DISEI - Università degli Studi di Firenze

Working Papers - Economics

Privacy and Persuasion: are we getting the best deal?

E. Carroni, L. Ferrari, S. Righi

Working Paper N. 09/2017

DISEI, Università degli Studi di Firenze
Via delle Pandette 9, 50127 Firenze, Italia
www.disei.unifi.it

Privacy and Persuasion: are we getting the best deal? *

Elias Carroni† Luca Ferrari‡ Simone Righi§¶

Abstract

We shed new lights on the desirability of privacy invasion by web–masters in online markets. We consider website's users uncertain about their need for a product and advertisers being offered banner spaces to show their commercials. The latter are designed as bayesian experiments. We show the emergence of different types of advertisement, ranging from fully informative to cheap talk. However, fully–informative banners are never showed if users can privately and costly acquire information about their state of necessity. As a result, when users' privacy is violated, they buy products they do not need and that they would not have bought if privacy were protected.

Keywords: Privacy, Bayesian Persuasion, Advertisement, Targeting.

JEL codes: D21, D80, L10, L51.

*We wish to thank Ennio Bilancini, Leonardo Boncinelli and the participants to the seminar at the University of Firenze, for their useful suggestions. Simone Righi acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 648693).

†Dipartimento di Scienze Economiche - Alma Mater Studiorum - Università di Bologna - 1, piazza Scaravilli, 40126 Bologna, Italy. Email: elias.carroni@unibo.it.

‡Department of Economics and Management, - Università di Firenze - 9, Via delle Pandette, 50127 Firenze, Italia. email: luca.ferrari.sdc@gmail.com.

§Department of Agricultural and Food Sciences - Alma Mater Studiorum - Università di Bologna - 50, viale Fanin, 40127 Bologna, Italy. Email: s.righi@unibo.it *and*

¶MTA TK "Lendület" Research Center for Educational and Network Studies (RECENS), Hungarian Academy of Sciences - H-1014 Budapest, Országház u. 30.

1 Introduction

Since the dawn of the digital era, the dominant business model of online platforms is fundamentally characterized by the provision of free services to users and of paying digital space to advertisers. More recently, the systematic collection and exploitation of users' private information –big data– allow platforms to offer targeted advertisements packages. From the advertisers' viewpoint, web channels are appealing as they make it easy to reach a high number of potential customers and to capitalize on the targeting power of web platforms. One key question –largely explored by both industrial and regulatory economics– is the impact of private information's availability to companies on surplus distribution and welfare.

To this regard, the answer of economic literature has evolved over time. Discussing the impact of privacy on economic interactions, Posner (1981) pointed out the inefficiency of privacy protection when the information is asymmetric. However, in the context of consumer–firm interaction, Varian (2002) highlights that the individual benefits of giving away private information can be outweighed by the lack of control on how firms use this information. This implies, in order to uncover the net welfare impact of different privacy setups, the need to analyze how firms' strategies vary in response to them. On the one hand, knowing more precisely users' preferences gives scope for price discrimination, which in a competitive setting induces firms to compete fiercely at the benefit of users.[1] On the other hand, data intermediaries collect information enabling an almost perfect matching between companies' products and users preferences (See Hagiu and Jullien 2011; Bergemann and Bonatti 2011, 2015). As pointed out by De Corniere and De Nijs (2016), this enhances efficiency creating new valuable trades for users who would have not bought otherwise. These results support the argument that invading privacy is economically desirable.

In this paper, a platform can be informed about each user's private characteristics, so to be possibly able to show to each user a personalized banner. The banner spaces are

[1] This is due to a "business–stealing" effect firstly discussed by Corts (1998). See Thisse and Vives (1988) for perfect price discrimination, Fudenberg and Tirole (2000) and Villas-Boas (1999) for pricing under customer recognition and Taylor (2004) for price discrimination in online markets.

auctioned to firms that want to advertise their products. In absence of private information about users' preferences, the platform is forced to show the same banner to all users. Hence, some users are not willing to click on the banner. Differently, when users' privacy is violated, the platform exploits users' private information, inducing everyone to click on a specific banner.

We consider advertisement as a Bayesian Experiment à la Kamenica and Gentzkow (2011) (hereatfer KG), allowing each advertiser to design its information structure as well as product's price. Each user is ex–ante uncertain about her state of necessity concerning the possible purchase of a good, which may turn out to be useful or not. Given this uncertainty, the Bayesian Experiment changes user's perception on her state of necessity.

The optimal pair ad–price trades off between two objectives, i.e., maximize probability of sale and surplus extraction. The trade–off can entail any possible type of advertisement between two extremes. On the one extreme, one finds fully informative advertisement in association with a high price. Namely, thanks to the ad, the user is revealed her state of necessity and she will buy the good only if needed. On the other extreme, it stands a Cheap–Talk advertisement associated to a lower price. In this case, the product will always be bought independently of users' needs. Gradually moving from the first to the second extreme, the price decreases together with advertisement informativeness.

From the user surplus viewpoint, the selection of the type of advertisement is not ex–post equivalent. Therefore, it is important to comprehend the characteristics of the equilibria that are most likely to emerge. We do this by introducing a cognitive cost that users face in order to individually get information about their state of necessity. We show that, since users prefer the Bayesian Experiment to individual information unravelling only if the price is sufficiently low, cognitive costs lead to the exclusion of fully informative equilibria.

Our result clearly shows that the fully informative targeting suggested by the economic literature is only a subset of a continuum of possible occurrences and it is actually the least robust one. When advertisement is not fully informative about state of necessity, the possibility of using users' private information may lead to the proliferation of wasteful transactions generating no additional social welfare. The advertisement sends signals that suggest the buyer to buy the product even when not needed, so that the transaction emerging perfectly offsets the loss of the user and the gain of the platform, with no additional social value. The number of these wasteful transactions dramatically increases going from privacy protection to privacy invasion, precisely as a consequence of improved targeting.

1.1 Related literature

Early approaches of the economists to the issue of privacy supported the view of Posner against privacy protection. As any friction reducing information in markets, Posner (1981) pointed out how less information generates inefficiency. Namely, in imperfect markets, allowing agents to hide information about themselves would induce them to hide negative traits, showing only positive one, with adverse selection consequences. More recently, the emergence of online markets and the so–called big–data, has made the economic analysis of consumers' privacy strictly interwined with the one on companies' targeting strategies. Indeed, the invasion of consumers' privacy by firms gives the latter the opportunity to enact pricing and advertising strategies.[2]

On the one hand, different degrees of knowledge may lead to segment the market through price discrimination. The seminal paper on perfect price discrimination of Thisse and Vives (1988) has been followed by the study of imperfectly discriminating strategies, e.g., loyalty schemes (Shaffer and Zhang (2000)), switching offers (Chen 1997; Fudenberg and Tirole 2000; Villas-Boas 1999) and online pricing (Taylor 2004). All these strategies require different levels of customer recognition. In sum, there is a consensus on a positive effect of the availability of information on consumer surplus, given by the fact that oligopolistic firms compete with lower prices to steal rival's business.

On the other hand, digital middlemen business model focuses on collecting and organizing private information of consumers in order to provide a highly effective matching between consumers' preferences and companies' products. At this regard, Hagiu and Jullien (2011) show that online intermediaries may have incentives to divert search, inducing more search than needed. With our view of advertisement, we look at a different kind of demand manipulation, focusing on a stochastic state of necessity rather than a match

[2] See Acquisti et al. (2016) for an exhaustive recent review on privacy in economics.

value. As pointed out by De Corniere and De Nijs (2016) and Bergemann and Bonatti (2011), targeting is efficient as it creates new opportunities for trade, by getting on–board consumers that would have been excluded otherwise. Furthermore, the introduction of a possibility of adopting a "hiding" technology may harm consumers, as shown by Belleflamme and Vergote (2016). Again, these results support the economic desirability of privacy invasion.

Our model interpretation of advertisement does not immediately match with the classification traditionally adopted by the literature, i.e., persuasive and informative advertisement.[3] The informative view (Stigler 1961; Nelson 1974) suggests the idea that individuals can lack information about products and that advertisement can the be seen as a mechanism to make individuals aware of products characteristics and thus on the relationship between them and consumers preferences. Another mechanism leading to informative advertisement can be the one of signaling (Nelson 1974; Milgrom and Roberts 1986), as advertisement - even when non directly informative - provide indirect experience to individuals about product characteristics and thus about its matching with individual preferences.

In its persuasive view (Robinson 1969; Kaldor 1950), advertisement induces the consumers to change their preferences and favor the advertiser's products. To put it differently, advertisement results in demand shifts. In this spirit, Friedman (1982) provides an oligopoly model in which increasing the intensity of advertisement is strategic choice to beat the rival. Similarly, in Schmalensee (1976) advertising is the only tool to compete in a market in which prices change infrequently and in Sutton (1991) ads create product differentiation.

In our paper, we re-interpret the informative advertisement approach in bayesian terms. Indeed, the bayesian advertisement has the only role of providing individuals with information (accurate or not) about their state of necessity. In this sense, the bayesian persuasion does not entail any ex-ante change in consumers preferences. However, there is a potential ex-post demand shift due to the change in perception about the state of necessity induced by a bayesian experiment.

[3] See Bagwell (2007) for a review on different classes of advertisement.

Our modelling choice is in the spirit of the recent stream of literature about information control, initiated by Kamenica and Gentzkow (2011) and Rayo and Segal (2010). In these models, a sender with commitment ability is able to design an experiment which reveals some information about decision maker's payoff relevant state. The sender designs the experiment in order to maximize the probability with which the decision maker takes his preferred action. Kamenica and Gentzkow (2011) study an application of their model in which an advertiser designs an experiment in order to inform consumers about the characteristics of the sponsored product whereas Rayo and Segal (2010) study a particular case in which the sender sells his experiment to an independent advertiser.

Relaxing the assumption on the commitment power of the sender, Hoffmann et al. (2014) study a model in which the latter may decide to acquire information about the personal characteristics of individuals and tailor messages that persuade them to take a particular action through selective information disclosure about horizontal aspects of a product. They find that the extent to which hyper-targeting[4] may harm consumers depends on the ability of firms to price-discriminate, on the competition between senders and on consumers' wariness.

The remainder of this paper is divided as follows. In Section 2 we formally present our model and discuss the optimal advertising behaviour. We further determine optimal platform behaviour under privacy protection and invasion in Section 2.2. We then assess comparatively the welfare characteristics of the two setups showing the increase of wasteful transactions (section 3). We then select select types of advertisements by introducing users' cognitive costs in Section 4. Section 5 concludes.

2 The Model

Consider a unit mass of users characterized by a state of necessity. Their utility from consumption is discovered only after buying and using a product. In this context, adver-

[4]Hoffmann et al. (2014) define hyper-targeting as "the collection and use of personally identifiable data by firms to tailor selective disclosure".

tisement may allow users to update their beliefs about their state of necessity.[5]

Users are showed advertisements through a website, owned and ran by a Web–Master. The latter offers banner spaces, each one containing up to one advertisement, to a set of potential sellers (advertisers). Whenever a user clicks on the banner, the advertiser filling the relative banner space must pay a per–click fee to the Web–Master.

The Web–Master wants to maximize profits by setting the per–click fee. At the beginning of the game each user is assigned by Nature an intensity of preference x and an unobservable state of necessity ω. Users who are satisfied with the banner click and receive a private shopping advice. Finally, purchase decisions are taken and payoffs are realized.

We consider two different privacy regimes. When *privacy is protected*, the Web–Master only knows the distribution of users' intensities of preference. This implies that the Web–Master can not discriminate between users. When *privacy is violated*, the Web–Master observes each user's preferences and can thus potentially discriminate with a personalised banner.

We study the model starting from the analysis of the purchase problem of a single user, thus deriving the optimal advertising rule (section 2.1). Then we close the model by studying the Web–Master's problem under the two different regimes (section 2.2). This allows us a welfare analysis performed in section 3. Finally, we provide an equilibrium selection argument that relies on users' ability to acquire perfect information about their true state of necessity by paying a cost in section 4.

2.1 Users and Advertisers

There is a unit mass of users characterized by intensity of preference x, assumed to be distributed according to an atom-less and continuous cumulative distribution function $G(x)$ with well defined density $g(x)$ on the support $[0, \bar{x}]$. Users are characterized by an unobservable state $\omega \in \{0, 1\}$. For $\omega = 1$, the user is in a state of necessity, while

[5]Users' uncertainty has to be intended completely unrelated with quality or intrinsic characteristics of the product. The latter are assumed to be perfectly communicated through advertisement, which is in this sense informative.

for $\omega = 0$ she is not. we denote as $\mu_0 \equiv \Pr(\omega = 0)$ the prior belief of not being in a state of necessity. The state of necessity affects the utility users experience from buying a product.[6] When the Web–Master has information on users, this has to be intended as information on the intensity of preferences, whereas the states of necessity remain ex-ante unknown by all players. Consider a product sold at price p. Then, the expected utility user x derives from buying the product is

$$U_B(x, p) = -p + (1 - \mu_0) \times x + \mu_0 \times 0. \tag{1}$$

In state $\omega = 0$ the user receives no utility from consuming the product, while she obtains a utility of x in the good state. The outside option of not buying is assumed without loss of generality to give zero utility, $U_{NB}(x, p) = 0$.

Surfing the web, users are exposed to a banner containing information about a product and its price. These information are used by the former to update their beliefs about the state of necessity. In particular, an advertisement is characterized by a pair of probability distributions π_0 and π_1 such that

$$\pi_0 \equiv \Pr(s = 0|\omega = 0) \text{ and } \pi_1 \equiv \Pr(s = 1|\omega = 1),$$

where $s \in \{0, 1\}$ denotes a shopping advice (for simplicity a message). Conditional on the observation of message s, users form the posterior belief $\Pr(\omega|s)$ using the Bayes rule.

Let us consider an advertiser y that optimal targets a generic user $\hat{x} \in [0, \bar{x}]$. Advertiser y would make profits whenever the user observes message $s = 1$ and buys the product at price $p(y)$. Thus, advertiser's profit is

$$\Pi(\pi, p(y)) = \tau(\pi) \times p(y) \tag{2}$$

where

$$\tau(\pi) \equiv \Pr(s = 1) = \mu_0(1 - \pi_0) + (1 - \mu_0)\pi_1$$

[6] In this model we do not explicitly account for tastes' and products' heterogeneity. This reflects the idea that the web–site is efficient in proposing relevant products. The only source of heterogeneity is the intensity of preferences, i.e., the user's pleasure from buying a product she needs.

is the probability with which the advertiser sends message $s = 1$, i.e., "buy". For user $\hat{x}$ to rationally follow the shopping advice, two conditions must be met:

$$U_{NB}(\hat{x}, p(y)|s = 0) > U_B(\hat{x}, p(y)|s = 0) \tag{IC0}$$

and

$$U_B(\hat{x}, p(y)|s = 1) \geq U_{NB}(\hat{x}, p(y)|s = 1). \tag{IC1}$$

These constraints imply that user $\hat{x}$ finds it optimal to act as suggested by the shopping advice. This is a Bayesian Persuasion problem *à la* KG in which a sender is able to decide, in addition to π, the price of the risky action "buy".

In state $\omega = 1$ users' and advertisers' preferences are aligned. Indeed, whenever the user in the state of necessity he obtains positive utility and thus wants to buy the product. Therefore, any $\pi_1 < 1$ is suboptimal as it would reduce the probability of buying when needed. As a result every optimal solution is such that $\pi_1 = 1$, which implies $\Pr(\omega = 0|s = 0) = 1$ so to trivially satisfy (IC0). Then, the optimal advertisement must leave the user with zero expected utility[7].

From the Bayes rule and the fact that $\pi_1 = 1$, (IC1) becomes

$$\frac{1 - \mu_0}{1 - \mu_0 \pi_0} \hat{x} = p(y). \tag{3}$$

The choice of the informational content of the advertisement uniquely defines an optimal price through relation (3). From the advertiser's viewpoint, all the pairs $(\pi_0, p(y))$ that satisfy (3) are optimal and payoff equivalent. This gives raise to a variety of advertisements that differ in terms of their information content. A Cheap–Talk (CT) advertisement always advices to buy ($s = 1$) so that $\tau = 1$. Hence, the optimal price will be lower than $\hat{x}$, i.e., from Equation (3), $p(y) = (1 - \mu_0)\hat{x}$. On the contrary, a Fully–Informative (FI) advertisement sends $s = 1$ only in the state of necessity i.e., $\tau = (1 - \mu_0))$. By giving full information to the use the advertiser can rise the price to the maximum one, i.e., $p(y) = \hat{x}$, from Equation (3). As the price of the sponsored good is a sender's strategic variable, the model endogenously generates a variety of optima advertisement-price pairs, summarised in the following proposition.

[7] If this were not the case, then the advertiser could either increase the price or decrease π_0 in order to rise the probability of sending message $s = 1$ through the banner.

Proposition 1. ***Advertisements and Information.***
The optimal Cheap–Talk advertisement targeting user $\hat{x}$ *is*

$$\pi_{CT} \equiv (\pi_0, \pi_1) = (0, 1) \text{ and } p_{CT}(y) = (1 - \mu_0)\hat{x}.$$

The optimal Fully–Informative advertisement targeting user $\hat{x}$ *is*

$$\pi_{FI} \equiv (\pi_0, \pi_1) = (1, 1) \text{ and } p_{FI}(y) = \hat{x}.$$

There also exist a continuum of optimal intermediate advertisements, which we call Persuasive, targeting user $\hat{x}$ *such that*

$$\pi_P \equiv (\pi_0, \pi_1) = (\pi_0, 1), \pi_0 \in (0, 1) \text{ and } p_P(y) \in (p(y)_{CT}, p(y)_{FI}).$$

For all optimal (π_0, π_1, p) *profits are always equal to*

$$\Pi = (1 - \mu_0)\hat{x} \tag{4}$$

Differently from KG in which uninformative signals may occur only when the receiver is more inclined towards sender's preferred action, here the advertiser can charge a smaller price in order to induce the user to always buy the product. Similarly, when the advertiser sets a fully informative experiment, this allows him to charge the maximum price the user is willing to pay. Finally, in between CT and FI, there exists a continuum of optimal prices associated with different informational contents, so that $\tau(\pi) \in (1-\mu_0, 1)$. In these equilibria, the price is lower with respect to the FI price, nevertheless, the probability of selling the product is higher as $\pi_0 < 1$. The probability of selling the product is maximized in the CT equilibria at the expense of a smaller price. However, equilibrium advertiser's profit do not depend on the informational content of the ad. All advertisement–price pairs leave the user, in expectation, with her reservation utility.

Although we solved the problem from the perspective of user $\hat{x}$, from Proposition 1 we can derive other, more general, insights. Indeed, if all users are exposed to the same banner, then user $\hat{x}$ will be indifferent between buying and not buying the product, whereas all users $x > \hat{x}$ are willing to buy and receive ex–ante positive utility (see equation (3)). To put it differently, they would be willing to pay a higher price regardless from ad's informational content. On the contrary, all users $x < \hat{x}$ would not even click on the banner, as purchase would give them negative ex–ante utility.

2.2 Privacy regimes

2.2.1 Privacy Protected

In section 2.1 we derived the optimal advertising targeting user $\hat{x}$. This simplifies Web–Master's problem as his choice is constrained to the set of optimal advertisements as described in Proposition 1.

In this section we consider the privacy–protected regime. That is to say, the Web–Master can not elicit each one private x and he is only informed about the distribution $G(x)$. Then, Web–Master's problem reduces to the choice of which user $\hat{x}$ to target. Recall from Proposition 1 that if the Web–Master targets user $\hat{x}$, then all users $x \geq \hat{x}$ will click on the banner. Formally, Web–Master's problem under the privacy protected regime is

$$\max_{\hat{x}} [1 - G(\hat{x})] f(\hat{x}) \tag{5}$$

where $f(\hat{x})$ is the maximum fee the Web–Master can set when hiring an advertiser targeting user $\hat{x}$ and $f(\hat{x}) = (1 - \mu_0)\hat{x}$. Clearly the fee can not be higher than advertiser's profit[8]. Intuitively, Web–Master's problem reduces to the choice of whom user to target. In particular, the Web–Master faces the following trade-off. By targeting a user with a lower preference intensity, the Web–Master increases the number of clicks and reduces the fee. On the contrary, by targeting a user with more intense preference, she is able to increase the fee at the expense of less clicks. Nevertheless, the choice of whom to target depends simply on the distribution $G(x)$ whereas the Web–Master is indifferent between any kind of optimal ad targeting user x^*.

Proposition 2. *When privacy is protected the Web–Master targets user*

$$x^* = argmax_{\hat{x}} [1 - G(\hat{x})] f(\hat{x})$$

and sets fee $f(x^*) = (1 - \mu_0)x^*$. *Finally, the Web–Master is indifferent between any optimal advertising targeting user* x^*.

[8] In other words, we implicitly assume that the advertisement markets is competitive and, in expectation, advertisers' profits are zero.

Proof. To see that a solution to Web–Master's problem exists notice that the function $[1-G(\hat{x})]f(\hat{x})$ is continuous in the closed and bounded interval $[0,\bar{x}]$, thus, by the extreme value theorem, it attains a minimum and a maximum which we denote with x^*. The fact that the Web–Master is indifferent between any kind of advertisements comes from the fact that the optimal ads targeting user x^* are payoff equivalent and led advertiser's profit to $(1-\mu_0)x^*$. ■

Clearly, the choice of whom user to target depends on the distribution.

Example 1. *Suppose that x is uniformly distributed on $[0,\bar{x}]$ so that the problem is concave in $\bar{x}$. From the first order conditions we have that*

$$\frac{1-G(x^*)}{g(x^*)} = x^*$$

and thus from the fact that $G(x) = \frac{x}{\bar{x}}$

$$\frac{\bar{x}-x^*}{\bar{x}}\bar{x} = x^* \implies x^* = \frac{\bar{x}}{2}.$$

Therefore, the Web–Master engages the advertiser targeting the median user and he gets the per-click fee $f(x^*) = (1-\mu_0)\frac{\bar{x}}{2}$.

Observe also that the Web–Master can not increase his payoff by randomizing his choice of which advertiser to hire as there exists an optimal user he can target. The fact that the Web–Master is indifferent between any optimal advertisement does not allow us to discriminate from CP to FI equilibria. Nevertheless, in terms of social welfare, the choice of which kind of advertisement to employ is essential. We carry out the welfare analysis in section 3 whereas now we turn on the privacy violated case.

2.2.2 Privacy Violated

We say that Privacy is violated whenever the Web–Master is able to elicit each one private preference intensity x. Nevertheless, as we stressed, the uncertainty beyond the state of necessity can not be unraveled. Whenever privacy is violated, the Web–Master is able to discriminate between users choosing for each $x \in [0,\bar{x}]$ the optimal advertisement that leaves her with zero expected utility. This maximizes each advertiser's profits and, in turn, the per-click fee paid to the Web–Master.

Proposition 3. *When privacy is violated, each user $x \in [0, \bar{x}]$ is shown an optimal advertisement that leaves her with zero expected utility. For each user x, the Web–Master is indifferent between any optimal advertisement.*

Proof. The proof is direct consequence of Proposition 1. ■

As in the privacy protected regime, the Web–Master is indifferent between which kind of advertisement to show. From his point of view, his profits do not depend on the choice of the informativeness of the ad.

Privacy violation intuitively increases Web–Master's profits. To see why, notice that with respect to the previous regime, now all users clicks on the banner. Each user is left with zero expected utility and the Web–Master extracts all surplus from the advertisers. Therefore, switching to a privacy violated regime reduces users' welfare through a mechanism which is analogous to the perfect-price discriminant monopolist. Nevertheless, according to the information content of the ad showed to users, the web-site may generate transactions that do not generate surplus as we explain in the following section.

3 Welfare

In the previous sections we discussed two simple scenario in which the Web–Master may or may not exploit users' private information in order to practice hyper-targeting strategies. We showed that the Web–Master is indifferent between the kind of advertisement to display as Cheap–Talk ads guarantee the same rent of a Fully–Informative one.

Social Welfare is composed of the sum of users' welfare and webmaster welfare, and shows properties which on the type advertisements chosen.

On the users' side, for each user $x \in [0, \bar{x}]$ the welfare is given by her ex–ante expected utility, which depends on the realization of the shopping advice $s \in \{0, 1\}$, i.e.:

$$\begin{aligned} U_B(x, \pi, p) &= \mu_0[-p(1-\pi_0) + 0 \times \pi_0] + (1-\mu_0)[\pi_1(x-p) + 0(1-\pi_1)] \\ &= \mu_0[-p(1-\pi_0)] + (1-\mu_0)\pi_1(x-p). \qquad (6) \end{aligned}$$

We can separate (6) in waste $\equiv \mu_0(1-\pi_0)(-p)$ and consumption $\equiv (1-\mu_0)\pi_1(x-p)$. Intuitively, the waste is the consumption of the product in the non-necessity state and vice-versa. Waste is created if and only if advertisement is not fully informative, i.e. if and only if $\pi_0 < 1$. In other terms, users' welfare is given by

$$W_u \equiv \int_i^{\bar{x}} (\text{waste} + \text{consumption}) di$$

where $i = 0$ when privacy is violated (since all users click on the banner) and $i = \hat{x}$ when privacy is protected since only users $x \geq \hat{x}$ click on the banner and behave as suggested.

On the other hand, Web–Master's welfare is simply given by his profits that are

$$W_\Pi \equiv \int_0^{\bar{x}} (1 - \mu_0) x dx,$$

if privacy is violated, and

$$W_\Pi \equiv \int_{x=\hat{x}}^{\bar{x}} (1 - \mu_0) \hat{x} dx,$$

when privacy is protected. Indeed, in this latter case the Web–Master is not able to discriminate between users and sets a fixed fee.

In the next proposition we underline the main effect switching from a privacy protected regime to a privacy violated one.[9]

Proposition 4. *Let $\pi^*(\hat{x})$ be the informational content of the optimal advertisement showed by the banner chosen in the privacy-protected regime. Then, when privacy is violated and all other hired advertisers adopt the same advertising strategy $\pi^*(x)$ then:*

1) Privacy violation always entails a reduction of users' surplus and an increase of Web–Master's profits.

2) Privacy violation weakly increases waste.

Proof. When privacy is protected the Web–Master targets user $\hat{x}$ and hires the advertiser supplying the ad $\pi^*(\hat{x})$ where the price $p(\hat{x})$ is determined by means of (1). Web–Master's profits are instead given by $[1 - G(\hat{x})](1 - \mu_0)\hat{x}$ regardless of the informational content

[9] We consider, for the sake of exposition, only the case in which switching from one regime to the other one, the advertising strategy is constant. However the results would extend to the cases where the Web–Master sell his banner spaces to advertises adopting different advertising strategies about which he is indifferent.

of π^*. Users' welfare is positive as from Proposition 1, only user $\hat{x}$ is made indifferent by the choice of the banner, while all users $x \geq \hat{x}$ are made ex-ante better off an. For each user $x \geq \hat{x}$ - in expectation

$$\text{waste} \equiv \mu_0(1-\pi_0)(-p(\hat{x})) \leq consumption \equiv (1-\mu_0)\pi_1(x - p(\hat{x}))$$

and therefore $W_u > 0$.

Point 1) then directly follows from Proposition 3. Since all users are left with zero expected utility when privacy is violated, then $W_u = 0$ and Web–Master's profits are

$$W_\Pi \equiv \int_0^{\bar{x}} (1-\mu_0)x dx = (1-\mu_0)\frac{\bar{x}^2}{2}$$

regardless from the informational content of π^*.

Point 2) then follows from Propositions 1 and 3: when privacy is protected $G(\hat{x})$ users do not click on the banner. When privacy is violated, waste increases exactly of $G(\hat{x})(1-\pi_0)\mu_0$ users, that is the fraction of users that will buy the product in state of non-necessity that did not click on the banner when privacy was protected. Waste increases if and only if $\pi_0 < 1$. ■

When privacy is violated the loss of users' welfare depends precisely on the ability of the Web–Master to engage in hyper-targeting strategies thus always extracting all surplus from advertisers and users. Nevertheless, surplus distribution among advertisers and users changes with the type of advertisement strategy adopted by the producer. It is possible to show the following:

Proposition 5. *Suppose that privacy is violated and advertisers are symmetric in the sense that they all play $\pi_{CT}(x), \pi_P(x)$ or $\pi_{FI}(x)$. Then:*

i) when $\pi = \pi_{CT}(x)$, all risk rests on users;

ii) when $\pi = \pi_P(x)$, users bear risk $1 - \pi_0(x)$ and advertisers bear risk $\pi_0(x)$;

iii) when $\pi = \pi_{FI}(x)$, all risk rests on advertisers.

Proof. Suppose that, for all $x \in [0, \bar{x}]$, $\pi = \pi_{CT}(x)$. Then $\tau = 1$, advertisers bear no risk as they sell with probability one whereas users make a mistaken purchase with probability μ_0, this proves i).

Suppose now that $\pi = \pi_{FI}(x)$. Users bear no risk since $\pi_1(x) = \pi_0(x) = 1$ and they only buy in state $\omega = 1$. On the contrary, in state $\omega = 0$ advertisers sell no product and suffer a loss of $-(1 - \mu_0)x$ which is exactly the amount due to the Web–Master for user's click on the banner. This proves iii).

Finally, when $\pi = \pi_P(x)$, $\pi_0 \in (0, 1)$, ii) follows from i) and iii). Indeed, $\pi_0(x)$ is the probability with which user x makes the correct choice in state $\omega = 0$. This is exactly the probability with which the user does not buy the product and therefore the probability with which the advertiser incurs a loss due to the fee he has to pay to the Web–Master for user's click. ■

Under any possible optimal advertising strategy π, at least one party between users and advertisers suffer a loss in state μ_0. This is due to the fact that, from advertisers point of view, truthful advertising is costly in the sense that they may end up paying the fee without having sold any product. On the contrary, on users' side, Cheap–Talk or Persuasive advertisements lead them to buy the product in the state of non necessity with positive probability. In the following section we argue that the risk is more likely to be a burden on users' side.

4 Equilibrium selection: Costly Information Acquisition

In the previous section we showed how surplus' distribution depends on the informativeness of the equilibrium advertisements showed to users. In particular, we showed that Fully–Informative advertisements do not generate any social waste (no users buy a product she does not need) and the risk of incurring a loss is moved onto advertisers. Although in all equilibria users are left with zero expected utility, Fully–Informative equilibria seem preferable from users' point of view since, in those equilibria, no user incurs in ex-post losses. In addition, recall that whenever a user buys a product she does not need, she would have been better off by investing her money in her outside option although we normalized the payoff deriving from consuming the outside option to zero.

One natural question however is the following. How do different advertisement equilibria arise? Which advertisements are more likely to be sent in equilibrium?

In order to answer to this equilibrium selection problem, we introduce a cost $c \in (0, \infty)$ that users can pay in order to learn their true state of necessity. Thus, at the beginning of the game, users face an information acquisition stage. If user x decides to pay the cost, she becomes perfectly informed about her state of necessity. The cost may have several interpretations and thus we discuss its generic role. For instance, it could be understood in psychological terms to underline that introspection about our necessities might be costly and so on. Whenever user x pays the cost, she buys the product if and only if her state of necessity is $\omega = 1$. Thus, whenever a user decides to pay the cost, her expected utility $U_C(\cdot)$ when the price of the advertised good is p is

$$U_C(x, p) = (1 - \mu_0)(x - p) - c.$$

Notice that the utility user x derives from paying the cost c does not depend on the information content of the advertisement because her decision ultimately depends on the knowledge of ω, which she acquired by paying the cost. If the user is going to pay the cost, she will buy the product only if $U_C(x, p) \geq U_{NB} = 0$, that is to say whenever

$$p \leq \frac{(1 - \mu_0)x - c}{1 - \mu_0}.$$

This means that an advertiser optimally targeting user x would set $p = \frac{(1-\mu_0)x-c}{1-\mu_0}$ and his expected profits would be

$$(1 - \mu_0) \times \frac{(1 - \mu_0)x - c}{1 - \mu_0} < (1 - \mu_0)x.$$

It is clear that, from advertisers' and Web–Master's point of view, it is not optimal for users to pay the cognitive cost. Both parties prefer users not to pay the cognitive cost so that advertisers can charge a higher price and, consequently, the Web–Master can set a higher per–click fee. This means that, in addition to the constraints derived in section 2.1, the optimal advertising mechanism requires one additional constraint, namely

$$U_B(x, p, \pi) \geq U_C(x, p)$$

and expanding the expected utilities

$$\mu_0(1 - \pi_0)(-p) + (1 - \mu_0)(x - p) \geq (1 - \mu_0)(x - p) - c. \tag{ICc}$$

The additional constraint simply says that user x must weakly prefer buying the product at price p when the banner characterized by information content π sends message $s = 1$ rather than acquiring perfect information at cost c. Observe that we have already considered that π_1 must be equal to one[10]. This also implies that IC0 is trivially satisfied and we can therefore study the optimal advertising rule targeting user x.

Proposition 6. *Consider a generic user $x \in [0, \bar{x}]$. Then, the optimal advertising rule satisfies*

$$\pi_0 = \min\left\{0, 1 - \frac{c}{p\mu_0}\right\} \text{ and } p(x) = \begin{cases} (1-\mu_0)x & \text{if } \pi_0 = 0 \\ x - \frac{c}{1-\mu_0} & \text{if } \pi_0 \in (0,1). \end{cases} \tag{7}$$

Proof. The advertiser maximizes his profits which are given by $\tau \times p(x)$ Intuitively, ICc must be binding. Otherwise the advertiser can either increase the price or reduce π_0 in order to increase his profits by raising the price or increasing the probability of selling the good. Thus, from ICc it directly follows that

$$\pi_0 = \min\left\{0, 1 - \frac{c}{p\mu_0}\right\}$$

since π_0 must be a probability. In particular, whenever $1 \geq \frac{c}{p\mu_0}$, the advertiser would set $\pi_0 = 0$. In order to pin down the price, recall that also the IC1 constraint must be satisfied with equality. Thus, from (3) we have that $p(x) = (1-\mu_0)x$ if $\pi_0 = 0$ and $p(x) = x - \frac{c}{1-\mu_0}$ if π_0 is interior. ■

To understand the implication of Proposition 6 notice that the choice of the information content of the ad depends essentially on the cost of acquiring exact information. Keeping all other parameters fixed, as $c \to \infty$ we have that $\pi_0 \to 0$. The intuition is straightforward. When c is extremely high, from users' point of view bearing the cost is not a valuable option. Nevertheless, what really drives the choice of the informative content of the ad is the knowledge of user's preference intensity x. In particular, the positive relation between information and price it is really a positive relationship between user's preferences and information. Intuitively, the higher x, the higher the willingness to pay and the information advertisers are willing to give up as the next proposition states.

[10]See section 2.1.

Proposition 7. *For each $x \in [0, \bar{x}]$ the optimal advertising rule $\pi_0(x)$ satisfies*

$$\pi_0(x) = \max\left\{\frac{\mu_0(1-\mu_0)x - c}{\mu_0(1-\mu_0)x - c\mu_0}, 0\right\}$$

which implies that for any $c > 0$, it holds that:

i) When privacy is violated, no user pays the cognitive cost in equilibrium;

ii) All users with $x \in \left[0, \frac{c}{(1-\mu_0)}\right]$ are informed by a Cheap–Talk advertising;

iii) All users $x \in \left(\frac{c}{\mu_0(1-\mu_0)}, \bar{x}\right]$ are informed by a Persuasive advertising;

iv) Fully–Informative advertisements are never sent in equilibrium.

Proof. Point i) follows from the fact that in equilibrium ICc must hold with equality (Proposition 6) which implies that when privacy is violated all users click on the banner and buy the product whenever the message sent by the banner is $s = 1$. From condition (7) it follows that when π_0 is interior, since $p(x) = x - \frac{c}{1-\mu_0}$,

$$\pi_0(x) = \frac{\mu_0(1-\mu_0)x - c}{\mu_0(1-\mu_0)x - c\mu_0}.$$

Simple algebra reveals that CT-advertisements (that is $\pi_0(x)$=0) are sent to all users

$$x \in \left[0, \frac{c}{1-\mu_0}\right].$$

On the contrary, all users $x > \frac{c}{1-\mu_0}$ receive a persuasive advertisement $\pi_0(x) > 0$. This proves ii) and iii). Finally, notice that $\pi_0(x)$ is strictly increasing in x and that

$$\lim_{x\to\infty} \frac{\mu_0(1-\mu_0)x - c}{\mu_0(1-\mu_0)x - c\mu_0} = 1$$

which means that FI advertisements require an infinite willingness to pay and therefore are never sent in equilibrium and proves iv). ■

Proposition 7 stresses the following idea. Users value advertisement precisely because it is a substitute of their effort to understand their state of necessity. Nevertheless, advertisement is designed in such a way to maximize advertisers' profits. In particular, advertisers are willing to increase the information content of the ad as users' willingness to pay increases. The intuition is the following. From advertiser's perspective, informative advertisements are costly in the sense that, when $\pi_0 > 0$, with positive probability the banner will send message $s = 0$ and the user will not buy the product. Clearly, the advertisers are willing to bear this risk of not selling the good only if the targeted users is willing to pay an high price when $s = 1$[11]. Indeed, Proposition 7 defines the cut–off user $x' = \frac{c}{1-\mu_0}$ so that all users $x \leq x'$ are exposed to Cheap–Talk advertisements. Notice that, in turn, the cut-off user depends on the cost: the higher the cost the higher the cut–off user (the larger x). This underlines advertisers' tendency to save on information keeping their balance onto equilibria in which users bear the risk. Nevertheless, for large values of x the advertisers are willing to share some information. In this case, all users $x > x'$ are targeted with Persuasive advertisements. Although the information advertisers are willing to share with users increases in x, Fully–Informative advertisements require infinite willingness to pay and thus are never sent in equilibrium. In general, this means that for any positive cost $c > 0$ Fully–Informative advertisements cease to be consistent with equilibrium behavior and that privacy violation can not but increase social waste.

5 Conclusion

This paper proposes a new framework to study the connection between online advertising and privacy when a Web–Master has the ability to elicit users' private information and the latter are uncertain about their state of necessity. First, we show that when only the distribution of users' preferences is known, advertising is only partially efficient in maximizing profits as advertisers are not sure about the characteristics of the users they are facing. On the contrary, when this private information is elicited, each user is optimally persuaded by the advertising she is shown.

Within our framework, a multiplicity of advertising equilibria emerges. The associated advertisements differ in their informational content, which can range from being purely

[11]Recall that whenever users click on the banner, regardless from the realization of s, the advertisers have to pay the fee to the Web–Master.

informative to be pure cheap talk. Nevertheless, when user are allowed to bear a cost in order to find out their state of necessity, advertisers are no longer willing to supply informative advertisements. Specifically, cheap–talk (partially–informative) advertisement is optimal when users are willing to pay a low (high) price for the product. This uncovers a positive relationship between advertisement informativeness and price, so that for given cognitive cost an advertiser finds it optimal to convey some information for individuals willing to pay a higher price.

In general, this implies that a switch from a privacy protected to a privacy violated regime increases social waste, i.e., the number of users who buy a product they do not need. Indeed, only fully–informative advertisements would eliminate social waste, although the risk in this case would be entirely burdened by the advertising side. In addition, when privacy is violated, the Web–Master can always extract all surplus as it is able to engage in hyper–targeting strategies.

To conclude, this paper underlines a new policy concern related to privacy. Differently from the conventional wisdom for which advertising is known to be informative, we show that this is not the case when users have the possibility to chose between costly acquiring information and being informed by commercials. In such a case, advertising allows users to save this cost in exchange of a lower price at the expense of a more risky choice. This behavior may generate social waste locking users into equilibria in which they are persuaded to buy product they do not need.

References

Acquisti, A., Taylor, C., and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2):442–92.

Bagwell, K. (2007). The economic analysis of advertising. *Handbook of industrial organization*, 3:1701–1844.

Belleflamme, P. and Vergote, W. (2016). Monopoly price discrimination and privacy: The hidden cost of hiding. *Economics Letters*, 149:141–144.

Bergemann, D. and Bonatti, A. (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3):417–443.

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

Chen, Y. (1997). Paying customers to switch. *Journal of Economics & Management Strategy*, 6(4):877–897.

Corts, K. S. (1998). Third-degree price discrimination in oligopoly: all-out competition and strategic commitment. *The RAND Journal of Economics*, pages 306–323.

De Corniere, A. and De Nijs, R. (2016). Online advertising and privacy. *The RAND Journal of Economics*, 47(1):48–72.

Friedman, J. (1982). Oligopoly theory. *Handbook of mathematical economics*, 2:491–534.

Fudenberg, D. and Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of Economics*, pages 634–657.

Hagiu, A. and Jullien, B. (2011). Why do intermediaries divert search? *The RAND Journal of Economics*, 42(2):337–362.

Hoffmann, F., Inderst, R., and Ottaviani, M. (2014). Persuasion through selective disclosure: Implications for marketing, campaigning, and privacy regulation. *JW Goethe University Frankfurt, UCL, and Bocconi.*

Kaldor, N. (1950). The economic aspects of advertising. *The review of economic studies*, 18(1):1–27.

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

Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of political economy*, 94(4):796–821.

Nelson, P. (1974). Advertising as information. *Journal of political economy*, 82(4):729–754.

Posner, R. A. (1981). The economics of privacy. *The American Economic Review*, 71(2):405–409.

Rayo, L. and Segal, I. (2010). Optimal information disclosure. *Journal of political Economy*, 118(5):949–987.

Robinson, J. (1969). *The economics of imperfect competition*. Springer.

Schmalensee, R. (1976). A model of promotional competition in oligopoly. *The Review of Economic Studies*, 43(3):493–507.

Shaffer, G. and Zhang, Z. J. (2000). Pay to switch or pay to stay: Preference-based price discrimination in markets with switching costs. *Journal of Economics & Management Strategy*, 9(3):397–424.

Stigler, G. J. (1961). The economics of information. *Journal of political economy*, 69(3):213–225.

Sutton, J. (1991). *Sunk costs and market structure: Price competition, advertising, and the evolution of concentration*. MIT press.

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *The RAND Journal of Economics*, 35(4):631–650.

Thisse, J.-F. and Vives, X. (1988). On the strategic choice of spatial price policy. *The American Economic Review*, pages 122–137.

Varian, H. R. (2002). Economic aspects of personal privacy. In *Cyber Policy and Economics in an Internet Age*, pages 127–137. Springer.

Villas-Boas, J. M. (1999). Dynamic competition with customer recognition. *The Rand Journal of Economics*, pages 604–631.