University of Groningen

Targeted advertising, platform competition and privacy

Kox, Henk; Straathof, Bas; Zwart, Gijsbert

Published in:
Journal of Economics & Management Strategy

DOI:
[10.1111/jems.12200](https://doi.org/10.1111/jems.12200)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Final author's version (accepted by publisher, after peer review)

Publication date:
2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Kox, H., Straathof, B., & Zwart, G. (2017). Targeted advertising, platform competition and privacy. *Journal of Economics & Management Strategy*, 26(3), 557-570. <https://doi.org/10.1111/jems.12200>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Targeted advertising, platform competition and privacy*

Henk Kox^{a,b}

Bas Straathof^a

Gijsbert Zwart^{c,d}

November 14, 2016

Abstract

Targeted advertising can benefit consumers through lower prices for access to websites. Yet, if consumers dislike that websites collect their personal information, their welfare may go down. We study competition for consumers between websites that can show targeted advertisements. We find that more targeting increases competition and reduces the websites' profits, but yet in equilibrium websites choose maximum targeting as they cannot credibly commit to low targeting. A privacy protection policy can be beneficial for both consumers and websites. If consumers are heterogeneous in their concerns for privacy, a policy that allows choice between two levels of privacy will be better. Optimal privacy protection takes into account that the more intense competition on the high-targeting market segment also benefits consumers on the less competitive segment. Consumer surplus is maximized by allowing them a choice between a high targeting regime and a low targeting regime which affords more privacy.

Keywords: platform competition, targeted advertising, privacy policy

JEL classification: D43, L13, L82, M38

*We are grateful to Paul Westra for research assistance. We thank Susan Athey, Leon Bettendorf, Michiel Bijlsma, Marco Haan, Marielle Non and Rianne van Dalen for comments and discussions.

^aCPB Netherlands Bureau for Economic Policy Analysis.

^bKVL Economic Policy Research, Den Bosch, The Netherlands

^cUniversity of Groningen, The Netherlands.

^dCorresponding author. Email: g.t.j.zwart@rug.nl.

1. Introduction

In 2015, internet advertising revenues in the US for the first time exceeded 17 billion dollars per quarter. The share of internet advertising has been rising steadily, currently totalling a third of total US advertising revenues, and almost equals that of broadcast television.¹

Internet advertising allows for more precise targeting, compared to traditional media.² Whereas a television commercial will be the same for all those watching a particular show, and therefore tailored to consumers' average interests conditional on their watching the show, the internet allows targeting on each individual's specific characteristics.

Targeting of advertisements is valuable to advertisers as it increases the probability that the advertisement leads to a purchase. This is also reflected in pricing schemes, which are increasingly based on click-through rates or other performance measures, rather than mere numbers of viewers. A website that can target its advertisements better will collect higher per-viewer revenues from advertisers, which may be partly reflected in lower subscription prices (if present) or higher quality to attract more visitors.

The advantages of targeting for advertisers induce firms to collect personal information on their consumers. Personal information used for targeting includes keywords entered in search engines, recent browsing history, previous web purchases or even the topics in their emails. Web companies such as Axiom or Bluekai collect information on individuals' web behaviour and use that to categorize consumers into profiles, which they then sell on to advertising sites.³

Better targeting has a potential drawback for consumers: consumers may care about the associated loss of privacy. Consumers have limited possibilities to verify what kind of personal information firms collect and how they use this information. How well do firms protect personal information against theft or manipulation by criminals? Do firms use personal information to

¹As reported by the Interactive Advertising Bureau's "Internet Advertising Revenue Report", 2016. See the annex to this paper, available at www.cpb.nl/en/publication/targeted-advertising-platform-competition-and-privacy for an overview of facts and figures on online advertising.

²Evans (2009) and Goldfarb (2014) provide overviews of the economics of online advertising.

³See e.g. "Who Do Online Advertisers Think You Are?", New York Times, November 30, 2012. In the annex to this paper, we describe the various types of players in this industry in more detail.

raise prices for some groups of consumers? In addition, some internet users may feel uneasy when they discover that their movements on the web may be recorded and reflected in the ads they are shown.⁴

These privacy costs are heterogeneous across consumers. A survey by Turow et al. (2009) shows that 66 percent of Americans does not want to have ads tailored to their personal characteristics. Goldfarb and Tucker (2012) find that consumers increasingly refuse to disclose sensitive information online, and Goldfarb and Tucker (2011) demonstrate that these costs are economically relevant, in the sense that they alter consumers' purchasing decisions in response to advertising. See Tucker (2012) for an overview.

One response to such consumer uneasiness is for web companies to offer consumers a choice on how much information can be collected on them. As an example, internet provider AT&T offered customers a 29 dollar reduction on their monthly subscription bill if the firm can use their information on browsing behaviour to better target the ads it shows them.⁵ Also, many websites allow consumers either to opt for signing in to the site or to browse anonymously. Signing in may increase the quality the site can offer, at the expense of the site storing previous browsing history. Alternatively, consumers may choose not to accept cookies, or may join industry "do-not-track" registers.

There may be a role for public intervention to protect online privacy. For one thing, many consumers may be ill-informed about websites' information gathering activities and privacy policies. It is costly or impossible for consumers to verify whether the websites they visit collect and use personal information. In the absence of verifiable contracts on the degree of privacy protection, these sites may have trouble committing to a strict privacy policy. Government intervention can help in providing a credible standard for privacy protection. Indeed, both in the EU and in the US stricter online privacy laws are being put in place.

In this paper, we show that competition among websites may also drive sites to choose

⁴There are various forms of targeting, see the annex. In the case of retargeting, an ad for a product which has been previously looked up on the internet may afterwards appear as a banner ad on completely unrelated webpages. Consumers may feel watched and stalked by the product as a result, see e.g. "Web trackers are totally out of control", ITworld.com, March 21, 2013.

⁵"AT&T says your privacy is worth \$ 29 a month", Techtimes.com, Dec 12 2013.

levels of privacy protection that are too low from their own perspective, as well as from their consumers' perspectives when these suffer costs from loss of privacy. The lower profitability may explain why websites make endeavours to commit to reduced targeting, through standards proposed by trade associations or private certification.⁶ Government regulation of privacy may lead to a Pareto improvement, increasing both websites' and consumers' surplus.

We analyze a model of websites that act as two-sided platforms, matching advertisers to consumers. The websites compete for consumers in a Hotelling fashion. Consumers visit only one platform (single-homing). This implies that the websites are competitive bottlenecks (Armstrong, 2006): websites are effectively gatekeepers for advertising access to consumers. This allows them to extract monopoly rents from advertisers. Websites can strategically choose the level of targeting of ads to their consumers. Advertisers have higher willingness to pay (per consumer) for an ad that is better targeted at a consumer's taste, which corresponds to his location on the Hotelling line. We assume consumers dislike being targeted, but their disutility from the amount of targeting that they are exposed to is heterogeneous.

We demonstrate that an increasing amount of ad targeting leads the web platforms to compete more vigorously with each other. Although a larger surplus is created by the better match between advertiser and consumer, this additional surplus is more than dissipated to consumers, reducing the platforms' profits. The intuition is that better ad targeting in particular renders marginal consumers more profitable to the website. Without targeting, these consumers are not very valuable to potential advertisers, as these will focus on the average consumer on the website. This changes when the ad can be better matched to the consumer. The higher ad revenues are channeled through to consumers, as is standard in the competitive bottleneck framework (see e.g. Anderson and Coate, 2005). But in addition, the higher value of the marginal consumer compared to the inframarginal ones spurs competition among the websites (Crampes, Haritchabalet and Jullien, 2009).

⁶Listokin (2015) studies two examples of such programs. TrustE provides private certification, see www.truste.com/consumer-resources/trusted-directory for a directory of clients, which include firms such as Apple or Cisco. The Digital Advertising Alliance, DAA, is a trade association, providing privacy standards for their members. Many large advertisers participate in their self-regulatory program, see www.aboutads.info/participating. Listokin (2015) argues, however, that such initiatives have only limited success, due to difficulties in private enforcement and adverse selection of firms into certification.

In spite of the reduced equilibrium profits, websites maximize targeting when choosing the level of targeting non-cooperatively, if consumers cannot observe that targeting level. Consumers benefit from the increased competition among websites and from the higher advertising revenues that decrease website prices, but if average costs associated with loss of privacy are high, the outcome is socially suboptimal as well. In that case a ban on targeting can be optimal.

Since consumers have heterogeneous costs of privacy loss, one may do better by allowing websites to differentiate the levels of targeting, as in the example of AT&T. We analyze equilibrium outcomes when privacy policy sets and enforces two maximum levels of targeting, one involving high privacy (low targeting) and one with lower privacy and thus more targeting. Websites now compete in menus of two vertically differentiated products: one with high and one with low targeting. We demonstrate that high privacy (low targeting) consumers now also benefit from the more intense competition in the low privacy (high targeting) segment. Prices in the high privacy segment are dragged down through the marginal effect of those consumers whose privacy costs are intermediate, such that they are indifferent between the high price, high privacy product and the low price, low privacy product. We find that a total surplus-maximizing regulator always allows some targeting on the low privacy segment. Consumer surplus maximization involves maximal targeting on the low privacy segment.

Our paper is related to the literature on advertising on two-sided platforms, drawing on Anderson and Coate (2005). We consider the model on a Salop circle, as in Crampes, Haritchabalet and Jullien (2009). Their result that with decreasing returns to scale competition is relaxed ties in with our result on the competitive effect of targeting. In contrast with that paper, in our model the level of targeting is endogenous.

The marketing literature on targeting and competition (e.g. Chen, Narasimhan and Zhang, 2001; Gal-Or and Gal-Or, 2005; Iyer, Soberman and Villas-Boas, 2005) focuses on targeting strategies by product suppliers themselves, and shows how this targeting can soften competition. We find an opposite effect when looking at targeting by intermediators, the website platforms.

Relatedly, recent literature on privacy focuses on its effects on price discrimination by firms (see Fudenberg and Villas-Boas, 2006, for an overview). If firms can customize not the adver-

tising, but the prices, based on observed client characteristics, a typical conclusion is that this leads to lower profits, as demonstrated first in Thisse and Vives (1988). The reason is that firms will tend to engage in Bertrand competition over each individual consumer. Chen and Iyer (2002) study a model with advertising and price discrimination, and also find that price discrimination tends to make those consumers that can choose among different providers to gain from lower prices. More recently, Taylor and Wagman (2014) analyze, using various models of competition, who gains and who loses from privacy regulation in models where websites can use information to price discriminate. In our setting, in contrast, we assume that firms do not price discriminate in their own product, e.g. because they can coordinate on a uniform pricing equilibrium which is more profitable for them, but do use consumer information to improve matching of advertisers to consumers, increasing surplus from the advertising side of the market.

Our paper incorporates the interaction of competition with privacy concerns and privacy legislation. Campbell, Goldfarb and Tucker (2015) is a recent study in this direction. They focus on entry barriers related to scale economies on the consumer side of having to familiarize oneself with the privacy policies, and having to consent with them.

2. Model

We consider a model of n horizontally differentiated internet firms (‘websites’), competing for consumers who can be homogeneously mapped to a preference space in the form of a circle, following Salop (1979). The utility consumers obtain from visiting a website depends on the distance on the circle between the consumer and the website, as well as on price and privacy policy. Websites’ revenues come from two sources. First, the websites offer content to consumers and compete in prices to attract consumers to their sites. In addition, websites also derive revenues from presenting advertisements to the consumers that visit their site. We consider a continuum of horizontally differentiated advertisers, uniformly distributed on the same Salop circle. Advertisers compete perfectly to have their advertisement shown to the

websites' consumers.

The focus of our model will be on the websites' ability to target advertisements to consumers. We assume that websites can freely choose the fraction of their subscribers $\rho \in [0, 1]$ for which they gather personal information. Higher ρ means more targeting. When a website has personal information on a consumer, it uses this information to match the consumer to an advertiser. Without that information, it shows consumers an ad of the average best match, which is that of the advertiser located at the same position as the website itself.⁷ Consumers derive disutility from the loss of privacy associated with the collection of personal data. Consumers are heterogeneous in the size of this disutility, θ , and we assume that the distribution of privacy preferences is independent of the consumers' location on the circle. Websites cannot observe the costs θ of individual consumers.

In terms of information, we assume that websites choose their targeting technology ρ and that this choice is verifiable for the regulator, but not observable to consumers. Advertisers are able to verify whether their advertisement was targeted or not, but they do not observe ρ directly.⁸

2.1. Consumers

We consider a unit mass of consumers uniformly distributed along the Salop circle of horizontal consumer preferences, parametrized by location $x \in [0, 1]$. Consumers subscribe to a single website (i.e. we consider single-homing on the part of consumers). The utility that a consumer located at x derives from subscribing to website i located at position x_i and charging price p_i equals⁹

$$u_i(x, \theta) = w - t|x - x_i| - p_i - \bar{\rho}_i \theta$$

⁷The use of Salop's circular city model rather than the Hotelling line greatly simplifies the location of average advertiser.

⁸The verification by advertisers can in practice be done by several mechanisms, like pay-per-view, pay-per-click, or by the conversion rate (share of advertisement exposures that result in a click-through to the advertisers own website). More details on such technology can be found in the annex of this paper.

⁹expressed in terms of some numeraire good

where $|x - x_i| \leq 1/2$ is the distance between the consumer and the website along the circle, t is a travel cost parameter measuring the disalignment between the consumer's location and the location of the website, w is the gross utility of consuming the website's service, which is assumed the same for all websites and consumers.

If a consumer expects that a website has collected personal information, his utility is reduced by θ .¹⁰ This privacy cost θ may measure discomfort of being targeted, but might also include expectations of monetary loss from price discrimination by advertisers when selling their products, or gains from finding better matched products.¹¹ We assume that θ is independently and identically distributed among all consumers, with distribution function $F(\theta)$ that has continuous density on $[\theta_L, \theta_H]$. The consumer expects that the website has collected his personal information with probability $\bar{\rho}_i$, so that $\bar{\rho}_i\theta$ is the consumer's expectation of the costs of privacy loss when subscribing to website i . We focus on rational expectations equilibrium, where the consumers' expectation of the websites' targeting choice corresponds with the actual equilibrium targeting choice of the website, i.e. $\bar{\rho}_i = \rho_i$ in equilibrium.

2.2. Advertisers

We have a continuum of price-taking advertisers, selling products that can be uniquely located at coordinate $y \in [0, 1)$ on the same preference circle as consumers and websites. Advertisers enjoy surplus when consumers are exposed to their advertisements on websites.¹² The size of this surplus, and hence the advertiser's willingness to pay, depends on the quality of the match between consumer and advertiser.

We parametrize the quality of a match by the arc distance between a consumer's location

¹⁰In addition, consumers could have extra direct (dis-)utility of viewing a targeted advertisement relative to viewing an untargeted advertisement. We assume that θ captures the net effect of privacy loss and viewing a targeted advertisement. Note that all consumers are shown exactly one advertisement, which is either targeted or non-targeted. Any generic disutility of advertisements therefore affects all consumers alike.

¹¹As is well known, price discrimination may have ambiguous effects on welfare (see e.g. Varian, 1989), which we cannot address here since we do not explicitly model the interactions of advertisers with consumers. We leave this as a topic for future research.

¹²We do not explicitly model sales of products. We assume there is some form of competition among suppliers in differentiated products, and an advertisement may give one supplier a competitive advantage, for instance by increasing the consumer's demand for his product, or by directing the consumer first to his product in a model of directed search.

x and the advertiser's location y , as in Wolinsky (1983), and more recently in Chandra (2009) and De Corniere (Forthcoming). An advertiser y 's perfect match is therefore with the consumer located at $x = y$, generating a match surplus ν . For a match with a less aligned consumer, the advertiser's surplus equals $a(|x - y|)$, with $a(\cdot) > 0$ a decreasing function of the arc distance between consumer and advertiser. This decrease represents the costs of a mismatch. We denote the value of a perfect match by $\nu = a(0)$.

Advertisers contract with the websites to have their advertisement shown to (some of) the platforms' consumers. Each time a consumer visits a website, the website has the opportunity to place one advertisement. We assume that this advertising space is sold at a price that extracts all rents of the advertiser - advertisers are price takers.

2.3. Websites

The n websites are located on the circle at equal distances $1/n$ from each other. Website i sets the subscription price p_i for consumers, and sells ad space aimed at each individual consumer to the advertisers. Since consumers single-home, websites are competitive bottlenecks and extract monopoly surplus from the advertisers (as in Armstrong, 2006). We focus on the case where all consumers subscribe to a website.

We model the choice of targeting intensity by allowing the websites to choose the proportion ρ_i of their customers for which they collect personal information. For simplicity, we assume that information collection is costless; adding a cost does not materially affect the analysis.¹³ When a website has personal information on a consumer it can identify his true position on the line x with probability one. We assume that the probability that a website has personal information on a customer is independent from his location x .¹⁴

For those customers whose position the website can exactly identify, it will sell the available advertisement space to the matching advertiser at $y = x$, at price ν , the advertiser's willingness to pay for an exact match.

For the $1 - \rho_i$ consumers of website i 's content that the website cannot target, the website

¹³In appendix B, we show how inclusion of a targeting cost affects the main results of section 3.

¹⁴Websites have no prior information that allows them to stratify customers before collecting information.

cannot find perfectly matching advertisers. It does have some information: those consumers' choice of visiting website i rather than another website signals that these consumers are likely to be close to website i . The website will sell advertising space to that advertiser that will bring the greatest match surplus contingent on consumers visiting website i . In a symmetric situation, this best match will be the advertiser located at the same position as the website, $y = x_i$. The price charged for this advertising space will be the advertiser's surplus averaged over all consumers that visit this website.

For a given choice of ρ_i , website i 's expected surplus from selling an advertisement to a consumer located at distance $|x - x_i|$ from the platform's location will then be equal to:

$$a(x - x_i, \rho_i) \equiv \rho_i \nu + (1 - \rho_i) a(x - x_i).$$

The timing of the model now is that first websites choose targeting intensities ρ_i and prices p_i . They contract on a per-consumer price with advertisers who can verify whether the advertisement was well-targeted. After all websites have set prices p_i and targeting intensities ρ_i , consumers observe prices p_i and form expectations over ρ_i . They then choose which website to visit. The marginal consumer who is indifferent between website i and $i + 1$ will be located at a distance

$$d_i = \frac{1}{2n} + \frac{p_{i+1} - p_i + (\bar{\rho}_{i+1} - \bar{\rho}_i)\theta}{2t}$$

from website i .

We assume that in the equilibrium, prices p_i are positive. In a free subscription model, websites would need another strategic parameter, such as quality, to compete for consumers. Depending on the nature of investments in quality, one might identify such quality choices with a negative contribution to prices.

3. Targeting and competition

We first consider the competitive equilibrium without regulation. We show that the degree of targeting chosen by the websites affects the intensity of competition among those websites: more precise targeting leads to more intense competition. Yet, *ex ante*, individual websites find it a dominant strategy to increase their level of targeting ρ_i .

To see this, consider a symmetric equilibrium with all websites choosing the same targeting accuracy ρ and price p . Since profits for an individual website i equal

$$\pi_i = \int_{-d_i}^{d_i} (p + a(|x|, \rho)) dx,$$

where d_i is the distance between the website's most remote customer and the website itself.¹⁵ In the symmetric equilibrium this leads to the first-order conditions for prices p_i ,¹⁶

$$p_i + a\left(\frac{1}{2n}, \rho\right) = \frac{t}{n}.$$

Substituting the expression for p_i into the website's profit we find

$$\pi_i = \frac{t}{n^2} + \frac{1}{n} \left(\bar{a}\left(\frac{1}{2n}, \rho\right) - a\left(\frac{1}{2n}, \rho\right) \right) \quad (1)$$

where

$$\bar{a}(d) = \frac{1}{2d} \int_{-d}^d a(|x|) dx, \quad \text{and} \quad \bar{a}(d, \rho) = \rho \nu + (1 - \rho) \bar{a}(d)$$

is the average advertising income over all consumers subscribing to website i , without targeting ($\rho = 0$) and with targeting ρ , respectively.

For $\rho < 1$, advertising income $a(d, \rho)$ is decreasing with the marginal consumer's distance d from the website, the average is higher than the marginal value $\bar{a}(d, \rho) - a(d, \rho) = (1 - \rho)(\bar{a}(d) - a(d)) > 0$, and profits are larger than in the case of perfect targeting $\rho = 1$. In terms of the analysis of Crampes, Haritchabalet and Jullien (2009), this is a case of decreasing returns to

¹⁵Normalizing the website's location $x_i = 0$; the absolute location of the website does not affect its profits.

¹⁶We do assume that prices are positive, i.e. t/n is large compared to advertising income.

audience size.

Finally, since the difference between the marginal and the average advertising revenues decreases with accuracy of targeting ρ , we have the result that

Proposition 1 *In the symmetric equilibrium, profits are decreasing in the degree of targeting ρ .*

The intuition is that an increase in targeting precision ρ increases the value to the website of marginal consumers. Without targeting, advertising impressions on marginal consumers are of low value, since advertisements are tailored to the average consumer. This means that in that case, an additional marginal consumer does not allow the website to increase his revenues from advertisers very much. With targeting, this changes: with perfect targeting, the marginal consumer is as valuable to the website as any inframarginal consumer, as each consumer is linked with its optimal advertiser, providing the website with advertising income ν . Targeting takes away the mismatch between consumer and advertiser, and this mismatch is greater for marginal consumers. This increase in value of the marginal consumer, in turn, heats up competition between adjacent websites for this consumer, and as a result profits go down.

If websites could coordinate on targeting, proposition 1 suggests that they might want to agree to keep targeting to a minimum. However, we next show that individually, websites win by increasing the accuracy of targeting over that of their competitors, so that in the non-cooperative equilibrium, maximal targeting results.¹⁷

Proposition 2 *Websites gain by increasing the accuracy of targeting ρ above that of their rivals. As a result, in the symmetric equilibrium all websites will choose maximum allowed targeting.*

The websites' problem is that they cannot commit toward their consumers on the level of targeting, ρ_i . As a consequence of the strategic interaction outlined in proposition 2, they are caught in a prisoners' dilemma. The only equilibrium is where each website chooses the

¹⁷The proof, and those of subsequent results, are in the appendix.

maximum level of targeting: more targeting is individually profitable as consumers do not observe changes in targeting, so cannot respond by changing their decision in response to such an individual increase of targeting.

Maximum targeting is undesirable from the point of view of the websites: from proposition 1, the websites' profits are smallest at maximum ρ , and they would be better off when targeting is not possible at all. But also social welfare may be suboptimal at this equilibrium. Although the improved matching of advertisers to consumers under increasing targeting raises advertiser surplus and hence websites' revenues from advertisers – and this benefit is fully passed through to consumers in the form of lower website prices – consumers who experience privacy costs with increased targeting may lose out.

Intervention by a planner may therefore be welfare improving if the average costs of privacy loss are sufficiently large.¹⁸ Let us here first consider the welfare effects of a simple privacy regulation that puts a maximum on the extent of targeting allowed by websites, $\rho_{max} \leq 1$. Such a regulation provides a credible commitment for websites to keep targeting accuracy ρ at this bound, and is therefore clearly beneficial for them.

The analysis of the effect on consumer surplus is also straightforward in the model. Since both price and privacy costs are linear in ρ , consumers prefer either full targeting, $\rho = 1$, or no targeting at all, $\rho = 0$. Average surplus per consumer is given by¹⁹

$$CS = w - p - \rho\tilde{\theta}$$

where $\tilde{\theta}$ is the average value of privacy costs θ . With price p depending on the marginal value of advertising income

$$p = \frac{t}{n} - a\left(\frac{1}{2n}; \rho\right), \quad (2)$$

and $a(\cdot; \rho) = \rho\nu + (1 - \rho)a(\cdot)$, we find that the optimal value of ρ depends only on the sign of $\nu - a(\frac{1}{2n}) - \tilde{\theta}$.

¹⁸In addition, there will be costs of enforcement, which we ignore in this analysis, but which in a full analysis would be traded off against the welfare benefits of regulation.

¹⁹ignoring average transportation costs, which are independent of ρ . Also, recall that we assume full coverage, or w high enough that all consumers participate.

Total surplus, in turn, depends not on the marginal value of advertising, but on the average value per consumer, and its trade-off with average privacy costs. Summarizing,

Corollary 1 1. *Producer surplus is maximized for $\rho = 0$.*

2. *Consumer surplus is maximized at $\rho = 0$ iff $\nu - a(\frac{1}{2n}) < \tilde{\theta}$. Otherwise, $\rho = 1$ maximizes CS.*

3. *Total surplus is maximized at $\rho = 0$ iff $\nu - \bar{a}(\frac{1}{2n}) < \tilde{\theta}$. Otherwise, $\rho = 1$ maximizes total surplus.*

If average privacy costs $\tilde{\theta}$ are sufficiently high, a welfare maximizing regulator takes into account consumer costs of privacy and restricts targeting. In this way it helps the websites circumventing their commitment problem, raising profits and increasing welfare at the same time. Consumers themselves are less averse to targeting: while they incur privacy costs, they benefit from the effect of targeting on competition intensity, lowering prices.²⁰

Increasing the number of websites n not only reduces prices and profits through the usual channel of increased competition among websites. It also improves advertising efficiency without targeting, as consumers' platform choice now more accurately reveals their location along the circle. Hence, both the average match value with advertisers, $\bar{a}(\frac{1}{2n})$, and the marginal match value $a(\frac{1}{2n})$, are larger, which increases advertising income and decreases prices. This means that the relative value of any targeting carried out by the websites (increasing ρ) decreases, and high-targeting will be optimal less often.

In our benchmark model, we have made the assumption that websites cannot commit to a targeting intensity ρ_i . Relaxing that assumption, let us briefly explore a model where consumers can reliably observe the ρ_i set by each website.²¹ In that case, the targeting intensity acts

²⁰The fact that the welfare maximizing regulator treats prices as welfare neutral relies on our assumption that all consumers participate in the market, so that there is no deadweight loss as a result of high prices. With price-elastic consumption, a welfare maximizing regulator would choose ρ weakly higher to reduce prices and increase consumption.

²¹Though for consumers it will be hard to directly verify the level of information collection by an individual website, websites might create a reputation for non-targeting. The internet search engine DuckDuckGo, for instance, emphasizes that it does not use historical search profiles, thus making itself more attractive to consumers that value their privacy more. Relatedly, websites might opt for private certification of their privacy policies, for instance by cooperating with do-not-track-me services.

as another strategic parameter that influences consumer decisions, apart from price p_i . The websites optimize their choice of targeting intensity taking into account consumer response, unlike in the case when ρ is not observable.

Solving the Nash equilibrium in ρ_i in addition to price p_i , we straightforwardly find that, in the symmetric equilibrium, websites choose ρ_i to maximize the difference of average advertizing surplus and average consumers' privacy costs,

$$\rho_i = \arg \max_{\rho \in [0,1]} \left[\bar{a}\left(\frac{1}{2n}, \rho\right) - \rho \tilde{\theta} \right].$$

This result is familiar from the theory of two-part tariffs: websites maximize joint surplus with their consumers, and use the fixed part of the tariff, in our case price p_i , to extract that surplus to the extent allowed by competition.

In our linear model, this entails an equilibrium choice $\rho_i = 1$ as long as $\nu - \bar{a}(\frac{1}{2n}) > \tilde{\theta}$, and $\rho_i = 0$ otherwise. Note that this exactly coincides with the total surplus maximizing choice, as both the website and the welfare maximizer are now aligned in maximizing total surplus from targeting. Consumers, as before, prefer higher ρ , as they are concerned with price levels, and take into account the higher surplus they obtain as higher targeting intensity drives competition up and prices down.

Likewise, even though websites now can commit vis-a-vis consumers, they still choose excessive targeting from a profit maximization point of view, as they ignore the effect their targeting choice has on the intensity of competition, and hence on profits. In this sense, the ordering of targeting preferences is preserved if we allow for commitment on ρ : websites would prefer lower targeting than the welfare optimal level, while consumers in contrast favour higher levels of targeting. It is the effect on competition intensity that drives this wedge in preferences, as in our base case analysis. If websites could also commit to a value of ρ vis-a-vis their rivals, they would always prefer coordination on $\rho = 0$, as this reduces competition and increases prices just as in our base case.

Finally, one might ask whether the websites might use the information they have on consumers to offer customized prices to consumers, instead of offering uniform prices as we assume

here. Such pricing based on consumer characteristics has been studied in a literature going back to the seminal paper by Thisse and Vives (1988), see also Fudenberg and Villas-Boas (2006) for a survey. As stressed in these papers, in price competition models such as the one we study, customised prices essentially make firms compete à la Bertrand over each individual consumer, as they do not have to trade off any losses on inframarginal consumers when they lower prices to attract the marginal consumers. Equilibrium prices and profits in such models then typically go down if firms use customer information to price discriminate. In this paper, we assume that the websites do succeed in avoiding that equilibrium, for instance because they can more easily observe each other's pricing policies and can collude on refraining from such customised pricing.²² In a model with customised prices, our main results would not continue to hold, as they hinge on the observation that more targeting intensifies competition on the marginal consumer and drag down prices for inframarginal consumers as a result. With customized prices, price competition would already be maximal, and no further gain on competition will arise from targeting the advertising.

4. Segmentation of the market

Consumers benefit from increased targeting via pass-through of higher advertising surplus (driven by the average value of advertising revenue per consumer), as well as from the more intense competition between websites (reflected in the difference between marginal and average advertising revenue). These benefits are equal for all consumers.

In contrast, consumers are heterogeneous in the costs they experience from loss of privacy that goes hand-in-hand with improved targeting. Those with high privacy costs θ may prefer a lower level of targeting, while those with lower θ will value the benefits higher than the costs.

It therefore makes sense to explore privacy regulations that allow for differentiation in targeting between low- and high-cost consumers. Again, the problem for the website is that

²²An alternative would be to use a model where consumer prices are zero, and websites compete on quality parameters that cannot easily be differentiated among consumers.

of commitment. The websites themselves cannot tailor their levels of targeting to consumer preferences directly, since that level is not observable for consumers. The planner, on the other hand, may help by setting and enforcing a menu of maximum levels for targeting that consumers can choose from.

Let us assume that the planner can set two levels of privacy: one where websites first ask consent (and which consumers can opt out of, e.g. cookies, do-not-track), ρ^{min} . And second, a maximum level of targeting for those consumers who are willing to give up some privacy in return for lower prices or better quality, ρ^{max} . Government enforced maximum targeting levels allow websites to credibly offer two vertically differentiated products to consumers, one with a government-enforced high level of privacy protection (low targeting), and one which allows a higher degree of targeting by the website.

Consistent with the websites' incentives to increase their level of targeting up to the bound set by the planner, from proposition 2, websites will thus compete in two, vertically differentiated offers, each with a different price. Consumers opting for the high privacy product experience minimal targeting ρ^{min} , and pay price p^h , while those opting for the low privacy subscription will pay the (lower) price p^l , and be exposed to greater targeting ρ^{max} . Hence, this leaves consumers with utility

$$u_i^h(\theta) = w - p_i^h - \rho^{min}\theta - \text{travel costs}, \quad u_i^l(\theta) = w - p_i^l - \rho^{max}\theta - \text{travel costs}.$$

We denote the marginal consumer privacy type indifferent between website i 's high and low quality products by $\bar{\theta}_i$. Clearly,

$$\bar{\theta}_i = \frac{p_i^h - p_i^l}{\rho^{max} - \rho^{min}}$$

and those consumers with low (or even negative) θ prefer the cheaper low privacy option to the more expensive high privacy one. Furthermore, we assume that the range of θ 's is sufficiently large to ensure that the equilibrium $\bar{\theta}$ is not a corner solution. Having $\theta_L < \nu - \bar{a}(\frac{1}{2n})$ and $\theta_H > \nu - a(\frac{1}{2n})$ makes sure of that.

Websites compete on both products, taking into account the elasticity of substitution of

consumers between the low and the high privacy products. We have total profits of website i

$$\begin{aligned}\Pi_i &= F(\bar{\theta}_i) \int_{-d_i^l}^{d_i^l} (p_i^l + a_i^l(|x|; \rho^{max})) dx + (1 - F(\bar{\theta}_i)) \int_{-d_i^h}^{d_i^h} (p_i^h + a_i^h(|x|; \rho^{min})) dx \\ &\equiv F(\bar{\theta}_i) \pi_i^l + (1 - F(\bar{\theta}_i)) \pi_i^h.\end{aligned}$$

These total profits are a weighted average of website i 's per-consumer profits π_i^l on the low-privacy segment (in which all consumers of type $\theta < \bar{\theta}_i$ will self-select) and profits π_i^h on the high-privacy segment.

In the high privacy segment h the website will target with intensity ρ^{min} . This means that on the high privacy segment, we have perfect matching with probability ρ^{min} , while with probability $1 - \rho^{min}$, the advertiser with the best average match is displayed to the user. The same holds on the low privacy segment with targeting intensity ρ^{max} , so that

$$a_i^h(|x|; \rho^{min}) = \rho^{min} \nu + (1 - \rho^{min}) a(|x|), \quad a_i^l(|x|; \rho^{max}) = \rho^{max} \nu + (1 - \rho^{max}) a(|x|).$$

In this two regime set-up we find that the competitive effects of targeting that we found before have spill-over effects among the two regimes:

Proposition 3 *When websites can offer both a high privacy (low targeting at ρ^{min}) product and a low privacy (high targeting at ρ^{max}) product, in the symmetric equilibrium with consumer segmentation, high-privacy consumers benefit from the presence of the low privacy market: prices are lower than without this second market. Similarly, prices for the low-privacy product are higher than they would be in the absence of the high-privacy product.*

With only a single targeting regime, we have profit optimization per consumer in that regime. Now, with two targeting regimes, we have an additional effect: changing prices on one segment not only leads to gains or losses of consumers to rival websites. It also causes some marginal consumers to switch from the low to the high privacy segment on the same website, or vice versa. Since these segments generate different profits per consumer $\pi^{h,l}$, this switching will affect price setting by the website: the website wants to reduce the incentives of

high privacy consumers to switch to the lower priced, and lower profit, low privacy segment. Without such switching, we had, by proposition 1, $\pi^l < \pi^h$: the high-privacy h market is less competitive. When the two markets are both present, and linked through the marginal $\bar{\theta}$ -consumer, we see that in equilibrium prices are reduced on the high-privacy segment, compared to the single targeting regime. And vice versa, prices are raised on the low-privacy segment. Hence, high-privacy users benefit from the stronger competition on the low-privacy segment, and vice versa.

As a next step, we solve explicitly for the resulting prices, given levels of targeting ρ^{min}, ρ^{max} ,

Lemma 1 *Equilibrium prices for high and low privacy products satisfy*

$$p^l + a(d, \rho^{max}) = 2td - (\bar{\theta} - \nu + a(d))\Delta\rho(1 - F) \quad (3)$$

$$p^h + a(d, \rho^{min}) = 2td + (\bar{\theta} - \nu + a(d))\Delta\rho F, \quad (4)$$

with $d = \frac{1}{2n}$ the distance from the website to its marginal consumer, and $\bar{\theta} = \frac{p^h - p^l}{\rho^{max} - \rho^{min}} = \frac{\Delta p}{\Delta \rho}$ satisfying

$$2td(\bar{\theta} - \nu + \bar{a}(d)) + \frac{F(\bar{\theta})(1 - F(\bar{\theta}))}{f(\bar{\theta})}\Delta\rho(\bar{\theta} - \nu + a(d)) = 0. \quad (5)$$

In particular,

$$\nu - \bar{a}(d) < \bar{\theta} < \nu - a(d). \quad (6)$$

Recall that with a single level of targeting, we had $p + a(d, \rho) = 2td$. From the expressions for prices with different levels of targeting, equations (3),(4), we see that p^l is indeed higher than when in isolation, and vice versa p^h is lower, since $\bar{\theta} - \nu + a < 0$. This is consistent with proposition 3.

We can now again turn to the regulator's choice of optimal $\rho^{min, max}$. For that, we first use lemma 1 to write down the explicit expressions for profits, total welfare and consumer surplus.

Lemma 2 *Total website profits, Π , equal*

$$\Pi = 2td + \bar{a}(d) - a(d) - \rho^{max} \int_{\theta_L}^{\bar{\theta}} (\bar{a}(d) - a(d))dF - \rho^{min} \int_{\bar{\theta}}^{\theta_H} (\bar{a}(d) - a(d))dF. \quad (7)$$

with $d = \frac{1}{2n}$. Total welfare is given by

$$TW = w + \bar{a}(d) + \rho^{max} \int_{\theta_L}^{\bar{\theta}} (\nu - \bar{a}(d) - \theta) dF + \rho^{min} \int_{\bar{\theta}}^{\theta_H} (\nu - \bar{a}(d) - \theta) dF \quad (8)$$

and consumer surplus

$$CS = w - 2td + a(d) + \rho^{max} \int_{\theta_L}^{\bar{\theta}} (\nu - a(d) - \theta) dF + \rho^{min} \int_{\bar{\theta}}^{\theta_H} (\nu - a(d) - \theta) dF \quad (9)$$

Now let us explore the combinations of targeting levels ρ^{min}, ρ^{max} that optimize these expressions. In doing so, we have to take into account that changing these targeting levels also changes the marginal consumer's privacy cost $\bar{\theta}$. That changes with $\Delta\rho$ according to equation (5).

As in the single-targeting analysis, we find that welfare maximization involves higher targeting intensity than the profit-maximizing choice for websites, while consumers benefit even more from targeting:

Proposition 4 *Websites optimize their profits when targeting is banned, $\rho^{min} = \rho^{max} = 0$. Total welfare maximization requires positive targeting, $\rho^{max} > 0$. Consumer surplus is maximized by allowing full targeting for the low-privacy segment, $\rho^{max} = 1$.*

The result for websites is immediate: they pass on any gains from advertising in the form of lower prices, and they suffer as before from higher competition as a result of targeting. Total welfare, on the other hand, benefits when adding a choice for some targeting in addition to the no-targeting product. As long as the marginal welfare gain from targeting, $\nu - \bar{a}(d)$, exceeds the privacy costs θ_L of the lowest type consumers, adding targeting is strictly optimal.

For consumers, the gain from targeting is higher: they also take into account the benefits in terms of fiercer competition, driving down prices. In fact, from equation (6), even for the marginal consumer, with privacy cost $\bar{\theta}$, the marginal consumer benefits $\nu - a(d)$ of increased targeting in the high-targeting regime outweigh the privacy cost. As a result, it is optimal to maximally increase the maximum targeting level.

We cannot make similar definitive statements on ρ^{min} . Whether raising ρ^{min} above zero can be optimal will depend on the actual distributions of θ and the values of a and $\bar{a}$, as the following example illustrates.

Example Let us consider a uniform distribution of privacy costs $\theta \in [0, B]$. In that case, equation (5) for the relation between the marginal consumer's privacy costs $\bar{\theta}$ and targeting intensity difference $\Delta\rho$ can be written as

$$\Delta\rho = 2td \frac{B}{\bar{\theta}(B - \bar{\theta})} \frac{\bar{\theta} - \nu + \bar{a}}{\nu - a - \bar{\theta}}$$

Maximizing consumer surplus now is equivalent to

$$\max B\rho^{max} \int_0^{\frac{\bar{\theta}}{B}} \left(\frac{\nu - a}{B} - \frac{\theta}{B} \right) d\left(\frac{\theta}{B} \right) + B\rho^{min} \int_{\frac{\bar{\theta}}{B}}^1 \left(\frac{\nu - a}{B} - \frac{\theta}{B} \right) d\left(\frac{\theta}{B} \right)$$

over ρ^{max} and $\Delta\rho$. Note that, as $\Delta\rho$ increases from 0 to 1, $\bar{\theta}$ increases monotonically from $\nu - \bar{a}$ to some intermediate value θ^* between $\nu - \bar{a} < \theta^* < \nu - a$. Hence optimization over $\Delta\rho$ is equivalent to optimization over the fraction of low privacy consumers $\frac{\bar{\theta}}{B} \in [\frac{\nu - \bar{a}}{B}, \frac{\theta^*}{B}]$.

By proposition 4, we have that $\rho^{max} = 1$ in the optimum. Writing $\rho^{min} = 1 - \Delta\rho$, and normalizing to $B = 1$, we need to optimize

$$\tilde{CS} = -\Delta\rho(\bar{\theta}) \int_{\bar{\theta}}^1 (\nu - a - \theta) d\theta.$$

It is convenient to do the equivalent maximization with respect to $\bar{\theta}$ (over the range $[\nu - \bar{a}, \theta^*]$), rather than over $\Delta\rho$ itself. Taking the derivative gives

$$\frac{d\tilde{CS}}{d\bar{\theta}} = -\frac{d\Delta\rho}{d\bar{\theta}} \int_{\bar{\theta}}^1 (\nu - a - \theta) d\theta + \Delta\rho(\nu - a - \bar{\theta}).$$

The second term is always positive. For the first, it depends on the sign of the integral. If $\nu - a$ is large, so that there are large benefits of targeting, the optimum will be at $\Delta\rho = 0$, i.e. $\rho^{min} = 1$ as well. Conversely, if $\nu - a$ is small, the optimum consumer surplus will be attained

at $\Delta\rho = 1$, with $\bar{\theta} = \theta^*$. In that case, consumer surplus is maximized by having a full privacy product, $\rho^{min} = 0$.

Let us also investigate the total welfare maximizing choices for ρ^{max}, ρ^{min} in this example. Looking at the total welfare expression (8), and writing down the first-order conditions for both ρ^{min} and ρ^{max} , it is immediately clear that these cannot both hold with equality: if $\nu - \bar{a} < \frac{1}{2}B$, the average consumer privacy costs, necessarily we have a corner solution $\rho^{min} = 0$. And vice versa, for $\nu - \bar{a} > \frac{1}{2}B$, we find $\rho^{max} = 1$. We can then solve for the other parameter to optimize welfare. With $B = 1$, and taking as an example the former case, $\nu - \bar{a} < \frac{1}{2}$ – which means that even without targeting a lot of advertising income can be realized – we have $\rho^{min} = 0$ and $\rho^{max} = \Delta\rho$ optimizes

$$TW = \Delta\rho(\bar{\theta}) \int_{\bar{\theta}}^1 (\nu - \bar{a} - \theta) d\theta.$$

With large difference between average and marginal advertising income, $\bar{a}$ and a , increasing ρ means increasing $\bar{\theta}$ by a lot, which means exposing many higher cost consumers to maximal targeting. In that case, optimally, $\rho^{max} = \Delta\rho$ is smaller than one, and we have positive but limited targeting ρ^{max} for the low type consumers. If, on the other hand, the difference between marginal and average advertising costs is small, optimization results in $\rho^{max} = 1$.

Summarizing, we find that allowing targeting on consumers that opt for targeting will be good for total welfare. From a consumer welfare perspective, a higher targeting intensity is preferred, since apart from better matches, also the increased competition among websites contributes to consumer welfare. This higher targeting on consumers who place low value on privacy also drives down prices on the high privacy segment, as competition effects spill over to those consumers.

5. Conclusions

In this paper we explored the interaction between competition among internet platforms and the degree of ad targeting they use. More targeting implies stronger competition. Yet, since

websites cannot commit to low targeting intensity, they are caught in a prisoners' dilemma: each firm individually benefits from increased targeting. In the equilibrium, websites will therefore drive up targeting. On the one hand, this reduces consumer prices, because of improved matching of consumers with advertisers. However, if consumers dislike the loss of privacy that is a consequence of targeting, privacy policy can lead to better outcomes than the laissez-faire outcome. In that case, also websites can benefit from the less intense competition that goes with this commitment to privacy protection.

In practice, consumers are heterogeneous in the costs they associate with loss of privacy. By allowing websites to offer multiple products, differing in the degree of targeting and price they offer, welfare can be increased. In this case, even those consumers that opt for the high privacy (and low targeting) product benefit: their prices are reduced as a result of the endogenously higher competition on the low privacy market segment.

Our paper provides a general discussion of welfare trade-offs in the presence of heterogeneous privacy concerns among consumers and websites with market power. Potential extensions could provide a more elaborate analysis of the welfare effects of private certification of targeting behaviour, the impacts of scale effects in consumer targeting, and the public costs of enforcing privacy policies.

References

- Anderson, Simon P and Stephen Coate. 2005. "Market provision of broadcasting: A welfare analysis." *The Review of Economic Studies* 72(4):947–972.
- Armstrong, Mark. 2006. "Competition in two-sided markets." *The RAND Journal of Economics* 37(3):668–691.
- Campbell, James, Avi Goldfarb and Catherine Tucker. 2015. "Privacy Regulation and Market Structure." *Journal of Economics and Management Strategy* 24(1):47–73.
- Chandra, Ambarish. 2009. "Targeted Advertising: The Role of Subscriber Characteristics in Media Markets." *The Journal of Industrial Economics* 57(1):58–84.
- Chen, Yuxin, Chakravarthi Narasimhan and Z John Zhang. 2001. "Consumer heterogeneity and competitive price-matching guarantees." *Marketing Science* 20(3):300–314.
- Chen, Yuxin and Ganesh Iyer. 2002. "Consumer addressability and customized pricing." *Marketing Science* 21(2):197–208.
- Crampes, Claude, Carole Haritchabalet and Bruno Jullien. 2009. "Advertising, Competition and Entry in Media Industries." *The Journal of Industrial Economics* 57(1):7–31.
- De Corniere, Alexandre. Forthcoming. "Search Advertising." *American Economic Journal: Microeconomics* .
- Evans, David S. 2009. "The Online Advertising Industry: Economics, Evolution, and Privacy." *The Journal of Economic Perspectives* 23(3):37–60.
- Fudenberg, Drew and J Miguel Villas-Boas. 2006. "Behavior-based price discrimination and customer recognition." *Handbook on economics and information systems* 1:377–436.
- Gal-Or, Esther and Mordechai Gal-Or. 2005. "Customized advertising via a common media distributor." *Marketing Science* 24(2):241–253.
- Goldfarb, Avi. 2014. "What is different about online advertising?" *Review of Industrial Organization* 44(2):115–129.

- Goldfarb, Avi and Catherine Tucker. 2011. "Online display advertising: Targeting and obtrusiveness." *Marketing Science* 30(3):389–404.
- Goldfarb, Avi and Catherine Tucker. 2012. "Shifts in Privacy Concerns." *American Economic Review: Papers and Proceedings* 102:349–353.
- Iyer, Ganesh, David Soberman and J Miguel Villas-Boas. 2005. "The targeting of advertising." *Marketing Science* 24(3):461–476.
- Listokin, Siona. 2015. "Industry self-regulation of consumer data privacy and security." *The John Marshall Journal of Information Technology and Privacy Law* 32(1).
- Salop, Steven C. 1979. "Monopolistic Competition with Outside Goods." *Bell Journal of Economics* 10(1):141–156.
- Taylor, Curtis and Liad Wagman. 2014. "Consumer Privacy in Oligopolistic Markets: Winners, Losers, and Welfare." *International Journal of Industrial Organization* 34:80–84.
- Thisse, Jacques-Francois and Xavier Vives. 1988. "On the Strategic Choice of Spatial Price Policy." *American Economic Review* 78(1):122–137.
- Tucker, Catherine E. 2012. "The economics of advertising and privacy." *International Journal of Industrial Organization* 30(3):326–329.
- Turow, Joseph, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley and Michael Hennessy. 2009. "Americans reject tailored advertising and three activities that enable it." *Working paper*.
- Varian, Hal R. 1989. "Price discrimination." *Handbook of industrial organization* 1:597–654.
- Wolinsky, Asher. 1983. "Retail Trade Concentration Due to Consumers' Imperfect Information." *Bell Journal of Economics* 14(1):275–282.

A. Proofs

Proof of proposition 2 Consider a hypothetical symmetric equilibrium characterized by ρ, p for all websites, and consumer expectations $\bar{\rho} = \rho$. In that case, since market share $2d_i$ does not directly depend on ρ_i (for consumers, only their expectation $\bar{\rho}$ determines their choice of website), we have

$$\frac{\partial \pi_i}{\partial \rho_i} = \frac{\partial}{\partial \rho_i} 2d_i(p + \bar{a}(d_i, \rho)) = 2d_i(\nu - \bar{a}(d_i)) > 0$$

and hence a symmetric equilibrium must have maximum allowed ρ . *Q.E.D.*

Proof of proposition 3 With profits

$$\Pi_i = F(\bar{\theta}_i)\pi_i^l + (1 - F(\bar{\theta}_i))\pi_i^h,$$

and the marginal consumer's privacy cost $\bar{\theta}$ given by

$$\bar{\theta}_i = \frac{p_i^h - p_i^l}{\rho^{max} - \rho^{min}},$$

we can compute first-order conditions for prices

$$\frac{\partial \pi_i^l}{\partial p_i^l} = \frac{f}{F} \frac{\pi_i^l - \pi_i^h}{\rho^{max} - \rho^{min}} \tag{10}$$

$$\frac{\partial \pi_i^h}{\partial p_i^h} = -\frac{f}{1 - F} \frac{\pi_i^l - \pi_i^h}{\rho^{max} - \rho^{min}}. \tag{11}$$

Following our previous analysis, competition on the high-privacy market h is less intense, so that $\pi^l < \pi^h$. Hence, in the two-regime market,

$$\frac{\partial \pi_i^l}{\partial p_i^l} < 0 \quad \frac{\partial \pi_i^h}{\partial p_i^h} > 0$$

so that $p^h < p^{h*} = t/n - a(d; \rho^{min})$ and $p^l > p^{l*} = t/n - a(d; \rho^{max})$, where stars denote the price for a single market with level of targeting ρ^{min} or ρ^{max} , given in equation (2). Hence, high-privacy users benefit from the stronger competition on the low-privacy segment, and vice

versa.

Proof of lemma 1 On each segment h, l we have per consumer profits for website i :

$$\pi_i^{l,h} = \int_{-d_i}^{d_i} dx (p_i^{l,h} + a_i^{l,h}(|x|, \rho)) = 2d_i(p_i^{l,h} + \bar{a}_i^{l,h}(d_i, \rho))$$

and the derivative,

$$\frac{\partial \pi_i^{l,h}}{\partial p_A^{l,h}} = 2d_i - \frac{1}{t}(p_A^{l,h} + a_A^{l,h}(d_i, \rho)).$$

Now, substituting into the first-order conditions for low and high prices, equations (10,11), we get in the symmetric equilibrium with $d_i = d = \frac{1}{2n}$,

$$\begin{aligned} \frac{\partial \pi_i^l}{\partial p_i^l} &= 2d - \frac{1}{t}(p^l + a^l(d)) = \frac{f}{F\Delta\rho} 2d(p^l - p^h + \bar{a}_l - \bar{a}_h) \\ \frac{\partial \pi_i^h}{\partial p_i^h} &= 2d - \frac{1}{t}(p^h + a^h(d)) = -\frac{f}{(1-F)\Delta\rho} 2d(p^l - p^h + \bar{a}_l - \bar{a}_h), \end{aligned}$$

and subtracting these, we find the condition for the privacy costs $\bar{\theta}$ of consumer who is indifferent between high and low privacy products,

$$2td(\bar{\theta} - \nu + \bar{a}) + \frac{F(1-F)}{f}\Delta\rho(\bar{\theta} - \nu + a) = 0.$$

Here we used that $p^h - p^l = \bar{\theta}\Delta\rho$, and $a^h = \rho^{min}\nu + (1 - \rho^{min})a$, $a^l = \rho^{max}\nu + (1 - \rho^{max})a$, with $\Delta\rho = \rho^{max} - \rho^{min}$.

From that condition for $\bar{\theta}$, and $\bar{a} > a$, it directly follows that

$$\nu - \bar{a}(d) < \bar{\theta} < \nu - a(d).$$

Finally, we solve for p^h and p^l separately, plugging in the condition for $\bar{\theta}$ into the first-order conditions. This gives equations (3,4). Q.E.D.

Proof of proposition 4 It is clear from the expression for total profits Π , (7), that from the point of view of the websites, banning targeting altogether ($\rho^{max} = 0 = \rho^{min}$) maximizes profits.

The reason is similar as before: this maximizes the wedge between the advertising income of the marginal consumer $a(d)$, and the average advertising income exceeds the marginal one, $\bar{a}(d) > a(d)$.

From a welfare point of view, allowing some targeting is always desirable. Consider the derivative of total welfare to the targeting intensity ρ^{max} ,

$$\frac{\partial TW}{\partial \rho^{max}} = \int_{\theta_L}^{\bar{\theta}} (\nu - \bar{a} - \theta) dF + \Delta \rho (\nu - \bar{a} - \bar{\theta}) f(\bar{\theta}) \frac{d\bar{\theta}}{d\Delta \rho}.$$

Note that at $\rho^{max} = \rho^{min} = 0$, we have from lemma 1 that $\bar{\theta} = \nu - \bar{a}$ since $\Delta \rho = 0$. Hence, the first term is positive at that point, while the second vanishes. The first-order effect of raising ρ^{max} above zero is therefore positive.

Since $\bar{\theta} < \nu - a$ for any value of $\Delta \rho$, the first integral expression in the corresponding equation for consumer surplus (9),

$$\frac{\partial CS}{\partial \rho^{max}} = \int_{\theta_L}^{\bar{\theta}} (\nu - a - \theta) dF + \Delta \rho (\nu - a - \bar{\theta}) f(\bar{\theta}) \frac{d\bar{\theta}}{d\Delta \rho}.$$

is always positive and raising ρ^{max} above zero certainly increases consumer surplus. In fact, since $\bar{\theta}$ increases with $\Delta \rho$, it is optimal for consumers to set ρ^{max} equal to one. The reason is that for all consumers opting for the low privacy product, the gain from lower prices outweighs the loss in privacy. *Q.E.D.*

B. Results with costly targeting

Here we explore the effects of costly targeting. Suppose websites incur costs $c\rho$ when choosing targeting intensity ρ , so that total profits are

$$\pi_i = \int_{-d_i}^{d_i} (p + a(|x|, \rho) - c\rho) dx.$$

First-order conditions for prices p_i in the symmetric equilibrium now become

$$p_i = \frac{t}{n} - a\left(\frac{1}{2n}, \rho\right) + c\rho,$$

so in equilibrium, costs are reflected in prices. The website's profits, given the level of targeting ρ , are therefore unaltered from the zero cost result,

$$\pi_i = \frac{t}{n^2} + \frac{1}{n} \left(\bar{a}\left(\frac{1}{2n}, \rho\right) - a\left(\frac{1}{2n}, \rho\right) \right)$$

and proposition 1 remains unaffected by the introduction of targeting costs. If a website increases its targeting intensity, keeping prices and consumer expectations fixed, there is now a cost in addition to the benefits of increased advertising income (as in proposition 2). As long as the difference between average advertising surplus with and without targeting, $\nu - \bar{a}(d)$, is higher than the marginal cost of targeting c , websites gain from increased targeting and proposition 2 remains intact. With high costs of targeting, an individual website is better off without any targeting, so $\rho = 0$ in the symmetric equilibrium and the prisoner's dilemma is not present.

On average, consumers compare the effects of price with loss of privacy, $\rho\tilde{\theta}$, where $\tilde{\theta}$ is the average privacy cost parameter. With targeting costs c included in the prices, consumers on the whole prefer full targeting as long as $\nu - a(\frac{1}{2n}) > \tilde{\theta} + c$, while total surplus is maximized at $\rho = 1$ as long as $\nu - \bar{a}(\frac{1}{2n}) > \tilde{\theta} + c$. Corollary 1 therefore remains valid as long as privacy costs θ are replaced with $\theta + c$.