ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

ONLINE ADVERTISING AND PRIVACY

Alexandre de Corniere and Romain De Nijs

Number 650
March 2013

Manor Road Building, Manor Road, Oxford OX1 3UQ

Online Advertising and Privacy*

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

March 18, 2013

Abstract

An online platform makes a profit by auctioning an advertising slot that appears whenever a consumer visits its website. Several firms compete in the auction, and consumers differ in their preferences. Prior to the auction, the platform gathers data which is statistically correlated with consumers' tastes for products. We study the implications of the platform's decision to allow potential advertisers to access the data about consumers' characteristics before they bid. On top of the familiar trade-off between rent extraction and efficiency, we identify a new trade-off: the disclosure of information leads to a better matching between firms and consumers, but results in a higher equilibrium price on the product market. We find that the equilibrium price is an increasing function of the number of firms. As the number of firms becomes large, it is always profitable for the platform to disclose the information, but this need not be efficient, because of the distortion caused by the higher prices. When the quality of the match represents vertical shifts in the demand function, we provide conditions under which disclosure is optimal.

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

JEL Classification: D4.

*For insightful discussions and suggestions we thank Bernard Caillaud, Chris Dellarocas, Gabrielle Demange, 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).

[†]Department of Economics and Nuffield College, University of Oxford. Email: adecorniere@gmail.com

[‡]Haas School of Business, UC Berkeley, and Ecole des Ponts ParisTech; Email: romaindenijs@gmail.com

1 Introduction

The online advertising industry has grown rapidly in the last decade.[1] With 33% of total revenue, display is the second most important type of online advertising after search. 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, but who can also help in the design of ads. Advertising networks (such as Google Adsense) and advertising exchanges (such as AdBrite) 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. 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 that we have in mind is the following (see Figure 1): web users (consumers) visit a website (a platform), which makes profit by selling a single advertising slot through an auction. Users are heterogeneous, in the sense that 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

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

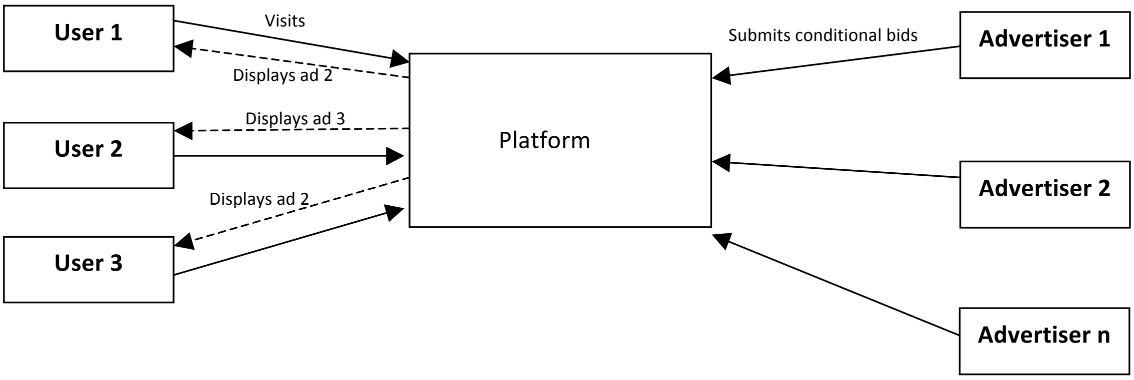

Figure 1: Market structure

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)).[2] 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), 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

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

of contracts) that he offers to the consumer. Thus, in all these papers, personal information is used 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. 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.

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 and on the number of advertising slots. With a large number of bidders (or a low number of slots) 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.

The model can also be interpreted as a model of 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 (2011) shows that when consumers actively search for products, targeting leads to more intense competition.

Another branch of the literature assume the prices of advertised goods are set exogenously and concentrates on other aspects of targeted advertising. 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. Finally, Van Zandt (2004), Anderson and de Palma (2009), and Johnson (2011) investigate the topic of privacy and congestion with consumers having limited attention.

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 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.[3] 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.

Hagiu and Jullien (2011) also study how intermediaries can use information on consumers characteristics in order to affect matching between firms and consumers. Among other results, they show that if the intermediary receives a fee every time a consumer visits an affiliated firm, the intermediary has an incentive to direct consumers towards firms that they would not have visited otherwise. Doing so, the platform manipulates the elasticity of the demands faced by its affiliated firms. In a different set-up, de Cornière (2011) shows that the optimal level of accuracy of the matching solves a trade-off between consumers participation and the level of firms' profit, which can be captured completely by the intermediary. In this paper, we suggest another rationale for voluntary imperfect matching by the platform, namely that implementing good matching (by disclosing information) might be too costly in terms of informational rent left to firms (See also Levin and Milgrom (2010)).

In section 2, we present the model in a rather general way. In section 3, we characterize symmetric equilibria under privacy and disclosure, and analyze the main implications of either regime. In section 4, we put more restrictions on the model and conduct a normative analysis. Section 5 presents some extensions of the basic model, and section 6 concludes. The appendix contains omitted proofs as well as several extensions of the model.

2 Model

There are n advertisers, who compete for a single slot on a platform's website. There is a continuum of consumers who 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 $[0, \overline{\theta}]$. 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 $\pi(p_i, \theta_i) = (p_i - c)D(p_i, \theta_i)$.[4] 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

[3]See also Bergemann and Pesendorfer (2007) and Board (2009).

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

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.

Welfare and a consumer's indirect utility, if the consumer is matched with firm i, are noted respectively $W(p_i, \theta_i)$ and $V(p_i, \theta_i)$. 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.[5] 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 : $\frac{\partial^2 \pi}{\partial p_i \partial \theta_i} \geq 0$*

We thus assume that 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) $D(p, \theta) = 1 - p^{\theta}$.

It is immediate to see that Assumptions 1-4 imply

$$\frac{\partial W}{\partial p_i} \leq 0, \quad \frac{\partial W}{\partial \theta_i} \geq 0, \quad \frac{\partial V}{\partial p_i} \leq 0, \quad \frac{\partial V}{\partial \theta_i} \geq 0$$

We assume that if a firm is matched with a consumer, it is in a monopoly situation with respect to that consumer. This strong assumption is made for expositional purpose: our model can accommodate some downstream competition among firms. However, it is essential that advertising attracts some consumers that would not have considered firm i without being exposed to an advertisement.

[5] 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))

It is also not essential whether profit is realized immediately or later on. Indeed we know that display advertising is less efficient than search advertising at generating clicks or immediate sales. Rather, display is often used as a brand-building device (See Manchanda, Dube, Goh, and K.Chintagunta (2006)).[6] Under an assumption of delayed sales, our analysis would carry through, provided that firms do not set a different price for consumers who click on an ad and for consumers who visit the website later on.

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.[7] 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.[8]

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.

3. Under *Disclosure*, each firm i learns θ_i. Under *Privacy*, firms do not learn the θ_i's.

[6]See Lewis and Reiley (2009), Chatterjee, Hoffman, and Novak (2003), Drèze and Hussher (2003), Rutz and Bucklin (2009) for more on the links between online advertising and sales.

[7]An alternative assumption is that the platform commits not to use this mapping to price its advertising slot.

[8]We discuss reserve prices in section 5.

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

The case of 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:

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

.

The case of 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.

Since 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, since 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,

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

where $f_{k:m}$ is the probability distribution function of the kth order statistic of θ_j among m.[9] 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

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

After some extra manipulations, we get :

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

$$\int_0^{\bar{\theta}} \frac{\partial \pi(p^{\mathcal{D}}(n), \theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i = 0. \tag{2}$$

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.*

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

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 $E[\pi(p^{\mathcal{D}}(n), \theta_{n:n})]$. Since $p^{\mathcal{D}}(n)$ 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. $\square$

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.

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 (see Appendix **??** for the detailed treatment of an example in which both phenomena may arise).

From a positive point of view, one would like to know under which conditions the platform is likely to adopt a disclosure policy. While we cannot predict whether the platform prefers privacy or disclosure for any n without imposing further restrictions on the demand function, the following polar case will prove useful in our normative analysis:

Proposition 3 *(i) When n goes to infinity, the platform's optimal policy is disclosure. (ii) The equilibrium price for the good tends to $p^*(\overline{\theta})$.*

The proof is in the appendix. Part (i) relies on the observation that as n goes to infinity, the informational rent of the winning bidder goes to zero, which implies that the platform captures the whole industry profit, which is greater under disclosure. The intuition for (ii) is that 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 $\overline{\theta}$.

4 Normative analysis

So far, the analysis does not allow us to determine the cases in which privacy must be favored over disclosure from a welfare point of view. In order to gain further insight, we specify the model by putting more structure on the family of demand functions $D(p, \theta)$. To facilitate the analysis, we also assume in the first part of this section that that $n = \infty$. We come back to the case with a finite number of firms in the second part of this section.

4.1 Infinite number of firms

In this subsection we assume that there is an infinite number of firms. Although we believe that the case with a large number of firms is the most relevant one when one thinks about advertising platforms such as Google or Facebook, let us stress the implications of this assumption. First, as shown in Proposition 3, the platform 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. Second, when $n = \infty$ firms charge a price $p^*(\overline{\theta})$ under disclosure. Therefore, our results will depend on the distribution F of types θ only through its mean m and the upper-bound of its support $\overline{\theta}$.

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 situation in which a consumer of type θ buys θ units of the product as long as the price is lower than the reservation price $\overline{p}$, and, in addition to these units, buys $q(p)$ units.

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 $2\eta(p^*(\theta),\theta) > \alpha(p^*(\theta))$).[10] Since 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 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,

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

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 θ. This leads to the following lemma, in which $\eta \equiv \eta(p^*(\theta), \theta)$ and $\alpha \equiv \alpha(p^*(\theta))$.

Lemma 4 *Disclosure (resp. Privacy) is welfare superior if, for every* θ, $\frac{p^*(\theta)}{\overline{p}-c} \leq ($ *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 4, 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 *If* q *is log-concave, disclosure is socially optimal.*

Proof: By Lemma 4, 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, since $\overline{p} > p$, $\frac{p-c}{p} = -\frac{Q}{pq'} \Rightarrow \overline{p} - c > -\frac{Q}{q'}$ 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.$\square$

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.[11] The log-concavity of the reliability function $1 - G$ is a property 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.

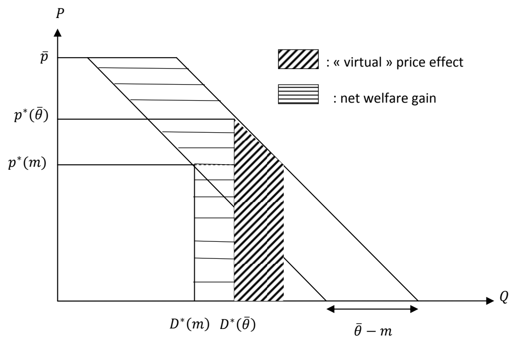

Figure 2: Linear demand

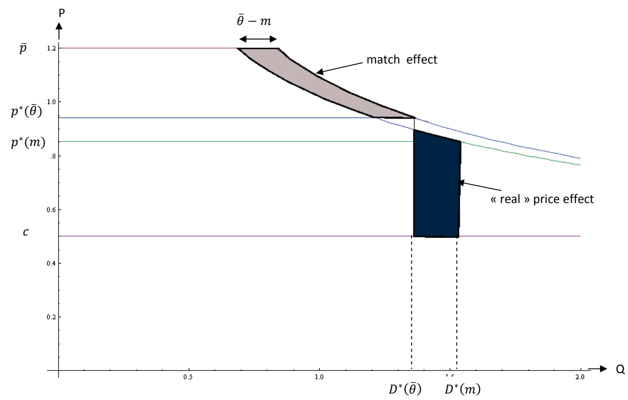

Figure 3: Constant elasticity

Examples: In Figure 2, 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. Figure 3 represents the inverse demand function when q has a constant price elasticity: $q(p) = p^{-\epsilon}$ if $p \leq \overline{p}$.[12] 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.

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

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

4.2 Finite number of firms

The above analysis relies on there being an infinite number of firms, which we think is a reasonnable approximation for large platforms. When n is finite, analytical results with the general formulation $D(p,\theta) = \theta + q(p)$ are hard to obtain, and so we use an even simpler formulation to gain some insights. We assume that the willingness to pay for product i is 1 with probability θ_i, and $v < 1$ with probability $1 - \theta_i$. The θ_i's are i.i.d. according to the cumulative distribution function $F(\theta) = \theta^\gamma$. The interested reader is referred to the appendix for a detailled treatment of this model. Let us just briefly mention the results that we obtain. First, as in the general analysis of section 3, switching from Privacy to Disclosure tends to lead firms to charge higher prices (1 instead of v), but also leads to better matches. Comparing the socially optimal policy and that chosen by the platform, we can determine when the platform discloses too much, too little or the right amount of information to advertisers. In particular, the platform discloses too little information when the number of advertisers is small and v is low, as in this case the informational rent left to the winning bidder would be too high under disclosure, whereas the low value of v makes it less socially costly to exclude some consumers through higher prices under disclosure.

5 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.

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 expected 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 his 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.[13]

Uniform multi-unit auctions. On many websites users can see several advertisements on each page, yet in the model we only allow for one advertisement to be displayed. What can we say if the platform was allowed to choose the number of ads? Suppose that the platform has K identical slots

[13]See Eso and Szentes (2007) for a model in which the auctioneer can sell information to bidders.

up for sale[14], 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.

Suppose that consumers have limited attention, and only inspect one advertisement, picked randomly among the K ads. This implies that the demand for product i is independent from the demand for product j.[15]

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, we show that 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 $\Theta = (\theta_1, ..., \theta_n)$.[16] 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

$$\pi_i(p_i, p_{K,n}) = \frac{1}{K} \int_{\hat{\theta}_K \in [0;\overline{\theta}]} \int_{\theta_i \in [\phi(\hat{\theta}_K, p_i, p_{K,n});\overline{\theta}} \Big(\pi(p_i, \theta_i) - \pi(p_{K,n}, \hat{\theta}_K) \Big) f_{n-K:n-1}(\hat{\theta}_K) f(\theta_i) d\hat{\theta}_K d\theta_i$$

The first order condition at a symmetric equilibrium is

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

Then there is a new trade-off, between the quantity sold (the number of slots) and the per-unit price (the highest losing bid). Interestingly, by choosing the number of slots the platform is able to fine tune the extent to which firms use the information under disclosure. To see this clearly, suppose that there is a large number of firms on the market. If the platform only sells one slot, firms expect that if they win the auction they will face a consumer with a high θ, and they charge a high price accordingly. Now, if the platform sells a large number of slots, winning the auction is less informative with respect to the consumer's type. Firms have an incentive to charge a lower price. We see then that selling more slots leads to a decrease in the equilibrium price of the goods, for reasons that have nothing to do with downstream competition on the product market.

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

[14] 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.

[15] Alternatively, one could assume that consumers inspect all the ads and that demands are independent.

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

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 indepenetly and identically uniformly distributed on $[0, 1]$, that there 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 $(\theta_i)_{i=1,\ldots,\infty}$, the platform generates a vector of signals $(s_i)_{i=1,\ldots,\infty}$, 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 ($\gamma = 1$).

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 4 and 5 depict the welfare as a function of γ.

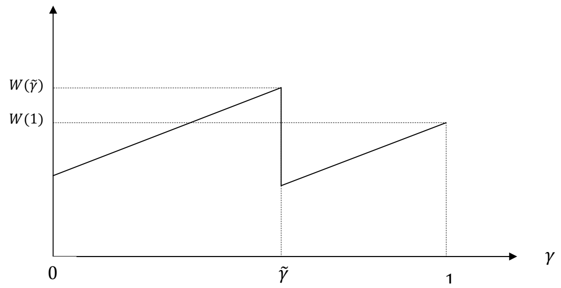

Figure 4: $\text{argmax}_\gamma W(\gamma) = \tilde{\gamma}$

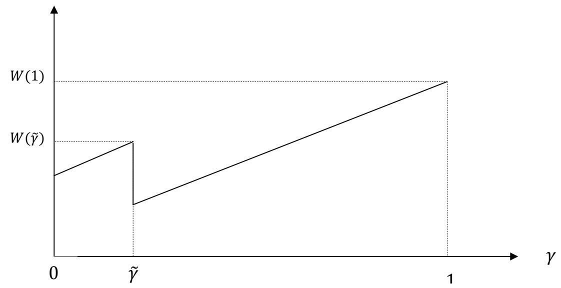

Figure 5: $\text{argmax}_\gamma W(\gamma) = 1$

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

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

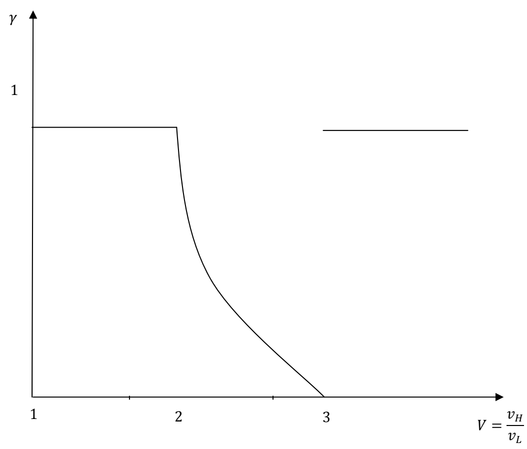

Figure 6: Socially optimal level of information revelation

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.

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 specific consumer's type they gain access to. While 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, we show in Appendix C 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 wore-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 (2011) 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.

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.

Competition between advertisers. In the current analysis, we have assumed that there is no competition once a firm has won the auction. Our insights can survive when competition is introduced. Consider for instance the case in which the ads of two competing products are displayed simultaneously on a consumer's screen. If the number of potential advertisers is large, those two products are very close substitutes. In addition, assume that these ads only convey information on the consumer valuations for each good but nothing about prices. If it is costly for consumers to search for prices by clicking and visiting the corresponding websites, then the well-known Diamond result tells us that in equilibrium firms charge their monopoly price and consumer visit only one firm. Consequently, our analysis with a monopolist firm could be directly translated in the above described duopolist competition with search costs. However, if prices were directly observable, the platform would have different incentives to reveal information. Indeed, if the platform reveals information, consumers are aware that the two displayed products are very similar. Then, Bertrand competition drives down final prices and decreases the platform's revenue. On the contrary, if the platform opts for privacy, the two products are more likely to be poor substitutes. This relaxes competition and ensures a greater revenue for the platform. However, unless the platform is willing to foster competition to boost consumers participation, it is clearly suboptimal for it to allocate two advertising slots for competing products if prices are observable.

6 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. Disclosing the information increases the sum of the platform's and advertisers' profits, but the platform has to leave an informational rent to the winning bidder. As the number of firms grows large, the later effect is dominated so that the platform always prefers disclosure. The disclosure of information by the platform induces a shift in the demand function, which leads to an increase in the equilibrium price of goods. When the shift is uniform along the demand curve, we give sufficient conditions on the elasticity of demand and of its slope for privacy or disclosure to be socially optimal.

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 (2011) 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, since in it we focus on the case in which the platform discloses the information.

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.

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," *Quarterly Journal of Economics*, 126.

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.

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

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.

CHATTERJEE, P., D. L. HOFFMAN, AND T. P. NOVAK (2003): "Modeling the click-stream: Implication for Web-based advertising effort," *Marketing Science*, 22(4), 520–541.

CHEN, Y., AND C. HE (2011): "Paid Placement: Advertising and Search on the Internet," *Economic Journal*, 121.

CONITZER, V., C. TAYLOR, AND L. WAGMAN (2011): "Hide and seek: costly consumer privacy in a market with repeated purchases," *mimeo*.

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. (2011): "Search Advertising," Discussion paper.

DRÈZE, X., AND F.-X. HUSSHER (2003): "Internet advertising: Is anybody watching ?," *Journal of Interactive Marketing*, 17(4), 8–23.

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.

Hagiu, A., and B. Jullien (2011): "Why do intermediaries divert search?," Discussion Paper 2.

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.

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, R., and D. Reiley (2009): "Retail Advertising Works! Measuring the Effects of Advertising via a Controlled Experiment on Yahoo!," *Yahoo! Research Technical Report*.

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

Manchanda, P., J.-P. Dube, K. Y. Goh, and P. K.Chintagunta (2006): "The Effect of Banner Advertising on Internet Purchasing," *Journal of Marketing Research*, pp. 98–108.

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

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

Rutz, O. J., and R. E. Bucklin (2009): "Does Banner Advertising Affect Browsing Paths? Clickstream Model Says Yes, For Some," *Working paper*.

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.

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*, vol. 1 of *MIT Press Books*. The MIT Press.

A Proofs

A.1 Proof of Proposition 1

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

$$\zeta^1(p) \equiv \int_0^{\overline{\theta}} \frac{\partial \pi(p, \theta_i)}{\partial p_i} f_i(\theta_i) d\theta_i$$

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

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

$$\int_0^{\overline{\theta}} \frac{\partial \pi(p^{\mathcal{P}}, \theta_i)}{\partial p_i} h(\theta_i) f_i(\theta_i) d\theta_i \geq 0 \tag{8}$$

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 (8) 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

$$\zeta^n(p) \equiv \int_0^{\overline{\theta}} \frac{\partial \pi(p, \theta_i)}{\partial p_i} F^{n-1}(\theta_i) f_i(\theta_i) d\theta_i$$

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 $p^{\mathcal{D}}(n) \geq p^{\mathcal{D}}(n-1)$. □

A.2 Proof of Proposition 3

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 since $(p^{\mathcal{D}}(n))_{n \geq 1}$ is non decreasing and bounded (by $\overline{p}$), it has a limit, that we note $p^{\mathcal{D}}_{\infty}$.

From (2), we have

$$\int_0^{\overline{\theta}} \frac{\partial \pi_i(p^{\mathcal{D}}(n), \theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i = 0$$

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

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

Let

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

and

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

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

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

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 $B_{n,\epsilon}$:

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

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

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

By taking n to infinity, one gets

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

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

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

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

The platform's profit is

$$E[\Pi^{\mathcal{D}}_{0,n}] = \int_0^{\overline{\theta}} \pi(p^{\mathcal{D}}(n), \theta_i) f_{n-1:n}(\theta_i) d\theta_i = E[\pi(p^{\mathcal{D}}(n), \theta_{n-1:n})]$$

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

By the monotone convergence theorem, we can conclude that

$$\lim_{n\to\infty} E[\Pi^{\mathcal{D}}_{0,n}] = E[\lim_{n\to\infty} \pi(p^{\mathcal{D}}(n), \theta_{n-1:n})] = \pi(p^*(\overline{\theta}), \overline{\theta})$$

B Finite number of firms

In this section we relax the assumption that there is an infinite number of firms. However, to carry out the analysis we need to further specify the demand function. This special case of the model allows us to study how the number of firms affects whether a monopolistic platform reveals too much or too little information.

More specifically, we assume that a consumer of type $(\theta_1, ..., \theta_n)$ has a probability θ_i to have a high willingness to pay for the good i ($v = v_H$), and a probability $1 - \theta_i$ to have a low willingness to pay for the same good ($v = v_L$). We also assume that the θ's are drawn independently from the c.d.f $F(\theta) = \theta^\gamma$ with $\theta \in [0, 1]$ and $\gamma > 0$. A higher value of γ means that the frequency of high values of θ increases. We also normalize v_H to 1 and denote $v_L = v \in [0,1]$. Let $m \equiv \frac{\gamma}{1+\gamma}$ be the expected value of θ_i.

B.1 Equilibrium under privacy

If the platform chooses not to disclose information to the firms, the equilibrium of the subgame is as follows:

Proposition 5 *If $m \leq v$, the equilibrium price is $p^{\mathcal{P}} = v$. The platform's revenue is $R^{\mathcal{P}} = v$. Consumers' surplus is $V^{\mathcal{P}} = m(1-v)$, and social welfare is $W^{\mathcal{P}} = (1-m)v + m$.*

If $m > v$, the equilibrium price is $p^{\mathcal{P}} = 1$. The platform's revenue is $R^{\mathcal{P}} = m$. Consumers' surplus is $S^{\mathcal{P}} = 0$, and social welfare is $W^{\mathcal{P}} = m$.

Proof: Under privacy, a firm expects to be matched with a consumer who has a probability m of having a high willingness to pay ($v_H = 1$). If $m \leq v$, firms prefer to serve every one rather than charge a high price and serve only the v_H consumers. If $m > v$, the opposite is true. The expressions of the platform's profit, consumers' surplus and social welfare are straightforward. □

B.2 Equilibrium under disclosure

When the platform chooses a disclosure policy, firms learn, for each consumer, the probability that he has a high valuation for its product. Firms are thus able to condition their bid on this information. The following lemma allows us to reduce the number of strategy profiles that can constitute an equilibrium.

Lemma 5 *When the platform discloses the information about consumers to firms, there is no equilibrium in which more than one firm set a price equal to v.*

Proof: Suppose by contradiction that at least two firms (firms 1 and 2) set a price $p = v$. Then they both bid v for every consumer, and thus they never make a positive profit. Now, if instead one of them (say, firm 1) deviates and sets a price 1 and bids θ_i, for every consumer there is a positive probability that it wins the auction and makes a profit equal to $\theta_1 - max_{i \neq 1} b_i^{\mathcal{P}}(\theta_i, p_i) > 0$. This is therefore a profitable deviation. □

Lemma 5 tells us that when the platform discloses the information about consumers, there can exist only two kinds of equilibrium:[17] a symmetric one in which all firms choose a high price 1, and an asymmetric one in which one firm chooses the low price equal to v while $n-1$ firms choose the high price 1. Let us first focus on the symmetric equilibrium.

Symmetric equilibrium Suppose that all firms choose the high price equal to 1. Their bidding function is therefore $b_i^{\mathcal{D}}(\theta_i, 1) = \theta_i$. Indeed, upon learning that a consumer is a good match with probability θ_i, firm i's expected profit conditional on being matched with the consumer is θ_i. The winner of the auction is thus the firm j such that $\theta_j = \max_l \theta_l$. The expected profit of firm i is therefore

$$\pi^{\mathcal{D}}(1) = \int_0^1 \int_0^{\theta_i} (\theta_i - \hat{\theta}_{-i})\gamma(n-1)\hat{\theta}_{-i}^{\gamma(n-1)-1}\gamma\theta_i^{\gamma-1} d\hat{\theta}_{-i} d\theta_i = \frac{\gamma}{(1+\gamma(n-1))(1+\gamma n)} \tag{9}$$

[17] This result depends on our normalization of $\overline{\theta}$ and of v_H to 1. If we had $v_L > \overline{\theta} v_H$, the only equilibrium would be for all the firms to charge v_L.

Indeed, firm i only wins the auction when θ_i is larger than all the θ_js, and firm i then pays the second highest bid. The term $\gamma(n-1)\hat{\theta}_{-i}^{\gamma(n-1)-1}$ is the pdf of the random variable equal to the highest realization of θ_j among $n-1$, while $\gamma\theta_i^{\gamma-1}$ is the pdf for θ_i.

It remains to check that there is no profitable deviation. The only possible deviation is to charge a low price v, which leads to a revenue of v per consumer who sees the ad. Firm i wins the auction as long as $v > \theta_j$ for all $j \neq i$. Its net profit is then:

$$\pi^{\mathcal{D}}(v) = \int_0^v (v - \hat{\theta}_{-i})\gamma(n-1)\hat{\theta}_{-i}^{\gamma(n-1)-1} d\hat{\theta}_{-i} = \frac{v^{\gamma(n-1)+1}}{1+\gamma(n-1)} \tag{10}$$

The deviation is not profitable if and only if

$$v < (\frac{\gamma}{1+\gamma n})^{1/(1+\gamma(n-1))} \tag{11}$$

The platform's revenue equals to the second highest bid:

$$R^{\mathcal{D}} = E[\theta_{n-1:n}] = \int_0^1 \theta dF_{n-1:n}(\theta) = \frac{n(n-1)\gamma^2}{(1+\gamma n)(1+\gamma(n-1))} \tag{12}$$

where $F_{n-1:n}(\theta) = \theta^{\gamma n} + n(1-\theta^{\gamma})\theta^{\gamma(n-1)}$ is the cdf of the second highest realization among all the θ_i's.

The expected surplus of consumers is zero: $V^{\mathcal{D}} = 0$. Indeed, the price is set so as to extract the whole surplus of buyers with a high willingness to pay, while other consumers do not buy.

The expected total welfare is the sum of the platform and firms' profits. It is equal to the expectation of the highest realization of θ_i:

$$W^{\mathcal{D}} = \int_0^1 \theta\gamma(n-1)\theta^{\gamma(n-1)-1} d\theta = \frac{\gamma(n-1)}{\gamma(n-1)+1} \tag{13}$$

The welfare is an increasing function of the number of firms, as a larger number of firms increases the probability for a consumer to be well matched. It converges to one when the number of firms tends to infinity.

Asymmetric equilibrium If (11) does not hold, the only equilibrium is an asymmetric one, in which one firm charges a low price v while the other firms charge the high price 1. Let us reorder the firms so as to have firm 1 charging the low price and $\theta_i < \theta_j$ for $j > i > 1$. The platform's revenue is the sum of three terms:

$$R^{\mathcal{D}}_{asym} = \int_0^v \theta dF_{n-1:n-1}(\theta) + \int_v^1 \int_0^v v dF_{n-2:n-1}(\tau)(\theta) dF_{n-1:n-1}(\theta) + \int_v^1 \int_v^1 \tau dF_{n-2:n-1}(\tau)(\theta) dF_{n-1:n-1}(\theta)$$

The first term corresponds to situations in which $\theta_{n-1} < \theta_n < v$. The firm who wins in this case is firm 1, and it pays the highest bid among the other firms (here, θ_n). If $\theta_{n-1} < v < \theta_n$, firm n wins the auction and the revenue equals v. Finally, if $v < \theta_{n-1} < \theta_n$, firm n wins and pays θ_{n-1}.

Using the appropriate expressions for $F_{n-1:n-1}$ and $F_{n-2:n-1}$, one gets

$$R^{\mathcal{D}}_{asym} = \frac{\gamma(-2+n)v^{(1+\gamma(-2+n))}}{(1+\gamma(-2+n))}$$

$$+v^{(1+\gamma(-2+n))}(1-n+(-2+n)v^{\gamma})(-1+v^{\gamma(-2+n)})$$

$$-\frac{(\gamma(-2+n)(-1+n)v^{-2\gamma}(-1+v^{\gamma(-2+n)})(\gamma v^{2\gamma} + v^{1+\gamma n}(-1+\gamma-\gamma n+(1+\gamma(-2+n))v^{\gamma}))}{((1+\gamma(-2+n))(1+\gamma(-1+n)))}$$

The social welfare equals

$$W_{asym}^{\mathcal{D}} = \int_0^v v dF_{n-1:n-1}(\theta) + \int_v^1 \theta dF_{n-1:n-1}(\theta)$$

i.e

$$W_{asym}^{\mathcal{D}} = \frac{\gamma(n-2) + v^{1+\gamma(n-2)}}{1+\gamma(n-2)}$$

B.3 Normative analysis

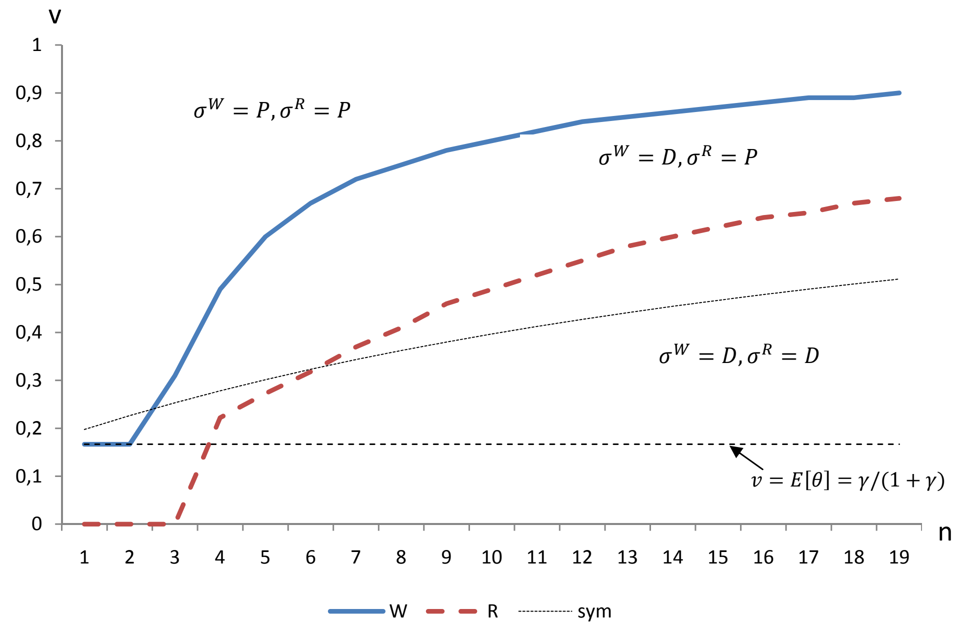

Figure 7: Socially vs privately optimal policy ($\gamma = 0.2$).

Figures 7, 8 and 9 illustrate how the optimal policies vary with v, n and γ. The socially optimal policy σ^W and the platform's optimal policy σ^R can be either Privacy ($\mathcal{P}$) or Disclosure ($\mathcal{D}$).

Above the thick dashed line, it is socially optimal to respect privacy ($\sigma^W = \mathcal{P}$) whereas under that line it is better to disclose the information ($\sigma^W = \mathcal{D}$). The platform's decision depends on the position of (n, v) with respect to the thick continuous line. Above it, the platform conceals the information ($\sigma^R = \mathcal{P}$), and it prefers disclosure under the line ($\sigma^R = \mathcal{D}$). The increasing thin dashed line represents the frontier between the symmetric and the asymmetric equilibrium (the symmetric equilibrium is played below the frontier).

The comparative statics with respect to n and v is the same whether we consider social welfare or the platform's revenue. As the number of bidders increases, the improvement in the matching quality under disclosure gets bigger, while the platform leaves a smaller informational rent to the winning bidder. Thus a higher n makes disclosure more desirable.

On the other hand, an increase in v leads to more privacy: as the population of consumers becomes more homogenous, excluding the low valuation consumers by disclosing the information (through a rise in the equilibrium price) is more costly, both socially and for the platform.

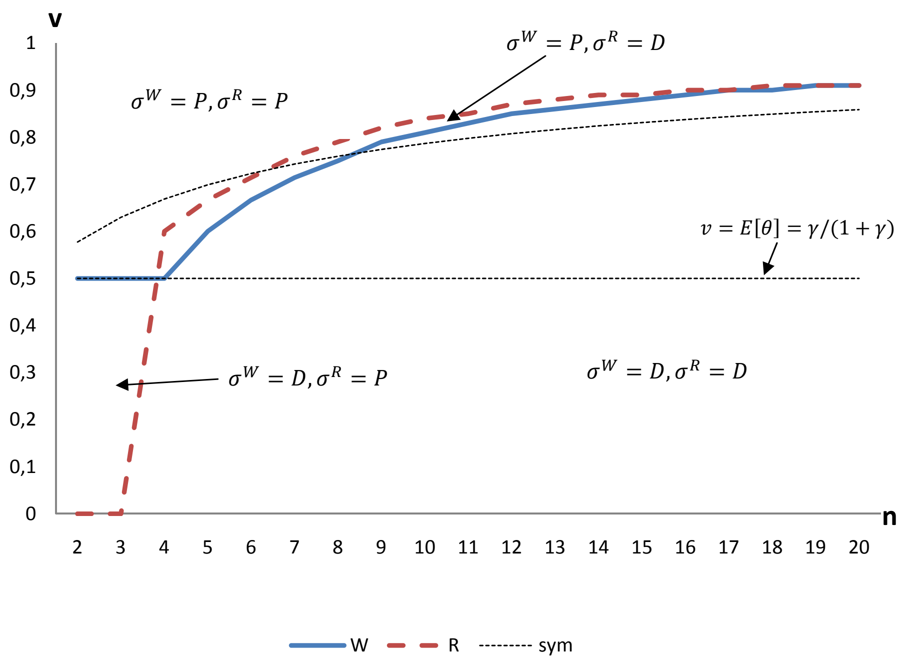

Figure 8: Socially vs privately optimal policy ($\gamma = 1$).

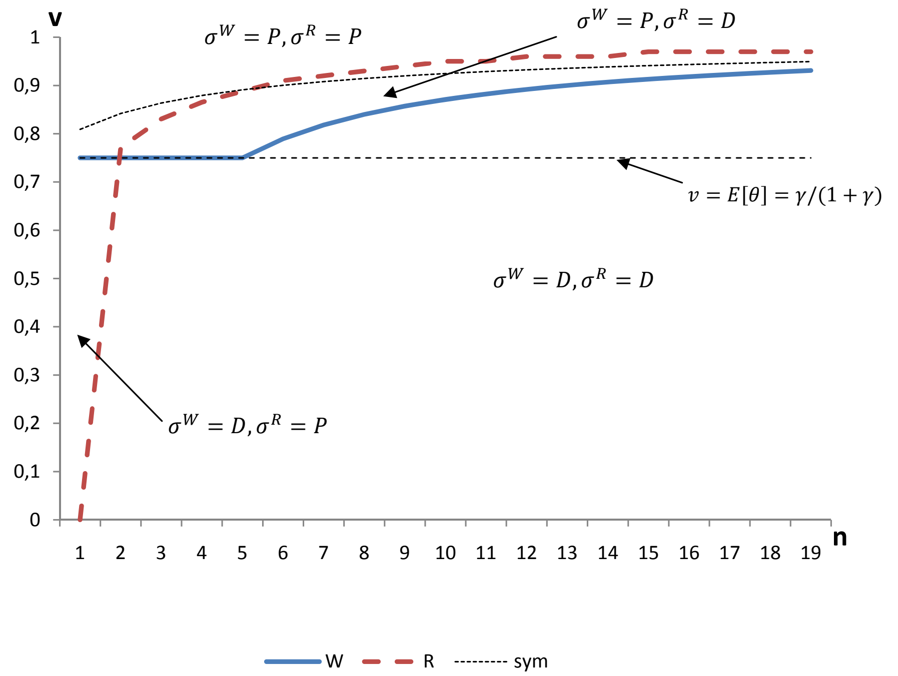

Figure 9: Socially vs privately optimal policy ($\gamma = 3$).

Notice that since the lines are not identical, there exists a set of parameters (n, v) such that the platform does not implement the socially optimal policy. It can disclose too much information ($\sigma^W = \mathcal{P}, \quad \sigma^R = \mathcal{D}$) when v and n take intermediate values, and γ is not too high ($\gamma \in \{0.2; 1\}$ in the example). The intuition is that the platform only takes v into account insofar as it determines its revenue under privacy. In particular, the platform does not fully internalize the social loss due to the exclusion of low valuation buyers. The number of firms must be large enough to reduce the informational rent, but low enough not to offset the socially negative effect of low valuation consumers' exclusion.

The platform may also disclose too little information ($\sigma^W = \mathcal{D}, \quad \sigma^R = \mathcal{P}$), for instance when both n and v are small. Indeed, if v is too small, firms always charge a high price, even under privacy (since $E[\theta] > v$), so that low valuation consumers are always excluded. It is thus clear that disclosure is socially optimal. However, if the number of bidders is small, leaving an informational rent is too costly for the platform, who therefore prefers privacy.

The "under provision" of information by the platform may also happen with larger values of v and n, provided γ is low enough (in our case, $\gamma = 0.2$). In such a situation, the ex ante probability that a match is good is low (0.16 for $\gamma = 0.2$). Switching to disclosure greatly enhances the efficiency of the matching: the probability of a good match would jump to 0.75 with 16 bidders. However, the informational rent left to the winning bidder is large for low values of γ (0.17 with $\gamma = 0.2$ and $n = 16$, i.e 45% of the value created by disclosure). Figure B.3 represents the informational rent as a function of γ, when $n = 16$. It is clear that if γ is small and v is not too small (so that it is not too tempting for the platform to give up a certain revenue of v under privacy), the platform prefers to conceal the information.

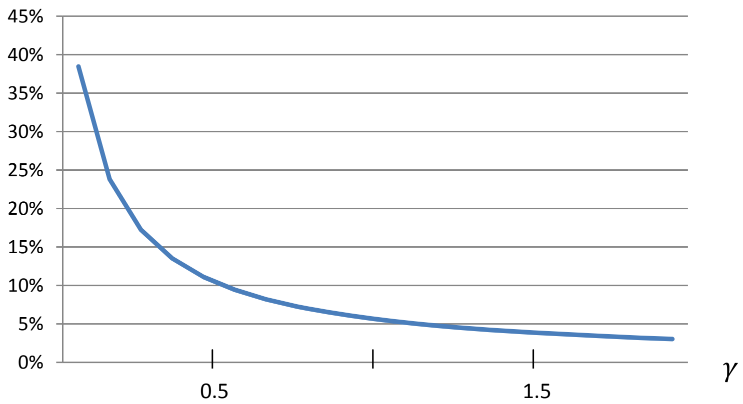

Figure 10: Relative information rent : $E[\frac{\theta_{n:n} - \theta_{n-1:n}}{\theta_{n:n}}]$, with $n = 16$.

The first order effect of an increase in γ is to make disclosure more desirable, both socially and for the platform. However, the private incentives to disclose the information grow more quickly than the public incentives, so that when γ is small there is a bias towards too much privacy provided by the platform, whereas for high values of γ the bias is towards too much disclosure. In the latter case, the improvement in the matching quality following disclosure is relatively modest, since there is a high probability that a consumer is of type v_H absent any information.

C The case of disclosure with customized prices

Suppose that firms are able to condition both their bid and their price on the consumer's type θ, whenever the platform chooses to reveal it. This means that the second and the third stage of the game in the basic model are inverted. To make the analysis simpler and to focus on an interesting case, we further assume that $\gamma = 1$ and that $v > 1/2$.

Under privacy, the situation is the same than without customized pricing: each firm bids v. Under disclosure, a firm i who learns that a consumer has a probability θ_i to be a good match will bid $b(\theta_i) = max\,(\theta_i, v))$. This bid is higher than the bid under privacy, and the bid under disclosure without customized pricing. As a result the platform always opts for a disclosure policy when firms can customize their prices. In addition, its revenue is always higher than without price discrimination.

As for firms, when customized pricing induces a shift for privacy to disclosure, they become better-off. Indeed the expected profit a firm under disclosure and customized pricing is strictly positive. It writes as follows:

$$\pi^{pd} = \int_v^1 (\int_0^v (\theta_i - v)\gamma(n-1)\hat{\theta}_{-i}^{\gamma(n-1)-1} d\hat{\theta}_{-i} + \int_v^{\theta_i} (\theta_i - \hat{\theta}_{-i})(n-1)\hat{\theta}_{-i}^{(n-2)} d\hat{\theta}_{-i}) d\theta_i \tag{14}$$

Firm i only wins the auction when θ_i is larger than all the θ_js, and firm i then pays the second highest bid. This second bid is equal to v when $\theta_{-i} \leq v$ and θ_{-i} otherwise. We can also deduce from this result that conditioning on being in the disclosure regime, firms are worse off when they engage in customized pricing. This is because the gap between their private valuation for the slot and the expected second highest bid shrinks when customized pricing is allowed. This result therefore delineates a new rationale for the non profitability of price discrimination in competitive environments.