoemer is $v/2$. Hence, the revenue is as characterized in (4).

Partial Disclosure on History. Consider first a *thin* market with a dominant advertiser. Consider advertiser j's bidding strategy for each targeted impression on publisher i. As in Section 4.2, we can characterize advertiser j's willingness-to-pay for each impression on publisher i as follows

$$w_{ij} = v^+ \left(1 - \frac{2mS_{i'j}}{1+m}\right), \quad i = 1, 2. \tag{34}$$

Similarly to the proof of Lemma 1 (see Appendix A.2), we can establish that the following configurations can potentially emerge in equilibrium

- A: $S_{2b} = 0$ and $S_{1b} = 0$. The willingness-to-pay for impressions are $w_{ia} = v^+ \left(1 - \frac{2m}{1+m}\right) > w_{ib} = v$. In this equilibrium, a outbids b for every targeted impression on every publisher.

- B: $S_{ib} = 0$ and $S_{i'b} = 1$. The willingness-to-pay for impressions are $w_{ia} = v^+ \left(1 - \frac{2m}{1+m}\right) < w_{ib} = v$. Hence, a and b single-home on different publishers. The same condition is necessary and sufficient for an equilibrium entailing the symmetric configuration, $S_{i'b} = 0$ and $S_{ib} = 1$.

All other configurations can be ruled out as follows. Consider any equilibrium candidate bidding strategy such that $S_{1b} > 0$ and $S_{2b} > 0$ (i.e. advertiser b multi-homes). Suppose $S_{ib} = 1$ for either $i = 1$ or $i = 2$. Then $w_{i'a} = v^+$ must exceed $w_{i'b} \leq v$. This implies that $S_{i'a} = 1$, which contradicts the assumption that both S_{1b} and S_{2b} are strictly positive. Suppose now that $1 > S_{1b} > 0$ and $1 > S_{2b} > 0$. These inequalities can hold if and only if all advertisers bid $w_{ia} = w_{ib}, \forall i$. However, if the latter equalities hold, both advertisers get a surplus equal to zero in the candidate equilibrium. Hence, advertiser a can deviate by bidding zero for each targeted impression on publisher i' and bidding $w_{ia} = w_{ib}$ for each targeted impression on i. As the advertiser acquires impressions on a single publisher, each such impression would be worth v^+. This deviation would be profitable because, given $S_{i'b} > 0$, the equilibrium price of impresions would be $w_{ib} < v < v^+$.

Summing up, the subgame that takes place conditional on the publishers not outsourcing to IN admits two possible equilibrium configurations. If and only if $v^+ \left(1 - \frac{2m}{1+m}\right) \geq v$ holds, the equilibrium is such that a multi-homes and buys all the targeted impressions on both outlets. Otherwise, the equilibrium is such that each advertiser single-homes on a different publisher. In the former equilibrium, the price of each targeted impression equals v, so the intermediary earns $vq\,(1+m)$. In the latter equilibrium, the price on publisher 1 is $v^+ \left(1 - \frac{2m}{1+m}\right)$, whereas it equals $v \left(1 - \frac{2m}{1+m}\right)$ on publisher 2. Hence, the intermediary earns $(v^+ + v)\,\tilde{q}\frac{1-m}{2}$.

Consider now a *thick* market. Similarly to the proof of Lemma 1 (see Appendix A.2), we can establish that, regardless of whether a dominant advertiser acquires all the impressions or not, the price of impressions is v, so the intermediary earns $vq\,(1+m)$.

C.5.2 Revenue comparison and choice of disclosure regime

Given (34) , one concludes that $v \leq v^+ \left(1 - \frac{2m}{1+m}\right)$ is necessary and sufficient for a dominant advertiser to acquire all the targeted impressions in its market with PH, i.e. for an equilibrium with multi-homing by advertisers to emerge. We assume such condition holds. Hence, in all markets with a dominant advertiser, we have $R_{IN}^{PH} = v\tilde{q}\,(1+m)$. In the main analysis, in either thin or thick markets without a dominant advertiser, we find that the revenues from F and PH are identical (see (2),(7) and (9)). With a dominant advertiser, however, $R^{PH} = v\tilde{q}\,(1+m)$ dominates the revenue with F in thin markets. Hence, under our assumptions, PH is strictly preferable to F. The total revenue earned by the intermediary with PH is

$$R_{IN}^{PH} = \tilde{q}\bar{v}\left(1 + m\left(1 - 2x\right)\right). \tag{35}$$

The revenue with PT is the same as in the baseline model, for all markets. Hence, if the intermediary adopts F on single-homers, and PT on multi-homers, its revenue if $R_{IN} = \tilde{q}\bar{v}$, as in (10). The comparison between this revenue and (35) reveals that IN will choose the PH if and only if $x \leq 1/2$, and otherwise adopt F on single-homers and PT on multi-homers, as in the baseline model. Overall, the intermediary earns the same revenue as in (10).

C.5.3 No intermediary

Consider thin market and evaluate advertiser j's bidding strategy for each targeted impression on publisher i. As in Section 4.2, we can characterize advertiser j's willingness-to-pay for each impression on publisher i as follows

$$w_{ia} = v^{+}\left(1 - q\frac{2mS_{i'a}}{1+m}\right), \quad i = 1, 2. \tag{36}$$

Similarly to the proof of Lemma 1 (see Appendix A.2), we can establish that the following configurations can potentially emerge in equilibrium

- A: $S_{2b} = 0$ and $S_{1b} = 0$. The willingness-to-pay for impressions are $w_{ia} = v^{+}\left(1 - \frac{2mq}{1+m}\right) > w_{ib} = v$. In this equilibrium, a outbids b for every targeted impression on every publisher.
- B: $S_{ib} = 0$ and $S_{i'b} = 1$. The willingness-to-pay for impressions are $w_{ia} = v^{+}\left(1 - \frac{2mq}{1+m}\right) < w_{ib} = v$. Hence, a and b single-home on different publishers. The same condition is necessary and sufficient for an equilibrium that entails the symmetric configuration, $S_{i'b} = 0$ and $S_{ib} = 1$.

All other configurations can be ruled out following a similar reasoning as above, for the case of PH.

Summing up, the subgame that takes place conditional on the publishers not outsourcing to IN admits two possible equilibrium configurations. If and only if $v^{+}\left(1 - \frac{2mq}{1+m}\right) \geq v$ holds, the equilibrium is such that a multi-homes and buys all the targeted impressions. Otherwise, the equilibrium is such that each advertiser single-homes on a different publisher. In the former equilibrium, the price of each targeted impression equals v, so each publisher earns $vq\frac{1+m}{2}$. In the latter equilibrium, the price on publisher 1 is $v^{+}\left(1 - \frac{2mq}{1+m}\right)$, whereas it equals $v\left(1 - \frac{2mq}{1+m}\right)$ on publisher 2. Hence, the publishers earn $v^{+}q\frac{1+m(1-2q)}{2}$ and $vq\frac{1+m(1-2q)}{2}$ respectively.

Consider now a thick market. Following a similar reasoning as above, one can establish that if and only if $v^{+}\left(1 - \frac{2mq}{1+m}\right) \geq v$ holds, the equilibrium is such that a multi-homes and buys all the targeted impressions. Otherwise, the equilibrium is such that each advertiser single-homes on a different publisher. In a thick market, in both cases the equilibrium price of the impressions is v, and each publisher earns a revenue equal to $vq\frac{1+m}{2}$.

Given the assumption that $v \leq v^{+}\left(1-\frac{2m}{1+m}\right)$ made in the previous subsection, the equilibrium must be such that the dominant advertiser multi-homes and buys all the impressions. Under this assumption, each publisher earns $R_i = vq\frac{1+m}{2}$, regardless of whether the market is thick or thin.

C.5.4 When do the publishers outsource?

We restrict our attention to the case where $v \leq v^{+}\left(1-\frac{2m}{1+m}\right)$, which is sufficient for the "high valuation" advertiser a to acquire all the impressions on each publisher when operating independently. Given the revenues (13) and (14) we calculated for markets without a dominant advertiser, the aggregate revenue earned by the publishers is

$$R_1 + R_2 = \bar{v}q\left(1 + m\left(1 - 2qx\right)\right),$$

which is identical to (15). Given the revenue earned by the intermediary is equal to (10), as established above, we conclude that Proposition 2 does not change.