Online Advertising and Privacy*

Alexandre de Cornière[†] and Romain De Nijs[‡]

October 7, 2015

Abstract

An online platform auctions an advertising slot. Several advertisers compete in the auction, and consumers differ in their preferences. Prior to the auction, the platform decides whether to allow advertisers to access information about consumers (disclosure) or not (privacy). Disclosure improves the match between advertisers and consumers but increases product prices, even without price-discrimination. We provide conditions under which disclosure or privacy is privately and/or socially optimal. When advertisers compete on the downstream market, disclosure can lead to an increase or a decrease in product prices depending on the nature of the information.

Keywords: online advertising, privacy, information disclosure, auctions.

JEL Classification: D4, L1, M37.

*We thank Bernard Caillaud, Chris Dellarocas, Gabrielle Demange, Peter Eso, Philippe Février, Benjamin Hermalin, Nabil Kazi-Tani, Frédéric Koessler, Philippe Jéhiel, Bruno Jullien, John Morgan, Régis Renault, Jean Tirole and seminar participants at Crest, the Paris School of Economics, the sixth bi-annual "Conference on The Economics of Intellectual Property, Software and the Internet" in Toulouse, the second "Workshop on the Economics of ICT" in Universidade de Evora, the 2011 IIOC conference, the workshop "Communication and Beliefs Manipulation" at PSE, and the Third Annual Conference on Internet Search and Innovation. The first author benefited from funding by the ANR (2010-BLANC-1801 01).

[†]Toulouse School of Economics. Email: adecorniere@gmail.com

[‡]Paris School of Economics, and Ecole des Ponts ParisTech; Email: romaindenijs@gmail.com

1 Introduction

The online advertising industry has grown rapidly in the last decade.[1] One of the two mains forms of online advertising, along with search, is display advertising. The basic organization of the display advertising market is the following: while surfing the internet, users visit various websites (also called *publishers*) who have an inventory of advertising slots to sell. This can be done directly by reaching to potential advertisers, or indirectly, through intermediaries whose main role is to aggregate supply and demand of advertising space and to act as matchmakers.Advertising networks (such as Google Adsense, AdBlade) and advertising exchanges (DoubleClick, OpenX) are such intermediaries.

Intermediaries and large publishers (such as Facebook or Google), which we will designate as *platforms*, have access to technologies that enable them to gather and analyze a considerable amount of data at a very high speed, making it possible to customize advertising using real-time auctions (see, e.g., The Economist (2014)). Advertisers may submit bids that depend, for instance, on the correspondence between the website's content and the advertisement, but also on data about the location of the consumer (obtained through the IP address), his past browsing history (obtained through cookies), or whatever information he gave to the platform or its partners (through subscription questionnaires for instance, or any information posted on his Facebook wall). These new opportunities give firms additional incentives to acquire and use personal information about consumers, which has led regulators and consumers to express worries, or at least to acknowledge some potential pitfalls. Among these are privacy breaches or fraudulent use of personal information, but also practices of behavioral targeting and pricing.

In this paper, we study the decision by a platform of whether to use the information it has gathered about consumers, in order to increase its revenue from advertising. Do such practices have social value? Who benefits most from them?

The situation we have in mind is the following (see Figure 1): consumers visit a platform, which auctions an advertising slot among n advertisers.[2] Consumers are heterogeneous, in the sense that

[1] See Evans (2008) and Evans (2009) for insightful discussions about this industry.

[2] The case of several slots is analyzed in section 5.

they do not derive the same value from consuming advertisers' products. Thanks to its technology, the platform gathers, for each consumer, information correlated to the consumer's willingness to pay for any product. The platform does not know how to interpret the information in terms of implied willingness to pay for different products, but advertisers are able to do it. For instance, the platform knows that the consumer is a young man living in a metropolitan area, but it is not able to infer his willingness to pay for good A or B. On the other hand, firm A knows that young men living in a metropolitan area are especially likely to have a high willingness to pay for its product, whereas firm B offers a product which is less likely to be a good match for such consumers.

In the economic literature, privacy has been defined a "the restriction of the collection or use of information about a person or corporation" (Stigler (1980)).[3] In our set-up, we define privacy as the platform's policy under which it does not disclose consumer information to advertisers. In order to make the analysis as transparent as possible, we assume that consumers do not exhibit intrinsic preferences over their privacy, but care about it insofar as it has economic effects.

We are thus concerned with two main questions: (1) what are the effects of the disclosure policy on market outcomes, that is on the interactions between consumers and advertisers?, and (2) when does the platform provide the efficient amount of privacy?

Regarding (1), we show that disclosure has both positive and negative consequences. On the positive side, when advertisers can condition their bids on information about consumers, the highest bidder, in equilibrium, is the firm that offers the best match, which is efficient. However, when good matches correspond to higher marginal revenues for advertisers, we show that disclosure of personal information leads to higher prices for consumers. Although reminiscent of results by Taylor (2004), Acquisti and Varian (2005), Hermalin and Katz (2006) or Calzolari and Pavan (2006), the latter effect stems from a different logic. The aforementioned papers have the following structure: (i) the seller observes a signal about the consumer's type, either by previous experimentation or by buying information from another firm, then (ii) the seller uses the signal to determine the price (or the menu of contracts) that he offers to the consumer. Thus, in all these papers, personal information is used

[3] See Png and Hui (2006) for a survey of the economics of privacy.

by sellers to price-discriminate among buyers. Such a mechanism is not very plausible in settings in which the identity of the buyer is not verifiable or when there are possibilities of arbitrage among consumers or across channels (offline and online). Moreover, in many instances firms are reluctant to use first- or third-degree price discrimination, for fear of a public relation backlash akin to what happened to Amazon in 2001.[4] Overall, "examples of personalized pricing remain fairly limited." (Council of Economic Advisors (2015)).

In our model, we rule out price-discrimination by assuming that firms choose their prices before learning the information. Once they learn it, they submit a bid and the winner of the auction has its advertisement displayed to the consumer. However we exhibit another channel through which disclosure can lead to higher prices: by allowing firms to condition their bid on consumers' characteristics, the disclosure of information leads to a situation in which firms expect their ads to reach only the consumers with a low price-elasticity of demand (the good matches). Firms then rationally set higher prices ex ante. The magnitude of the price increase depends on the number of bidders. With a large number of bidders winning the auction is more informative than with fewer bidders, so that the products prices (as well as the ad slot price) will be higher when there are many bidders. An increase in the number of slots plays the same role as a decrease in the number of bidders when they do not compete on the product market.

We also study the case in which the platform sells more than one slot and advertisers compete to sell their product (the "downstream competition case"). With downstream competition, disclosure can lead to higher or lower equilibrium product prices. In a discrete choice model, we show that disclosure leads to higher prices when the distribution of willingness to pay conditional on good signals is more concentrated than the unconditionnal distribution.

Regarding our second question, namely whether the platform provides the right amount of information from a social welfare perspective, our analysis partially relies on insights formulated by Ganuza (2004) and Ganuza and Penalva (2010). As in those papers, disclosing information increases

[4]Anderson and Simester (2010) also provide empirical evidence on why firms should be cautious in engaging in price discrimination that may antagonize consumers who then react by making fewer purchases.

the total profits of the industry (platform and advertisers) but comes at the price, for the platform, of leaving an informational rent to the winning bidder. This effect is particularly important when the number of bidders is small.[5] We show that when the number of firms is large, the platform always prefers to disclose information about consumers. Indeed, in that case, the rent left to the winning bidder vanishes. However, and in contrast to Ganuza (2004), such a policy is not necessarily efficient: some consumers are excluded from the market following the increase in the equilibrium price of the goods. Following the approach of Cowan (2007), we give conditions under which privacy or disclosure is optimal when the quality of the match determines a vertical shift of the demand function.

On top of the papers already mentionned, our paper contributes to the following strands of the literature:

Targeted advertising: By disclosing information about the consumer, the platform ensures that a consumer will see the most relevant advertisement, whereas when no information is disclosed (the privacy case) ads are displayed randomly. Some papers in the literature on targeted advertising also find that targeting leads to higher prices, but for different reasons: in Roy (2000), Iyer, Soberman, and Villas-Boas (2005), or Gaelotti and Moraga-Gonzalez (2008), targeting allows firms to segment the market, thereby softening price competition. In Esteban, Gil, and Hernandez (2001), targeting leads to higher prices for cost-related reasons. On the other hand, de Cornière (2013) shows that when consumers actively search for products, targeting leads to more intense competition.

Another branch of the literature on targeted advertising assume the prices of advertised goods are set exogenously and concentrates on other aspects of targeting. Athey and Gans (2010), Bergemann and Bonatti (2011) and Athey, Calvano, and Gans (2012) study how targeting affects advertising markets and competition between online and offline media. Spiegel (2013) consider the trade-off between targeted advertising and loss of privacy in the software market. Van Zandt (2004), Anderson and de Palma (2009), and Johnson (2011) investigate the topic of privacy and congestion with consumers having limited attention.

[5] See also Bergemann and Pesendorfer (2007), Board (2009) and Levin and Milgrom (2010).

Auction design: One of our contributions is to study how the design of an auction, in particular the amount of information revealed to participants, affects pre-auction strategic decisions by participants, here their pricing strategy. Other papers following a related approach include Bergemann and Välimäki (2002), Arozamena and Cantillon (2004) and Piccione and Tan (1996), although the choice of the designer in these papers concerns the format of the auction (e.g., first-price versus second-price) rather than the decision to disclose information.

In section 2, we present the model. In section 3, we characterize symmetric equilibria under privacy and disclosure, and analyze the main implications of either regime as well as the platform's decision. In section 4, we put more restrictions on the model and conduct a normative analysis. In section 5 we extend our results to cases in which the platform has several slots to auction. Section 6 presents some extensions of the basic model, and section 7 concludes. The appendix contains some omitted proofs.

2 Model

There are n advertisers, who compete for a single slot on a platform's website. A continuum of consumers visit the website. A consumer's type is a vector $(\theta_1, ..., \theta_n)$. The θ_i are independent and identically distributed according to a continuous cumulative distribution function F over an interval set $[\underline{\theta}, \overline{\theta}]$. The corresponding density function is f.

If a consumer of type $(\theta_1, ..., \theta_n)$ is matched with firm i, which charges price p_i, firm i's profit is $\pi(p_i, \theta_i) = (p_i - c)D(p_i, \theta_i)$.[6] Depending on the context, $D(p, \theta_i)$ can either be interpreted as the number of units bought or, for markets with unit demand, as the probability that the willingness to pay is higher than p_i, given θ_i. In the latter case, a consumer may see an ad and not buy the product.

Let $P(q, \theta_i)$ be the inverse demand function for a consumer of type θ_i when he faces firm i. Welfare and a consumer's surplus are respectively $W(p_i, \theta_i) = \int_0^{D(p_i,\theta_i)} (P(q, \theta_i) - c)dq$ and $V(p_i, \theta_i) =$

[6] The assumption of a constant marginal cost is not essential but simplifies the notations.

$\int_0^{D(p_i,\theta_i)}(P(q,\theta_i)-p_i)dq$. We make the following assumptions:

Assumption 1 *The demand function D is twice continuously differentiable in both arguments.*

There exists $\overline{p}$ such that for all θ_i, and for all $p_i \geq \overline{p}$, $D(p_i,\theta_i)=0$.

Assumption 2 *π is strictly concave in p_i over $[0,\overline{p}]$. For every θ_i, there exists $p^*(\theta_i) \in [0;\overline{p}]$ such that $\frac{\partial \pi(p^*(\theta_i),\theta_i)}{\partial p_i}=0$.*

Assumptions 1 and 2 are made for analytical simplicity. In particular, they ensure that for any ϕ, the function $\int_0^{\overline{p}} \pi(p,\theta)\phi(\theta)d\theta$ is concave in p, which will allow us to only look at the first-order conditions of the profit-maximization problem. The price $p^*(\theta_i)$ is the price that firm i would charge if it could use first degree price-discrimination, knowing that the consumer's (relevant) type is θ_i. We will refer to it as the *complete information* price.

The following assumptions bear more economic significance.

Assumption 3 *For every price $p < \overline{p}$, $D(p,\theta_i) \geq D(p,\theta_i')$ if and only if $\theta_i \geq \theta_i'$*

Given Assumption 3, a better match corresponds to an upward move of the demand function.[7] This implies $\frac{\partial W}{\partial \theta_i} \geq 0$ and $\frac{\partial V}{\partial \theta_i} \geq 0$. In section 4, we make the stronger assumption that the effect of an increase in θ_i on the demand function is independent of the price p_i, but for the time being we only impose the following condition:

Assumption 4 *The profit function exhibits increasing differences : $\frac{\partial^2 \pi}{\partial p_i \partial \theta_i} \geq 0$*

Therefore for any price p_i the marginal revenue of firm i is larger for high values of θ_i. The following parameterizations of demand functions satisfy Assumptions 1-4: (i) $D(p,\theta)=\theta+q(p)$ if $p \leq \overline{p}$, and zero otherwise, with $2q'(p)+pq''(p) \leq 0$ for all p. (ii) $D(p,\theta)=1-p^\theta$. For instance, let $\theta \in [0,1]$, and, using specification (i), $q(p)=A-p$, with $A>1$ and $\overline{p}=A$. Then the complete information price is $p^*(\theta)=\frac{A+\theta}{2}$ and the corresponding profit is $\pi^*(\theta)=\left(\frac{A+\theta}{2}\right)^2$.

[7] In particular, we rule out situations in which an increase in θ_i corresponds to a rotation of the demand function with an interior rotation point (see Johnson and Myatt (2006))

We assume that if a firm is matched with a consumer, it is in a monopoly situation with respect to that consumer. We relax this assumption in section 5.

An important implication of Assumptions 2 and 4 is the following:

Lemma 1 *The complete information price $p^*(\theta_i)$ is non-decreasing in θ_i.*

Proof : The proof is a standard result of monotone comparative statics (see for instance Vives (2001)). Let $\theta_i > \theta_i'$, and $p' > p^*(\theta_i)$. From Assumption 4 we have $\pi_i(p', \theta_i) - \pi_i(p^*(\theta_i), \theta_i) \geq \pi_i(p', \theta_i') - \pi_i(p^*(\theta_i), \theta_i')$. But, by Assumption 2, $\pi_i(p', \theta_i) - \pi_i(p^*(\theta_i), \theta_i) < 0$. Therefore p' cannot maximize $\pi_i(p, \theta_i')$. $\square$

We assume that the consumer's type is private information, but the platform observes a signal about it. The platform does not know the mapping from the signal to the actual value of the type.[8] It can choose to reveal the value of the signal to advertisers. In that case, each firm i privately learns the value of θ_i. One can imagine that θ_i is the score that firm i would assign to the consumer. The platform knows the age, gender, address of the consumer, as well as some other information related to his valuations for the different goods, but is not able to compute the score, because it lacks some information about the firm. Still, if the platform reveals these characteristics to advertisers, they are able to compute the score. If the platform decides to reveal the information, we say that it follows a *disclosure* policy. If not, we say that it follows a *privacy* policy. Anytime a consumer visits the website, the platform runs a second price auction in order to determine which firm will appear on the consumers' screen. For simplicity, we assume that the platform cannot set a reserve price for the auction.[9]

The timing of the game is the following:

1. The platform commits to a policy $\sigma \in \{\mathcal{D}, \mathcal{P}\}$, where $\mathcal{D}$ stands for *disclosure* and $\mathcal{P}$ for *privacy.*

2. Firms choose independently and simultaneously their prices p_i.

[8] An alternative assumption is that the platform commits not to use this mapping to price its advertising slot.
[9] We discuss reserve prices in section 6.

3. Under disclosure, each firm i learns θ_i. Under privacy, firms do not learn the θ_i's.

4. Under disclosure, firms can submit bids which depend on the realization of θ_i: $b_i^{\mathcal{D}}(\theta_i, p_i)$. Under privacy, they submit a single bid $b_i^{\mathcal{P}}(p_i)$. The auction is a second price auction with no reserve price.

5. The consumer is matched with the winning firm, say firm j. Total welfare, consumer's surplus and firm j's profit are given by $W(p_j, \theta_j)$, $V(p_j, \theta_j)$, and $\pi_j(p_j, \theta_j)$. The platform's revenue R is given by the highest losing bid.

In the auction we only consider equilibria in undominated strategies, that is in which firms bid truthfully.

3 Equilibrium under privacy and disclosure - the general case

Firm pricing and bidding

Privacy. Suppose that the platform chooses not to disclose information. Let $P \equiv (p_1, ..., p_n)$ be the vector of prices, and P_{-i} be the vector of prices of firms other than i. If it sets a price p_i, firm i's profit is

$$E[\pi_i^{\mathcal{P}}(p_i, P_{-i})] = \max\{\int_0^{\overline{\theta}} \pi(p_i, \theta_i) f(\theta_i) d\theta_i - T_i(P_{-i}), 0\}$$

where $T_i(P_{-i}) = \max_{j \in N-i} \int_0^{\overline{\theta}} \pi(p_j, \theta_j) f(\theta_j) d\theta_j$ is firm i's payment if it wins the auction. Notice that this payment does not depend on the realization of the consumer's type, because firms do not learn the θ_i's before they bid. Maximizing this profit with respect to p_i leads to the following result:

Lemma 2 *When the platform chooses to implement a privacy policy, a symmetric equilibrium is such that the price $p^{\mathcal{P}}$ satisfies:*

$$\int_0^{\overline{\theta}} \frac{\partial \pi(p^{\mathcal{P}}, \theta_i)}{\partial p_i} f(\theta_i) d\theta_i = 0. \tag{1}$$

Given that firms cannot infer anything from the fact that they win the auction, they set a price corresponding to the monopoly case when they have no information about consumers.

Also, in a symmetric equilibrium under privacy, firms bid the same amount for every consumer, and therefore the platform extracts all the profits of the industry:

$$R^{\mathcal{P}} = E[\pi(p^{\mathcal{P}}, \theta_i)].$$

Disclosure. Now we assume that firms privately learn the consumer's type before bidding (but after having chosen their price). We look for a symmetric equilibrium, in which firms charge a price $p^{\mathcal{D}}(n)$ and bid truthfully for every realization of the consumer's type.

Because firms bid truthfully, firm i's bid is $\pi(p_i, \theta_i)$. Suppose that all the firms other than i set a price $p^{\mathcal{D}}(n)$. Let $\hat{\theta}_{-i}$ be the highest realization of θ_j for $j \in N - i$. Let j_0 be the identity of the corresponding firm. By Assumption 3, i will win the auction if it bids more than firm j_0, as j_0 outbids all the other firms. Let $\phi(\hat{\theta}_{-i}, p_i, p^{\mathcal{D}}(n))$ be the smallest value of θ_i such that i wins the auction. Notice that by Assumption 3, $\phi(\hat{\theta}_{-i}, p, p) = \hat{\theta}_{-i}$ for every p. Firm i's expected profit is, therefore,

$$E[\pi_i^{\mathcal{D}}(p_i, p^{\mathcal{D}}(n))] = \int_{\hat{\theta}_{-i} \in [\underline{\theta}, \overline{\theta}]} \int_{\theta_i \in [\phi(\hat{\theta}_{-i}, p_i, p^{\mathcal{D}}(n)), \overline{\theta}]} \left[\pi(p_i, \theta_i) - \pi(p^{\mathcal{D}}(n), \hat{\theta}_{-i})\right] f_{n-1:n-1}(\hat{\theta}_{-i}) f(\theta_i) d\theta_i$$

where $f_{k:m}$ is the probability distribution function of the kth order statistic of θ_j among m.[10] At a symmetric equilibrium, we must have $\frac{\partial E[\pi_i^{\mathcal{D}}(p_i, p^{\mathcal{D}}(n))}{\partial p_i}|_{p_i = p^{\mathcal{D}}(n)} = 0$, by concavity of the profit function. This first-order condition can be rewritten as

$$\int_{\hat{\theta}_{-i} \in [\underline{\theta}, \overline{\theta}]} \left\{ \int_{\theta_i \in [\hat{\theta}_{-i}, \overline{\theta}]} \frac{\partial \pi(p^{\mathcal{D}}(n), \theta_i)}{\partial p_i} f(\theta_i) d\theta_i - \frac{\partial \phi(\hat{\theta}_{-i}, p_i, p^{\mathcal{D}}(n))}{\partial p_i} \left(\pi(p^{\mathcal{D}}(n), \hat{\theta}_{-i}) - \pi(p^{\mathcal{D}}(n), \hat{\theta}_{-i})\right) \right\} f_{n-1:n-1}(\hat{\theta}_{-i}) d\hat{\theta}_{-i} = 0$$

After some extra manipulations, we get :

[10] $f_{m:m}$ corresponds to the highest realization, $f_{m-1:m}$ to the second highest, and so on.

Lemma 3 *Under disclosure, a symmetric equilibrium price is given by*

$$\int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial \pi(p^{\mathcal{D}}(n), \theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i = 0. \tag{2}$$

The difference between (1) and (2) comes from the term $F^{n-1}(\theta_i)$ in the integrand. Under privacy, winning the auction for a consumer does not bring any information about the consumer's type. Under disclosure, on the other hand, firm i wins the auction only when all the θ_j's are smaller than θ_i, which occurs with probability $F^{n-1}(\theta_i)$. As we show in the next proposition, the equilibrium price is then higher under disclosure than under privacy.

Proposition 1 *(i) For every n, the equilibrium price under disclosure is larger than the equilibrium price under privacy: $p^{\mathcal{D}}(n) \geq p^{\mathcal{P}}$. (ii) Under disclosure, the equilibrium price of the good increases with the number of firms.*

Omitted proofs are in the appendix. The intuition for Proposition 1 is the following: under disclosure, conditional on winning the auction, firm i expects to face consumers with a higher θ_i than under privacy, and therefore the optimal strategy is to charge a higher price. This effect is all the more important as the number of firms is large, because being the winner among a large set of bidders is a stronger signal that the value of θ_i is high.

The following proposition summarizes the results, and describes the effects of the disclosure policy on profits.

Proposition 2 *(i) The adoption of a disclosure policy by the platform leads to a better expected match between advertisers and consumers, but also to a higher price of the advertised good. (ii) The profits of the industry (advertisers and platform) are higher under disclosure. (iii) Advertisers' share of the total profits is also higher under disclosure.*

Proof: (i) The expected value of the match is $E[\theta_{n:n}]$ under disclosure, against $E[\theta_i]$ under privacy. The fact that disclosure leads to higher prices has been shown in Proposition 1. (ii) The industry's profit under disclosure is $E[\pi(p^{\mathcal{D}}(n), \theta_{n:n})]$. Because $p^{\mathcal{D}}(n)$ is the price that maximizes $E[\pi(p, \theta_{n:n})]$, we have $E[\pi(p^{\mathcal{D}}(n), \theta_{n:n})] \geq E[\pi(p^{\mathcal{P}}, \theta_{n:n})]$. Using the fact that, for every p and every $\theta > \theta'$, $\pi(p, \theta) \geq \pi(p, \theta')$, we get $E[\pi(p^{\mathcal{P}}, \theta_{n:n})] \geq E[\pi(p^{\mathcal{P}}, \theta)]$, the last term being the industry's profit under privacy. (iii) The expected profit of an advertiser is $\frac{1}{n}\left(E[\pi(p^{\mathcal{D}}(n), \theta_{n:n})] - E[\pi(p^{\mathcal{D}}(n), \theta_{n-1:n})]\right) > 0$ under disclosure, whereas it is zero under privacy. □

The fact that disclosure leads to better matches is intuitive, and in line with empirical findings (see, e.g., Goldfarb and Tucker (2011a) and Goldfarb and Tucker (2011b)). Proposition 2 also reveals that advertisers have strong incentives to lobby for more disclosure by intermediaries who possess some information about consumers. We saw that under privacy, the platform extracts all the industry profits, while this is not the case under disclosure, because of the informational rent left to the winning bidder.

Platform policy

From a positive point of view, one would like to know under which conditions the platform is likely to adopt a disclosure policy. As in Ganuza (2004), the main trade-off for the platform is between efficiency (increasing the industry's profit through disclosure) and rent-extraction. However, given that the product price is endogenous to the disclosure policy, the analysis is more intricate.

The lemma below is helpful to prove Proposition 3. It is also interesting in itself.

Lemma 4 *The equilibrium price for the good tends to $p^*(\overline{\theta})$ when n goes to infinity.*

Intuitively, when the number of firms is very large, firm i knows that it will win the auction only when θ_i is very close to $\overline{\theta}$, and so it charges the complete information price corresponding to a consumer of type $\overline{\theta}$.

In the case in which consumers can consume several units of the good, the assumption of linear pricing entails a loss in generality. Indeed, in the case with $n = \infty$, a two-part tariff $p = c$ and $T = W(c, \overline{\theta})$ leads to a higher profit. However the same does not hold in a unit-demand case in which $D(p, \theta) = Pr[v \geq p|\theta]$. Thereafter we shall thus favor this interpretation.

We can now state the main result regarding the platform's optimal policy as a function of the number of bidders.

Proposition 3 *There exists $\overline{n}$ such that, for all $n > \overline{n}$, the platform's revenue is higher under disclosure than under privacy.*

Intuitively, as the number of bidders increase, the informational rent of the winner decreases. For n sufficiently large, the share of the industry's profit captured by the platform is large enough that it prefers to implement disclosure.

Example In order to obtain some further qualitative insights, we solve the model numerically using the following specification: $D(p,\theta) = \theta + \frac{e^{-\frac{p-\mu}{s}}}{1+e^{-\frac{p-\mu}{s}}}$ for $p \leq \overline{p}$, and $D(p,\theta) = 0$ for $p > \overline{p}$. As in Cowan (2007), a good signal corresponds to a vertical shift in the demand function. The demand function, net of the vertical shift, corresponds to a logistic distribution of consumers' willingness to pay. The density of the distribution is single-peaked, with a mean μ and a variance $\frac{\pi s^2}{3}$.[11] An increase in s corresponds to a decrease in the price-elasticity of the demand function. We also assume that θ can take only two values: $\theta_i = \overline{\theta}$ with probability α, and $\theta_i = \underline{\theta}$ with probability $1-\alpha$, which allows us to discuss how the shape of the distribution of types affects the platform's incentives to disclose information.

We choose the following baseline parameters: $\mu_0 = 1$, $s_0 = 1/3$, $c_0 = 1/2$, $\overline{p}_0 = 2$, $\underline{\theta}_0 = 0$, $\overline{\theta}_0 = 1/5$, $\alpha_0 = 1/6$ and $n_0 = 5$. The complete information prices in this benchmark are (approximately) $p^*(\underline{\theta}) = 1.3$ and $p^*(\overline{\theta}) = 1.9$. We solve the model under privacy and disclosure using equations (1) and (2). We obtain $p^{\mathcal{P}} = 1.11$ and $p^{\mathcal{D}} = 1.19$, i.e a 7% increase in the equilibrium price. Under these parameters, the platform is roughly indifferent between privacy and disclosure (its profit under disclosure is 0.2% higher than under privacy).

Figure 2 shows how changes in parameter values affect the relative profitability of disclosure for the platform. For instance, a 100% increase in n (i.e. $\frac{\Delta n}{n_0} = 1$) leads disclosure to be 15% more

[11] Similar results obtain for a Pareto distribution, i.e. such that the demand function, net of the vertical shift, has a constant price elasticity.

profitable than privacy. If α increases by 100%, disclosure is 8% more profitable. The platform is more likely to choose privacy when n is small or when good signals are unlikely (α small), so as not to leave too much of an informational rent to the winning bidder. The effects of μ and s are more modest.

Figures 3 to 6 show the platform's preferred policy as a function of two parameters, *ceteris paribus*. In the range of parameters we considered, the decision of the platform is not very sensitive to the values of s, μ and $\overline{\theta}$, the characteristics of the demand function, as opposed to n and α. Figure 3, for instance, shows that, for any "reasonnable" value of s,[12] the decision of the platform is based on a thereshold $\overline{n} \in (9/2, 11/2)$, when the other parameters take their baseline values. We obtain analogous figures for μ and $\overline{\theta}$. On the other hand, the value of α still plays a role when it varies jointly with n: for lower values of α, the threshold $\overline{n}$ over which the platform chooses disclosure increases, especially when α is small. The baseline values we chose are not crucial to obtain these qualitative results.

Even though it is of course delicate to draw definitive conclusions from such a numerical exercise, our results suggest that the decision of the platform is somehow robust to the details of the demand function: to a first-order approximation, the platform's decision only depends on n and, to a lesser extent, on α.

As in Bergemann and Bonatti (2011), we have the prediction that more targeting (going from privacy to disclosure) can lead the price of advertising to increase or to decrease. Among other things, Bergemann and Bonatti (2011) find that, while targeting increases the value of advertising, its equilibrium price first increases and then decreases in the targeting ability. In our model, the price of an ad is the highest losing bid. Whether more targeting leads to an increase in the price of advertising relies on the trade-off between the increase in the value of advertising and the information rent. Roughly speaking, our model predicts that when the number of advertisers is small (resp. large) more targeting is more likely to decrease (resp. increase) the price of advertising.

[12] i.e. such that solutions to (1) and (2) are interior.

4 Normative analysis with a large number of firms

In the numerical application of section 3, welfare is always higher under disclosure. This illustrates the first potential distortion in the platform's decision, namely that, to minimize informational rent, the platform can disclose too little information. The contribution of this section is to show that the platform can disclose *too much* information under certain circumstances.

We focus on cases with a large number of firms: $n = \infty$. As shown in Proposition 3, the platform then captures the whole industry profits, and always prefers to implement disclosure. Thus, we assume away cases in which the platform discloses strictly less information that what would be optimal . The second implication of having $n = \infty$ is that firms charge a price $p^*(\overline{\theta})$ under disclosure. Therefore, our results depend on the distribution of types F only through its mean m and the upper-bound of its support $\overline{\theta}$, which considerably simplifies the analysis.

Following Cowan (2007), let us assume that for every $p \leq \overline{p}$, $D(p,\theta) = \theta + q(p)$, with $q' < 0$ and $D(p,\theta) = q(p) = 0$ for all $p > \overline{p}$. Let $P(x,\theta)$ be the inverse demand, defined as $D(P(x,\theta),\theta) = x$, or $P(D(p,\theta),\theta) = p$. Such a formulation corresponds to a conditional distribution of the willingness to pay v_i $Pr[v_i = \overline{p}|\theta_i] = \theta_i$ and $Pr[v_i \in [p,\overline{p})|\theta_i] = q(p)$.

The elasticity of demand is $\eta(p,\theta) = -pq'(p)/(\theta + q(p))$, and the elasticity of the slope of the demand is $\alpha(p,\theta) = -pq''(p)/q'(p)$. Notice that the latter expression does not depend on θ, so that we can drop θ from the arguments of α.

For any θ, the optimal price $p^*(\theta)$ is given by the Lerner formula

$$\frac{p^*(\theta) - c}{p^*(\theta)} = \frac{1}{\eta(p^*(\theta),\theta)} \tag{3}$$

The second-order condition is $(p - c)q''(p) + 2q'(p) \leq 0$ at $p = p^*(\theta)$ (i.e $2\eta(p^*(\theta),\theta) > \alpha(p^*(\theta))$).[13] Throughout this section we assume that $p^*(\overline{\theta}) < \overline{p}$, so that we only have interior solutions. As the type θ enters linearly in the demand function, one can see that the optimal price under privacy is $p^{\mathcal{P}} = p^*(m)$, i.e the complete information price corresponding to the average type. Thanks to

[13]Notice that we no longer require the profit to be concave everywhere.

Proposition 3, we also know that the price under disclosure is $p^*(\overline{\theta})$, the complete information price corresponding to the best type.

Define:

$$W(\theta) = \int_0^{D(p^*(\theta),\theta)} (P(x,\theta) - c)dx$$

The function W measures the social welfare when a firm faces a consumer of type θ and charges the complete information price $p^*(\theta)$. Given the previous observations, we can see that disclosure is preferable to privacy from a social welfare point of view when $W(\overline{\theta}) > W(m)$.

We are interested in sufficient conditions over q such that privacy or disclosure is better for welfare. To do so, let us consider the derivative of W:

$$W'(\theta) = (p^*(\theta) - c)[1 + p^{*'}(\theta)q'(p^*(\theta))] + \int_0^{D(p^*(\theta),\theta)} \frac{\partial P(x,\theta)}{\partial \theta} dx \tag{4}$$

Notice that $P(x,\theta) = q^{-1}(x - \theta)$, so that $\frac{\partial P(x,\theta)}{\partial \theta} = -\frac{\partial P(x,\theta)}{\partial x}$. Therefore, (4) simplifies to

$$W'(\theta) = \underbrace{(p^*(\theta) - c)p^{*'}(\theta)q'(p^*(\theta))}_{\text{price effect}} + \underbrace{\overline{p} - c}_{\text{match effect}} \tag{5}$$

The right-hand side of (5) is made of two terms. The first term, the “price effect” measures the loss in utility due to a higher price: an increase in θ leads to an increase in the price, and thus a lower quantity from the elastic part of the demand $q(p)$. The second term, the “match effect” corresponds to the shift in the inverse demand that results from the increase in θ.

At this point it may be useful to discuss the relationship between Cowan (2007) and the present paper. Cowan (2007) uses the parametrization $D(p,\theta) = \theta + q(p)$ to study whether third degree price discrimination is desirable. He assumes that there are two sub populations, one with demand $q(p)$ (the weak market) and one with demand $\theta + q(p)$ (the strong market). However, contrary to what we do, he assumes that the utility of a consumer in the strong market is $U(q(p))$, so that θ does not directly enter the utility function. This remark sheds some light on the differences between our approach and the well-known analysis of price-discrimination. While third-degree price discrimination leads to a higher price for consumers in the strong market, disclosure of information by the platform leads to a higher price *for everyone*. However, disclosure improves the match, so that all consumers face products for which they have a *strong* demand.

A sufficient condition for disclosure (privacy) to be socially optimal is for $W'(\theta)$ to be positive (negative) for all $\theta \geq m$. This leads to the following lemma, in which $\eta \equiv \eta(p^*(\theta), \theta)$ and $\alpha \equiv \alpha(p^*(\theta))$.

Lemma 5 *Disclosure (resp. privacy) is socially optimal if, for every* $\theta \geq m$, $\frac{p^*(\theta)}{\overline{p}-c} \leq ($ *resp.* $\geq) 2\eta - \alpha$

Proof: For notational simplicity, we drop the arguments of the functions. We have

$$W' \geq 0 \Longleftrightarrow \frac{\overline{p} - c}{p^* - c} \geq -p^{*'} p q', \tag{6}$$

$p^{*'}$ may be obtained from the profit maximizing condition:

$$(p^*(\theta) - c) q'(p^*(\theta)) + \theta + q(p^*(\theta)) = 0$$

Differentiating with respect to θ gives

$$p^{*'} = -\frac{1}{2q' + (p^* - c) q''}.$$

Plugging this expression into (6), and rearranging terms, one gets the desired condition.□

Given Lemma 5, we would like to know which characteristics of the demand function $D(p, \theta)$ make it more desirable to disclose the information. The first one of such characteristics is the reservation price $\overline{p}$: loosely speaking, an increase in $\overline{p}$ makes disclosure more desirable. Regarding the shape of the elastic part of the demand, q, we have the following:

Proposition 4 *(i) If q is log-concave and $p^*(\overline{\theta}) < \overline{p}$, disclosure is socially optimal.*

(ii) If $q(p) = p^{-\epsilon}$ (constant elasticity) then privacy is socially optimal if $\epsilon \in [\underline{\epsilon}, \frac{2\overline{p}-c+\sqrt{c(4\overline{p}-3c)}}{2(\overline{p}-c)}]$, where $\underline{\epsilon}$ solves $\frac{\partial \pi(\overline{p},\overline{\theta})}{\partial p} = 0$.

Proof: (i) By Lemma 5, we know that disclosure is optimal when $\frac{p^*(\theta)}{\overline{p}-c} \leq 2\eta - \alpha$ for all θ. For notational simplicity, let us drop the arguments θ and p. The previous condition is equivalent to

$$\frac{2q'}{Q} + \frac{1}{\overline{p}-c} \leq \frac{q''}{q'}$$

with $Q = \theta + q$. Multiplying the previous expression by $Qq' < 0$ and rearranging terms, we find that disclosure is optimal when

$$(q')^2 \left(2 + \frac{Q}{q'(\overline{p}-c)}\right) - q''Q \geq 0$$

Now, as $\overline{p} > p$, $\frac{p-c}{p} = -\frac{Q}{pq'}$ implies $\overline{p} - c > -\frac{Q}{q'}$ and so $\frac{Q}{q'(\overline{p}-c)} > -1$.

Thus

$$(q')^2 \left(2 + \frac{Q}{q'(\overline{p}-c)}\right) - q''Q \geq (q')^2 - q''Q$$

To conclude, notice that $(q')^2 - q''Q \geq 0$ when Q (and therefore q) is log-concave.

(ii) The condition $\epsilon \geq \underline{\epsilon}$ ensures that $p^*(\theta) < \overline{p}$ for all θ, so that the previous analysis is valid.[14] By Lemma 5, privacy is optimal if $\frac{p^*(\theta)}{\overline{p}-c} \geq 2\eta - \alpha$. Using Lemma 1, we have $p^*(\theta) \geq p^*(0)$ for all $\theta \geq 0$. Moreover, $p^*(0)$ maximizes $(p-c)p^{-\epsilon}$, i.e. $p^*(0) = \frac{\epsilon c}{\epsilon - 1}$. We also have $\epsilon \geq \eta(p^*(\theta), \theta)$ for all $\theta \geq 0$, as η is the elasticity of $\theta + p^{-\epsilon}$, and $\alpha = 1 + \epsilon$. Therefore, if $\frac{\epsilon c}{(\epsilon-1)(\overline{p}-c)} \geq \epsilon - 1$, then the condition $\frac{p^*(\theta)}{\overline{p}-c} \geq 2\eta - \alpha$ holds for every θ, and privacy is optimal. It is straightforward to check that $\epsilon = \frac{2\overline{p}-c+\sqrt{c(4\overline{p}-3c)}}{2(\overline{p}-c)}$ is the only solution which is larger than 1 ($\epsilon < 1$ is not consistent with interior solutions, as it leads to $p^*(0) = \overline{p}$). $\square$

To interpret Proposition 4, let $q(p) = Pr[v \geq p]$, where v is interpreted as the willingness to pay of an individual with $\theta = 0$. If G is the cumulative distribution function of v, then q is log-concave if and only if $1 - G(v)$ is log-concave.[15] The log-concavity of the reliability function $1 - G$ is a property

[14] Note that there are cases in which $\underline{\epsilon} > \frac{2\overline{p}-c+\sqrt{c(4\overline{p}-3c)}}{2(\overline{p}-c)}$, i.e. the sufficient condition is never satisfied.

[15] Another equivalent condition is $\frac{g}{1-G}$ non-decreasing, where g is the density of v.

shared by a number of usual distributions (uniform, normal, logistic among others, see Bagnoli and Bergstrom (2005)). To get an intuition for this result, notice that from the first-order condition we have

$$\frac{dp^*(\theta)}{d\theta} = \frac{1}{g(p)\left(1 - \phi'(p) + g'(p)/g(p)^2\right)}$$

where $\phi = (1-G)/g$. Heuristically, when ϕ' is small, the optimal price does not vary much with θ. When q is log-concave, $\phi' < 0$, and so we are precisely in the case in which the price effect is small in magnitude compared to the match effect. Proposition 4 (ii) tells us that when q has a constant elasticity, privacy is optimal when the elasticity is intermediate. For too low an elasticity, $p^{\mathcal{P}} = p^{\mathcal{D}} = \overline{p}$ and disclosure is optimal. When the elasticity is high, even though $p^{\mathcal{P}} < p^{\mathcal{D}}$, they are both very close to the marginal cost c, so that the match effect dominates. When ϵ is intermediate, $p^{\mathcal{D}}$ is sufficiently above $p^{\mathcal{P}}$ to offset the match effect.

Examples: In Figure 7, we represent the case of a linear demand $q(p) = a - bp$. Switching from privacy to disclosure leads to a vertical shift of $\overline{\theta} - m$. Notice that here the quantity consumed under disclosure $D^*(\overline{\theta})$ is greater than that under privacy $(D^*(m))$. This is why we speak of a "virtual" price effect: under disclosure, consumers buy a smaller quantity than if the price remained the same as under privacy. As we see, the welfare gain is always positive, as implied by Proposition 4 (i). Figure 8 represents the inverse demand function when q has a constant price elasticity: $q(p) = p^{-\epsilon}$ if $p \leq \overline{p}$.[16] Here we see that disclosure leads to a lower quantity (the price effect) that is partially compensated by the increase in the utility generated by the infra-marginal units (the match effect). In this specific example, privacy leads to a higher welfare than disclosure.

5 Multi-slot case

So far in our analysis we have only considered situations of downstream monopoly, i.e such that winning the auction gives a firm the exclusive right to appear on the consumer's screen. However,

[16] We take $\overline{\theta} = 0.2, m = 0.05, \overline{p} = 1.2, \epsilon = 2.5$ and $c = 0.5$.

many platforms are designed so as to display several advertisements to each consumer. The purpose of this section is to investigate the link between upstream information disclosure (at the auction stage) and downstream interactions (prices in particular) when there is more than one advertising slot.

We tackle two separate cases. In the first one, we shut down downstream competition by assuming that consumers look at the $K > 1$ ads they see on their screen, and that the K purchasing decisions are independent. In the second case, $K = 2$ but consumers view the products as imperfect substitutes.

Multi-unit auction with independent demands

Suppose that the platform has K identical slots up for sale.[17], and that it allocates them through a uniform price auction, in which the price paid by all the firms whose ads are displayed is equal to the highest losing bid. Consumers inspect the K ads they see, and final demands are independent.

Whether the platform prefers privacy or disclosure still depends on the trade-off between increasing the value of trade (with disclosure) and eliminating the informational rent of the winner (privacy). However, the equilibrium price charged by advertisers will decrease as the number of slots increases, *even though we explicitly rule out competition on the product market.*

To see this, consider a symmetric configuration in which all firms (except i) charge a price $p_{K,n}$ and bid their profit $\pi(p_{K,n}, \theta_j)$ after learning the consumer's type $\Theta = (\theta_1, ..., \theta_n)$.[18] Let $\hat{\theta}_k$ be the k-th highest value of θ_j, for $j \neq i$. Then firm i wins a slot in the auction if and only if $\pi(p_i, \theta_i) \geq \pi(p_{K,n}, \hat{\theta}_K)$. Let $\phi(\hat{\theta}_K, p_i, p_{K,n})$ be the smallest value of θ_i such that i wins a slot. Notice that in equilibrium $\phi(\hat{\theta}_K, p_{K,n}, p_{K,n}) = \hat{\theta}_K$. Firm i's profit is

$$\pi_i(p_i, p_{K,n}) = \int_{\hat{\theta}_K \in [\underline{\theta}; \overline{\theta}]} \int_{\theta_i \in [\phi(\hat{\theta}_K, p_i, p_{K,n}); \overline{\theta}} \left(\pi(p_i, \theta_i) - \pi(p_{K,n}, \hat{\theta}_K) \right) f_{n-K:n-1}(\hat{\theta}_K) f(\theta_i) d\hat{\theta}_K d\theta_i$$

[17] A more realistic assumption would be to also introduce heterogeneity among slots whose values would depend on their prominence (See Athey and Ellison (2011) and Chen and He (2011)) For the sake of brievity we do not pursue that avenue in the current paper and leave the question for future research.

[18] Such a strategy is weakly dominating in the uniform auction with unit-demand.

The first order condition at a symmetric equilibrium is

$$\int_{\theta\in[\underline{\theta},\bar{\theta}]} \frac{\partial\pi(p_{K,n},\theta)}{\partial p_i} F_{n-K:n-1}(\theta) dF(\theta) = 0 \tag{7}$$

Proposition 5 *With independent demands, the equilibrium product price is a decreasing function of the number of advertising slots.*

Proof: The logic is the same as in Proposition 1 (ii), by noting that $F_{n-K:n-1}(\theta) \geq F_{n-K':n-1}(\theta)$ for any $K' \leq K$.$\square$

On top of the standard trade-off between the number of slots and the price at which each slot is sold, another trade-off exists between the number of slots and the final price charged to consumers. The intuition for Proposition 5 is that as K increases, winning the auction becomes less informative with respect to the expected elasticity of demand.

How does having the ability to display several advertisements affect the platform's incentives to disclose information? Under privacy, the revenue of the platform is, for $K < n$, $\mathcal{R}_K^{\mathcal{P}} = KE[\pi(p^{\mathcal{P}}, \theta)]$. Under disclosure, it is $\mathcal{R}_K^{\mathcal{D}} = KE[\pi(p_{K,n}^{\mathcal{D}}, \theta_{n-K:n})]$, i.e. K times the profit of the highest losing bidder. When $D(p,\theta) = \theta + q(p)$ as in section 4, privacy is more profitable for the platform if and only if $(p^{\mathcal{P}} - c)\left(E[\theta] + q(p^{\mathcal{P}})\right) > (p_{K,n}^{\mathcal{D}} - c)\left(E[\theta_{n-K:n}] + q(p_{K,n}^{\mathcal{D}})\right)$. Given that the θ_i are i.i.d., there exists $\hat{K}$ such that for all $K > \hat{K}$, $E[\theta_{n-K:n}] < E[\theta]$. We then have, for all $K > \hat{K}$,

$$(p_{K,n}^{\mathcal{D}} - c)\left(E[\theta_{n-K:n}] + q(p_{K,n}^{\mathcal{D}})\right) < (p_{K,n}^{\mathcal{D}} - c)\left(E[\theta] + q(p_{K,n}^{\mathcal{D}})\right) < (p^{\mathcal{P}} - c)\left(E[\theta] + q(p^{\mathcal{P}})\right),$$

the last inequality following from revealed preferences. We have showed the following Proposition:

Proposition 6 *When $D(p,\theta) = \theta + q(p)$, if K is large enough, the platform prefers to implement privacy.*

When K is small, a numerical analysis using the same setup as in section 3 indicates that increasing K reduces the incentives to disclose information (see Figure 9).

As K increases, the highest losing bidder, whose bid determines the platform's revenue, tends to receive more negative signals under disclosure. This leads the platform to implement privacy more frequently.

Notice that the decision to disclose information to bidders in turn affects how many slots the platform is willing to have. Indeed, under privacy, selling an extra slot doesn't affect bidders' willingness to pay, so that the optimal value of K is $n-1$. Under disclosure however, selling an extra slot reduces the willingness to pay of each bidder, because it becomes more likely that, conditional on winning a slot, the value of θ_i is small. Therefore, under disclosure the platform sells (weakly) fewer slots than under privacy.

Downstream competition

The general case with two slots.

Suppose that there are two advertising slots, and that the consumer views the products as imperfect substitutes. Let $D_i(p_i, p_j, \theta_i, \theta_j)$ the probability that a consumer buys from firm i when he sees ads by firms i and j. As in the downstream monopoly case, we assume that $\frac{\partial^2 \pi_i}{\partial p_i \partial \theta_i} \geq 0$: a higher θ_i gives incentives to increase p_i. To make things interesting, let us also assume that $\frac{\partial^2 \pi_i}{\partial p_i \partial \theta_j} \leq 0$: competing against a "stronger" rival makes a firm willing to cut its price.[19] The timing of the game is unchanged, with firms choosing their prices after a regime (privacy or disclosure) has been announced, but before observing any signal. The auction format is a uniform auction: the highest two bidders appear on the consumer's screen, and pay the bid of the third highest bidder. We focus on equilibria with truthful bidding, which is a weakly dominating strategy. We denote $p^{\mathcal{P}}$ and $p^{\mathcal{D}}$ the equilibrium prices under privacy and disclosure respectively.

[19]The case in which $\frac{\partial^2 \pi_i}{\partial p_i \partial \theta_j} > 0$ is straightforward to analyze: disclosure leads to higher prices, as in the downstream monopoly case.

Privacy Firms do not observe any signal before bidding, and thus all bid the same amount. The first-order condition which determines the equilibrium price is thus

$$\int_{[\underline{\theta},\overline{\theta}]^2} \frac{\partial \pi(p^{\mathcal{P}}, p^{\mathcal{P}}, \theta_i, \theta_j)}{\partial p_i} f(\theta_i) f(\theta_j) d\theta_j d\theta_i = 0 \tag{8}$$

Disclosure In a symmetric equilibrium under disclosure, firm i wins one of the two slots if it receives one of the two highest signals. Let $\theta_{n-1:n-1}$ and $\theta_{n-2:n-1}$ be the highest and second highest signals among the firms other than i. Firm i's first-order condition is thus

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(\int_{\underline{\theta}}^{\theta_i} \frac{\partial \pi(p, p, \theta_i, \theta_{n-1:n-1})}{\partial p_i} f(\theta_i) f_{n-1:n-1}(\theta_{n-1:n-1}) d\theta_{n-1:n-1} \right.$$
$$\left. + \int_{\theta_i}^{\overline{\theta}} \frac{\partial \pi(p, p, \theta_i, \theta_{n-1:n-1})}{\partial p_i} f(\theta_i) f_{n-1:n-1}(\theta_{n-1:n-1}) F_{n-2:n-1}(\theta_i | \theta_{n-1:n-1}) d\theta_{n-1:n-1} \right) d\theta_i = 0$$

The first term in the the parenthesis corresponds to the cases in which $\theta_{n-1:n-1} < \theta_i$, whereas the second one corresponds to the cases in which $\theta_{n-2:n-1} < \theta_i < \theta_{n-1:n-1}$.

Using the fact that $f_{n-1:n-1}(x) = (n-1) f(x) F(x)^{n-2}$ and $f_{n-1:n-1}(y) F_{n-2:n-1}(x|\theta_{n-1:n-1} = y) = (n-1) f(y) F(x)^{n-2}$ (see Krishna (2009)), we obtain, after some simple manipulations and the relabelling $\theta_{n-1:n-1} = \theta_j$, the first-order condition

$$\int_{[\underline{\theta},\overline{\theta}]^2} \frac{\partial \pi(p^{\mathcal{D}}, p^{\mathcal{D}}, \theta_i, \theta_j)}{\partial p_i} f(\theta_i) f(\theta_j) F(\min\{\theta_i, \theta_j\})^{n-2} d\theta_j d\theta_i = 0 \tag{9}$$

Unlike the downstream-monopoly case, the comparison between (8) and (9) is ambiguous. Indeed, through the term $F(\min\{\theta_i, \theta_j\})^{n-2}$ in the integrand of (9), larger realizations of *both* θ_i and θ_j are given a larger weight than in (8). Given that $\frac{\partial^2 \pi_i}{\partial p_i \partial \theta_i} > 0 > \frac{\partial^2 \pi_i}{\partial p_i \partial \theta_j}$, we cannot say in general how the equilibrium price varies with disclosure. Firms expect to only win when they have received a good signal, but they also expect to face stronger rivals than under privacy. In order to say more, below we study the special case of a discrete choice model.

Discrete choice model

Suppose that consumers have use for only one product. The gross utility v_i that a consumer gets from consuming good i, given a type θ_i, is distributed according to a c.d.f. $G(.|\theta_i)$, with $G(v_i|\theta_i) \leq G(v_i|\theta_i')$ if and only if $\theta_i \geq \theta_i'$ (first-order stochastic dominance). Let g be the associated density. Following Perloff and Salop (1985), a consumer facing firm i and firm j buys from i if and only if $v_i - p_i \geq v_j - p_j$.[20] Under privacy, the two firms which win a slot are randomly selected. Let $G_{\mathcal{P}}(v)$ be the probability that $v_i \leq v$ given that the firm is randomly drawn. Using equations (12) and (13) in Perloff and Salop (1985) , the equilibrium price under privacy is given by

$$p^{\mathcal{P}} = c + \frac{1}{2\int (g_{\mathcal{P}}(v))^2 dv} \tag{10}$$

Under disclosure, the firms who win the auction are those with the two highest signals. Let $G_{\mathcal{D}}(v) \equiv G(v|\theta_i > \theta_{n-2:n-1})$ be the c.d.f. of v_i conditional on i winning a slot. The equilibrium price under disclosure is then

$$p^{\mathcal{D}} = c + \frac{1}{2\int (g_{\mathcal{D}}(v))^2 dv} \tag{11}$$

Notice that $\int (g(v))^2 dv$ is the continuous analog to the Herfindahl index $\sum_v (g(v))^2$ which measures the concentration of the distribution G. Using this definition of concentration, we have

Proposition 7 *The price of the product is higher under disclosure if and only if $G_{\mathcal{D}}$ is less concentrated than $G_{\mathcal{P}}$.*

The intuition for Proposition 7 is as follows: when high signals allow firms to infer the willingness to pay with high precision (i.e $G_{\mathcal{D}}$ is more concentrated than $G_{\mathcal{P}}$), the two winning firms are relatively homogenous in the eyes of the consumer. Expecting to always compete against such a rival, each firm sets its price close to marginal cost to begin with (in the spirit of Bertrand competition). If, on the

[20]This formulation implies that consumers always buy one product. It allows us to obtain closed-form solution, thus facilitating the exposition.

other hand, $G_{\mathcal{D}}$ is less concentrated than $G_{\mathcal{P}}$, disclosure leads to a relaxation of competitive pressure because winning firms are more differentiated. For example, if the distribution of v_i conditional on θ_i is $\mathcal{U}[\underline{v}, \theta_i]$, then $G_{\mathcal{P}}$ is more concentrated than $G_{\mathcal{D}}$, and disclosure increases prices. If on the other hand the distribution is $\mathcal{U}[\theta_i, \overline{v}]$, disclosure leads to a more concentrated distribution and therefore to a higher price.

Even when disclosure increase prices, it does not follow that it harms consumers, because they benefit from better matches.[21] When it decreases prices, consumers necessarily benefit from disclosure. Unlike with downstream monopoly, the platform may no longer prefer to disclose information, even when n is large, if $G_{\mathcal{D}}$ is more concentrated than $G_{\mathcal{P}}$. Indeed, firms' profits are then dissipated through price competition, which lowers the bids. Therefore we see a new tension arise: when consumers benefit the most from disclosure, the platform tends to prefer privacy.

The platform's decision to offer 1 vs 2 advertising slots depends on the intensity of competition, which in turn depends on the concentration of the distributions $G_{\mathcal{D}}$ and $G_{\mathcal{P}}$. Under disclosure, the higher the concentration of $G_{\mathcal{D}}$, the less likely the plaform is to offer two slots, because of the low equilibrium profits. When $G_{\mathcal{D}}$ is not very concentrated, profits remain high under competition, and having two slots is more likely to be profitable. The same goes for the privacy case. Therefore, unlike the case with independent demands across firms, privacy does not necessarily lead the platform to put more slots for sale than disclosure, as the decision depends on the relative concentration of $G_{\mathcal{D}}$ and $G_{\mathcal{P}}$.

6 Other extensions

In this section we discuss some extensions to our basic model. We first consider the mechanism through which the platform allocates the advertising slot, and then introduce the possibility of partial disclosure.

[21] In this model with inelastic total demand, disclosure always increases total welfare.

Optimal mechanism with one slot. In our model we focus on the platform's choice of information revelation rather than on the optimal mechanism. Here the optimal mechanism is straightforward to design. The platform should sell the access to information at a price T equal to the expected net profit of a firm, and forbid firms to participate if they do not pay T. This way, the platform can extract the profit of the whole industry $E[\pi(p_n^{\mathcal{D}}, \theta_{n:n})]$, which is maximized under disclosure. For legal reasons or reputational concerns the platform may not be able to sell detailed information on its customers. However it may be authorized to share this information with its commercial partners. In such a case information revelation to bidders may then be seen a means for the platform to still monetize his information without selling it.[22]

Partial disclosure. So far we have only considered two types of information revelation policy for the platform. Under the disclosure policy, the platform reveals all the available information about consumers prior the auction, whereas under the privacy policy it reveals nothing. One can imagine that it could be better to adopt an information revelation policy between these two extremes. The question of the optimal provision of information by a principal has been investigated in different contexts in the literature. Most related to this work are the papers that study this question in the context of auctions (e.g.Ganuza and Penalva (2010)), and in the context of a monopolistic seller (e.g.Lewis and Sappington (1994), Johnson and Myatt (2006), Anderson and Renault (2006)).

Consider the following demand function:

$$D(p_i, \theta_i) = \begin{cases} 1 + \theta_i & \text{if } p_i \leq v_L \\ \theta_i & \text{if } p_i \in (v_L, v_H] \\ 0 & \text{if } p_i > v_H \end{cases}$$

This demand function is such that a higher value of θ_i corresponds to a vertical shift, as in section 4. We assume that the θ_i's are independently and identically uniformly distributed on $[0, 1]$, that there

[22] See Eso and Szentes (2007) for a model in which the auctioneer can sell information to bidders. See also Bergemann and Bonatti (2013) for a model in which the platform can sell different signals.

is an infinite number of firms, and that the marginal cost of production is zero.

Hereafter, we focus on the *truth or noise* technologies of information revelation introduced by Lewis and Sappington (1994). More specifically, we assume that, given a type $(\theta_i)_{i=1,...,\infty}$, the platform generates a vector of signals $(s_i)_{i=1,...,\infty}$, such that $s_i = \theta_i$ with probability γ, and s_i is an independent uniform draw from $[0,1]$ otherwise. Firm i only observes s_i and γ, so that its belief over θ_i is summarized by $\beta_i \equiv E[\theta_i|s_i,\gamma] = \gamma s_i + \frac{1-\gamma}{2}$. The platform's only strategic tool is the choice of γ, and this specification encompasses the cases of privacy ($\gamma = 0$) and disclosure ($\gamma = 1$). [23]

Given γ, the highest possible signal that a firm can receive is $s_i = 1$, which translates into $\beta_i = \frac{1+\gamma}{2}$. Because there is an infinite number of firms, by the same logic as in Proposition 3 (ii), firms behave as if they faced consumers with $\beta_i = \frac{1+\gamma}{2}$ with probability 1. They thus maximize $pD(p, \frac{1+\gamma}{2})$, which leads to the following equilibrium price:

$$p_\gamma = \begin{cases} v_H & \text{if } \frac{v_H}{v_L} \geq \frac{3+\gamma}{1+\gamma} \\ v_L & \text{otherwise} \end{cases}$$

Just as in Proposition 3 (i), the platform's revenue is maximized with full disclosure ($\gamma = 1$). To see this, notice that the expected industry profit, which is entirely captured by the platform, is $\max\left(v_H\frac{1+\gamma}{2}, v_L(\frac{1+\gamma}{2}+1)\right)$, increasing in γ.

Now let us turn to the socially optimal policy, under the constraint that advertisers choose their own prices. Let $V \equiv v_H/v_L > 1$. For any $\gamma \leq \tilde{\gamma} \equiv \frac{3-V}{V-1}$, advertisers choose $p = v_L$, and welfare is given by $W(\gamma) = v_H\frac{1+\gamma}{2} + v_L$. For $\gamma > \tilde{\gamma}$, the equilibrium price is v_H, the units that are valued at v_L by consumers are not sold, and welfare is $W(\gamma) = v_H\frac{1+\gamma}{2}$. Figures 10 and 11 depict the welfare as a function of γ.

It is clear that there can be two solutions: $\gamma^* = \tilde{\gamma}$ or $\gamma^* = 1$. Straightforward computations

[23] The truth or noise technology does not require the platform to observe anything about θ. Other possible technologies (e.g. pooling of types above a threshold à la Anderson and Renault (2006) - Kamenica and Gentzkow (2011)) can lead to partial disclosure, but they require the plaform to observe signals about θ.

reveal that

$$\gamma^* = \begin{cases} \frac{3-V}{V-1} & \text{if } V \in [2,3] \\ 1 & \text{otherwise} \end{cases}$$

When V is too small or too large, the policy chosen by the platform has no impact on firms' pricing. In the former case, they always choose $p = v_L$ whereas in the latter case they set $p = v_H$. In these cases, it is optimal to disclose the information, as it improves the quality of the matching. However, for intermediate values of V, disclosing too much information leads firms to charge a higher price. The socially optimal policy involves disclosing as much information as possible without triggering this pricing distortion.

The bottom line is that partial disclosure can be welfare-maximizing, in which case the platform reveals too much information.

Reserve price. It is well known that in the independent and private valuation model of auctions, the optimal auction for the seller consists in a second-price auction with an optimal reserve price. In our model, allowing the platform to use a reserve price would provide more incentives to disclose information. Indeed, using a reserve price allows the platform to reduce the expected informational rent of the winning bidder under the disclosure policy. If the platform is unable to establish a reverse price directly, an indirect means to implement it is not to reveal information to all the advertisers. Indeed, under a disclosure regime, the platform could increase its revenue by keeping two advertisers uniformed. If the number of bidders is relatively large, this does not have a great statistical impact on the highest advertisers' valuations, but guarantees a minimum revenue to the platform.

Customized pricing. In the baseline model we assume that firms choose their prices before submitting their bids, thus preventing any form of price customization based on the type of the consumer they gain access to. Although this assumption is reasonable in many settings, in view of the relative technological ease with which firms that compete online could price discriminate, an important issue is how the optimal policy disclosure changes when firms customize their price for each consumer they reach. In a simplified version of our model, available upon request, we show that price customization gives the platform more incentives to reveal information. This is because firms are now better equipped to take advantage on the information to extract consumers surplus. More interestingly, we also show that, holding constant the fact that the platform discloses the information, firms become worse-off when they can customize their prices. The reason is that the gap between their private valuation for the slot and the expected second highest bid shrinks when customized pricing is allowed. Firms engaging in price customization are therefore 'instrumentalized' by the platform who further exploits them to extract consumer surplus. This result is related to Conitzer, Taylor, and Wagman (2012) who study, in an environment with dynamic price discrimination, how a platform controls the level of privacy by choosing the cost of anonymization for consumers who want to avoid being price discriminated against. They show that such an information gatekeeper chooses free anonymization for consumers, which enables the seller to credibly commit to not price discriminate, thus increasing its profits. In our context, firms would also be better-off by finding a means to not engage in price discrimination.

7 Conclusion

In this paper, we study the incentives of an ad-supported platform to disclose information about its users to advertisers prior to an auction. The choice of a disclosure regime affects firms' pre-auction pricing decision: disclosure can lead to higher prices even without price-discrimination. This effect, absent from the standard literature on information disclosure in auctions, opens the possibility that the platform discloses *too much* information, in particular when the demand function has an intermediate (constant) price-elasticity. However, because disclosure improves the quality of the match between consumers and advertisers, the price effect can be compensated by an increase in the amount of trade. When advertisers compete downstream, the equilibrium price can go up or down. In the latter case, the trade-off between price effect and match effect disappears, so that disclosure is socially optimal, but the platform may prefer privacy.

Our model ignores several issues, that we discuss below. Although introducing these features would enrich the analysis by making it more realistic, most of the effects that we highlight would not be fundamentally altered.

Elastic participation. One assumption of the model is that the number of advertisers and the number of consumers is fixed. An interesting avenue for future research would be to study a model in which the platform can influence the number of users. For instance, if consumer participation is determined by the surplus that they expect to obtain through their interactions with advertisers, the platform could be tempted to foster competition by displaying several ads from competing advertisers. On the other hand, in order to attract advertisers the platform would need to give them some market power. These questions are related to the two-sided markets literature (see Armstrong (2006), in particular pp. 686-687, or de Cornière (2013) for an application to the search engine industry). If users exhibit intrinsic privacy concerns, the platform could also use the degree of privacy as an instrument to generate more traffic.

Competition between platforms. One important assumption of the model is that the platform is a monopoly. What would change in a model with several platforms? We think that the main consequence would be that the platforms would tend to internalize to a lesser extent the impact of their policy on advertisers' pricing strategy, with the implication that they would implement more disclosure. This does not change our analysis of the effects of disclosure on the product market, nor does it affect our normative analysis, because in it we focus on the case in which the platform discloses the information.[24]

[24] See Casadesus-Masanell and Hervas-Drane (2015) for a recent model of competition between platforms with privacy considerations.

A Proofs

Proof of Proposition 1

(i) The proof is based on a comparison of the first order conditions (1) and (2). Let

$$\zeta^1(p) \equiv \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial \pi(p,\theta_i)}{\partial p_i} f_i(\theta_i) d\theta_i$$

From (1), we have $\zeta^1(p^{\mathcal{P}}) = 0$.

Also, as by Assumption 4 we have $\frac{\partial^2 \pi_i}{\partial p \partial \theta_i} \geq 0$, then, for any increasing function h,

$$\int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial \pi(p^{\mathcal{P}},\theta_i)}{\partial p_i} h(\theta_i) f_i(\theta_i) d\theta_i \geq 0 \tag{12}$$

Now let

$$\zeta^n(p) \equiv \int_{\{\theta_i \in [0,\overline{\theta}]\}} \frac{\partial \pi(p,\theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i$$

From (2), we have $\zeta^n(p^{\mathcal{D}}(n)) = 0$. Using (12) with $h \equiv F^{n-1}$, one gets $\zeta^n(p^{\mathcal{P}}) \geq 0$.

Moreover, ζ^n is non increasing by concavity of the profit function, and so we obtain $p^{\mathcal{P}} \leq p^{\mathcal{D}}(n)$.

(ii) The proof of the second point obeys a similar logic. Let

$$\zeta^n(p) \equiv \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial \pi(p,\theta_i)}{\partial p_i} F^{n-1}(\theta_i) f_i(\theta_i) d\theta_i$$

For every n, $p^{\mathcal{D}}(n)$ is such that $\zeta^n(p^{\mathcal{D}}(n)) = 0$. By choosing $h_n \equiv F^{n-1}/F^{n-2} = F$, which is increasing, we get $p^{\mathcal{D}}(n) \geq p^{\mathcal{D}}(n-1)$. $\square$

Proof of Lemma 4

Let $p^{\mathcal{D}}(n)$ be the price if the platform chooses to implement a disclosure policy when n firms are on the market. First, notice that because $(p^{\mathcal{D}}(n))_{n\geq 1}$ is non decreasing and bounded (by $\overline{p}$), it has a limit, that we note $p^{\mathcal{D}}_{\infty}$.

From (2), we have

$$\int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial \pi(p^{\mathcal{D}}(n),\theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i = 0$$

Therefore, for any n and $\epsilon \in (0,\overline{\theta})$,

$$n\left(\int_0^{\overline{\theta}-\epsilon} \frac{\partial \pi(p^{\mathcal{D}}(n),\theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i + \int_{\overline{\theta}-\epsilon}^{\overline{\theta}} \frac{\partial \pi(p^{\mathcal{D}}(n),\theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i\right) = 0$$

Let

$$A_{n,\epsilon} \equiv n \int_0^{\overline{\theta}-\epsilon} \frac{\partial \pi(p^{\mathcal{D}}(n), \theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i$$

and

$$B_{n,\epsilon} \equiv n \int_{\overline{\theta}-\epsilon}^{\overline{\theta}} \frac{\partial \pi(p^{\mathcal{D}}(n), \theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i$$

Because $\frac{\partial^2 \pi}{\partial p_i \partial \theta_i} \geq 0$, one can write

$$\frac{\partial \pi(p^{\mathcal{D}}(n), 0)}{\partial p_i} \int_0^{\overline{\theta}-\epsilon} nF^{n-1}(\theta_i) f(\theta_i) d\theta_i \leq A_{n,\epsilon} \leq \frac{\partial \pi(p^{\mathcal{D}}(n), \overline{\theta}-\epsilon)}{\partial p_i} \int_0^{\overline{\theta}-\epsilon} nF^{n-1}(\theta_i) f(\theta_i) d\theta_i$$

The integral on the left side and on the right side is equal to $F^n(\overline{\theta}-\epsilon)$, and goes to zero as n goes to infinity. Therefore $\lim_{n\to\infty} A_{n,\epsilon} = 0$

By the same argument, we can provide a lower and an upper bound on $B_{n,\epsilon}$:

$$\frac{\partial \pi(p^{\mathcal{D}}(n), \overline{\theta}-\epsilon)}{\partial p_i} [F^n(\overline{\theta}) - F^n(\overline{\theta}-\epsilon)] \leq B_{n,\epsilon} \leq \frac{\partial \pi(p^{\mathcal{D}}(n), \overline{\theta})}{\partial p_i} [F^n(\overline{\theta}) - F^n(\overline{\theta}-\epsilon)]$$

Using the fact that $F^n(\overline{\theta}) = 1$, and that $A_{n,\epsilon} + B_{n,\epsilon} = 0$, we obtain

$$A_{n,\epsilon} + \frac{\partial \pi(p^{\mathcal{D}}(n), \overline{\theta}-\epsilon)}{\partial p_i} [1 - F^n(\overline{\theta}-\epsilon)] \leq 0 \leq A_{n,\epsilon} + \frac{\partial \pi(p^{\mathcal{D}}(n), \overline{\theta})}{\partial p_i} [1 - F^n(\overline{\theta}-\epsilon)]$$

By taking n to infinity, one gets

$$\frac{\partial \pi(p^{\mathcal{D}}_{\infty}, \overline{\theta}-\epsilon)}{\partial p_i} \leq 0 \leq \frac{\partial \pi(p^{\mathcal{D}}_{\infty}, \overline{\theta})}{\partial p_i}$$

If $\epsilon \to 0$, and by continuity of the derivative of the profit, we finally get

$$\frac{\partial \pi(p^{\mathcal{D}}_{\infty}, \overline{\theta})}{\partial p_i} = 0$$

i. e $p^{\mathcal{D}}_{\infty} = p^*(\overline{\theta})$.

Proof of Proposition 3

The platform's profit is

$$R^{\mathcal{D}}(n) = \int_{\underline{\theta}}^{\overline{\theta}} \pi(p^{\mathcal{D}}(n), \theta_i) f_{n-1:n}(\theta_i) d\theta_i = E[\pi(p^{\mathcal{D}}(n), \theta_{n-1:n})]$$

Notice that $\theta_{n-1:n}$ converges almost surely to $\overline{\theta}$. Then, by continuity of π and using Lemma 4, $\pi(p^{\mathcal{D}}(n), \theta_{n-1:n})$ converges to $\pi(p^*(\overline{\theta}), \overline{\theta})$ almost surely.

By the monotone convergence theorem, we can conclude that

$$\lim_{n\to\infty} R^{\mathcal{D}}(n) = E[\lim_{n\to\infty} \pi(p^{\mathcal{D}}(n), \theta_{n-1:n})] = \pi(p^*(\overline{\theta}), \overline{\theta}) > E[\pi(p^{\mathcal{P}}, \theta) = R^{\mathcal{P}},$$

which implies that there exists $\overline{n}$ such that $R^{\mathcal{D}}(n) > R^{\mathcal{P}}$ for all $n > \overline{n}$.

References

Acquisti, A., and H. Varian (2005): "Conditioning Prices on Purchase History," *Marketing Science*, 24(3), 367–381.

Anderson, E., and D. Simester (2010): "Price Stickiness and Customer Antagonism," *Quarterly Journal of Economics*, 125(2), 729–765.

Anderson, S. P., and A. de Palma (2009): "Information congestion," *RAND journal of Economics*, 40(4), 688–709.

Anderson, S. P., and R. Renault (2006): "Advertising Content," *American Economic Review*, 96(1), 93–113.

Armstrong, M. (2006): "Competition in Two-Sided Markets," *RAND Journal of Economics*, 37(3), 668–691.

Arozamena, L., and E. Cantillon (2004): "Investment incentives in procurement auctions," *The Review of Economic Studies*, 71(1), 1–18.

Athey, S., E. Calvano, and J. S. Gans (2012): "The impact of the Internet on advertising market for news media," *Working paper*.

Athey, S., and G. Ellison (2011): "Position Auctions with Consumer Search," *The Quarterly Journal of Economics*, 126(3), 1213–1270.

Athey, S., and J. S. Gans (2010): "The impact of targeting technology opn advertising markets and media competition," *American Economics Review*, 100(2), 608–613.

Bagnoli, M., and T. Bergstrom (2005): "Log-concave probability and its applications," *Economic Theory*, 26(2), 445–469.

Bergemann, D., and A. Bonatti (2011): "Targeting in Advertising Markets: Implications for Offline vs. Online Media," *RAND Journal of Economics*, 42(2), 417–443.

——— (2013): "Selling Cookies," Discussion paper, Cowles Foundation for Research in Economics, Yale University.

Bergemann, D., and M. Pesendorfer (2007): "Information structures in optimal auctions," *Journal of Economic Theory*, 137(1), 580–609.

Bergemann, D., and J. Välimäki (2002): "Information acquisition and efficient mechanism design," *Econometrica*, 70(3), 1007–1033.

Board, S. (2009): "Revealing Information in Auctions: The Allocation Effect," *Economic Theory*, 38(1), 123–135.

Calzolari, G., and A. Pavan (2006): "On the optimality of privacy in sequential contracting," *Journal of Economic Theory*, 130(1), 168–204.

Casadesus-Masanell, R., and A. Hervas-Drane (2015): "Competing with Privacy," *Management Science*, 61(1), 229–246.

Chen, Y., and C. He (2011): "Paid placement: Advertising and search on the internet*," *The Economic Journal*, 121(556), F309–F328.

Conitzer, V., C. R. Taylor, and L. Wagman (2012): "Hide and seek: Costly consumer privacy in a market with repeat purchases," *Marketing Science*, 31(2), 277–292.

Council of Economic Advisors (2015): "Big Data and Differential Pricing," Report.

Cowan, S. (2007): "The welfare effects of third-degree price discrimination with non-linear demand functions," *RAND Journal of Economics*, 38(2), 419–428.

de Cornière, A. (2013): "Search Advertising," Discussion paper.

Eso, P., and B. Szentes (2007): "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74(3), 705–731.

Esteban, L., A. Gil, and J. M. Hernandez (2001): "Informative Advertising and Optimal Targeting in a Monopoly," *Journal of Industrial Economics*, 49(2), 161–80.

Evans, D. S. (2008): "The Economics of the Online Advertising Industry," *Review of Network Economics*, 7(3), 359–391.

——— (2009): "The Online Advertising Industry: Economics, Evolution, and Privacy," *Journal of Economic Perspectives*, 23(3), 37–60.

Gaelotti, A., and J. L. Moraga-Gonzalez (2008): “Advertising, segmentation and prices,” *International Journal of Industrial Organization*, 26(5), 1106–1119.

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

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

Goldfarb, A., and C. Tucker (2011a): “Online Display Advetising: Targeting and Obtrusiveness,” *Marketing Science*, 30(3), 389–404.

——— (2011b): “Privacy Regulation and Online Advertising,” *Management Science*, 57(1), 57–71.

Hermalin, B., and M. Katz (2006): “Privacy, property rights and efficiency: The economics of privacy as secrecy,” *Quantitative Marketing and Economics*, 4(3), 209–239.

Iyer, G., D. Soberman, and J. M. Villas-Boas (2005): “The Targeting of Advertising,” *Marketing Science*, 24(3), 461–476.

Johnson, J. (2011): “Targeted Advertising and Advertising Avoidance,” *Working Paper*.

Johnson, J. P., and D. P. Myatt (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96(3), 756–784.

Kamenica, E., and M. Gentzkow (2011): “Bayesian Persuasion,” *American Economic Review*, 101(6), 2590–2615.

Krishna, V. (2009): *Auction theory*. Academic Press, San Diego.

Levin, J., and P. Milgrom (2010): “Online Advertising: Heterogeneity and Conflation in Market Design,” *American Economic Review, Papers and Proceedings*, 100(2), 603–607.

Lewis, T. R., and D. E. M. Sappington (1994): “Supplying Information to Facilitate Price Discrimination,” *International Economic Review*, 35(2), 309–27.

Perloff, J. M., and S. C. Salop (1985): “Equilibrium with product differentiation,” *The Review of Economic Studies*, 52(1), 107–120.

Piccione, M., and G. Tan (1996): “Cost-reducing investment, optimal procurement and implementation by auctions,” *International Economic Review*, 37(3), 663–686.

Png, I. P. L., and K.-L. Hui (2006): "The Economics of Privacy," *Handbooks in Information Systems*, 1(9), 471–497.

Roy, S. (2000): "Strategic segmentation of a market," *International Journal of Industrial Organization*, 18(8), 1279–1290.

Spiegel, Y. (2013): "Commercial software, adware, and consumer privacy," *International Journal of Industrial Organization*, 31(6), 702–713.

Stigler, G. J. (1980): "An Introduction to Privacy in Economics and Politics," Discussion paper.

Taylor, C. (2004): "Consumer Privacy and the Market for Customer Information," *RAND Journal of Economics*, 35, 631–650.

The Economist (2014): "Data: Getting to Know You," .

Van Zandt, T. (2004): "Information Overload in a Network of Targeted Communication," *RAND Journal of Economics*, 35(3), 542–560.

Vives, X. (2001): *Oligopoly Pricing: Old Ideas and New Tools*. The MIT Press, Cambridge.

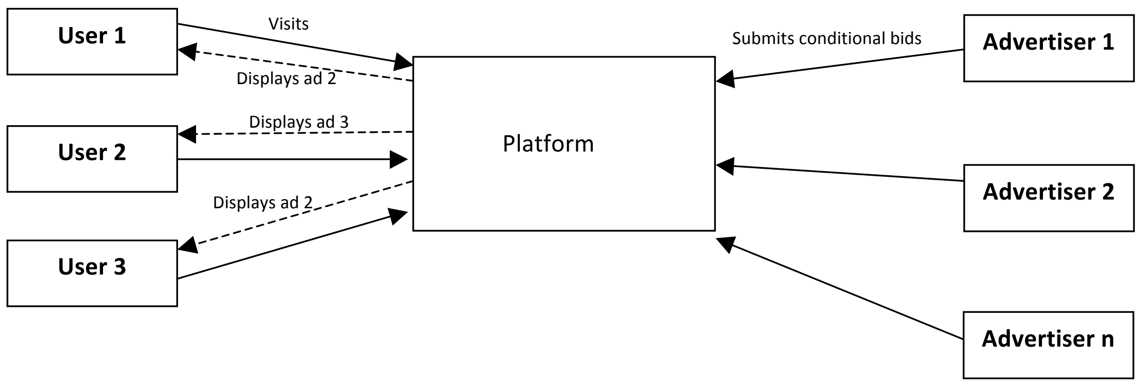

Figure 1: Market structure

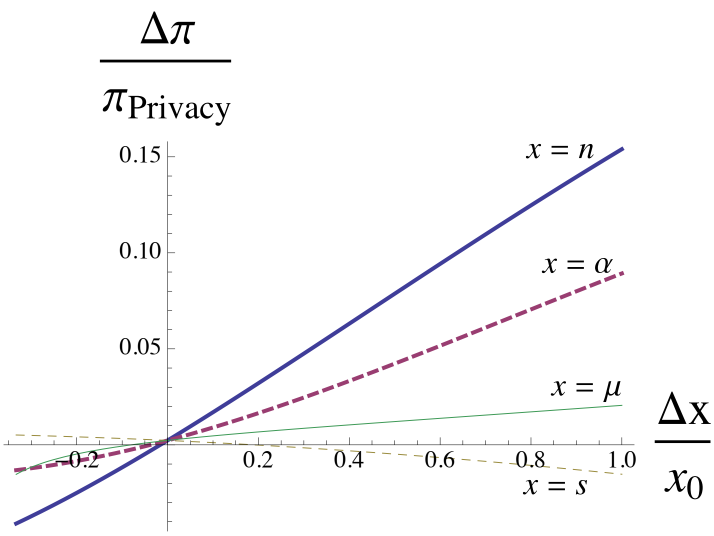

Figure 2: Effect of relative changes in parameter values on relative profitability of disclosure.

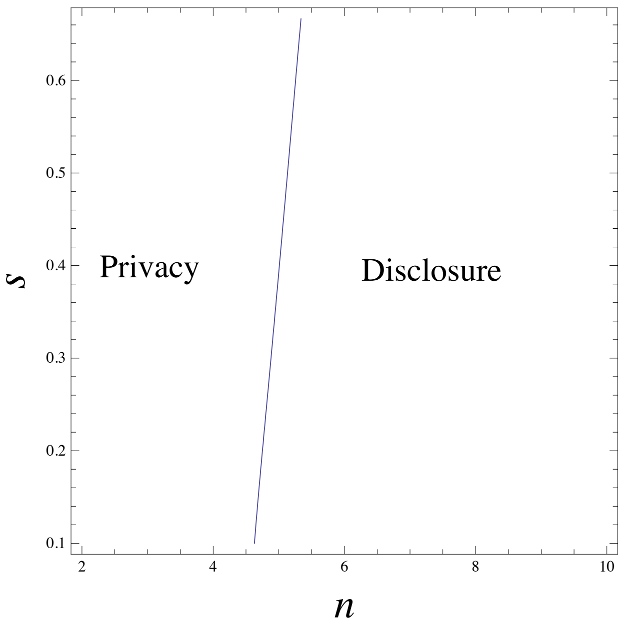

Figure 3: Policy in (n, s) plane

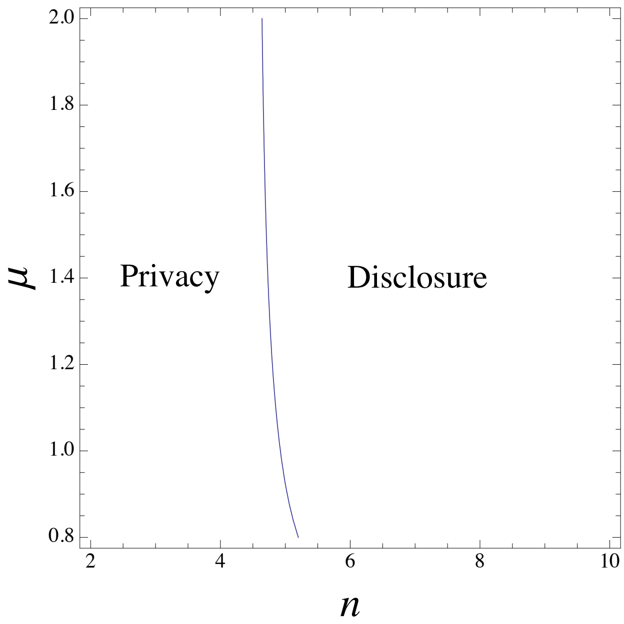

Figure 4: Policy in (n, μ) plane

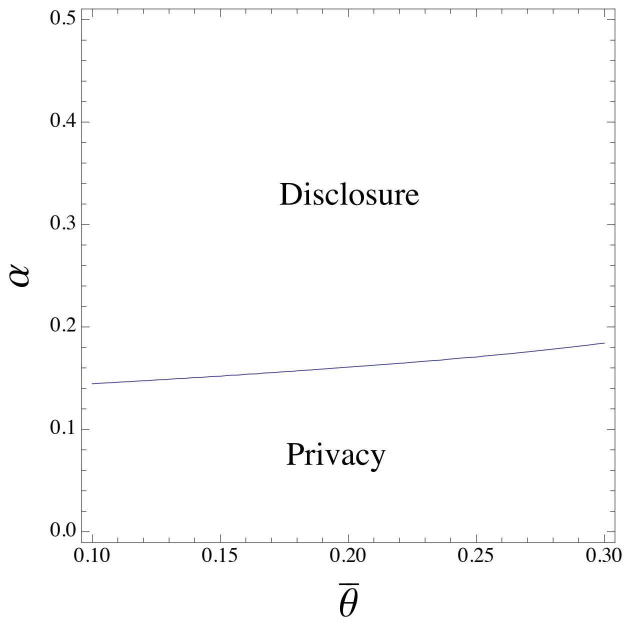

Figure 5: Policy in $(\bar{\theta}, \alpha)$ plane

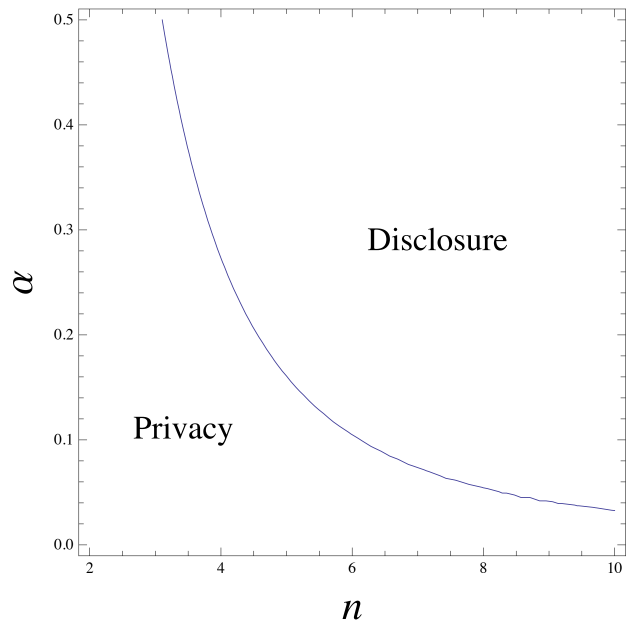

Figure 6: Policy in (n, α) plane

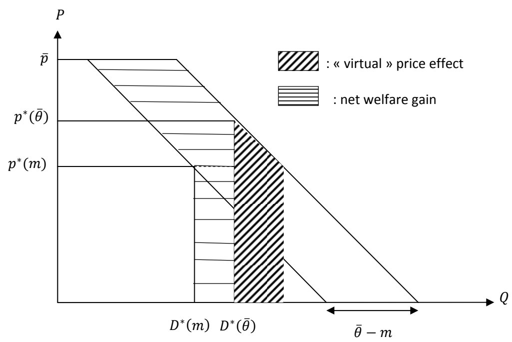

Figure 7: Linear demand

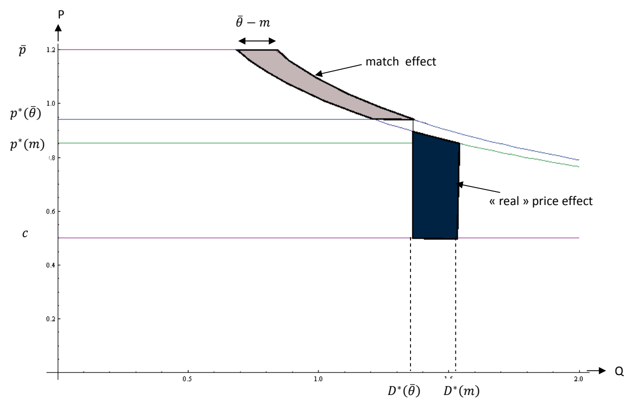

Figure 8: Constant elasticity

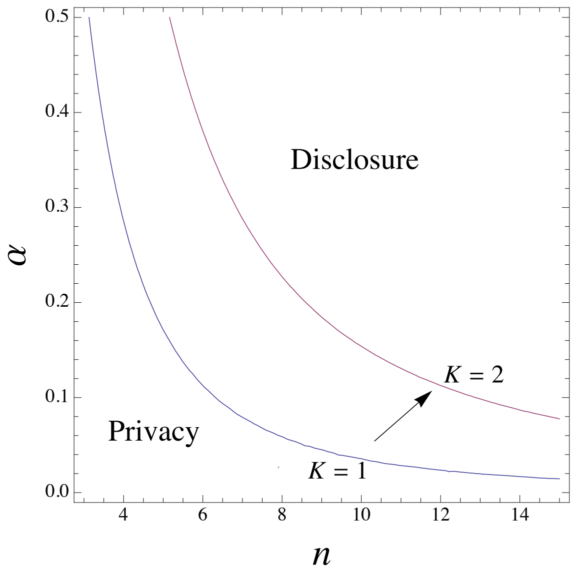

Figure 9: Platform policy in (n, α) plane after an increase in K

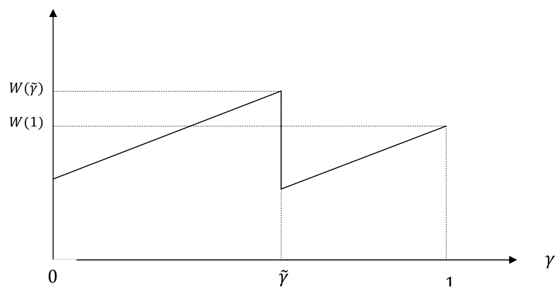

Figure 10: $\text{argmax}_\gamma W(\gamma) = \tilde{\gamma}$

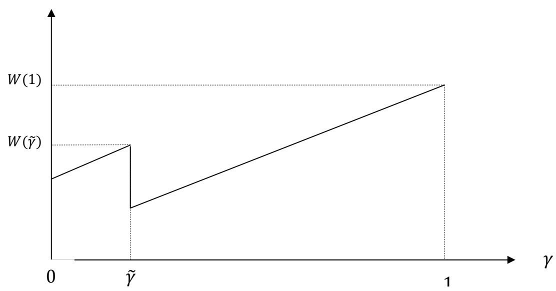

Figure 11: $\text{argmax}_\gamma W(\gamma) = 1$