



This is a repository copy of *Intermediaries in the online advertising market*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/195453/>

Version: Accepted Version

Article:

D'Annunzio, A. and Russo, A. orcid.org/0000-0002-8837-1265 (2023) Intermediaries in the online advertising market. *Marketing Science*. ISSN 0732-2399

<https://doi.org/10.1287/mksc.2023.1435>

© 2023 INFORMS. This is an author-produced version of a paper subsequently published in *Marketing Science*. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Intermediaries in the Online Advertising Market*

Anna D'Annunzio

TBS Business School

Antonio Russo

University of Sheffield and CESifo

January 16, 2023

Abstract

Most of the ads displayed by digital publishers are sold via intermediaries, which 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. When markets are thin, retaining information on consumers' type is superior to retaining information on exposure to ads. Even though 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 and 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

*We gratefully acknowledge the contribution by Christian Peukert to the first stages of this work. We also thank the Senior Editor, the Associate Editor, two referees, Doh-Shin Jeon, Bruno Jullien, Leonardo Madio, Mark Tremblay 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 online advertising market plays an increasingly central role in modern economies. In addition to being a sizable and rapidly growing market,¹ online advertising is a key source of revenue to many digital content providers and publishers, such as newspapers, blogs and review websites, with significant implications for the rest of society.² Furthermore, advertising affects the prices consumers pay for goods and services (Bagwell, 2007). For all these reasons, understanding how this market works is important.

Digital publishers typically provide “display” ads.³ A striking feature of the market for this particular kind of advertising is the major role played by intermediaries: for instance, intermediaries transacted more than 60 percent of display ad spend in the EU in 2017 (IAB, 2017). Furthermore, one firm (Google) has a dominant position in every link of the chain of intermediaries that connects advertisers to publishers (see Figure 1).⁴ In this paper, we study the transparency choices of a monopolist advertising intermediary when selling ad impressions, and consider the implications for publishers, advertisers and regulators. We also examine how transparency and privacy regulation affect the display advertising market.

Conceivably, a large intermediary presents several attractive features to digital publishers, particularly when consumers multi-home, i.e., visit multiple online publishers in a short time frame. The intermediary centralizes the sale of ads from multiple publishers and can typically achieve a more precise profiling of consumers. Moreover, the intermediary can coordinate consumers’ exposure to ads on multiple publishers. This feature makes it possible, for example, to limit wasteful cross-publisher ad repetition that, according to previous literature (e.g., Ambrus et al., 2016; Athey et al., 2018) and to several market operators, hinders the efficiency and effectiveness of ad campaigns.⁵

¹Global 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/>).

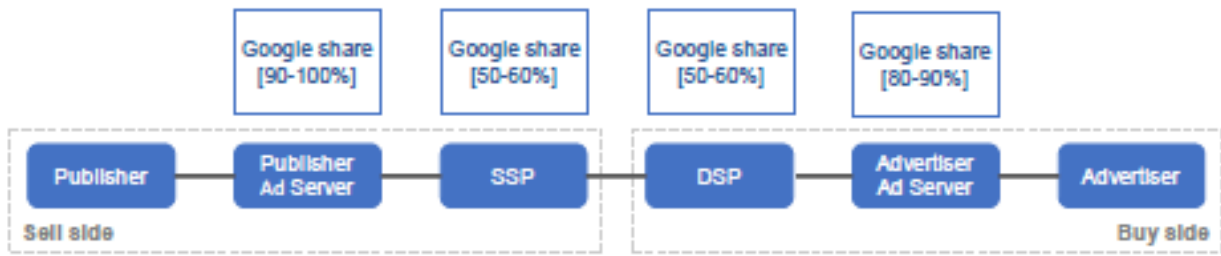
²For instance, advertising revenue drives investment in content quality in crucial domains such as journalism. On the relation between the online advertising market - with particular regard to the role of large platforms - and the viability of high quality journalism, see chapter 4 of the Cairncross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

³Display ads are 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, whereas the search advertising market was worth GBP 7.3 billion (CMA, 2020).

⁴The 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).

⁵See, e.g., <https://digiday.com/marketing/ad-techs-frequency-cap-problem/> and the concerns expressed by industry bodies (<https://www.thedrum.com/news/2019/03/05/the-industry-s-five-step-plan-improve-trust-advertising>), smaller intermediaries (<https://www.thedrum.com/news/2019/03/05/the-industry-s-five-step-plan-improve-trust-advertising>), smaller intermediaries (<https://www.thedrum.com/news/2019/03/05/the-industry-s-five-step-plan-improve-trust-advertising>).

Figure 1: Chain of intermediaries in online display advertising. Source: CMA (2020)



In practice, however, the functioning of intermediaries is quite complex and obscure. Regulators and market operators have pointed to the lack of transparency in, for instance, how intermediaries run advertising auctions and how they allocate the impressions. One of the main issues is that these platforms strategically retain valuable information from advertisers, making it difficult to assess the effectiveness and reach of their campaigns (CMA, 2020). On the other hand, there are concerns about consumer privacy when the intermediary tracks consumers online (Goldfarb and Tucker, 2011).

The observations above raise several interesting questions. What drives the incentives of the intermediary to retain information from advertisers when auctioning ad impressions? Which sort of information should the intermediary retain? How does transparency (or lack thereof) by the intermediary affect the advertising market? Do digital publishers gain by selling their advertising space via the intermediary in presence of multi-homing? How does the intermediary respond to transparency and privacy regulation?

To shed some light on the above questions, we consider a simple setting with two publishers and an intermediary. Consumers either visit one or both publishers, 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 intends to reach that type. Advertising markets differ in their thickness (i.e., the number of advertisers belonging to that market) and the returns to advertising. Impressions are sold via auctions, run by the intermediary if the publishers outsource their ad inventories. The intermediary can gather more accurate information about consumers than the publishers, which allows to sell a larger volume of targeted ads. Moreover, the intermediary observes a consumer's ad exposure on different publishers.

[//www.appnexus.com/sites/default/files/whitepapers/whitepaper-futureoftrading.pdf](http://www.appnexus.com/sites/default/files/whitepapers/whitepaper-futureoftrading.pdf), p.12) and ad agencies (<https://www.kantar.com/inspiration/advertising-media/the-digital-explosion-how-do-people-feel-about-online-ads>). Google's own campaign evaluation tools emphasize unique users and impression repetition (https://support.google.com/google-ads/answer/2472714?hl=en&ref_topic=3123050).

In the baseline model, we assume diminishing returns to advertising to the same consumer. Consequently, disclosing a consumer’s ad exposure entails a key trade-off for the intermediary: bids increase from advertisers that are not sending an impression to the consumer on another publisher, but decrease from the other advertisers. Hence, disclosure effectively “thins the market out”. Unless the market is thick (i.e., the number of advertisers is large enough), the price of impressions on multi-homers drops sharply. Therefore, if many consumers belong to thin markets, the intermediary is better off disclosing less information to the advertisers.

Motivated by the current debate, we explore two dimensions along which the intermediary can retain information from advertisers. One is to simply not disclose consumers’ ad exposure, preventing the advertisers from capping the frequency of their messages.⁶ Somewhat surprisingly, we find that retaining this information is not optimal. The intuition is that, with no control on frequency, the advertisers single-home on different publishers to avoid the risk of repetition. If the market is thin, each advertiser thus reaches all multi-homers on one publisher, and heavily discounts impressions on the other. Therefore, the price of impressions drops, as with full information disclosure. Alternatively, the intermediary can “thicken” the set of advertisers interested in a multi-homer by retaining information about her/his preferences, making targeting less granular. When markets are thin, less granular targeting is optimal because it increases the equilibrium price of impressions on multi-homers. Recent empirical studies have highlighted the incentive to retain information about consumer preferences by advertising intermediaries (Lu and Yang, 2020; Rafeian and Yoganarasimhan, 2021). Our result establishes a previously unexplored link between this incentive and the potential for wasteful repetition of ads.

We then investigate whether the publishers gain by outsourcing the sale of their ad inventories to the intermediary. This choice is profitable to the publishers if their audience mostly belongs to thick advertising markets. In such markets, the intermediary can sell a higher volume of targeted impressions (given its superior ability to profile consumers), but also at a higher price. However, if a substantial share of the audience belongs to thin advertising markets, outsourcing to the intermediary may reduce the equilibrium price of targeted impressions, particularly when many consumers multi-home.

We also analyze how the intermediary affects the size and distribution of surplus in the advertising market. In our model, the publishers outsource if and only if the intermediary

⁶Some operators have questioned the transparency of frequency capping tools provided by intermediaries like Google, for instance regarding the imminent phasing-out of third party cookies. See, e.g., <https://www.adexchanger.com/data-driven-thinking/why-we-cannot-miss-the-mark-on-ad-frequency-capping-in-the-post-cookie-era/>.

increases the revenue generated from their ad inventory. However, the advertisers may be worse off, because competition among the publishers weakens. Moreover, despite its superior ability to target consumers and allocate ads, the intermediary may reduce the total surplus in the advertising market if it reduces the granularity of targeting on multi-homing consumers.

Next, we turn to the effects of regulation on transparency. Imposing Full Disclosure to the intermediary (see, e.g., the proposed remedies in [CMA, 2020](#), p. 395) may either increase or reduce total 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, this policy may induce the publishers not to outsource, reducing the extent to which the industry benefits from the intermediary's superior tracking capabilities.

To evaluate the effects of privacy regulation, we extend the model allowing consumers to block tracking by the intermediary (e.g., rejecting third-party cookies).⁷ When choosing its degree of transparency towards advertisers, the intermediary must consider not only the effect on the revenue from ads on tracked consumers, but also how consumers react (e.g., because they find targeted ads to be intrusive). Thus, our analysis suggests, regulation such as the GDPR ([European Parliament, 2016](#)) affects not just the intermediary's ability to collect information about consumers, but also its incentives to share such information with the advertisers. The policy reduces the number of tracked consumers, but may either increase or decrease the incentives to disclose to advertisers the data of consumers who do not opt out of tracking.

Finally, we present several extensions to the baseline model. These include reserve prices in advertising auctions, increasing returns to advertising (retargeting), competition among advertisers in the product market and heterogeneous advertising returns within each advertising market. Our main results are fundamentally robust to these modifications.

The remainder of the paper is organized as follows. Section 2 discusses previous literature. Section 3 presents the model, that we solve in Section 4. We analyze the distribution of profits and evaluate 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 Appendices (Appendices B, C, and D are for online publication).

⁷Advertising 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>).

2 Literature review

A growing literature studies how multi-homing affects the advertising market ([Ambrus et al., 2016](#); [Athey et al., 2018](#); [Affeldt et al., 2021](#); [Gentzkow et al., 2021](#)). This literature highlights the inefficiencies that arise whenever consumers can be exposed to multiple ads on different outlets. Our model builds on these insights and contributes to this literature by introducing online advertising intermediaries that manage the allocation of ads on multi-homing consumers.

Our work connects the above literature on multi-homing to a recent literature that studies how the granularity of ad targeting affects the revenue of ad financed platforms. In line with an early intuition by [Levin and Milgrom \(2010\)](#), this literature emphasizes the market-thinning effect of targeting and discusses the incentives for platforms to conflate ad impressions. [Sayedi \(2018\)](#) studies how a publisher should allocate impressions among real-time bidding (RTB) and reservation contracts, considering a single publisher and two horizontally differentiated advertisers. With RTB, information about consumer preferences has a market-thinning effect because each advertiser bids only for consumers that strongly prefer its product. [Lu and Yang \(2020\)](#) use data from Alibaba’s ad network to show that narrow targeting based on consumer interests reduces ad revenue by weakening competition among advertisers. [Rafeian and Yoganarasimhan \(2021\)](#) consider an ad network delivering ads on mobile apps that can adopt behavioral targeting (i.e., based on consumer attributes and preferences) and/or contextual targeting (i.e., based on which app they are using and when). In the empirical analysis, the authors estimate counterfactual click-through rates for ads on multiple apps, and find that the ad network should restrict to contextual targeting to avoid market-thinning. However, their theoretical model does not allow to rank these two orthogonal targeting regimes in terms of platform revenue and total surplus, because it only establishes a relation among targeting regimes that differ by their granularity.

Differently from the above papers, we focus on consumers’ exposure to repeated ads depending on how many publishers they visit, which matters because the marginal returns from impressions on the same consumer are not constant. Therefore, we contribute to this literature in several ways. First, we explore an additional dimension of the market-thinning effect of disclosure: the repetition of ads across publishers. We provide a theory that establishes a previously unexplored link between the incentive to retain information on consumer preferences by advertising intermediaries (highlighted by [Rafeian and Yoganarasimhan, 2021](#), and [Lu and Yang, 2020](#)), and multi-homing. Consequently, we study forms of information disclosure by the

intermediary that are relevant in this context, concerning the granularity of information about consumer preferences and, differently from previous papers, the frequency of ad exposure. Differently from [Rafeian and Yoganarasimhan \(2021\)](#), our theoretical analysis allows to order different information disclosure regimes (that cannot be interpreted as being more or less granular) in terms of their effect on platform revenue, surplus of advertisers and publishers, as well as total surplus. Furthermore, we endogenize the outsourcing decision by the publishers, showing that, despite the intermediary’s superior tracking capabilities, the publishers may not benefit from outsourcing their ad inventories. This aspect is novel in itself, and it also allows to provide novel policy and managerial implications. Also, we study the effects of transparency and privacy regulation not only for a given market structure, but also considering the possible changes in market structure determined by policy.

Few other papers study intermediaries in online advertising. [Marotta et al. \(2021\)](#) consider how a platform that can share consumer information with advertisers affects competition on the product market, but ignore digital publishers. [Sharma et al. \(2019\)](#) study the contractual arrangements between digital publishers and two differentiated ad networks, and how the publishers sort across the two networks. [D’Annunzio and Russo \(2020\)](#) consider an ad network that centralizes the sale of ads in presence of multi-homing consumers and advertisers. The authors focus on how the ad network influences the advertising intensities on the publishers, and on the implications of consumers avoiding third-party tracking. [Peitz and Reisinger \(2020\)](#) show that, by centralizing the sale of ads on multi-homing consumers, an ad network can have a negative effect on the equilibrium price of impressions. Unlike these papers, we study the disclosure of information about consumers visiting different publishers. We also connect the transparency choice of the intermediary to the choice of digital publishers to outsource the sale of ads.

Our analysis also contributes to the literature on the effects of privacy policy in online advertising markets ([Goldfarb and Tucker, 2011, 2012](#); [Acquisti et al., 2016](#)). Recent work in this literature points to the externalities related to consumers disclosing information that enables platforms to profile other consumers ([Acemoglu et al., 2019](#); [Bergemann et al., 2019](#); [Choi et al., 2019](#); [D’Annunzio and Russo, 2020](#)). Empirical papers have studied the response by consumers, platforms and advertisers to regulation such as the GDPR ([Peukert et al., 2022](#); [Jia et al., 2021](#); [Johnson et al., 2020](#)). [Aridor et al. \(2020\)](#) show that, while the GDPR induced more consumers to opt out of third party tracking, advertisers’ bids for impressions on those who do not opt out increased, suggesting that platforms can profile such consumers more effectively. Rather than focusing on how privacy policy affects the ability of a platform

to gather consumer information, we study how the policy affects the incentive to disclose such information to the advertisers, which the literature has ignored so far.

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) share of multi-homers and $\frac{1-m}{2}$ the share of single-homers on each publisher. Each consumer is characterized by a type, θ , summarizing a set of characteristics, such as interests (culture, sports, etc.), demographics and geographic location, which determine her/his relevance to the advertisers. We let θ be distributed among consumers according to a uniform distribution with support $[0, 1]$, independently of the allocation of consumers on the publishers. Each publisher exposes a consumer to one impression. Therefore, single- and multi-homing consumers receive, respectively, one and two impressions in total.

Advertisers. Ads inform consumers about products. Let $k(\theta)$ be the set of advertisers that intends 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. Each consumer belongs to one and only one market. 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 non-targeted impressions is zero.

Each advertising market is characterized by two parameters. First, the number of advertisers, n , i.e., the market *thickness*. We refer to markets as “thin” if $n = 2$, “intermediate” if $n = 3$, and “thick” if $n \geq 4$.⁸ Let x , y and $1 - x - y$ be the share of thin, intermediate and thick markets, respectively. The thickness of a market may depend, for example, on how specific the preferences of the respective consumers are, or on the number of sellers operating in

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

a geographical area. To simplify the exposition, we set $y = 0$ in the baseline model, relegating the analysis with intermediate markets to Appendix C.

The second parameter characterizing advertising markets is the marginal return from informing a relevant consumer. We denote this return by v and assume it is distributed according to a distribution $G(v)$ on $[0, v_H]$, with smooth density $g(v)$, and mean $\bar{v} \equiv \int_0^{v_H} v dG(v)$. The advertising return may depend, e.g., on product margins and on the probability that informing a consumer converts into a sale. In the baseline model, v is homogeneous among advertisers within a market. 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.

Each advertiser maximizes its expected return from ads, net of the prices paid to the platforms. Impressing a consumer with one ad is enough to inform her/him. Sending the same ad twice to the consumer is thus wasteful. Therefore, as we shall see, 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, the marginal return from informing a consumer does not depend on the consumer being exposed to ads from other advertisers in the same market.

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.

When selling its impressions directly, each publisher generates a signal σ for each consumer (e.g., using first-party cookies), that conveys information about the consumer's type. 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. To streamline the exposition, we assume the publishers always reveal σ to the advertisers when selling an impression. 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 its ad inventory. If a publisher accepts, the intermediary sells the impressions in simultaneous first-price auctions. The intermediary generates a signal about each 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 the same probability, $\tilde{q} = q$. If both publishers outsource, instead, 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 as well as which ads she/he is exposed to during each visit.

We assume that, at the bidding stage, the intermediary informs the advertisers about which publisher delivers each impression. Furthermore, when selling each targeted impression, the intermediary decides whether to fully reveal information on (i) the consumer's type and (ii) her/his ad exposure, i.e. whether she/he multi-homes and which other ad she/he is exposed to on another publisher, if any. The first piece of information allows the advertisers to target only relevant consumers. The latter piece of information allows the advertisers to control the frequency of their messages across publishers. Specifically, the intermediary decides among the following disclosure regimes for each impression on a profiled consumer:

1. *Full Disclosure (F)*: IN discloses θ and the consumer's ad exposure.
2. *Partial Type Disclosure (PT)*: IN discloses the consumer's ad exposure and that σ^{IN} takes one among a finite set Θ of $t \geq 2$ values, including θ .
3. *Partial Exposure Disclosure (PE)*: IN discloses the consumer's θ , but not her/his ad exposure.

Under *PT*, the intermediary limits the granularity of information about consumer preferences and attributes, by conflating impressions on different consumer types. For example, instead of revealing the consumer's exact location, IN could only reveal her/his postcode. Under *PE*, IN prevents the advertisers from managing the frequency of ads on the same consumer.

Timing. The timing of moves is as follows:

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 impression opportunities are generated simultaneously.
 - (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 $r \in \{F, PT, PE\}$, for each impression. The platforms sell the impressions in simultaneous first-price auctions. The advertisers bid for each impression separately and simultaneously.
- 4. Impressions take place and all payoffs are realized.

Note that the choice of disclosure regime is conditional on the information available to the intermediary when selling each impression, i.e. θ (if observed) and which publisher(s) the consumer visits. Recall that n and v are not observed at that stage. We focus on pure strategy Subgame-Perfect Nash Equilibria.

3.2 Discussion of the setup

We briefly discuss some of our assumptions. Consistently with Google’s dominant position in the “open display” digital market, we assume the intermediary is a monopolist. We focus on first-price auctions because most digital intermediaries in the display market run auctions based on this format. For example, Google’s ad exchange switched to first-price auctions in 2019 (CMA, 2020). We ignore reserve prices in the baseline model, but include them in an extension (see Section 6.1).⁹

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.¹⁰ To economize on notation, we set to zero the marginal return from sending an ad to a consumer more than once, so each advertiser values only the first impression on a consumer. This assumption is not crucial: what matters for our results is that returns decrease with duplicated ads. In an extension (Section 6.4), we assume increasing returns to advertising on the same consumer. By normalizing the number of impressions on each publisher to one, we rule out repetition on the same outlet. This repetition is not a major concern in reality, since digital publishers have the means to manage the frequency of ads within their own domains, e.g. using first-party cookies. Cross-outlet repetition is a much more relevant challenge in the management of ad campaigns.

⁹Allowing for the impressions on multi-homing consumers to take place (and be auctioned) sequentially would not change our results (see Appendix D.6.1 for a proof of this claim).

¹⁰According to the advertising literature, online advertisers typically have a target range for the number of impressions per consumer (Yuan et al., 2013 find it to be between three and eight).

The baseline model assumes that the return from informing consumers is homogeneous across advertisers within the same market. We relax this assumption in Section 6.5. In another extension, we allow the return from informing a consumer to depend on whether the consumer sees messages from competing advertisers (Section 6.2).

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 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 fundamentally affected.

We model the contractual arrangements between the publishers and the intermediary in a stylized way, but this assumption is not crucial. 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 discuss in Section 4.3.

We also 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 $\tilde{q} > q$ regardless of how many publishers outsource. This assumption would make the analysis more involved without yielding qualitatively different results.

Assuming the choice of disclosure regime takes place when each impression is sold is consistent with the assumption that the intermediary uses all the available information when making this choice. Little would change if we let r be chosen at an earlier stage. Restricting the intermediary to a single regime for all impressions would not change our results significantly.

4 Solving the model

We begin by considering the subgame where the publishers outsource to the intermediary and characterizing its profit-maximizing disclosure regime. Next, we consider 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, each (first-price) auction for an impression is won by the advertiser with the highest willingness-to-pay, and the equilibrium price equals the second-highest willingness-to-pay. The outcome is therefore identical to that of a second-price auction, given that the Revenue Equivalence principle (Myerson, 1981) applies in our setting. Second, all non-targeted

impressions sell at zero, because no advertiser is willing to place a positive bid. Therefore, we focus on targeted impressions in the following.

The subgames we consider may 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 Intermediary

Let $w_{i,j}^r$ be advertiser j 's willingness-to-pay for an impression delivered on publisher i , under the disclosure regime $r \in \{F, PT, PE\}$.

4.1.1 Full Disclosure ($r = F$)

In the Full Disclosure regime, the advertisers in a given market bid for each targeted impression knowing θ . Only the advertisers in $k(\theta)$ are thus willing to place positive bids. Furthermore, the advertisers know the consumer's ad exposure and can control the frequency of their messages. If the consumer is a single-homer, the willingness-to-pay for a targeted impression is $w_{i,j}^F = v$, for any i , since the impression is not repeated. Hence, all advertisers in the market bid v , and the equilibrium price of any impression on a single-homer is v .

Suppose now the consumer is a multi-homer. An advertiser is willing to pay v for an impression if and only if it is not already acquiring the other impression available on the same consumer. Consequently, the equilibrium price of each targeted impression under Full Disclosure depends on the thickness of the advertising market. If the market is thick ($n \geq 4$), there are $n - 1$ advertisers willing to pay v for impressing a multi-homer on publisher j and one advertiser willing to pay 0 (because only one advertiser impresses the multi-homer on the other publisher). Hence, the equilibrium price of both impressions on a multi-homer in a thick market is v . By contrast, if the market is thin, only one advertiser is willing to pay v for each impression on a multi-homer, because the other one has already impressed the multi-homer on the other publisher. In a first-price auction, each advertiser makes a bid equal to the second-highest willingness-to-pay, implying that the price of the impressions on multi-homers drops to zero (see Table 1). Letting p denote the equilibrium price of an impression, we have

$$p_{n=2,MH}^F = 0, p_{SH}^F = p_{n \geq 4,MH}^F = v. \quad (1)$$

Full Disclosure maximizes the advertisers' willingness-to-pay for impressions that do not

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

fall on already exposed consumers, but, if the market is thin, also results in a sharp reduction in the price. The expected revenue (R) earned by the intermediary from single- and multi-homers from targeted impressions is, respectively

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

so the total expected revenue is

$$R^F = \bar{v}\tilde{q}(1 + m(1 - 2x)). \quad (3)$$

4.1.2 Partial Type Disclosure ($r = PT$)

In this regime, the advertisers bid for each targeted impression knowing the consumer’s ad exposure. Therefore, as with Full Disclosure, they can control the frequency of their ads on different publishers. However, information about consumer preferences 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 these conflated markets have the same return from informing consumers (we relax this assumption in Appendix A.1, showing that there is no significant change in the results).¹¹ Given 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.¹²

By conflating the consumer’s type, θ , with other $t - 1$ “wrong” types, the intermediary

¹¹A 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/his home. Suppose also that, instead of disclosing the consumer’s precise location, the intermediary only discloses her/his postcode. The Thai restaurants within this enlarged geographical area sell similar products and have therefore similar returns from informing the consumer.

¹²In our setting, there would be no gain to the intermediary in disclosing to the advertisers that the consumer belongs to one of the conflated markets with higher probability than to another market (see Appendix B for the proof of this claim).

ensures that there are at least two advertisers with positive willingness-to-pay for *any* targeted impression. Hence, even if a multi-homer belongs to a thin market, the equilibrium price of a targeted impression never drops to zero, as it does under Full Disclosure. The drawback is that the advertisers are not willing to pay as much as they would for a non-repeated impression targeted more granularly. At equilibrium, the intermediary sets $t = 2$ because conflating more than two markets would only reduce its revenue. The equilibrium price, regardless of which publishers the consumer visits and of the thickness of the market, is equal to the second highest willingness-to-pay of advertisers, that is,

$$p^{PT} = \frac{v}{2}. \quad (4)$$

Intuitively, the intermediary dilutes the quality of the information about consumer attributes just enough to “thicken” the set of advertisers that bids for impressions on multi-homers in thin markets. Hence, the expected revenue on single- and multi-homers is, respectively

$$R_{SH}^{PT} = \frac{\bar{v}\tilde{q}}{2} (1 - m), \quad R_{MH}^{PT} = \bar{v}\tilde{q}m. \quad (5)$$

4.1.3 Full vs Partial Type Disclosure

Comparing (2) and (5), we see that the intermediary may want to conflate the impressions on multi-homing consumers. The PT regime avoids the market-thinning effect that occurs for multi-homers under F in thin markets, generating more revenue on such consumers. We can therefore state the following

Lemma 1. *Full Disclosure results in higher advertising revenue on single-homers than Partial Type Disclosure. However, the latter results in higher revenue on multi-homers if and only if $x > \frac{1}{2}$.*

We shall now analyze the Partial Exposure Disclosure regime. As we will show, this regime does not generate more revenue than the other two.

4.1.4 Partial Exposure Disclosure ($r = PE$)

In this regime, the advertisers bid for each targeted impression knowing the consumer’s θ , but not whether the consumer receives another impression on a different publisher. The advertisers thus cannot control the frequency of exposure to ads across publishers. The willingness-to-pay, $w_{i,j}^{PE}$, 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 $\frac{m}{1+m}$, whereas the probability it falls on a single-homer is $\frac{1-m}{1+m}$. Second, the probability that the targeted impression the consumer receives on the other publisher is from the same advertiser, j . This probability is $S_{i'j}$, i.e. the share of targeted impressions advertiser j acquires on i' . We have

$$w_{i,j}^{PE} = v \left(\frac{\frac{1-m}{2} + m(1 - S_{i'j})}{\frac{1+m}{2}} \right) = v \left(1 - \frac{2mS_{i,j}}{1+m} \right), \quad i, i' = 1, 2; i' \neq i. \quad (6)$$

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

The above 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 impression on publisher i if and only if b acquires a larger share of impressions on publisher i' than a . Consequently, there can only be two sets of bidding strategies in equilibrium. The first equilibrium candidate 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 set of equilibrium candidate bidding strategies is such that all advertisers in a market place equal bids (given by (6)) 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 $w_{i,j}^{PE}$ 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 2. *Under Partial Exposure Disclosure, all advertisers in each market single-home on a different publisher.*

Therefore, the advertisers avoid cross-outlet repetition by placing ads on a single publisher.

Table 2: Advertiser’s willingness-to-pay, equilibrium prices and revenue in thin and thick markets under PE .

	Thin ($n = 2$)		Thick ($n = 4$)	
	Publisher 1	Publisher 2	Publisher 1	Publisher 2
w_{ia}	v	$v \frac{1-m}{1+m}$	w_{ia}	v
w_{ib}	$v \frac{1-m}{1+m}$	v	w_{ib}	v
			w_{ic}	$\frac{v}{1+m}$
			w_{id}	$\frac{v}{1+m}$
$p_{n=2}^{PE}$	$v \frac{1-m}{1+m}$	$v \frac{1-m}{1+m}$	$p_{n=4}^{PE}$	v
$R_{n=2}^{PE}$	$v\tilde{q}(1-m)$		$R_{n=4}^{PE}$	$v\tilde{q}(1+m)$

This finding is in line with [Athey et al. \(2018\)](#). We are now in a position to characterize the equilibrium prices, the allocation of impressions and the revenue earned by the intermediary with PE .

Thin markets ($n = 2$). In a thin market, given Lemma 2 and expression (6), advertiser a (resp. b)’s willingness-to-pay is v for each targeted impression on publisher 1 (resp. 2).¹³ Furthermore, advertiser a (resp. b)’s willingness-to-pay is $v(1 - \frac{2m}{1+m}) = v \frac{1-m}{1+m}$ for each 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 impression on publisher 2 is worthless if it hits a multi-homer. The equilibrium price of impressions is therefore equal to the second highest willingness-to-pay, that is

$$p_{n=2}^{PE} = v \frac{1-m}{1+m}. \quad (7)$$

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. Hence, when determining how much to bid for the impressions on the other publisher, the advertiser heavily discounts those falling on multi-homers. The equilibrium price is thus equal to the incremental value of ads on the other publisher. Therefore, the intermediary’s revenue equals $v\tilde{q}(1-m)$. We summarize the equilibrium bids and revenue in a given thin market in Table 2.

Thick markets ($n \geq 4$). Suppose now that $n = 4$. Given Lemma 2, we focus without loss of generality on the equilibrium such that two advertisers (say, a and b) single-home on

¹³To ease exposition, we present only one of the symmetric equilibria in each market, since the equilibrium prices and profits for all parties are identical.

publisher 1 while the other two (say, c and d) single-home on 2.¹⁴ 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 (6) and $S_{ij} = 1/2, \forall i, j$, the willingness-to-pay by a and b (resp. c and d) for each impression on publisher 2 (resp. 1) equals $v(1 - \frac{m}{1+m}) = \frac{v}{1+m}$. The equilibrium price of impressions is therefore

$$p_{n \geq 4}^{PE} = v, \quad (8)$$

and each advertiser receives half the impressions supplied by the respective publisher. In a thick market, therefore, the intermediary can extract the full value of the targeted impressions under PE . Therefore, the intermediary's revenue equals $v\tilde{q}(1+m)$. We summarize the bids and revenue for a given thick market in Table 2. Note that 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 on each publisher.

Table 2 states the intermediary's revenue in given thin and thick markets, from which we obtain the following total expected revenue under PE :

$$R^{PE} = \bar{v}\tilde{q}[x(1-m) + (1-x)(1+m)] = \bar{v}\tilde{q}(1+m(1-2x)). \quad (9)$$

4.1.5 Revenue-maximizing disclosure regime

Comparing (3) and (9), we see that F and PE yield the same expected revenue to the intermediary. Under both regimes, the full value of each impression can be extracted if the market is thick, whereas only the impressions on single-homers generate value if the market is thin. Hence, there is no gain in not disclosing consumer ad exposure and impeding frequency capping. Accordingly, we shall assume that, when indifferent among full and partial disclosure, the intermediary adopts the former, more transparent regime.¹⁵ Given this finding and Lemma 1, we can conclude the following:

Proposition 1. *The intermediary sells all targeted impressions on single-homing consumers under Full Disclosure. Furthermore, it sells the impressions on multi-homing consumers under*

¹⁴Equilibria 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.

¹⁵The indifference among regimes F and PE breaks in favor of F if we include intermediate markets (i.e., markets with $n = 3$ advertisers) in the model, as we show in Appendix C.

Partial Type Disclosure 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 advertising markets and the intermediary’s disclosure of information to the advertisers. By tracking consumers, the intermediary can avoid wastefully repeating impressions on multi-homers. However, the ensuing market-thinning effect reduces the price of these impressions in thin markets, creating an incentive to reduce the quality of the information disclosed to the advertisers.¹⁶ Interestingly, although information about consumer ad exposure has a market-thinning effect, we find that the intermediary should not retain this information from the advertisers (i.e., make cross-outlet frequency capping less effective). Rather, the intermediary should reduce the granularity of targeting of consumer preferences. This result is consistent with the findings of recent empirical literature showing that reducing the granularity of targeting can increase the revenue of ad financed platforms (Lu and Yang, 2020; Rafeian and Yoganarasimhan, 2021). Repetition across outlets when consumers multi-home, our result suggests, could provide an additional incentive to make targeting less granular.

Given Proposition (1) and equations (2) and (5), the intermediary’s total expected revenue in equilibrium is:

Lemma 3. *When both publishers outsource, the intermediary’s total revenue is*

$$R_{IN} = \begin{cases} \bar{v}\tilde{q}(1 + m(1 - 2x)), & \text{if } x \leq \frac{1}{2}, \\ \bar{v}\tilde{q}((1 - m) + m) = \bar{v}\tilde{q}, & \text{if } x > \frac{1}{2}. \end{cases} \quad (10)$$

4.2 No intermediary

We now consider the subgame where no publisher outsources to IN . Note that this is the same subgame as if only one publisher outsourced, 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.

4.2.1 Advertisers’ willingness-to-pay for impressions

In this scenario, the information available to the advertisers when bidding on a targeted impression is essentially the same as in the case where the intermediary adopts PE : advertisers

¹⁶This 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 (Ganuzo, 2004; Ganuzo and Penalva, 2010; Bourreau et al., 2017).

know the consumer's type, θ , and on which publisher the impression takes place, but not the consumer's ad exposure on the other publisher.

Consider an advertiser $j \in k(\theta)$ and let w_{ij} be its willingness-to-pay for a targeted impression delivered by publisher i . Each targeted impression on i is worth v to advertiser j if the consumer is a single-homer (probability $\frac{1-m}{1+m}$). Instead, if the consumer is a multi-homer (probability $\frac{2m}{1+m}$), she/he 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 j . Hence, we have

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

This expression is similar to (6), the willingness-to-pay under the *PE* regime. Just like in that scenario, w_{ij} decreases in the share of targeted impressions acquired by advertiser 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 *PE* than when the publishers do not outsource. In the latter case, a multi-homer can receive the same ad twice only if she/he is identified by *both* publishers independently. By contrast, when the intermediary profiles a multi-homer, both impressions are targeted. 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 4.1.4 to claim the following:

Lemma 4. *If neither or only one of the publishers outsources to the intermediary, all advertisers in each market single-home on different publishers.*

We can thus follow the same steps as in Section 4.1.4 to characterize the equilibrium prices, the allocation of impressions and the profits earned by the publishers when neither outsources to the intermediary. Specifically, in a thin market, the equilibrium price of a targeted impression on each publisher is

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

Each publisher therefore earns the following expected revenue

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

If the market is thick, the equilibrium price of impressions on both publishers is v , and each publisher earns

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

Based on the results above, we can compute the total expected revenue of each publisher:

Lemma 5. *If neither publisher joins the intermediary, each publisher collects the following total revenue*

$$\begin{aligned} R_i &= x \left(\frac{\bar{v}q}{2} (1 + m(1 - 2q)) \right) + (1 - x) \left(\frac{\bar{v}q}{2} (1 + m) \right) = \\ &= \frac{\bar{v}q}{2} (1 + m(1 - 2xq)), \quad i = 1, 2. \end{aligned} \quad (15)$$

4.3 The publishers' decision to join the intermediary

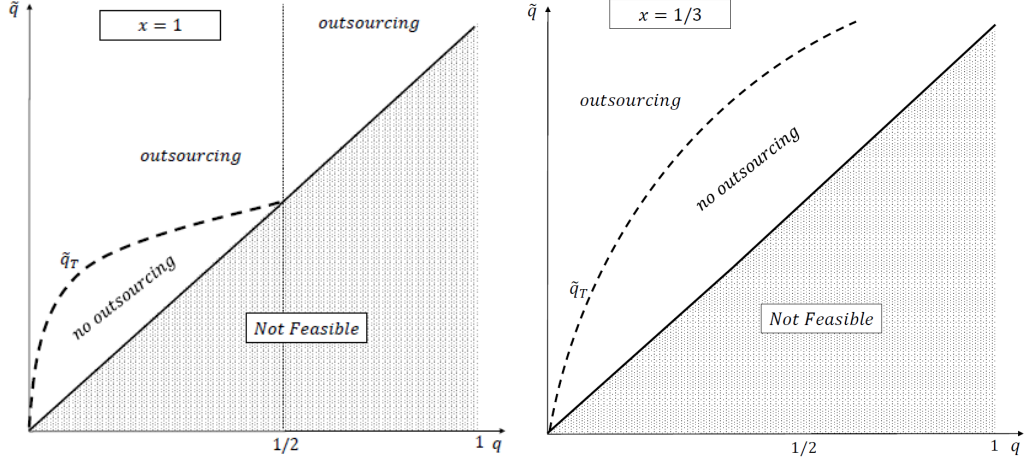
We now consider the first stage of the game and investigate the following question: is it profitable for the publishers to outsource the sale of their ad inventories to the intermediary? At the beginning of the game, the intermediary offers a transfer, T_i , to each publisher i . The publisher outsources if and only if this transfer is 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.¹⁷ The intermediary chooses T_i to maximize its profit $\pi_{IN} = R_{IN} - \sum_{i=1,2} T_i$, subject to $T_i \geq R_i, i = 1, 2$. Letting T_i^* denote the solution to this problem, we have $T_i^* = R_i, i = 1, 2$. Therefore, the intermediary can profitably induce each publisher to outsource if and only if

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

Using Lemmas 3 and 5 (that is, equations (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

¹⁷This observation also implies that we can rule out the situation where only one publisher outsources on the equilibrium path, since the intermediary cannot earn enough revenue to compensate the publisher while making a profit in that subgame.

Figure 2: Threshold \tilde{q}_T , variation with respect to q (for $x = 1$ on the left and $x = \frac{1}{3}$ on the right). Figures obtained setting $m = 1/2$.



as follows


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

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. $x \leq \frac{1}{2}$. 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 $\tilde{q} \geq \tilde{q}_T$. 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. $x > \frac{1}{2}$. The publishers outsource if their ability to identify consumers is large enough, i.e. $q > \frac{1}{2x}$. Otherwise, there exists a threshold $\tilde{q}_T \equiv q(1 + m(1 - 2qx))$, 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 .*

It is useful to start from some polar cases to get the intuition. First, suppose there are only thick markets ($x = 0$). In this case, the intermediary's disclosure regime is F and we have $\tilde{q}_T = q$. 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 advantageous to the publishers. Suppose now the share of thin markets, x increases marginally (so that IN keeps choosing F). In each thin market the intermediary extracts zero revenues from multi-homers under F , while the publishers get positive revenues: the average price of each target impression for independent publishers increases with m . Hence, the higher x , the lower is the incentive to outsource.

Figure 3: Threshold \tilde{q}_T , variation with respect to m (for $x = 1$ on the left and $x = \frac{1}{3}$ on the right). Figures obtained setting $q = 1/3$.



Focus now on the opposite polar case, $x = 1$. The intermediary chooses $r = PT$ for multi-homers and \tilde{q}_T boils down to $q(1 + m(1 - 2q))$, so $\tilde{q} < \tilde{q}_T$ can hold only if $q < 1/2$. Although the intermediary sells a larger volume of targeted impressions than the publishers (given $\tilde{q} > q$), it may have to sell them at a lower price. The difference in price can be seen by comparing (12) to (4). When markets are thin, the intermediary conflates impressions on multi-homers from different markets, so the advertisers only pay half of their value. However, when selling independently, the price the publishers are able to extract for each targeted impression decreases with q , and this effect is more pronounced the larger is m . When q is low enough ($q < 1/2$), the probability of repetition across publishers is small enough that the expected value of a targeted impression to the advertisers is higher than it would be with the intermediary. 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. Consider now a marginal decrease in x (so that IN keeps choosing PT for multi-homers). Outsourcing becomes less profitable, because each impression in thick markets is sold at $\frac{v}{2}$ under PT and at v by independent publishers (see (4) and (12)). Hence, as x decreases, the incentive to outsource weakens.

Finally, consider 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 all impressions. However, in thin markets, the publishers sell each targeted impression on multi-homers 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 is large enough, the intermediary adopts *PT* for multi-homers. Again, unless $q > \frac{1}{2x}$, the publishers sell each impression at a higher average price, implying that \tilde{q}_T increases with m . The upshot is that consumer multi-homing does not imply that digital publishers necessarily benefit from relying on the intermediary, despite its superior tracking capabilities.

5 Welfare analysis and the effects of regulating the intermediary

We now evaluate how the intermediary affects the total surplus in the advertising market and its distribution. 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 total surplus and its distribution?

In our model, the publishers outsource to the intermediary if and only if the total profit on the supply side of the advertising market increases. However, the advertisers may not benefit from outsourcing. 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, advertisers get zero surplus in thick markets. Hence, advertiser surplus (*AS*) without outsourcing is (given $(1+m)q$ targeted impressions)

$$AS^{noIN} = 2\bar{v}xm q^2. \quad (18)$$

If the publishers outsource, the intermediary extracts the whole surplus from impressions on single-homers, while the surplus left to advertisers from impressions on multi-homers depends on the disclosure regime. Specifically, under Full Disclosure, the advertisers retain all the surplus from each targeted impression on multi-homers in thin markets, since the price of impressions 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 (4)). The advertisers also get zero net surplus in thick markets. Hence, advertiser surplus with the

intermediary is

$$AS^{IN} = \begin{cases} 2xm\bar{v}\tilde{q} & \text{if } x \leq \frac{1}{2}, \\ 0 & \text{if } x > \frac{1}{2}. \end{cases} \quad (19)$$

Consequently, advertiser surplus increases with outsourcing if and only if $x \leq \frac{1}{2}$. Despite selling a larger volume of targeted impressions ($q < \tilde{q}$), the intermediary does not necessarily make the advertisers better off, in particular if advertising markets are thin.

Finally, total surplus in the advertising market, defined as the sum of advertisers surplus and platforms' profits, may either increase or decrease when the publishers outsource to the intermediary. If the publishers operate independently, each targeted impression produces a gross surplus equal to v , since the impression is 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, but the volume of targeted impressions is higher with outsourcing, given $\tilde{q} > q$, implying that total surplus is higher with outsourcing and Full Disclosure. However, if the intermediary adopts Partial Type Disclosure, some impressions are wasted because they are not targeted in a granular way. 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 intermediaries to implement greater transparency. See, for example, the Digital Market Act (European Commission, 2022) and the remedies discussed by the UK Competition and Markets Authority (CMA, 2020, pag. 395).¹⁸ We study some 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 otherwise. Hence, we focus on the effect of regulation when $x > 1/2$.

Transparency reduces the profits earned by the intermediary to the benefit of advertisers. Indeed, if the publishers outsource to the intermediary, imposing Full Disclosure makes the price of impressions on multi-homers drop to zero in thin markets. Hence, the region of

¹⁸In 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 data necessary for advertisers and publishers to carry out their own independent verification of the advertisements inventory, including aggregated and non-aggregated data.”

parameters such that the publishers outsource shrinks. Also, the regulation makes the advertisers better off if it deters the publishers from outsourcing, because the advertisers earn a positive surplus (see (18)), whereas they would earn nothing with the intermediary (see (19)).

Consider now the effects of transparency regulation on total surplus. 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 total surplus can be either positive or negative, because there are fewer targeted ads without the intermediary, but more are targeted in a precise way (the conditions are the same as described in Proposition 3). Hence, imposing transparency does not necessarily increase efficiency in the online advertising market. 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. Total surplus increases if the publishers outsource to 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

We extend the model allowing privacy sensitive consumers to block third-party tracking, e.g., by installing browser extensions that block third-party cookies or anonymize online activity.¹⁹ This extension provides us with a setting to explore the implications of privacy regulation, such as the European Union’s GDPR (European Parliament, 2016), and similar laws adopted in some US States (e.g., California), Chile, Japan, Brazil, and South Korea, for the online advertising market. One of the main provisions of these laws is to increase consumer control over personal data and avoid tracking by third parties.²⁰ We are interested in understanding how these provisions 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-publisher tracking. If a consumer blocks, the intermediary cannot observe

¹⁹In 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 anti-tracking tools only block third-party tracking.

²⁰The GDPR mandates opt-in policies for collecting consumers’ consent regarding third-party cookies, as opposed to opt-out ones.

her/his ad exposure, and can profile her/him only with probability q . If a consumer does not block, the intermediary collects the same information as in the baseline model. To simplify the analysis, we assume that all consumers multi-home ($m = 1$) in this extension. Also, 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, we assume that sufficient conditions for the publishers to outsource hold.²¹

Let c be the idiosyncratic cost (or effort) of blocking tracking, distributed uniformly among consumers on the $[0, \bar{c}]$ interval. Let b_r , with $r = \{F, PT\}$, be the private benefit from opting out.²² This benefit captures consumers' enhanced privacy, but also the perceived relevance and/or intrusiveness of ads. For instance, under Full Disclosure consumers are more likely to see more precisely targeted ads than under PT . Consumers may appreciate such ads or perceive them as exceedingly intrusive (Goldfarb and Tucker, 2012; Turow et al., 2009). Accordingly, we let b_F be bigger or smaller than b_{PT} .

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/his net cost of blocking, $c - t$, is smaller than b_r . Hence, the fraction of consumers blocking is $\frac{b_r+t}{\bar{c}}$ (assumed to be smaller than one for consistency).

Consider now IN 's choice between F and PT . The latter regime results in a higher expected revenue for the intermediary if and only if

$$\left(1 - \frac{b_{PT}+t}{\bar{c}}\right) \tilde{q} + \left(\frac{b_{PT}+t}{\bar{c}}\right) 2q(1-xq) \geq \left(1 - \frac{b_F+t}{\bar{c}}\right) 2\tilde{q}(1-x) + \left(\frac{b_F+t}{\bar{c}}\right) 2q(1-xq). \quad (20)$$

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

$$x > x_T \equiv \frac{1}{2} - \frac{b_F - b_{PT}}{2m\tilde{q}(\bar{c} - b_F - t)} (\tilde{q} - 2q(1-xq)). \quad (21)$$

Given the last term in brackets is positive by assumption, we get

$$x_T < \frac{1}{2} \iff b_F > b_{PT}.$$

The intermediary adopts a two-sided logic when choosing its disclosure regime, considering

²¹As established in Section 4.3, outsourcing occurs if the revenue earned by the intermediary under each regime is at least equal to the revenue earned by the publishers independently. Formally, we assume that $\max(\tilde{q}, 2\tilde{q}(1-x)) > 2q(1-xq)$.

²²We assume b_r is the same for all consumers for simplicity.

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., $b_F > b_{PT}$ 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 privacy policy. Expression (21) implies that

$$\frac{\partial x_T}{\partial t} < 0 \iff b_F > b_{PT}.$$

Hence, a tighter privacy policy results in a lower probability that IN chooses F whenever $b_F > b_{PT}$, and viceversa. 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 $b_F > b_P$ (which implies that $x = x_T < \frac{1}{2}$), the loss is greater under F than under P . Therefore, the set of values of x such that the intermediary adopts F shrinks. Hence, if and only if $b_F > b_P$, a tighter privacy policy reduces the likelihood that the intermediary adopts granular targeting on those who do not block.

Proposition 5. *Suppose fewer (resp. more) consumers block third-party tracking with Full Disclosure than with Partial Disclosure. Privacy policy that facilitates blocking makes the intermediary more (resp. less) likely to adopt Full Disclosure.*

This result shows that privacy policy does not just reduce the intermediary's ability to track consumers, but also changes its incentives to share the information collected with the advertisers. Thus, there are two possibly unintended consequences of the policy: on the one hand, if consumers are less likely to allow tracking when ads are targeted, privacy policy induces the intermediary to be less transparent towards advertisers. On the other hand, if consumers are more likely to allow tracking with targeted ads, the policy may induce the intermediary to disclose their data to advertisers more often.

6 Extensions and robustness checks

6.1 Reserve prices in advertising auctions

In this extension, we allow the intermediary to introduce reserve prices in auctions.²³ We concentrate on regimes F and PT that emerge at equilibrium in the baseline model (Proposition 1). In our setting, reserve prices are redundant in auctions under PT , because the equilibrium price of the impressions equals the willingness-to-pay of advertisers. Under Full Disclosure, reserve prices are redundant for impressions on single-homers and on multi-homers in thick markets, because the intermediary can extract the whole surplus from advertisers even without such prices. Instead, reserve prices can limit the price drop for impressions on multi-homers in thin markets. Note that, however, the reserve price cannot be conditioned on v , since the realization of this parameter in each market is unobservable to the intermediary. Hence, the drawback of setting a higher reserve price is that the share of markets where advertisers drop out from the auction because their return from informing consumers, v , is too low gets also higher. Due to this trade-off, we show in Appendix D.1 that, for impressions on multi-homers in thin markets, the F regime with a reserve price does not necessarily outperform the PT regime. For instance, Full Disclosure results in weakly lower total revenue given some notable distributions of v (e.g., uniform and continuous Bernoulli distributions). We conclude that the possibility to adopt reserve prices does not necessarily induce the intermediary to be more transparent when auctioning ad impressions.²⁴

6.2 Advertiser competition on the product market

We assumed that the advertisers' marginal return from informing a relevant consumer does not depend on whether the consumer is exposed to ads from competitors. However, if competing advertisers provide substitute products, ads may intensify competition on the product market (see, e.g., de Cornière and de Nijs, 2016).

To account for the effects of advertising on price competition, we modify the model

²³Reserve 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>).

²⁴Relaxing the assumption that impressions on multi-homers take place simultaneously, we could let the intermediary implement a reserve price in the second auction based on the price at which the impression sells in the first one. As we argue in Appendix D.6.2, such reserve price would not allow to capture all advertiser surplus in thin markets under Full Disclosure, and would bring to qualitatively similar results as in the baseline model.

assuming that, if a consumer is exposed to two ads from different advertisers in the same market, each advertiser gets a return equal to αv , with $\alpha \in [0, 1]$. The analysis is in Appendix D.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 informed by a rival. On the other hand, there is also an incentive for advertisers who inform a consumer to prevent competitors from informing the same consumer.

If the effect of competition on advertising returns is relatively small, i.e. $\alpha \geq \frac{1}{2}$, we find no qualitative change in the results. However, if $\alpha < \frac{1}{2}$, some of our results change. In this case, when the publishers operate independently and when the intermediary adopts Full Disclosure, we find that in each market one advertiser acquires all the targeted impressions on multi-homers. The reason is that 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 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 competitor, 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 their ad inventories. In Appendix D.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.²⁵ Let this share be $\rho \in (0, 1)$. When choosing the disclosure regime for each impression, the intermediary would have the same incentives as in our baseline model, since it would retain a fraction $1 - 2\rho$ of the revenue and because the revenue sharing agreement does not affect the willingness-to-pay by advertisers. Therefore, Proposition 1 does not change. Similarly, Proposition 2 holds: the intermediary induces the publishers to outsource (by setting ρ sufficiently large) if and only if

²⁵Google’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>).

it earns at least as much total revenue as the publishers can earn independently, i.e. condition (17) holds.

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 D.4 we assume each advertiser gets a higher return from the second impression on a consumer than from the first one. Under this assumption, we find the intermediary should adopt Full Disclosure 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 publishers benefit from outsourcing to the intermediary.

6.5 Heterogeneous advertising returns within markets

In Appendix D.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.²⁶ More precisely, we assume there is a share of markets where one of the advertisers has a larger return from informing consumers.

First we consider the case where the dominant advertiser a has a return v^+ large enough compared to v . In this case, when the publishers do not outsource their ad inventories, advertiser a outbids the others for each targeted impression on *both* publishers. This is in contrast to what occurs with homogeneous returns, where all advertisers in a market single-home (see Lemma 4). 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 effect when the intermediary adopts *PE*. There is instead no increase

²⁶Although one could consider a more general distribution of values of v in thick markets, this would add little to the analysis since most of the interesting trade-offs faced by the intermediary and the publishers involve thin markets.

in revenue with F and PT , because the second-highest willingness-to-pay remains the same. Therefore, in this setting the intermediary prefers PE to F . Furthermore, the intermediary chooses PT for multi-homers and F for single-homers, if and only if the share x of thin markets with homogeneous advertisers is high enough, i.e. $x \geq 1/2$. The results of Proposition 1 do not change significantly. We also find no qualitative change in Proposition 2.

Consider now the case where the dominant advertiser a has a return from informing consumers close to v . In this case, when the publishers do not outsource, advertisers single-home, as in the baseline model. Again, the presence of an advertiser with higher returns has a positive effect on ad prices and may make PE preferable to F to the intermediary when markets tend to be thick. However, Propositions 1 and 2 do not change in a qualitative sense.

Overall, this analysis suggests that in a setting with heterogeneous advertising returns within markets the intermediary may be even more likely to retain information from the advertisers. However, we do not find significantly different results compared to our baseline model.

7 Concluding remarks

We have studied the incentives of a monopolist intermediary to disclose consumer information to advertisers when auctioning ad impressions, focusing on information that allows for targeting of consumer preferences and for managing the frequency of exposure to ads. We conclude by providing a summary of the policy and managerial implications of our results.

Policy implications. The paper contributes to the debate on regulating transparency in the digital advertising market. Transparency requirements to advertising gatekeepers are discussed, for instance, in the Digital Market Act (European Commission, 2022). We have found that the intermediary prefers not to disclose all the available information on consumers when auctioning impressions in thin advertising markets (Proposition 1). Imposing Full Disclosure redistributes surplus from the supply side (publishers and intermediary) to the demand side (advertisers) of the market. However, total surplus from advertising may decrease, because some publishers may find it unprofitable to outsource to the intermediary if the latter cannot maximize the revenue from ads (see Proposition 4).

Privacy regulations, such as the GDPR (European Parliament, 2016) and the California Consumer Privacy Act of 2020, aim to ensure that consumers have more control over their own personal information. We found that these policies may affect the advertising market

not only by changing the intermediary's ability to collect consumer information, but also its willingness to share such information with advertisers. Quite interestingly, if consumers are less likely to allow tracking when ads are highly targeted (e.g., because they find such ads intrusive), the intermediary is less likely to disclose their data to the advertisers. On the other hand, if consumers like personalized ads and these ads make them less likely to block tracking, the regulation may induce the intermediary to disclose consumer data to advertisers more often (Proposition 5). For an exhaustive analysis of privacy policy, however, a more detailed model of consumers' behavior is necessary. Such an analysis is beyond the scope of our investigation.

Managerial implications. Our results suggest that the pervasiveness of consumer multi-homing 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. If there are diminishing returns to advertising and advertisers want to avoid excessive repetition on the same consumer (frequency capping), the thickness of the advertising market and the extent of multi-homing should be key parameters 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 publishers 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 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 weakened when the publishers outsource to the same intermediary. The net effect of these forces on advertisers may be negative or positive. Nevertheless, the advertisers benefit from regulation that encourages greater transparency. Interestingly, the advertisers may also benefit from privacy regulation, if it makes the intermediary more likely to disclose consumer information when selling the impressions (Proposition 5).

Our model also provides some insights concerning the intermediary's optimal disclosure of consumer data to advertisers (Proposition 1). We have found that, if advertising markets are

thin, reducing the granularity of information regarding consumer preferences can increase the equilibrium price of impressions, whereas retaining information that allows frequency capping does not. By contrast, the disclosure of both types of information is advantageous when advertising markets tend to be thick.

There are many other aspects in the behavior of intermediaries that seem worthy of consideration. For example, recent reports by antitrust authorities (e.g., [CMA, 2020](#)) and the case led by the Texas AG ([US District Court of NY, 2021](#)) raise multiple concerns about alleged anti-competitive practices by Google, including the possible manipulation of advertising auctions and charging of hidden fees to advertisers. We plan to tackle these intriguing issues in future research.

References

- Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2019). Too Much Data: Prices and Inefficiencies in Data Markets. NBER Working Papers 26296, National Bureau of Economic Research, Inc. [2](#)
- Acquisti, A., Taylor, C., and Wagman, L. (2016). The Economics of Privacy. *Journal of Economic Literature*, 54(2):442–492. [2](#)
- Affeldt, P., Argentesi, E., and Filistrucchi, L. (2021). Estimating demand with multi-homing in two-sided markets. Workingpaper, CentER, Center for Economic Research. CentER Discussion Paper Nr. 2021-025. [2](#)
- 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. [1](#), [2](#), [3.2](#), [4.1.4](#)
- Anderson, S. P. and Coate, S. (2005). Market provision of broadcasting: A welfare analysis. *Review of Economic Studies*, 72(4):947–972. [3.2](#)
- Aridor, G., Che, Y.-K., and Salz, T. (2020). The effect of privacy regulation on the data industry: Empirical evidence from gdpr. Working Paper 26900, National Bureau of Economic Research. [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](#), [2](#), [3.2](#), [4.1.4](#), [4.1.4](#)
- Bagwell, K. (2007). The economic analysis of advertising. volume 3 of *Handbook of Industrial Organization*, pages 1701–1844. Elsevier. [1](#)
- Bergemann, D., Bonatti, A., and Gan, T. (2019). The Economics of Social Data. Cowles Foundation Discussion Papers 2203, Cowles Foundation for Research in Economics, Yale University. [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. [16](#)
- Cairncross, F. (2019). The Cairncross Review: a sustainable future for journalism. UK Department for Digital, Culture, Media and Sport. [2](#)
- Choi, J. P., Jeon, D.-S., and Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173:113–124. [2](#)
- CMA (2020). Online platforms and digital advertising market study. Market study final report, Competition and Markets Authority. [3](#), [4](#), [1](#), [1](#), [3.2](#), [5.2](#), [7](#)
- 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. [6.2](#)
- European Commission (2022). Digital Markets Act. Regulation 2022/1925. [5.2](#), [7](#)
- European Parliament (2016). General Data Protection Regulation. Regulation 2016/679. [1](#), [5.3](#), [7](#)
- Ganuzza, J.-J. (2004). Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations. *RAND Journal of Economics*, 35(3):583–598. [16](#)
- Ganuzza, 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. [16](#)

- Gentzkow, M., Shapiro, J. M., Yang, F., and Yurukoglu, A. (2021). Advertising prices in equilibrium: Theory and evidence. Brown University Working Paper. [2](#)
- Goldfarb, A. and Tucker, C. (2012). Shifts in privacy concerns. *American Economic Review*, 102(3):349–53. [2](#), [5.3](#)
- Goldfarb, A. and Tucker, C. E. (2011). Privacy regulation and online advertising. *Management science*, 57(1):57–71. [1](#), [2](#)
- IAB (2017). IAB Europe Report: European Programmatic Market Sizing 2017. available at <https://iab europe.eu/research-thought-leadership/iab-europe-report-european-programmatic-market-sizing-2017/>. [1](#)
- Jia, J., Jin, G. Z., and Wagman, L. (2021). The short-run effects of the general data protection regulation on technology venture investment. *Marketing Science*, 40(4):661–684. [2](#)
- Johnson, G. A., Shriver, S. K., and Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1):33–51. [2](#)
- Levin, J. and Milgrom, P. (2010). Online Advertising: Heterogeneity and Conflation in Market Design. *American Economic Review*, 100:603–07. [2](#)
- Lu, S. and Yang, S. (2020). Targeting Breadth in Behaviorally Targeted Display Advertising. Working paper, USC Marshall School of Business. [1](#), [2](#), [4.1.5](#)
- 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](#)
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73. [4](#)
- Ostrovsky, M. and Schwarz, M. (2016). Reserve prices in internet advertising auctions: A field experiment. Working Paper, Stanford University. [23](#)
- Peitz, M. and Reisinger, M. (2020). Advertising networks and consumer tracking. Mimeo. [2](#)
- Peukert, C., Bechtold, S., Batikas, M., and Kretschmer, T. (2022). Regulatory spillovers and data governance: Evidence from the gdpr. *Marketing Science*, 41(4):746–768. [2](#)

- Rafeian, O. and Yoganarasimhan, H. (2021). Targeting and privacy in mobile advertising. *Marketing Science*, 40:193–218. 1, 2, 4.1.5
- Sayed, A. (2018). Real-time bidding in online display advertising. *Marketing Science*, 37:553–568. 2
- Sharma, P., Sun, Y., and Wagman, L. (2019). The Differential Effects of New Privacy Protections on Publisher 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
- US District Court of NY (2021). Google Digital Advertising Antitrust Litigation. Civil Action No.: 1:21-md-03010-PKC. 7
- Yuan, S., Wang, J., and Zhao, X. (2013). Real-time bidding for online advertising: Measurement and analysis. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*, ADKDD '13, New York, NY, USA. Association for Computing Machinery. 10

Appendix

A Proofs of Lemmas and Propositions

A.1 Proof of the claims in Section 4.1.2

We show that no significant change in the main results occurs if we relax the assumption that v is identical in the two conflated markets. If the consumer is a single-homer, the intermediary cannot do better than sell the impression under F , earning a revenue equal to v . Consider now a multi-homer that is profiled and of type θ . Let v_θ be the return that characterizes the advertisers in the consumer’s “true” market, and v_k be the return for the advertisers in the conflated, “wrong”, market. Note that we retain the assumption that v is unobservable to the platform and, hence, the pairing of conflated markets is not conditional on v .

If both the conflated markets are thick, for each of the two available impressions there are at least two bidders who did not yet acquire an impression on the given consumer. Hence, each impression is won by an advertiser in market θ and the equilibrium price is $v_\theta/2$ if and only if $v_\theta \geq v_k$. Otherwise, the impression is won by an advertiser in k and the price is $v_k/2$. Hence, the revenue from each impression is $\frac{\max(v_\theta, v_k)}{2}$.

If both the conflated markets are thin, both impressions are won by advertisers in market θ if and only if $v_\theta \geq v_k$, but the equilibrium price is $v_k/2$ since only one of these advertisers is not already acquiring an impression on the same consumer. By contrast, the impressions are won by advertisers in market k if $v_\theta < v_k$, and the equilibrium price is $v_\theta/2$. Therefore, the equilibrium price of each targeted impression on a multi-homer is $\frac{\min(v_\theta, v_k)}{2}$.

Suppose now that only one of the conflated markets is thin and that this market is k . If $v_\theta \geq v_k$, both impressions are won by advertisers in θ , and the price of each impression is $v_\theta/2$, because there are at least two advertisers in such market that do not acquire one impression already. If $v_\theta < v_k$, the impressions are won by advertisers in k , but the price is still $v_\theta/2$ since only one of the advertisers in θ is not acquiring another impression on the consumer already. Therefore, the equilibrium price of the impressions must be $v_\theta/2$. In a similar way, one can show that the equilibrium price is $v_k/2$ if the thin market is θ .

We can now calculate the expected revenue to IN when adopting PT . The price on each single-homer in a given market is v , so the expected revenue from such consumers is $\tilde{q}(1-m)\bar{v}$. The probability that conflated markets are thick (since the markets to be conflated are chosen randomly) is $(1-x)^2$. The price of each impression in that case is $\frac{\max(v_\theta, v_k)}{2}$. The expected revenue from each profiled multi-homer is therefore $\frac{2 \int_0^{v^H} v \cdot h(v) dv}{2}$, given that there are two impressions on such consumers, where $h(v) = \frac{\partial H(v)}{\partial v}$ and $H(v) = pr [\max(v_\theta, v_k) \leq v]$. Since v_θ, v_k are i.i.d. and drawn from the distribution $G(v)$, we have $H(v) = pr [v_\theta \leq v] \cdot pr [v_k \leq v] = G(v)$. Hence, $h(v) = 2G(v)g(v)$, and the expected revenue from a multi-homer is $2\tilde{q} \int_0^{v^H} vG(v)g(v)dv$. If both conflated markets are thin (probability x^2), the price is $\frac{\min(v_\theta, v_k)}{2}$. The expected revenue is therefore $\frac{2 \int_0^{v^H} v \cdot g(v) dv}{2}$, where $g(v) = \frac{\partial G(v)}{\partial v}$ and $G(v) = pr [\min(v_\theta, v_k) \leq v]$. Since v_θ, v_k are i.i.d. and drawn from the distribution $G(v)$, we have $L(v) = pr [v_\theta \leq v] + pr [v_k \leq v] - pr [v_\theta \leq v] \cdot pr [v_k \leq v] = 2G(v) - G(v)$. Therefore the expected revenue from a profiled multi-homer is $2\tilde{q} \int_0^{v^H} v(1-G(v))g(v)dv$. If only one of the markets is thin (probability $2x(1-x)$), the equilibrium price is $v_\theta/2$ or $v_k/2$, but since these variables are i.i.d., the expected revenue from a profiled multi-homer is $\int_0^{v^H} vg(v)dv = \bar{v}$.

Summing up, when adopting PT , the IN earns the following revenue

$$R_{IN} = (1 - m)\bar{v} + 2m \left[(1 - x)^2 \int_0^{v_H} vG(v)g(v)dv + x^2 \int_0^{v_H} v(1 - G(v))g(v)dv + x(1 - x)\bar{v} \right] = \\ (1 - m)\bar{v} + 2m \left[(1 - x) \int_0^{v_H} vG(v)g(v)dv + x \left(\bar{v} - \int_0^{v_H} vG(v)g(v)dv \right) \right]$$

It is straightforward to establish that $0 < \int_0^{v_H} vG(v)g(v)dv < \bar{v} \equiv \int_0^{v_H} vg(v)dv$ and that $0 < \int_0^{v_H} v(1 - G(v))g(v)dv < \bar{v}$. Therefore, although the above revenue can be larger or smaller than (5) depending on the distribution, it compares in a similar way to the revenues under F and PT . Specifically, if markets are thick, the revenue with F dominates. However, if markets are thin, the revenue with PT dominates. Hence, qualitatively, Proposition 1 does not change. Furthermore, there are two different thresholds on \tilde{q} such that the publishers outsource, which depend on the size of x , with similar properties. Proposition 2 is therefore also qualitatively unchanged.

A.2 Proof of Lemma 2

Consider the scenario where publishers outsource ads under regime PE . Consider advertiser j 's bidding strategy for each relevant impression delivered on publisher i . The advertiser gets an expected return of (6) from each such impression. The equilibrium bidding strategies of the advertisers must be such that the 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 bidding 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 impressions acquired by an advertiser in this market on publisher 2 and the value of $w_{i'j}$ for impressions on publisher 1. Consider the bidding strategy of the advertisers on such publisher. There are two possible cases:

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

- Case B: if $S_{2a} = S_{2b}$, the advertisers' willingness-to-pay for each relevant impression on 1 are such that $w_{1a} = v \left(1 - \frac{2mS_{2a}}{1+m}\right) = w_{1b} = v \left(1 - \frac{2mS_{2b}}{1+m}\right)$. 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 (6), in a first-price auction the equilibrium winning bid is (7). 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 (7) 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 acquiring. There exists also the symmetric equilibrium obtained from swapping a and b .

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 (6) for impressions on publisher 1, and consider the bidding strategy of the advertisers on this publisher. The following cases are possible:

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

- Case B: if $\min(S_{2d}, S_{2c}) > S_{2b} = S_{2a}$, advertisers' willingness-to-pay for each impression on 1 are such that $\max\left(v\left(1 - \frac{2mS_{2d}}{1+m}\right), v\left(1 - \frac{2mS_{2c}}{1+m}\right)\right) < w_{1b} = v\left(1 - \frac{2mS_{2b}}{1+m}\right) = w_{1a} = v\left(1 - \frac{2mS_{2a}}{1+m}\right)$. Hence, a and b bid equally and outbid c and d for for all such impressions, i.e. $S_{1a} = S_{1b} = 1/2 > S_{1c} = S_{1d} = 0$.
- Case C: if $S_{2d} > S_{2c} = S_{2b} = S_{2a}$, advertisers' willingness-to-pay for each impression on 1 are such that $w_{1d} = v\left(1 - \frac{2mS_{2d}}{1+m}\right) < w_{1c} = v\left(1 - \frac{2mS_{2c}}{1+m}\right) = w_{1b} = v\left(1 - \frac{2mS_{2b}}{1+m}\right) = w_{1a} = v\left(1 - \frac{2mS_{2a}}{1+m}\right)$. Hence, a, b and c bid equally and outbid d for for all such impressions, i.e. $S_{1a} = S_{1b} = S_{1c} = 1/3 > S_{1d} = 0$.
- Case D: if $S_{2a} = S_{2b} = S_{2c} = S_{2d} = 1/4$, advertisers' willingness-to-pay for each 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 (6), 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 impressions on publisher 1 is thus v . By the same token, the price of 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 (6), 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 impressions on publisher 1 is thus v . On publisher 2, d 's willingness-to-pay is v while a, b and c 's equals $v\left(1 - \frac{2m}{3(1+m)}\right)$. 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 (6) each advertiser's willingness-to-pay is $v\left(1 - \frac{m}{2(1+m)}\right)$ 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 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, since all 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 in the advertising market when the publishers do not outsource. Given Lemma 4, 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

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

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

$$\begin{cases} (1+m)\tilde{q}\int_0^{v^H} vdv & \text{if } x \leq \frac{1}{2}, \\ (1-m)\tilde{q}\int_0^{v^H} vdv + 2m\tilde{q}\frac{\int_0^{v^H} vdv}{2} = \tilde{q}\int_0^{v^H} vdv & \text{if } x > \frac{1}{2}. \end{cases} \quad (23)$$

Therefore, if $x \leq \frac{1}{2}$, total surplus increases with outsourcing given $\tilde{q} > q$. 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, i.e. $2m\tilde{q}\frac{\int_0^{v^H} vdv}{2}$, because half the impressions are sold to the “wrong” advertiser. Hence, total surplus increases when the publishers outsource if and only if $\tilde{q} > \tilde{q}_w \equiv q(1+m)$.

A.4 Proof of Proposition 4

Given a regulation that restricts IN to adopting $r = F$, the intermediary’s revenue is $\bar{v}\tilde{q}[1+m(1-2x)]$ if the 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.

Given $x \geq 1/2$, according to Proposition 1 the intermediary would adopt $r = PT$ for impressions on multi-homers without the regulation. Total surplus would be $(1-m)\tilde{q}\int_0^{v^H} vdv + 2m\tilde{q}\frac{\int_0^{v^H} vdv}{2} = \tilde{q}\int_0^{v^H} vdv$ in that case. If outsourcing does not occur, total surplus is $(1+m)q\int_0^{v^H} vdv$. Hence, if the publishers do not outsource with the regulation in place, but would have outsourced without it, we conclude that the regulation reduces total surplus if and only if $\tilde{q}_w \equiv q(1+m) < \tilde{q}$ holds.

B Partial Disclosure with uneven disclosure probabilities

We show that, in our setting, the intermediary would never want to implement PT with uneven weights. From Lemma 1 we know that the intermediary always adopt Full Disclosure for single-homers, because the intermediary always finds it profitable to disclose all available information in this case. Hence, we concentrate here on auctions on multi-homers. Suppose the intermediary adopts PT . In the current version of the paper, we restrict attention to the case where consumers belong to each of the two conflated markets with equal probability. We now relax that hypothesis by letting the intermediary disclose to the advertisers that a consumer belongs to the true market, say market “ a ”, with probability $z \in [\frac{1}{2}, 1]$, and to market “ b ” with probability $(1-z) \in [0, \frac{1}{2}]$. For concreteness (and without loss), we suppose that the

intermediary gives to the true market the consumer belongs to a weakly higher weight (e.g., because of reputational concerns with respect to publishers and advertisers). Also, observe that the case with $z = 1$ corresponds to the Full Disclosure regime.

Suppose first the multi-homing consumer belongs to a thin market. Let v be the return that the advertisers in these markets get from informing a relevant multi-homer. Thus, for a targeted and non-repeated impression, the advertisers in market a expect a return of zv , whereas the advertisers in market b expect a return of $(1 - z)v$. Because $z \geq \frac{1}{2}$, the advertisers in market a have a (weakly) higher willingness-to-pay for unrepeated impressions than the advertisers in market b . However, whenever an advertiser in market a has already impressed the consumer, its willingness-to-pay for another impression on the same consumer is zero. Hence, in a first-price auction, the highest bid would be placed by an advertiser in market a , so both impressions would be acquired by advertisers in this market. However, the equilibrium price is equal to the second-highest willingness-to-pay, that is, the willingness-to-pay of an advertiser in market b , i.e. $(1 - z)v$ (see Table 3 below).

Suppose now the multi-homing consumer belongs to a thick market. The equilibrium price is then equal to the highest willingness-to-pay zv because there are at least three advertisers in market a willing to bid that amount on each multi-homer.

Summing up, conditional on adopting PT , the total revenue on multi-homers is $R_{MH}^{PT} = 2m\bar{v}\tilde{q}[x(1 - z) + (1 - x)z]$. This function is linear in z , meaning that its maximum is at either $z = \frac{1}{2}$ or $z = 1$. Specifically, we find that the maximum is at $z = 1$ if and only if $x \leq 1/2$, which is the same threshold as in Lemma 1. In other words, if and only if $x \leq 1/2$, PT does not outperform Full Disclosure (we assume that, when indifferent, the intermediary adopts the lost transparent regime). If $x > 1/2$, the intermediary should set $z = \frac{1}{2}$. Hence, as claimed above, there is no loss in assuming the intermediary gives equal weights to the two markets when adopting PT , as in the current version of the model.²⁷

C Intermediate markets (for online publication)

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

²⁷In the above discussion we assumed that the intermediary conflates two markets. Using a similar reasoning as in the paper, one can easily show that it is never profitable to conflate more than two markets.

Table 3: Equilibrium willingness-to-pay and price for impression on a multi-homer in a thin market under PT .

	Publisher 1	Publisher 2
$w_{i,a1}^F$	zv	0
$w_{i,a2}^F$	0	zv
$w_{i,b1}^F$	$(1-z)v$	$(1-z)v$
$w_{i,b2}^F$	$(1-z)v$	$(1-z)v$
$p_{n=2,MH}^F$	$(1-z)v$	$(1-z)v$

Full Disclosure ($r = F$). Following the same steps as in 4.1.1, we find that if a 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 is

$$p_{SH}^F = p_{n=3,MH}^F = v.$$

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

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

Partial Type Disclosure ($r = PT$). 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

$$p^{PT} = \frac{v}{2}. \quad (25)$$

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

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

Partial Exposure Disclosure ($r = PE$). 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 $v(1 - \frac{m}{1+m})$. This expression is derived from (11), noting that a and b each win half the available targeted impressions on publisher 1, so $S_{1a} = S_{1b} = 1/2$.

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

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	$v \left(1 - \frac{m}{1+m}\right)$
$R_{n=3}$	$v\tilde{q}\frac{(1+m)}{2} + v \left(1 - \frac{m}{1+m}\right)$	$\tilde{q}\frac{(1+m)}{2} = v\tilde{q} \left(1 + \frac{m}{2}\right)$

Then, the willingness-to-pay by advertiser c for each targeted impression on publisher 2 is v , whereas on publisher 1 it equals (7), 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.

Hence, considering all markets, total revenues are

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

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.

Proof of the statements regarding equilibrium in the PE regime. Let $\{a, b, c\}$ be the set of advertisers in the market. Focus, without loss of generality, on the relation between the share of impressions acquired by an advertiser in this market on publisher 2 and the value of (6) for impressions on publisher 1, and consider the bidding strategy of the advertisers on such publisher. The following cases are possible

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

- Case C: if $S_{2a} = S_{2b} = S_{2c}$, advertisers' willingness-to-pay for 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 (6), 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 because 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.

Choice among disclosure regimes. First, we compare F and PE . Revenues in intermediate markets are higher under F than under PE , because the intermediary are able

Table 5: Advertiser's willingness-to-pay, equilibrium price and publishers' 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	$v \left(1 - \frac{mq}{1+m}\right)$
$R_{i,n=3}$	$\frac{vq}{2} (1 + m)$	$\frac{vq}{2} (1 + m (1 - q))$

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 PE 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 (5)), we see that revenues are not affected. Indeed, revenues in intermediate markets are as in thick markets. Hence, Proposition 1 and Lemma 3 are not affected.

Market equilibrium without the intermediary. Following the same steps as in 4.2 and above for the regime PE 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

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

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.

D Proofs of robustness checks (for online publication)

D.1 Reserve price

We establish that adopting F with a revenue-maximizing reserve price does not necessarily increase IN 's revenue, compared to adopting the PT regime for multi-homers. We proceed as follows. We characterize the equilibrium bidding strategies of advertisers conditional on the reserve price p_R . Next, we compute the equilibrium value of p_R chosen by IN . Note that p_R cannot be made conditional on v and n , since the intermediary does not observe such parameters. Finally, we compare the revenue earned with the revenue-maximizing reservation price under F to the revenue earned with PT with a reserve price. For concreteness, we focus only on impressions on multi-homers and in thin advertising markets, where, as demonstrated in the baseline model, there is a strong drop in ad prices on multi-homers with full disclosure. For the sake of brevity, we do not consider the PE regime.

Consider an impression on a profiled consumer and assume IN adopts $r = F$. If the consumer is a multi-homer, the equilibrium willingness-to-pay are as characterized in Table 1 for the couple of impressions available. Hence, the equilibrium price of each such impression is zero if there is no reserve price. With a reserve price p_R , the equilibrium price is $p = p_R$, but the impressions are sold only if the market is such that $v \geq p_R$.

Let us now compute IN 's expected revenue from multi-homers in each (thin) market when adopting $r = F$ and setting a reserve price p_R . This revenue is $\tilde{q}(2mp_R)$ if $v \geq p_R$, and 0 if $v < p_R$, since in the latter case IN does not sell the impressions. Because impressions on multi-homers in thin markets are sold at zero in the baseline model, it follows that, conditional on adopting $r = F$, the intermediary is better off imposing a reserve price such that $v_H \geq p_R > 0$.

Let us now calculate the revenue-maximizing p_R . The revenue on multi-homers is

$$R^F(p_R) = \tilde{q} \left(2mp_R \int_{p_R}^{v_H} dG(v) \right) = 2\tilde{q}mp_R(1 - G(p_R)). \quad (28)$$

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

$$\frac{dR^F(p_R)}{dp_R} = 1 - G(p_R) - p_R g(p_R) = 0 \quad (29)$$

and therefore

$$p_R^* = \frac{1 - G(p_R^*)}{g(p_R^*)}. \quad (30)$$

Suppose now IN adopts PT . There is no scope for a reserve price to increase revenues under PT , because the equilibrium price equals the advertisers' willingness-to-pay (see Section 4.1.2). Hence, conditional on $r = PT$, IN does not impose any reserve price. Given the price of an impression under PT , in (4), the total revenue from multi-homers in thin markets under PT is $\tilde{q}\bar{v}m$.

Given the above findings, IN adopts $r = F$ with the reserve price p_R^* in (30) for each impression on multi-homers if and only if $R^F(p_R^*) > \tilde{q}\bar{v}m$. Otherwise, IN adopts $r = PT$ without a reserve price for each such impression. We find that

$$R^F(p_R^*) > \tilde{q}\bar{v}m \iff \tilde{q}m \left(2 \frac{(1 - G(p_R^*))^2}{g(p_R^*)} - \bar{v} \right) > 0. \quad (31)$$

Whether the above condition holds depends on the distribution of advertising returns, $G(v)$. The strict inequality does not hold for two notable distributions:

- Suppose that $G(v)$ is a uniform distribution with support $[0, v_H]$. Hence, $E(v) = \frac{v_H}{2}$, $G(p_R) = \frac{p_R}{v_H}$ and $g(p_R) = \frac{1}{v_H}$. Replacing these values in (30), we obtain that $p_R^* = \frac{v_H}{2}$. Replacing further in (31), we obtain that $R^F(p_R^*) = \tilde{q}\bar{v}m$.
- Suppose that $G(v)$ is a continuous Bernoulli distribution with $\lambda = \frac{1}{2}$. Hence, $E(v) = \frac{1}{2}$, $G(p_R) = p_R$ and $g(p_R) = 2$. Replacing these values in (30), we obtain $p_R^* = \frac{1}{3}$. Replacing further in (31), we obtain that $R^F(p_R^*) < \tilde{q}\bar{v}m$.

D.2 Competition between advertisers

No intermediary. Let w_{ij} be the willingness-to-pay of advertiser $j \in k(\theta)$ for a targeted impression on a 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:

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

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 impression 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 $\alpha \geq \frac{1}{2}$. Therefore, as in our baseline model, impressions on the two publishers are substitutes for the advertisers. 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 6. *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 lemma, 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's willingness-to-pay is v on one publisher and $v \left(\frac{(1-m)+2m(1-\alpha q)}{1+m} \right)$ on the other. Hence, the equilibrium 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 $R_{i,n=2} = \bar{v} \left(\frac{1-m}{2} + m(1-\alpha q) \right)$. 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 $R_{i,n \geq 4} = \bar{v} \left(\frac{1-m}{2} + m(\alpha q + 1 - q) \right)$.

Suppose now that $\alpha < \frac{1}{2}$. A single advertiser multi-homes and wins all impressions. The second-highest willingness-to-pay is by 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 $R = \bar{v} \left(\frac{1-m}{2} + m(\alpha q + 1 - q) \right)$ respectively, in thin and thick markets.

Intermediary and Full Disclosure. Assume the intermediary chooses $r = F$ 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. 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 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 $R^F = \bar{v}(1 - m + 2m(1 - \alpha))$. In thick markets, the price is $p = v\alpha$ and $R_{n=2}^F = \bar{v}(1 - m + 2m\alpha)$. 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 $R^F = \bar{v}(1 - m + 2m\alpha)$.

Intermediary and Partial Exposure Disclosure. Suppose now that the intermediary adopts $r = PE$. 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

$$w_{i,j}^{PE} = v \frac{\left(\frac{1-m}{2}\right) + m((1-\alpha)S_{i'j} + \alpha(1 - S_{i'j}))}{\frac{1+m}{2}}, i' \neq i.$$

Observe that $w_{i,j}^{PE}$ decreases in $S_{i'j}$ 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 $p_{n=2} = v \frac{(1-m)+2m(1-\alpha)}{1+m}$ and $R_{n=2}^{PE} = \bar{v}(1 - m + 2m(1 - \alpha))$ respectively. In thick markets, the equilibrium price and revenues are $p_{n \geq 4} = v \frac{(1-m)+2m\alpha}{1+m}$ and $R_{n \geq 4}^{PE} = \bar{v}(1 - m + 2m\alpha)$ 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 $p = v \frac{(1-m)+2m\alpha}{1+m}$ and $R^{PE} = \bar{v}(1 - m + 2m\alpha)$ respectively.

Intermediary and Partial Type Disclosure. Suppose now that the intermediary adopts $r = PT$ for targeted impressions on multi-homers. Recall that impressions on single-homers 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 competitors. If the consumer receives an impression from an advertiser in the conflated market, i.e. not a competitor, the return is still $v/2$. The return is instead $\alpha v/2$ if the consumer receives another impression from a competitor. 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 n , total revenues are $R^{PT} = (1 - m)\bar{v} + \frac{\bar{v}}{2}\tilde{q}2m = \tilde{q}\bar{v}$.

Choice of disclosure regime. As in the baseline model, Full Disclosure dominates Partial Exposure Disclosure. The intermediary thus sells all impressions on single-homers under full disclosure. As for impressions on multi-homers, IN chooses F over PT if and only if

$$\begin{cases} \bar{v}\tilde{q}m2(x(1-\alpha) + (1-x)\alpha) > \bar{v}\tilde{q}m, & \text{if } \alpha \geq \frac{1}{2}, \\ \bar{v}\tilde{q}m2\alpha > \bar{v}\tilde{q}m, & \text{if } \alpha < \frac{1}{2}. \end{cases}$$

Hence, if $\alpha \geq \frac{1}{2}$, there is F if and only if $x(1-2\alpha) + \alpha > 1/2$. Instead, if $\alpha < \frac{1}{2}$, PT dominates F .

D.3 Proof of the claim in Section 6.3

Let ρ 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 ρ , the intermediary's objective, $(1 - 2\rho)R_{IN}$, is maximized by the same choice of disclosure regime for any ρ . Let us now establish how ρ should be set in order to induce the publishers to outsource. Each publisher outsources if and only if the amount transferred by the intermediary, ρR_{IN} , is at least as large as the revenue the publisher could earn independently, R_i (recall, as established in the main text, that this revenue does not depend on whether the other publisher outsources). Since the intermediary's objective is decreasing in ρ , in equilibrium we will have $\rho = R_1/R_{IN} = R_2/R_{IN}$, R_1 and R_2 being identical in our setting. Replacing for ρ in the intermediary's profit, $(1 - 2\rho)R_{IN}$, we obtain that this profit is positive if and only if $R_{IN} \geq \sum_{i=1,2} R_i$, which is identical to condition (16).

	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 $n = 2$.

D.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

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

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 . 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 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 $\frac{vq}{2} (1 + m)$.

	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 $n \geq 4$.

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:

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

Although we do not repeat the entire analysis here, one can follow similar steps as in our baseline model, it is fairly easy to see that this revenue is weakly larger than the revenue the platform could earn under either PT or PE . 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.

D.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. 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^+ .

D.5.1 Intermediary

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

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 (2).

Partial Type Disclosure. 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-highest willingness-to-pay for each targeted impression on a multi-homer and a single-homer is $v/2$. Hence, the revenue is as characterized in (5).

Partial Exposure Disclosure. 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

$$w_{ia} = v^+ \left(1 - \frac{2mS_{i'a}}{1+m} \right), \quad w_{ib} = v \left(1 - \frac{2mS_{i'b}}{1+m} \right), \quad i = 1, 2. \quad (34)$$

Similarly to the proof of Lemma 2 (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 impressions 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 $(v^+ + v) \tilde{q} \frac{1-m}{2}$ in total.

Consider now a *thick* market. Similarly to the proof of Lemma 2 (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 $vq(1+m)$.

D.5.2 Revenue comparison and choice of disclosure regime

Case $v^+ \left(1 - \frac{2m}{1+m}\right) \geq v$. As explained above, if $v^+ \left(1 - \frac{2m}{1+m}\right) \geq v$ a dominant advertiser acquires all the targeted impressions in its market with PE . Hence, in all markets with a dominant advertiser, we have $R_{IN}^{PE} = v\tilde{q}(1+m)$. In the baseline model, in either thin or thick markets without a dominant advertiser, we find that the revenues from F and PE are identical (see (2), and (9)). With a dominant advertiser, however, $R^{PE} = v\tilde{q}(1+m)$ dominates the revenue with F in thin markets. Hence, under our assumptions, PE is strictly preferable to F . The total revenue earned by the intermediary with PE is

$$R_{IN}^{PE} = \tilde{q}\bar{v}(1+m(1-2x)). \quad (35)$$

Observe that in thin and thick markets with a dominant advertiser and in thick markets without a dominant advertiser, the intermediary obtains the same revenue, $v\tilde{q}(1+m)$, so the shares x' and z drop out from the above expression.

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 is $R_{IN} = \tilde{q}\bar{v}$. The comparison between this revenue and (35) reveals that IN will choose the PE if and only if $x \leq 1/2$. Otherwise, it adopts F on single-homers and PT on multi-homers, as in the baseline model. The revenue earned by the intermediary is isomorphic to (10).

Case $v^+ \left(1 - \frac{2m}{1+m}\right) < v$. Under this condition, the dominant advertiser only acquires the impressions on one publisher under PE . Hence, in all thin markets with a dominant advertiser, we have $R_{IN}^{PE} = (v^+ + v)\tilde{q}\frac{1-m}{2}$. In the baseline model, in either thin or thick markets without a dominant advertiser, the revenues from F and PE are identical (see (2) and (9)). With a dominant advertiser, however, $R^{PE} = (v^+ + v)\tilde{q}\frac{1-m}{2}$ dominates the revenue with F in thin markets. Hence, under our assumptions, PE is strictly preferable to F . The total revenue earned by the intermediary with PE is

$$R_{IN}^{PE} = \tilde{q}\bar{v}(1 + m(1 - 2(x + x'))) + (\bar{v}^+ - \bar{v})\tilde{q}\frac{1-m}{2}x', \quad (36)$$

where \bar{v}^+ is the average of the valuation of dominant advertisers computed over all markets. Observe that in thick markets, with and without a dominant advertiser, the intermediary obtains the same revenue under PE , so the shares z drop out from the above expression. 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 is $R_{IN} = \tilde{q}\bar{v}$. The comparison between this revenue and (36) reveals that IN will choose the PE if and only if the share of homogeneous thin markets, x , is smaller than a threshold, $T_{IN}^+ \equiv \frac{m(1-2x')}{2} + \frac{(\bar{v}^+ - \bar{v})(1-m)x'}{4\bar{v}}$. If $x \geq T_{IN}^+$, the intermediary should adopt F on single-homers and PT on multi-homers. Again, this suggests that the intermediary should be less likely to adopt this policy, with respect to our baseline setting. Overall, the intermediary earns therefore

$$R_{IN} = \begin{cases} \tilde{q}\bar{v}(1 + m(1 - 2(x + x'))) + (\bar{v}^+ - \bar{v})\tilde{q}\frac{1-m}{2}x', & \text{if } x \leq T_{IN}^+, \\ \bar{v}\tilde{q}, & \text{if } x > T_{IN}^+. \end{cases} \quad (37)$$

D.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

$$w_{ia} = v^+ \left(1 - q \frac{2mS_{i'a}}{1+m} \right), \quad w_{ib} = v \left(1 - q \frac{2mS_{i'b}}{1+m} \right), \quad i = 1, 2. \quad (38)$$

Similarly to the proof of Lemma 2 (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 PE .

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 $vq \frac{1+m}{2}$.

D.5.4 When do the publishers outsource?

We proceed again considering two cases:

Case $v \leq v^+ (1 - \frac{2m}{1+m})$. This condition 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

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

which is identical to (15). Given the revenue earned by the intermediary is equal to (10), as established above, we obtain the same threshold on \tilde{q} that is isomorphic to (17) provided in Proposition 2. Hence, there is no qualitative change in the proposition.

Case $v > v^+ (1 - \frac{2qm}{1+m})$. Under this condition, the dominant advertiser a only acquires the impressions available on one of the publishers (single-homing). The total revenue earned by the publishers is

$$R_1 + R_2 = \bar{v}q(1 + m(1 - 2q(x + x'))) + \frac{\bar{v}^+ - \bar{v}}{2}q(1 + m(1 - 2q))x'.$$

Comparing this revenue to (37), it is possible to characterize a threshold on \tilde{q} , denoted as \tilde{q}_T , that is similar to (17). More specifically,

$$\tilde{q}_T \equiv \begin{cases} q \left(\frac{\bar{v}(1+m(1-2q(x+x')))+\frac{\bar{v}^+-\bar{v}}{2}(1+m(1-2q))x'}{\bar{v}(1+m(1-2(x+x')))+(\bar{v}^+-\bar{v})\frac{1-m}{2}x'} \right), & \text{if } x \leq T_{IN}^+, \\ q \left(\frac{\bar{v}(1+m(1-2q(x+x')))+\frac{\bar{v}^+-\bar{v}}{2}(1+m(1-2q))x'}{\bar{v}} \right), & \text{if } x > T_{IN}^+. \end{cases}$$

Again, there is no qualitative change in Proposition 2.

D.6 Sequential auctions

D.6.1 Proof of the claim in footnote 9 (sequential actions)

In this Appendix we show that considering sequential auctions would not affect the results of the baseline model. Let us concentrate on the Full Disclosure case, and allow consumers to visit different publishers in different moments, meaning that the two impressions opportunity on a multi-homer do not occur simultaneously. Consider thin markets. Under Full Disclosure, the intermediary discloses the type of the consumer and whether the impression is the first or the second one on a given consumer. When an auction for an ad impression occurs, if the intermediary discloses that the auctioned impression is the second one on a given consumer, the

equilibrium bid will be equal to 0, because one advertiser has already impressed the consumer, and the other advertiser makes a bid equal to this second-highest willingness-to-pay on that consumer, that is zero. Then, if the intermediary discloses that the auctioned impression is the first one on a consumer, neither the advertiser nor the intermediary know whether a consumer is a single- or a multi-homer. Hence, all advertisers in the market will have a willingness-to-pay equal to v if this impression is on a single-homer (that occurs with probability equal to $1 - m$), and equal to 0 if this impression is on a multi-homer (that occurs with probability equal to m). Indeed, if the consumer is a multi-homer, each advertiser knows that another impression opportunity on the same consumer will occur, and that it will be sold at zero. Hence, the bid on the first impression on a given consumer would be equal to $v(1 - m)$, and there are a total of 1 first impressions. At equilibrium, under regime F , the revenues in thin markets are $\bar{v}\bar{q}(1 - m)$, that is, only impressions on single-homers are sold at a positive price. Following the same reasoning, it is easy to prove that under Full Disclosure, if markets are thick, even when auctions are not simultaneous, the intermediary extracts full revenues, because there are always two advertisers willing to bid v for each impression. Hence, revenues will be exactly the same as in equation (3). Following the same reasoning as in the paper, in all other disclosure regimes revenues would not change.

D.6.2 Proof of the claim in Footnote 24 (reserve price in sequential auctions)

In this Appendix, we briefly explore the implications of a reserve price imposed by the intermediary if impressions on multi-homers are not sold simultaneously. For concreteness, we focus on the Full Disclosure regime and, since the intermediary can extract the full value v from each impression on single-homers (regardless of market thickness) and on multi-homers in thick markets (even without a reserve price), we concentrate on impressions on multi-homers in thin markets ($n = 2$). Our objective is to establish that introducing a dynamic reserve price would not change our results fundamentally.

Suppose the auctions for impressions on a multi-homer take place sequentially. Assume the intermediary sets a reserve price p_R in the second of the two auctions for impressions on a given multi-homer, and no reserve price in the first auction. Let us also assume that $0 \leq p_R \leq v$ to focus on the relevant set of parameters (otherwise, the impression would not be sold). Finally, assume that p_R is set equal to the price the impression is sold in the first auction. Obviously, the equilibrium price in the second auction, p_2 , is equal to p_R : only one advertiser has a positive willingness-to-pay for the consumer, and the lowest price at which it can win the auction is p_R . Consider now the equilibrium in the first auction. This equilibrium must

also be such that the price, p_1 , equals p_R by construction. Hence, in any thin market, there is (in principle) an infinity of equilibria such that $p_1 = p_2 = p_R$, with $0 \leq p_R \leq v$. If any of these equilibria can emerge with equal probability, the expected revenue from each impression would be $v/2$, and so the total revenue from the two impressions sold on a multi-homer would be v . Therefore, the intermediary would not be able to extract the full revenue. In fact, the expected revenue would be the same as under PT (see Lemma 1 and Proposition 1). Moreover, among the set of equilibria described above, the equilibrium where $p_1 = p_2 = p_R = 0$ is the *only* Pareto efficient (and coalition-proof) one from the perspective of the advertisers. Under this selection criterion, the equilibrium would thus be the same as in our baseline model under Full Disclosure.