



The price of discovering your needs online[☆]

Elias Carroni^{a,*}, Luca Ferrari^b, Simone Righi^{c,d}

^a Dipartimento di Scienze Economiche - Alma Mater Studiorum - Università di Bologna - 1, piazza Scaravilli, 40126 Bologna, Italy

^b Department of Economics and Communication, University of Modena and Reggio Emilia, Italy

^c Department of Computer Science, University College London, United Kingdom

^d MTA TK "Lendület" Research Center for Educational and Network Studies (RECENS), Hungarian Academy of Sciences, Hungary



ARTICLE INFO

Article history:

Received 7 January 2019

Revised 27 May 2019

Accepted 10 June 2019

Available online 19 June 2019

JEL classification:

D80

D82

D83

L10

M37

Keywords:

Targeting

Platforms

Digital economy

Information design

ABSTRACT

Web users, even when not actively looking for a product, are continuously targeted by offers that potentially fit their interests. Does this kind of targeting always promote transactions with high social value? We consider a model in which a two-sided platform is visited by sellers seeking to promote their products and by users with state-contingent preferences about those products. Depending on its targeting ability, the platform could be able to target users by tailored advertising banners. We show that persuasion is stronger when targeting ability is maximal, fueling purchases that rational individuals may regret due to the persuasive nature of banners.

© 2019 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license.

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Introduction

Ann is surfing the web, she just started to play the guitar and suddenly, from her online social network feed, new advertisements appear:¹ they are related to music, music equipment and a software offering guided guitar lessons. Ann was not actively looking for such a product, but the advertising platform has access to private information collected during her browsing activity and thus exposes Ann to a banner promoting a product that might fit her needs. Nevertheless, she is uncertain about her need for the software and thus she buys the software only if she believes that the expected utility from the purchase will be positive.

[☆] We would really like to thank Friederike Mengel and two anonymous referees for insightful comments on previous versions of this paper. We are also particularly grateful to Ennio Bilancini, Leonardo Boncinelli, Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, J. C. Ely, Bruno Jullien and the participants of the seminars at the University of Firenze, Telecom Paris Tech, DSE (Bologna), the IIBEO workshop (Alghero), and the EARIE 2017 conference (Maastricht). 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).

* Corresponding author.

E-mail addresses: elias.carroni@unibo.it (E. Carroni), luca.ferrari1987@unimore.it (L. Ferrari), s.righi@ucl.ac.uk (S. Righi).

¹ Realistically speaking, Ann's browser history is partly collected in the form of cookies, text files stored in Ann's computer that can be accessed by other web servers and used to show Ann content more likely to be close to her actual needs.

What could a seller possibly do in order to convince Ann to buy the software? More precisely, once a seller is matched with a user who is potentially interested in the sponsored product, what are the seller's options? By designing a banner, the seller conveys some information about Ann's payoff-relevant state, that is, her needs. In other words, the seller designs a signal over purchase recommendations.² The more precise the information conveyed by the banner, the more Ann will be willing to pay when she receives a positive signal recommending her to buy the product. For instance, the seller might offer Ann a free trial of the product, which may or may not convince her to buy it. The possible recommendations to "buy" and "not buy" are sent with a probability conditional on Ann's true needs.³ Given this opportunity, Ann is willing to proceed with the free trial. When Ann likes the outcome of the experiment, i.e., she observes the recommendation "to buy", her expected payoff is larger than zero due to the change in her beliefs that results from the outcome of the experiment. Oppositely, whenever Ann does not like the trial, she prefers not to buy the product.

The example described above is a typical interaction that might occur thanks to an online platform with targeting abilities. Indeed, in the reality of digital markets, advertisers outsource the job of reaching and matching users interested in their products to Internet platforms in exchange for a fee. For instance, Facebook advertising platform explains to its users that:

"Our ad products let businesses and organizations connect with the people who are most likely to be interested in their products and services. We believe the ads you see across the Facebook family of apps and services should be useful and relevant to you."⁴

This paper inquires the following issue: does this matching always promote transactions with high social value? We show that the value of this matching activity is strongly related to the accuracy of information conveyed by online banners.

The information online banners convey may take different forms. In general, a banner links to a webpage or to media content (e.g. a video) advertising a product.

The choice of which content to include in the banner (wording, media, etc.), might result persuasive for some users rather than some others, in line with the recent evidence about psychographic persuasion (See for instance [Hirsh et al. 2012](#); [Kosinski et al. 2013](#); [Matz et al. 2017](#)). The users interpret the banner as an information structure whereas the realization is privately drawn by the user and cannot be manipulated by the seller.

Building on these observations, we consider a situation in which a two-sided platform is visited by sellers seeking to promote their products and by users with state-contingent preferences about those products. We consider both a situation in which the platform is able to perfectly distinguish the user's willingness to pay (*perfect targeting*) and one where it can only inform sellers about its distribution (*imperfect targeting*). Each seller pays a fee to the platform in exchange of its matching services, he decides his banner's content and the price of his product. We show, along the lines of [Rayo and Segal \(2010\)](#), that the more uncertainty is reduced by a banner, the larger the price each seller can set. In addition to this intuitive relationship, we underline how the cost of producing the good affects information provision. Since designing information is costly, expensive goods are advertised by more informative banners. However, the platform monetizes advertisement through clicking volumes, so that it might choose to show banners different from the most informative ones. This implies a social cost of persuasion. Users learn something about their needs thanks to the availability of their digital footprints, but at the cost of accepting transactions with negative value, i.e., sometimes buying non-needed products. Whenever a platform maximally exploits user's information (e.g., in online social networks) the persuasive effect is stronger with respect to other situations. For instance, when banners are auctioned in relationship to keywords (e.g., in online search engines), targeting is imperfect but the persuasive effect is mitigated. This exemplifies the distinction between *perfect* and *imperfect targeting*. The usual redistributive problems linked to price discrimination are enhanced by the fact the information is costly to design. Whereas targeting increases matching efficiency, the manipulative nature of banners erodes the welfare gains of passing from *imperfect* to *perfect targeting*. Therefore, if its main interest is consumer protection, the regulator should carefully evaluate targeting when providing valuable information is costly.

The remainder of this paper is divided as follows. In [Section 2](#) we present the review of the related literature. [Section 3](#) introduces the model and presents the optimal behavior of a single seller (3.1) as well as the optimal choices of the platform (3.2). [Section 4](#) discusses how our model can explain some stylized facts about online advertising and [Section 5](#) gives some welfare implications. Finally, [Section 6](#) concludes.

2. Related literature

The emergence of online markets and of the so-called big data have an important impact on companies' strategies. Indeed, the access to user-specific information by online middlemen makes them very attractive for firms, which have the opportunity to enact tailored pricing and advertising strategies and use these platforms to reach an enormous number of potential customers.

² See [Bergemann and Morris \(2019\)](#).

³ Formally: $\Pr(\text{not buy}|\text{no need})$ and $\Pr(\text{buy}|\text{need})$ respectively.

⁴ See https://www.facebook.com/ads/about/?entry_product=ad_preferences.

The huge mass of information accessible online is frequently used for targeting. In particular, different degrees of knowledge may lead to segment the market through price discrimination (Thisse and Vives 1988), loyalty schemes (Shaffer and Zhang 2000), switching offers (Chen 1997; Fudenberg and Tirole 2000; Villas-Boas 1999) and online pricing (Taylor 2004). Tailoring is thus prominent in online markets, where targeted advertising of products may be efficient as it creates new opportunities for trade, by getting on board consumers that would have been excluded otherwise (Bergemann and Bonatti 2011).

Digital middlemen have a key role, because of their ability to provide a two-sided matching service between users and advertisers, and their impact on welfare is multifaceted. At this regard, De Corniere and De Nijs (2016) have demonstrated how online intermediaries may have a negative impact on welfare, as their disclosure of consumer-specific information to advertising firms may lead to excessive prices. Differently from their paper, we show that while targeting is efficient but it also generates social waste and overspending in persuasion. Moreover, many scholars have recently focused on the sale by search engines of sponsored links to advertisers to reach online searchers. Hagi and Jullien (2011) show how platforms may have incentives to divert search, inducing more search than needed. Athey and Ellison (2011) analyze and discuss the effects of different *position auctions* on welfare and consumers surplus and Gomes (2014) highlights a trade-off between rent extraction and clicking volumes, as advertisers with the highest willingness to pay do not necessarily offer the most relevant advertisement and long-run clicks depend on the relevance of sponsored links. In a recent paper, Board and Lu (2018) study product's information disclosure in a context in which consumers who are actively searching for products would stop searching only when satisfied with the information provided by the signals released by a seller. They find that, even in a competitive market, seller may behave as monopolists, giving no information to the buyers. In particular, if consumers' beliefs are perfectly known by sellers, the monopoly disclosure strategy is the unique equilibrium, as buyers expect to get no more information with an additional search.

Differently from these contributions, in our model the platform is a website in which *users surf* to enjoy contents (e.g. newspapers, blogs) or to interact with friends (social media) and are exposed to banners even if they are not actively looking for products. Hence, for their attention to be captured by banners, the latter must display a relevant product. When users can be perfectly targeted with tailored banners, clicking volumes are always maximized and then the platform only focuses on rent extraction. Differently, when users are not perfectly targeted, the trade-off is between clicking volumes and the per-click fee paid by the seller to which the banner space is sold. We borrow from the literature of targeting the matching role of online intermediaries and their discriminating power, but we look at online targeting from a different perspective.

In a context in which users have state-dependent needs, advertisement is interpreted as a statistical experiment that is useful for users to discover their needs or the match of their preferences with a product. Therefore, our interpretation of advertisement does not immediately match with the classification traditionally adopted by the literature, i.e., persuasive (Kaldor 1950; Robinson 1969) and informative (Nelson 1974; Stigler 1961) advertisement.⁵ The experiment has the only role of providing individuals with information (more or less accurate) about their state of necessity. In this sense, it does not entail any ex-ante change in consumer preferences but only an ex-post demand shift due to the change in perception about the state of necessity induced by the observation of the experiment's outcome. Our paper is closer to Anderson and Renault (2006) who discuss the informative content of advertisement, analyzing type (price and preference-attribute match) and amount of information that sellers decide to reveal to consumers through ads. In line with us, they find that companies prefer to reveal a limited amount of information. Further, their study confirms the centrality of the degree of information provision for the communication strategies of firms. We innovate by connecting the limited informativeness of advertisements to their information structure.

The relationship between individual users and sellers is modelled in the spirit of the recent stream of literature on information design, initiated by Rayo and Segal (2010) and Kamenica and Gentzkow (2011). In these models, a sender with commitment power 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 her 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) consider a case in which a platform displays different experiments to maximize its profits exploiting their position across a webpage. 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-targeting⁶ may harm consumers depends on the ability of firms to price-discriminate, on the competition between senders and on consumers' wariness. In our setup the experiment is interpreted as a banner and the user-seller relationship is embedded in an online platform.

3. The model

An internet platform runs a website and offers banner spaces to sellers in exchange for a per-click fee. Users that visit the website are exposed to banners, which can be tailored whenever possible. Whenever a user clicks on a banner, she

⁵ See Bagwell (2007) for a review on different classes of advertisement.

⁶ Hoffmann et al. (2014) define hyper-targeting as "the collection and use of personally identifiable data by firms to tailor selective disclosure".




Fig. 1. Timeline.

receives an experiment whose outcome can be interpreted as a purchase recommendation. The information structure of the experiment is designed by the seller. Thus, a seller is characterized by a banner containing the price of the good and its information structure. When the outcome of the experiment is positive, the user buys the product. The timing of the model is summarized in Fig. 1.

In stage 1, the platform sets per-click fees that sellers have to pay in order to advertise their products through banners. This decision depends on the targeting abilities of the platform. We say that the platform has perfect targeting abilities whenever it is informed about each user's willingness to pay. In this case, individual banners can be displayed to users. This gives scope for two-sided price discrimination: a seller may be offered a different fee depending on the user reached with the tailored banner. When the platform has imperfect targeting abilities, it is only informed about the distribution of the willingness to pay of its users. In this case, individual banners cannot be displayed to users, with no possibility of discriminating fees. Given the fee offered by the platform, in stage 2 sellers decide the content of their banners and enter the platform by paying the fee. Users decide whether to click on a banner based on their prior beliefs and - conditional on the information received through the banner - whether to purchase the advertised product in stage 3. Finally, payoffs of all agents are realized in stage 4.

We will solve the model by backward induction. Thus, we first study in Section 3.1 a generic seller-user interaction, providing the optimal banner of each seller given the targeting ability of the platform and its decisions in the previous stage. Then, Section 3.2 is devoted to the analysis of the fee setting of stage 1.

3.1. Offers through banners

In this section, we analyze stage 3 and, going backward, stage 2. We start by detailing the objective functions of the agents in both sides of the market and to define how the user-seller interaction takes place.

Formalizing the example proposed in the introduction of the paper, consider one seller targeting a generic user with willingness to pay $x \in [\underline{x}, \bar{x}]$, distributed according to a cumulative $G(x)$ with p.d.f. $g(x)$. We further assume that the hazard rate function is monotonically non-decreasing, i.e., $\frac{\partial \frac{g(x)}{1-G(x)}}{\partial x} \geq 0$. In terms of our example, the seller is targeting Ann and all the users who share the same x . The seller's product is advertised through a banner and sold at price $p \geq 0$. While the user is informed by the banner about all the product's characteristics, she is uncertain about whether the product will indeed be needed or not, i.e. she is characterized by an unobservable state $\omega \in \{0, 1\}$, describing her necessities. As in Johnson and Myatt (2006), the user is characterized by state-dependent preferences $u(\omega, x, p)$. In our setup $\omega = 1$ denotes the state of the world in which the good is needed, while $\omega = 0$ denotes the state of the world in which the good is not needed.⁷ The user has a prior belief about being in the state 0, denoted as $\mu \equiv \Pr(\omega = 0)$. Thus, given her prior, the user's expected utility from buying the product is:

$$U_B(x, p) = -p + (1 - \mu)x + \mu \cdot 0.$$

In the "bad" state $\omega = 0$, the user will find the product not useful to her purposes whereas in the good state, $\omega = 1$, she gets a payoff equal to her willingness to pay minus the price. The outside option of not buying is assumed without loss of generality to give zero payoff, $U_{NB}(x, p) = 0$. Note that the willingness to pay x and the probability of needing the good $1 - \Pr(\omega = 0)$ are considered as uncorrelated since, a priori, there is no reason to assume that a good of higher quality (i.e. higher willingness to pay) is more (or less) likely to be needed.⁸

We say that a seller targets user x whenever she produces a banner β_x , where x refers to the willingness to pay of the user. We assume that the banner contains: (i) the price of the offered good and (ii) an experiment whose outcomes are statistically correlated with user's state of necessity. This is in line with the advertising policy of leading online social networks and search engines, as the banners cannot contain contents that are manipulative or deceptive per se.⁹ A banner is thus

⁷ In Johnson and Myatt (2006), ω denotes the value to the user of the match with the characteristics of seller's product. The same interpretation fits our setup.

⁸ Introducing a positive (negative) correlation would simply increase (decrease) the amount of information the seller is willing to send to consumers, without changing qualitatively the results.

⁹ In our setup, this is guaranteed by the fact that advertisers cannot manipulate the outcome of their experiments. However, this reputation-driven advertising policies do not ensure fully informative banners: Best and Quigley (2016) show that "the desire to persuade in the future can generate credibility, and hence persuasion, today".

defined by $\beta = (\pi, p)$, with π representing an experiment characterized by a pair of conditional probability distributions

$$\begin{aligned} \pi_0 &\equiv \Pr(s = 0 | \omega = 0), \\ \pi_1 &\equiv \Pr(s = 1 | \omega = 1) \end{aligned}$$

where $s \in \{0, 1\}$ denotes a shopping advice (for simplicity the *message*).

Through the choice of the banner, the seller designs what the user could possibly learn from privately interpreting the experiment. In our leading example, Ann downloads the free trial of the software. As long as the free trial might inform her about her needs, she could also discover something that prevents her from buying the product. Thus, the experiment can have two possible outcomes: when $s = 1$, it succeeds in persuading the user that, in expected terms, she is weakly better off buying the product, whereas when $s = 0$, the experiment fails to persuade her to buy it. In addition to [Kamenica and Gentzkow \(2011\)](#), we allow the seller to choose the price of the sponsored product which in our setup represents the cost of the risky action *buy*.

The banner β_x realized by the seller will be the one that maximizes profits, i.e. $\Pi = \Pr(\text{sale} | \beta_x)(p - c)$, where c is the marginal cost of producing the product.

Given the platform decisions in the first stage of the game, which in turns depend on its' targeting ability, the subgame composed by stages 2 and 3 develops as follows:

- i) Nature draws the user's x state of necessity ω which is unknown to both players;
- ii) The seller targeting user x designs his banner β_x ;
- iii) The user observes β_x and decides whether to click or not on the banner;
- iv) If the user does not click on the banner, she gets her reservation payoff, 0. If she clicks, she receives information and reacts accordingly.

To solve this user-seller interaction problem we rely on the Bayes Correlated Equilibrium (BCE) ([Bergemann and Morris 2016](#)) with the Sender-Preferred refinement. Thus, the equilibria are those BCE in which the *Sender* (in our case the seller) payoff is maximized, as in [Kamenica and Gentzkow \(2011\)](#).

3.1.1. Stage 3: clicking and purchase decisions

In stage 3 users decide whether to click on a banner and, if they do, whether to purchase the advertised product. In particular, we look to situations in which the banner induces the user to click on it, i.e., to situations that are incentive compatible and in which the user is obedient to the action recommended in the banner.¹⁰ Notice that an incentive compatible banner always leads to a click. Indeed a consumer is willing to click as long as the experiment π gives her a reliable purchase recommendation, so to make the user more informed. When this is the case, the utility of clicking is always higher than the outside option of not clicking (and not buying), which is assumed without loss of generality to be zero.

Conditional on the observation of message s , the user forms the posterior belief $\Pr(\omega = 0 | s) \equiv \mu_s$ using Bayes rule:

$$\begin{aligned} \mu_0 &= \frac{\mu\pi_0}{\mu\pi_0 + (1 - \mu)(1 - \pi_1)}, \\ \mu_1 &= \frac{\mu(1 - \pi_0)}{\mu(1 - \pi_0) + (1 - \mu)\pi_1}. \end{aligned}$$

In order for s to express a shopping advice of the kind “buy” or “not buy”, the experiment $\pi \equiv (\pi_0, \pi_1)$ must be incentive compatible so that the observation of s represents a recommended action. In particular, if the message suggests to buy the product, then it must be the case that, conditional on message $s = 1$, user x weakly prefers to buy the product, i.e. $U_B(x, p, \pi | s = 1) \geq U_{NB}(x, p, \pi | s = 1) = 0$, that is:

$$\frac{(1 - \mu)\pi_1}{(1 - \mu)\pi_1 + \mu(1 - \pi_0)}x \geq p. \tag{IC1}$$

Furthermore, for the experiment to be incentive compatible, conditional on $s = 0$, user x needs to strictly prefer avoiding the purchase, that is $U_{NB}(x, p, \pi | s = 0) = 0 > U_B(x, p, \pi | s = 0)$ thus

$$p > \frac{(1 - \mu)(1 - \pi_1)}{(1 - \mu)(1 - \pi_1) + \mu\pi_0}x. \tag{IC0}$$

When the IC conditions hold, the user behavior is determined by the observation of the result of the experiment, i.e., she is “obedient” to the latter. Notice that a user would click on a banner only if it is incentive compatible. Otherwise she would receive recommendations she does not agree with, and therefore she would not follow them.

¹⁰ In our leading example, Ann, captivated by the structure of the banner, clicks on it whenever she believes she can find valuable information.

3.1.2. Stage 2: sellers' decision

In stage 2, each seller decides the content of a banner, anticipating the clicking and the purchase behavior of consumers. Noticeably, the objective of the seller is to maximize the expected profit. The maximization problem is constrained by the fact that the seller has to produce a banner that induces the user to click on it, otherwise trade cannot occur and the seller would be left with zero profit. It is important to notice that, at this stage of the game, each seller observes the trading ability of the platform and the fee relative to each banner space. In particular, under imperfect targeting, each seller decides who to target knowing that the other side of the market is reached imperfectly. Differently, under perfect targeting, the platform is perfectly able to make each seller reach the perfect match on the other side of the market. For the sake of exposition, in what follows we first consider the optimal decisions of a seller deciding to target a generic user x on the other side of the market and then the optimal target under perfect and imperfect targeting.

In order to study the relationship between user's and seller's decisions we proceed by steps. First, we study the relationship between the information structure of the experiment and the price of the offered good. We start by abstracting from the information-design costs the seller incurs. We then relax the analysis by studying how these frictions affect seller's optimal strategy.

In equilibrium, the banner $\beta_x = (p, \pi_0, \pi_1)$ is such that the incentive-compatibility constraints hold. In this case the seller's profit is given by

$$\Pi = \Pr(s = 1)(p - c) - f_x,$$

where f_x is the fee that has to be paid to the platform for the banner to reach user x . From the seller's viewpoint, the fee is just a fixed cost, so it is relevant in the decision about whether to produce a banner, but irrelevant in the choice of the optimal banner.¹¹ The probability with which the outcome of the experiment results in a recommendation to buy the product is

$$\Pr(s = 1) = \mu(1 - \pi_0) + (1 - \mu)\pi_1,$$

which depends on the strategic choice π made by the seller. Therefore, equilibrium is an optimal banner

$$\beta_x^* = \arg \max_{\beta_x} \Pr(s = 1)(p - c) \text{ s.t. } IC0 \text{ and } IC1.$$

Proposition 1 The Optimal Banner. Consider the seller targeting user x . The optimal banner $\beta_x^* = (p^*, \pi_0^*, \pi_1^*)$ is characterized by $\pi_1^* = \pi_0^* = 1$ and $p^* = x$.

Proof. See Appendix A. \square

Proposition 1 is important because it highlights a crucial relationship between the accuracy of the experiment proposed by the seller and the correspondent price. Off equilibrium, the accuracy of the experiment can be defined as

$$Y := \mu\pi_0 + (1 - \mu)\pi_1, \quad Y \in [0, 1] \tag{Accuracy}$$

that is the probability with which the purchase recommendation matches user's needs.¹² Given that in equilibrium $\pi_1^* = 1$, from the relationship in Eq. (6) we obtain

$$Y^* = 1 - (1 - \mu) \frac{x}{p^*} - 1 - \mu.$$

Thus, when targeting user x , more accurate banners result in higher prices. It is worth to stress that, keeping the price constant, the accuracy of the experiment decreases in x . Indeed, if the user willingness to pay increases but the price remains constant, the seller can increase the probability of sales just by designing a less accurate, though incentive compatible, experiment. Nevertheless, **Proposition 1** underlines that if the seller incurs in a marginal cost of production, full information disclosure has to occur, as already presented by [Rayo and Segal \(2010\)](#). Indeed, since the seller can condition the price to the actual realization of the experiment, he can always extract the entire buyer surplus.¹³ As a consequence, the seller provides sufficient information to exhaust all gains from trade.¹⁴

Nevertheless, experiments are rarely fully informative as gathering information is usually costly. Designing an experiment requires the seller to sustain relevant costs to study the drivers that can inform a user about her needs. Those costs affect the seller's optimal behavior in a straightforward way as they create an additional trade-off the seller has to consider when designing his optimal banner.

Coherently with the information-design literature, the cost of designing an experiment is convex in the space of signals ([Kamenica and Gentzkow, 2014](#)). We assume that all sellers face a cost of providing an experiment $kC(\pi_0, \pi_1)$ with $C(\cdot)$

¹¹ In [Section 4](#) we discuss the effect of a per-sale fee, which, on the contrary, affects the strategic choice of how much information to convey.

¹² Accuracy measures how effectively s matches ω . It increases as the probability of false positives and false negatives decreases. This definition of accuracy is shared by the computer science literature on machine learning ([Fawcett 2006](#)) and by the economic literature on information acquisition problems (see for instance [Persico \(2000\)](#)).

¹³ Because of incentive compatibility the buyer buys only when she observes the realization $s = 1$.

¹⁴ For a related result in the context of a search model see [Anderson and Renault \(2006\)](#).

symmetric and convex in both arguments and $C(\pi_0, \pi_1) = 0$ whenever the experiment is not informative.¹⁵ The parameter k measures the technology used by the seller to provide a more accurate experiment.

Let us now consider how the introduction of the information-design cost affects the maximization problem of the seller. Exploiting the (IC1), the price must be equal to the posterior times the willingness to pay x . Thus, the seller’s problem reduces to:

$$\max_{(\pi_0, \pi_1) \in [0,1]^2} (1 - \mu)\pi_1 x - [(1 - \mu)\pi_1 + \mu(1 - \pi_0)]c - kC(\pi_0, \pi_1).$$

Given the assumptions on the cost of designing the experiment, FOCs are sufficient for a maximum. Therefore, the optimal signal is implicitly defined by the following conditions:

$$\underbrace{(1 - \mu)(x - c)}_{\text{MB of sale } \omega=1} = k \underbrace{\frac{\partial C(\pi_0, \pi_1)}{\partial \pi_1}}_{\text{MC of } \pi_0} \tag{1}$$

$$\underbrace{\mu c}_{\text{MB of no sale in } \omega=0} = k \underbrace{\frac{\partial C(\pi_0, \pi_1)}{\partial \pi_0}}_{\text{MC of } \pi_1}. \tag{2}$$

This directly leads to the following proposition:

Proposition 2 Costly experiment. Consider the seller targeting user x . When the seller faces marginal cost $c > 0$ and the cost of designing the experiment is $C(\pi) = kC(\pi_0, \pi_1)$ as described above, then the optimal banner satisfies $p_x^* = \mu_1(1)x$ as well as the FOCs (1) and (2). Furthermore, $\pi_1^* > \pi_0^*$.

To understand that $\pi_1^* > \pi_0^*$ observe from (1) and (2) that the marginal benefit of selling the product in the good state will always be larger than the marginal benefit of not selling the product in the bad state, whenever there are gains from trade. To our purposes, it is without loss of generality to assume that the value of the product is sufficiently large to lead to a corner solution with $\pi_1^* = 1 > \pi_0^*$.¹⁶ Therefore, hereafter we let $C(\pi_0, 1) = C(\pi_0)$, with $\pi_1^* = 1$ and $\pi_0^* = C^{-1}(\frac{\mu c}{k})$. Notice that when $\pi_0 = 0$ the signal is not informative and has zero cost. On the contrary, when $\pi_0 = 1$, all the uncertainty is resolved and the signal has the maximal cost. Therefore, the equilibrium accuracy is simply pinned down by the equilibrium value of π_0 .

When the choice of the experiment is not costly, the seller would opt for a fully informative banner whenever he faces a positive marginal cost. However, given that designing an experiment is costly, the seller reduces its informativeness. In addition, since the user understands that she is going to buy the product with some probability in the bad state of the world, the seller has to insure her by setting a lower price.

One can think about the seller being possibly efficient along two different dimensions: production and information design. Intuitively, a seller with a smaller marginal cost c is more efficient in producing or shipping the good, whereas a seller with a lower k is able to provide the same level of accuracy at a smaller cost. Condition (2) identifies this relationship and the following lemma formalizes the directions of the trade-off in terms of accuracy.

Lemma 1 Costs and Accuracy.

- i) A decrease (increase) of the cost of designing the experiment relative to a decrease (increase) in the marginal cost results in a higher (lower) equilibrium accuracy;
- ii) As $c/k \rightarrow 0$, the equilibrium experiment provides no additional accuracy, i.e., $\pi_0^* \rightarrow 0$.

Proof. See Appendix A. □

Since the accuracy of the experiment and the price of the offered good are bounded by the equilibrium relationship found in Proposition 1, the trade-off between the efficiency dimensions are reflected by the price of the offered good as shown in Fig. 2. A change in $\frac{\mu c}{k}$ induces a change in the price depending on the cost of designing the experiment the seller is facing. Whereas the magnitude of the price change depends on the shape of cost function $C(\cdot)$, lower marginal costs always translate into less accurate experiments, i.e., a lower equilibrium value of π_0^* .

Consider the example of beauty products. In such market, sellers of expensive products typically offer free samples, which are more informative than purely persuasive marketing strategies followed by sellers of cheap products. The results of Lemma 1 are also applicable to digital goods, such as software and online services, where it is true that marginal costs of production are very small, but still higher than the negligible cost of delivering the free trial of a software.

To sum up, we derived the optimal banner a seller would design in order to reach a potential consumer. The banner contains the price of the good being sold and an experiment. If the experiment-price pair is incentive compatible, then the

¹⁵ A π is not informative if it is of the form $(p, 1 - p)$ with $p \in [0, 1]$.

¹⁶ Notice that a common formulation in information design problems is to assume that $C(\pi_0, \pi_1)$ is the mutual-information function (Shannon, 1948) which, as pointed out by Caplin and Dean (2013), has unbounded derivatives at corners, leading to interior solutions. Nevertheless, taking into account interior solutions would just obfuscate the main focus of our work without altering the mechanics of the model.




Fig. 2. The equilibrium price $p^*(\pi_0^*)$ is determined by the optimal choice of how much information to convey in state 0, i.e., π_0 . The picture shows that the optimal level of π_0^* solves the trade-off between the cost of producing the good, c , and the marginal cost of increasing π_0 , $C'(\cdot)$. Then, the optimal price is determined by the equilibrium relationship $p^* = \frac{1-\mu}{1-\mu\pi_0^*}x$. A decrease of c , reduces the marginal cost of selling in the bad state, π_0^* decreases and to have the offer being incentive compatible, the price has to decrease accordingly.

user clicks on the banner, she receives a purchase recommendation and behaves accordingly. We showed that the efficiency dimensions of the seller affect the optimal equilibrium experiment-price pair. In general, the accuracy of the seller’s experiment depends on his efficiency trade-off, whereas the price of the offered good must satisfy the incentive-compatibility constraints. Therefore, the less accurate the experiment, the smaller the price. Having characterized the optimal banner targeting user x , we are now able to finish the analysis of the second stage of the game by distinguishing the two cases of perfect and imperfect targeting.

Perfect targeting. In this case, each seller takes into account that the platform is able to offer a perfect match on the other side of the market. Therefore, provided that the profits that can be made are sufficiently high, i.e., $(p_x^* - c) \Pr(s = 1 | \pi_0^*, \pi_1^*) - kC(\pi_0^*) \geq f_x$, a seller who aims to target user x on the other side of the market is able to do it and willing to enter the platform paying the fee f_x .

Imperfect targeting. Differently from the previous case, here the platform is informed only about the distribution. Hence, rather than offering a perfect match, it can only offer access to the distribution of users. Therefore, each seller knows that offering a banner β_x^* is incentive compatible for user x , so that it would generate a click also for all users with $x' > x$ but it will never generate a click from users with $x'' < x$.¹⁷ Therefore, a seller chooses its banner facing the tradeoff between clicking volumes and per-click profits. Formally, he solves

$$\max_x \underbrace{[1 - G(x)]}_{\text{click volume}} \underbrace{\Pi(\beta_x^*)}_{\text{per-click profits}} .$$

Therefore, when targeting is imperfect, all sellers show the same banner $\beta_{x^*}^*$, where:

$$x^* = \arg \max_x [1 - G(x)] \Pi(\beta_x^*)$$

To see that a solution to the maximization problem exists, notice that the function $[1 - G(x)] \Pi(\beta_x^*)$ is continuous in the closed and bounded interval $[x, \bar{x}]$, thus, by the extreme value theorem, it attains a maximum at x^* . Moreover, monotonicity of the hazard rate function guarantees that this solution is unique.

Note that the optimal experiment is the same no matter what the targeting technology is, and that the difference between perfect and imperfect targeting rests on price. This follows from the fact that, whenever perfect targeting is possible, the seller will engage in price discrimination. However, different platforms offer different possibilities of targeting. For example, search engines auction search keywords to sellers, while online social networks sell much more tailored advertisement services, relying on a wealth of information on consumers. The case of keywords searches configures a less precise targeting than the one based on idiosyncratic preferences of users. Following our results, one could predict that companies relying on platforms allowing more precise targeting would engage more often in price discrimination. Detecting price discrimination in online platforms is notoriously difficult as vendors keep their pricing algorithms secret. However, Mikić et al. (2012) report evidence that price discrimination is realized in platforms (including bestbuy.com and amazon.com) which collect data on the habits of consumers, and that many of the websites they investigate price discriminate based on the source URL of their incoming users. A further indirect confirmation of price discrimination in platforms allowing for better targeting is given in Ellison and Ellison (2018), who observe that prices in online markets tend to be higher than their off-line counterparts (where arguably there is less scope for price discrimination).

¹⁷ See Kolotilin et al. (2017) for an equivalence result between private persuasion mechanisms and experiments.

3.2. Stage 1: platform's decisions

In the previous section, we characterized the optimal banner targeting user x and found the solution of all subgames. Anticipating what would be the reaction of sellers in the second stage, the platform decides the optimal fee in stage 1.

In order to analyze this stage, it is important to recall that, regardless the targeting abilities of the platform, a banner β_x^* would generate a click for all users with $x' \geq x$. Moreover, the seller providing banner β_x^* is the one willing to pay the most to reach user x .¹⁸ The per-click profit made by the seller providing banner β_x^* will be:

$$\Pi(\beta_x^*) = (1 - \mu)x - (1 - \mu\pi_0^*)c - C(\pi_0^*).$$

In order to focus our attention on the most interesting cases, we assume the following:

Assumption 1. The marginal cost of production is sufficiently high: $c > \frac{kc'(1)}{\mu}$.

Assumption 2. \underline{x} is sufficiently high: $\Pi(\beta_{\underline{x}}^*) > 0$.

Under [Assumptions 1](#) and [2](#), it is worth targeting all agents and none of them is perfectly informed by the experiments proposed in the banners. Relaxing the first assumption would lead to corner solutions in which the experiments are perfectly informative and, as we will show, it will entail that the welfare gains of targeting are maximal. Relaxing the second one would only make some people never tailored by banners, as providing them with information would not generate any profit.

In principle, fees can be discriminating, so that the banner shown to user x is offered to a seller who is willing paying an amount f_x . The role of the platform is to target each user with a banner and it is precisely for this reason that sellers are willing to pay a fee for its services. However, the ability of the platform in targeting users depends on the precision of information the platform has about the individual x 's.

Perfect targeting. Under perfect targeting, individual x 's are known to the platform. In this case, it is easy to show how perfect targeting on the user side is always associated with perfect price discrimination on the seller side. Formally, this means that:

Lemma 2. When targeting is perfect, fees are discriminating, with $f_x^* = \Pi(\beta_x^*)$. As a consequence, each user $x \in [\underline{x}, \bar{x}]$ is shown an optimal banner β_x^* .

Proof. From [Proposition 2](#), we know that a seller who targets agent x produces a banner β_x^* and obtains profits equal to $\Pi(\beta_x^*) = (1 - \mu)x - (1 - \mu\pi_0^*)c - C(\pi_0^*)$. In order to fully extract the surplus of this seller, the platform has to set a fee f_x that is just sufficient to induce the seller to produce the relative banner, i.e., the optimal fee $f_x^* = \Pi(\beta_x^*)$. \square

The result in [Lemma 2](#) is a tale of two-sided perfect discrimination. In particular, since banner β_x^* is intended to target precisely user x , showing to the latter this perfectly targeted banner allows to maximize both the clicking volumes (as all user would click on a banner) and the seller's profit associated to each click. As a consequence, fully discriminating fees allow the platform to achieve the best payoff, as they allow for full surplus extraction.

Imperfect targeting. Differently, when the targeting is imperfect, any seller would choose banner $\beta_{x^*}^*$, so that the following holds:

Lemma 3. When targeting is imperfect, the platform sets a non-discriminating fee, so that $\forall x f^* = \Pi(\beta_{x^*}^*)$.

Proof. The fee is the one fully extracting surplus as in [Lemma 2](#). \square

Differently from the case discussed in [Lemma 2](#), here the platform still extracts the entire surplus of sellers, but this surplus is clearly lower as some consumers do not click on the banner $\beta_{x^*}^*$. Because of the impossibility to perfectly discriminate on the sellers' side, the banner trades-off between clicking volumes and per-click profits. This tradeoff does not exist when the platform can perfectly tailor consumers, as in that case volumes can always be maximized.

Another way to interpret the difference between targeting regimes is that there is much more variety of banners under perfect targeting. Indeed, when perfect targeting is not viable, the platform has no better option than placing the same banner $\beta_{x^*}^*$ to all users. On the contrary, if individual realizations of x are known by the platform, each user receives a personalized offer: a banner that contains a price and an experiment. This result follows from the platform's tracking abilities, which leads to a boost in the variety of banners. The consequences of such a difference will be better understood in the [Section 4](#).

4. Discussion

This section discusses the implications of our model, showing that it can give a rationale to some stylized facts regarding online advertising.

¹⁸ Notice that all banners $\beta_{x'}$ with $x' < x$ would not generate any profit when shown to user x , who does not click on them. Differently, all banners $\beta_{x'}$ with $x > x'$ will generate a lower profit than β_x^* , which is the optimal banner that can be shown to user x , as it entirely extracts her surplus.

The marginal cost-accuracy trade off. First of all, it is worth noticing that the platform has an influence on the informativeness of the banners that users are exposed to. Indeed, the platform's objective is to maximize clicking volumes and per-click rent extraction. Moreover, since clicking volumes essentially depend on the ability of the platform to target users, per-click rent extraction becomes the unique objective within each targeting regime. Therefore, the platform has always incentives to hire the seller making the highest profits. As a consequence, there exist situations in which the platform hires sellers offering less accurate experiments, as expressed in the following proposition.

Proposition 3. Consider the platform deciding which seller to hire for a banner β_x targeting user x . The platform may prefer to hire sellers offering less accurate experiments.

Proof. See Appendix A. \square

The intuition behind the result stated in Proposition 3 is that a more efficient seller can make higher profits when producing a banner tailored for user x . This is what ultimately matters to a platform. Given that the fee, discriminating or not, allows for surplus extraction, the optimal solution for the platform is to hire the seller generating the highest possible surplus, i.e., the most efficient one. As a consequence, the platform prefers a seller proposing less accurate experiments. Indeed, combining the results shown in Lemma 1 with the one in Proposition 3, we show that competition for banner spaces would not result in better accuracy, since a decrease of the sellers' marginal cost leads, at the limit, to *non-informative* experiments. On the one hand, more efficient sellers have weaker incentives to provide information about the bad state. On the other hand, efficient sellers are more appealing to the platform, as they achieve higher profits and consequently pay higher fees.

Persuasive advertising and sales. In Section 3.1.2 we described the optimal behavior of sellers advertising their products through banners within the web-space provided by a platform with different targeting abilities. In particular, we derived a positive relationship between the quantity of information conveyed by a banner and the price of the sponsored good. The latter also depends on the marginal cost of production, indeed sellers with lower marginal costs tend to design less informative banners. Overall, this affects sellers' probability of selling the good. Nevertheless there are two ways through which a seller can expand its sales. Consider the imperfect targeting case and suppose that the seller wants to increase the number of sales with his banner. Keeping everything else constant, which could have the largest effect on sales: an increase of the quantity of information revealed by the banner or conversely a decrease of the price?

Let us rewrite the IC1 as

$$x(p, \pi_0) = p^* \frac{1 - \mu\pi_0^*}{1 - \mu}.$$

Then, we have

$$\frac{\partial x(p, \pi_0)}{\partial p} = \frac{1 - \mu\pi_0^*}{1 - \mu}, \quad \frac{\partial x(p, \pi_0)}{\partial \pi_0} = -\frac{\mu}{1 - \mu} p^*.$$

When $-\frac{\partial x(p, \pi_0)}{\partial p} < \frac{\partial x(p, \pi_0)}{\partial \pi_0}$, then a decrease of the price is more effective in expanding the base of users interested in the product. In particular, this is true if and only if

$$1 > \mu(p^* + \pi_0^*). \quad (3)$$

Notice that, if condition (3) is satisfied, then a drop of the price boosts sales more than an increase of information conveyed by the banner. Also, the condition is more likely to be satisfied for products sold at a low price whose banners are not particularly informative. On the contrary, the fact that condition (3) does not hold does not directly imply that raising the information content will increase sales. Indeed, although more user with smaller x_s becomes potentially interested, an increase of π_0 results in a lower probability of selling the good to each of them and therefore there is not a clear cut solution. This trade-off may give a theoretical explanation to the lack of evidence about the positive effect of advertising in increasing sales. For instance, DellaVigna and Gentzkow (2010) survey recent empirical work from which no clear evidence about the effect of advertisement on sales can be advanced. Our model offers a novel interpretation since the probability of trading the good is affected by the quantity of information conveyed in equilibrium. Thus, in line with Bertrand et al. (2010), we show that varying the content of the advertisement, its signal structure, sometimes may have a larger effect on sales than on prices.

The determinants of accuracy. In our leading example, our model predicts that Ann should receive banners whose information content increases with the ratio $\mu \frac{c}{k}$. On the one hand, keeping the ratio c/k constant, the seller has less incentives to reveal information when the user is more likely to be in the good state. On the other hand, fixing the prior, persuasive and cheap-talk banners are expected as c/k drops. In other terms, the model predicts that different products will be advertised through different strategies, from persuasive and invasive banners that rely on aggressive pricing practices, to strategically designed banners that aim to increase user's willingness to pay by reducing their relevant uncertainty. For example, considering a software, the cost of designing a free-trial is reasonably small. According to our model, this induces sellers to give in their banner the possibility to try the digital good, as in Ann's example.

Per-sale fixed fees. To conclude the discussion, notice that a platform concerned with the kind of advertisements its users are exposed to, may switch to a different fee pricing scheme. For instance, the introduction of a per-sale fixed-fee

Table 1
Individual payoff and total surplus in each state ω from each click.

State	Prob	User x	seller	platform	TS(x)
1	$1 - \mu$	$x - p$	$p - c - f$	f	$x - c$
0	μ	$-p(1 - \pi_0^*)$	$(p - c)(1 - \pi_0^*) - f$	f	$-(1 - \pi_0^*)c$

$\alpha > 0$, would add up to seller’s marginal cost. Formally, the seller’s profit becomes $\Pi = [p - (c + \alpha)]\Pr(s = 1) - kC(\pi_0)$. As a result, sellers will reveal more information through their banners.

5. Welfare

This section discusses the welfare implications of our model. The desirability of perfect targeting is essentially linked to consumers privacy¹⁹ and it depends on the objective of the regulator. In the digital economy, the access and management of big data made the economic analysis of consumers privacy strictly intertwined with the one of companies’ targeting strategies. As a consequence, a regulator should somehow balance between two effects. On the one hand, people potentially interested to a product are reached by information regarding that product. On the other hand, targeting also entails a surplus-extraction effect, which goes at the detriment of consumers. On top of that, in our model banners are never perfectly informative and their accuracy becomes very important to understand the welfare generated by the platform. In particular, the regulator might be worried about the effect of persuasive advertisement as it reduces the maximum welfare gain from targeting practices and, at the same time, it increases the frequency of purchases that users ex-post regret.

For the sake of exposition, let us consider a user x who clicks on the banner. When the banner gives her a positive recommendation, she buys the product receiving a payoff of $x - p$ in the good state and $-p$ in the bad state. The seller receives the price and faces a marginal cost whenever the good is sold and always pays a fee whenever the user clicks. The platform receives the fee independently of the result of the experiment included in the banner. As the price and the fee are transfers among agents, the total surplus generated by each transaction user-seller is given in the good state of the world by $x - c$ and in the bad state by $-(1 - \pi_0^*)c$, as reported in Table 1. The latter refers to a user-seller interaction through a clicked banner. The viability of perfect targeting affects both the clicking volumes and the per-click surplus distribution.

If aggregate welfare is the only concern of the regulator, it will sum the surplus generated by each banner while also considering also the cost of designing the experiments. The aggregate welfare can be then decomposed as follows:

$$W = \overbrace{(1 - \mu) \int_{\underline{x}}^{\bar{x}} xg(x)B(x)dx}^{\text{surplus}} - \overbrace{(1 - \mu\pi_0^*)c \int_{\underline{x}}^{\bar{x}} g(x)B(x)dx}^{\text{production cost}} - \overbrace{C(\pi_0^*)}^{\text{information cost}} \tag{4}$$

where $B(x) = 1$ if agent x clicks on a banner and zero if she does not. When targeting is perfect $B(x) = 1 \forall x$, while when targeting is imperfect $B(x) = 1$ for $x \geq x^*$ and $B(x) = 0$ below x^* . If providing any level of information were free, users would buy only in the good state, or simply $\pi_1^* = \pi_0^* = 1$ and $C(1) = 0$, with no transactions in the bad state and maximal number of surplus-generating transactions. Welfare-related results can be summarized as follows.

Proposition 4. *Welfare is strictly larger under perfect targeting than under imperfect targeting. The welfare gain due to perfect targeting drops as the banners become less accurate.*

Proof. See Appendix A. □

Proposition 4 highlights two main results. On the one hand, it shows that perfect targeting entails a shift of consumer surplus to the platform, as all users buy a product at its maximum price. Moreover, the discriminating fee is used by the platform to fully take over the value created by banners. On the other hand, perfect targeting enhances total surplus, as it entails a massive boost in the clicking volumes and purchases. However, because experiments are costly, fully informative banner do not exist, i.e., $\pi_0 < 1$ for all x s. Hence, although perfect targeting guarantees a boost in the overall supply of information, banners are biased towards persuasion, thus generating undesirable transactions in the bad state.

As in any situation in which there is price discrimination, the extent of targeting implies redistributive issues, typically from the demand to the supply side. This is true even in our setup, as targeting implies a transfer of consumer surplus to the platform. Therefore, the regulator might be interested in protecting consumers whose privacy is violated. Calling α the weight that the regulator gives to consumer surplus with the resulting $1 - \alpha$ related to the platform, the welfare can be rewritten as follows:

$$W = \alpha \underbrace{\int_{\underline{x}}^{\bar{x}} g(x)B(x)[(1 - \mu)(x - p_x^*) - \mu p_x^*(1 - \pi_0^*)]dx}_{\text{Consumer Surplus}} + (1 - \alpha) \underbrace{\int_{\underline{x}}^{\bar{x}} f_x^*g(x)B(x)dx}_{\text{Platform profit}} \tag{5}$$

¹⁹ See Acquisti et al. (2016) for an exhaustive recent review on privacy in economics.

Comparing the case of perfect targeting with the case of imperfect targeting, we can conclude the following:

Proposition 5. For any given α , if banners become less accurate, a regulator prefers imperfect to perfect targeting.

Proof. See Appendix A. \square

The latter proposition brings to light new policy considerations related to the online flow of information across demand and supply – mediated by the online platform. Although targeting improves the market efficiency in matching the two sides of the market, the persuasive nature of banners raises concerns about how additional information is ultimately used in the market. Thus, from Proposition 5 follows that the regulator should be more attentive to consumer protection whenever banners are inherently persuasive, as persuasion undermines the social benefits of targeting.

6. Conclusion

In recent years, the evolution of Internet-marketed activities started to worry policy makers insofar as digital and global transactions do not constitute an easy target for regulators. In addition, the availability of users' private information has created a market which platforms exploit to match sellers and users. Although a platform effectively improves the match between users and sellers, its incentives do not generally imply more informed purchase decisions. Regardless from its targeting ability, which determines clicking volumes, the objective of the platform is essentially to extract a rent from sellers' profits. The higher the ability of sellers to produce profits, the more they become appealing to the platform.

Interestingly, whenever the increase in sellers' profits is due to an improvement in production efficiency, then sellers' willingness to inform users about their needs decreases. Indeed, an increase in sellers' production efficiency is transferred to consumers in form of lower prices but at the cost of more uncertain decisions. Thus, the enhancement in efficiency goes along with a deterioration of experiments' accuracy.

The strategic use of information firms make when advertising their products online constitutes an issue the policy makers have to deal with. In particular, the potential surplus created from efficient matching is undermined by the cost of providing information. Given that this leads to persuasive banners and, in turns, wasteful transactions, the regulator should carefully give more weight to redistributive issues when providing information gets more costly.

Appendix A

Proof of Proposition 1, page 13. User's expected payoff increases in the probability with which the experiment sends the correct recommendation. In addition, in state $\omega = 1$ the seller's and the user's interests are perfectly aligned. Thus, because of the linearity of the problem, $\pi_1^* = 1$. Then, at π_1^* we have posterior belief $\mu_0 = 1$ since message $s = 0$ can only be observed in state $\omega = 0$. This means that IC0 is always trivially satisfied. Substituting $\pi_1^* = 1$ into IC1 we are left with:

$$\frac{1 - \mu}{1 - \mu\pi_0}x \geq p.$$

The seller's and the user's interests are now perfectly opposed which means that, in equilibrium, IC1 is binding, so that:

$$\frac{1 - \mu}{1 - \mu\pi_0}x = p^*. \quad (6)$$

Plugging the relationship between p and π_0 in Eq. (6) into the profit function, the seller's problem reduces to:

$$\max_{\pi_0} \Pr(s = 1)[p(\pi_0) - c] = \max_{\pi_0} \left[(1 - \mu\pi_0) \left(\frac{1 - \mu}{1 - \mu\pi_0}x - c \right) \right]$$

which is maximized at $\pi_0^* = 1$. \square

Proof of Lemma 1, page 15. Since $\pi_0^* = C'^{-1}\left(\frac{\mu c}{k}\right)$, the equilibrium accuracy of the experiment can be written as

$$Y = \mu C'^{-1}\left(\frac{\mu c}{k}\right) + (1 - \mu)\pi_1^*.$$

As $C(\cdot)$ is strictly convex, it directly follows that $C'^{-1}(\cdot)$ increases in k and decreases in c . Then, since $C'(\cdot)$ is monotonic, $C'^{-1}(\cdot)$ reaches its minimum, $\pi_0^* = 0$, at $c = 0$. Finally, when no experiment is available, the accuracy of user's decision is $1 - \mu$, i.e., the probability of buying in the good state. Therefore, the change in accuracy due to the observation of the experiment is

$$Y - (1 - \mu) = \mu\pi_0^*$$

which goes to 0 when π_0^* is equal to 0. \square

Proof of Proposition 3, page 21. When deciding which banner to show to user x , the platform just sets a fee that takes the entire surplus of the corresponding seller. Since the price will always be set to make the IC1 binding and to induce a click,

the profit of a seller is the only concern for the platform. Since $\pi_1^* = 1$, the profit of a seller with cost c can be directly written as:

$$\Pi(\pi_0|c) = (1 - \mu)x - (1 - \mu\pi_0)c - C(\pi_0),$$

where π_0 is any conditional probability between zero and one. Now, consider a shock to a lower marginal cost $c' < c$. It will hold that: $\Pi(\pi_0|c') \geq \Pi(\pi_0|c)$ for all π_0 , so that the objective function $\Pi(\pi_0|c')$ lies above $\Pi(\pi_0|c)$ for any π_0 chosen by the seller. This also implies $\max_{\pi_0} \Pi(\pi_0|c') \geq \max_{\pi_0} \Pi(\pi_0|c)$. Hence, the platform will always choose the most efficient seller as it will guarantee a higher per-click fee. Combining this result with the one in Lemma 1, hiring a more efficient seller results in less accurate experiments. □

Proof of Proposition 4, page 25. When targeting is perfect, each user clicks on a personalized banner. This banner will entirely extract the tailored user's surplus, as the price would precisely be the one that makes her indifferent between buying and not buying. Therefore, there will be full surplus extraction. Differently, when targeting is imperfect, users with $x < x^*$ do not click on the unique banner $\beta_{x^*}^*$, whereas users with $x > x^*$ receive a positive expected payoff, as $U(x > x^*, \beta_{x^*}^*) > 0$. In terms of total welfare, when individual x 's are not known, we have:

$$W^{IT} = (1 - \mu) \int_{x^*}^{\bar{x}} xg(x)dx - c(1 - \mu\pi_0^*)[1 - G(x^*)] - C(\pi_0^*) \tag{7}$$

With perfect targeting, the total welfare is given by:

$$W^{PT} = (1 - \mu) \int_x^{\bar{x}} xg(x)dx - c(1 - \mu\pi_0^*) - C(\pi_0^*) \tag{8}$$

Comparing Eqs. (7) and (8) we find:

$$W^{PT} - W^{IT} = (1 - \mu) \int_x^{x^*} xg(x)dx - c(1 - \mu\pi_0^*)G(x^*) > 0,$$

since Assumption 2 implies that the first term is larger than the second one. □

Proof of Proposition 5, page 26. Consider a situation in which the regulator protects privacy, so that targeting is not viable. In that case, we have $f_x^* = f^* = (1 - \mu)x^* - (1 - \mu\pi_0^*)c - C(\pi_0^*)$ and $p_x^* = p^* = \frac{(1-\mu)x^*}{1-\mu\pi_0^*}$. Moreover, consumers click on the banner only if $x \geq x^*$, so that $B(x) = 0$ for all $x < x^*$ and $B(x) = 1$ when $x \geq x^*$. Plugging these values into Eq. (5) and rearranging terms, we get:

$$W^{IT} = \alpha(1 - \mu) \int_{x^*}^{\bar{x}} g(x)x dx - [1 - G(x^*)](1 - \mu)x^* - (1 - \alpha)[(1 - \mu\pi_0^*)c - C(\pi_0^*)]$$

When privacy is not protected, perfect targeting is viable. In that case, we have $f_x^* = (1 - \mu)x - (1 - \mu\pi_0^*)c - C(\pi_0^*)$ and $p_x^* = \frac{(1-\mu)x}{1-\mu\pi_0^*}$, and all consumers click on a targeted banner, so that $B(x) = 1$ for all x . Plugging these values into Eq. (5) and rearranging terms, we get:

$$\begin{aligned} W^{PT} &= \alpha \times 0 + (1 - \alpha) \left((1 - \mu) \int_x^{\bar{x}} g(x)(x - (1 - \mu\pi_0^*)c) - C(\pi_0^*) \right) dx \\ &= (1 - \alpha)(1 - \mu) \int_x^{\bar{x}} g(x)x dx - (1 - \alpha)[(1 - \mu\pi_0^*)c - C(\pi_0^*)] \end{aligned}$$

Notice that $W^{IT} > W^{PT}$ if:

$$\alpha > \alpha_{min} \equiv 1 - \frac{\int_{x^*}^{\bar{x}} g(x)x dx - (1 - G(x^*))x^*}{\int_{x^*}^{\bar{x}} g(x)x dx + \int_x^{\bar{x}} g(x)x dx - \frac{(2-G(x^*))(C(\pi_0^*)+c(1-\mu\pi_0^*))}{1-\mu}}$$

To conclude the proof, notice that α_{min} increases with π_0^* , given that $sign \left| \frac{\partial \alpha_{min}}{\partial \pi_0^*} \right| = -sign \left| \frac{\partial (C(\pi_0^*) - c\mu\pi_0^*)}{\partial \pi_0^*} \right| > 0$ by Assumption 1. □

References

Acquisti, A., Taylor, C., Wagman, L., 2016. The economics of privacy. *J. Econ. Lit.* 54 (2), 442–492.
 Anderson, S.P., Renault, R., 2006. Advertising content. *Am. Econ. Rev.* 96 (1), 93–113.
 Athey, S., Ellison, G., 2011. Position auctions with consumer search. *Q. J. Econ.* 126 (3), 1213–1270.
 Bagwell, K., 2007. The economic analysis of advertising. *Handb. Ind. Organ.* 3, 1701–1844.
 Bergemann, D., Bonatti, A., 2011. Targeting in advertising markets: implications for offline versus online media. *RAND J. Econ.* 42 (3), 417–443.
 Bergemann, D., Morris, S., 2016. Bayes correlated equilibrium and the comparison of information structures in games. *Theor. Econ.* 11 (2), 487–522.
 Bergemann, D., Morris, S., 2019. Information design: a unified perspective. *J. Econ. Lit.* 57 (1), 1–57.
 Bertrand, M., Karlan, D., Mullainathan, S., Shafir, E., Zinman, J., 2010. What's advertising content worth? Evidence from a consumer credit marketing field experiments. *Q. J. Econ.* 125 (1), 263–306.

- Best, J., Quigley, D., 2016. Persuasion for the Long Run. Working paper - Oxford University.
- Board, S., Lu, J., 2018. Competitive information disclosure in search markets. *J. Polit. Economy* 126 (1), 1965–2010.
- Caplin, A., Dean, M., 2013. Behavioral Implications of Rational Inattention with Shannon Entropy. Technical Report. National Bureau of Economic Research.
- Chen, Y., 1997. Paying customers to switch. *J. Econ. Manage. Strategy* 6 (4), 877–897.
- De Corniere, A., De Nijs, R., 2016. Online advertising and privacy. *RAND J. Econ.* 47 (1), 48–72.
- DellaVigna, S., Gentzkow, M., 2010. Persuasion: empirical evidence. *Annu. Rev. Econ.* 2 (1), 643–669.
- Ellison, G., Ellison, S.F., 2018. Match Quality, Search, and the Internet Market for used Books. Technical Report. National Bureau of Economic Research.
- Fawcett, T., 2006. An introduction to ROC analysis. *Pattern Recognit. Lett.* 27 (8), 861–874.
- Fudenberg, D., Tirole, J., 2000. Customer poaching and brand switching. *RAND J. Econ.* 31 (4), 634–657.
- Gomes, R., 2014. Optimal auction design in two-sided markets. *RAND J. Econ.* 45 (2), 248–272. doi:10.1111/1756-2171.12050.
- Hagiu, A., Jullien, B., 2011. Why do intermediaries divert search? *RAND J. Econ.* 42 (2), 337–362.
- Hirsh, J.B., Kang, S.K., Bodenhausen, G.V., 2012. Personalized persuasion: tailoring persuasive appeals to recipients? *Personality traits. Psychol. Sci.* 23 (6), 578–581. doi:10.1177/0956797611436349.
- Hoffmann, F., Inderst, R., Ottaviani, M., 2014. Persuasion through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation. JW Goethe University, Frankfurt, UCL, and Bocconi.
- Johnson, J.P., Myatt, D.P., 2006. On the simple economics of advertising, marketing, and product design. *Am. Econ. Rev.* 96 (3), 756–784.
- Kaldor, N., 1950. The economic aspects of advertising. *Rev. Econ. Stud.* 18 (1), 1–27.
- Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. *Am. Econ. Rev.* 101 (6), 2590–2615.
- Kamenica, E., Gentzkow, M., 2014. Costly persuasion. *Am. Econ. Rev. (Papers and Proceedings)* 104, 457–462.
- Kolotilin, A., Mylovannov, T., Zapechelnyuk, A., Li, M., 2017. Persuasion of a privately informed receiver. *Econometrica* 85 (6), 1949–1964.
- Kosinski, M., Stillwell, D., Graepel, T., 2013. Private traits and attributes are predictable from digital records of human behavior. *Proc. Natl. Acad. Sci.* 110 (15), 5802–5805.
- Matz, S., Kosinski, M., Stillwell, D., Nave, G., 2017. Psychological framing as an effective approach to real-life persuasive communication. In: Gneezy, A., Griskevicius, V., Williams, P. (Eds.), *NA - Advances in Consumer Research*, 45. Association for Consumer Research, Duluth, MN, pp. 276–281.
- Mikians, J., Gyarmati, L., Erramilli, V., Laoutaris, N., 2012. Detecting price and search discrimination on the internet. In: *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*. ACM, pp. 79–84.
- Nelson, P., 1974. Advertising as information. *J. Polit. Economy* 82 (4), 729–754.
- Persico, N., 2000. Information acquisition in auctions. *Econometrica* 68 (1), 135–148.
- Rayo, L., Segal, I., 2010. Optimal information disclosure. *J. Polit. Economy* 118 (5), 949–987.
- Robinson, J., 1969. *The Economics of Imperfect Competition*. Springer.
- Shaffer, G., Zhang, Z.J., 2000. Pay to switch or pay to stay: preference-based price discrimination in markets with switching costs. *J. Econ. Manage. Strategy* 9 (3), 397–424.
- Shannon, C.E., 1948. A mathematical theory of communication. *Bell Syst. Tech. J.* 27 (3), 379–423.
- Stigler, G.J., 1961. The economics of information. *J. Polit. Economy* 69 (3), 213–225.
- Taylor, C.R., 2004. Consumer privacy and the market for customer information. *RAND J. Econ.* 35 (4), 631–650.
- Thisse, J.-F., Vives, X., 1988. On the strategic choice of spatial price policy. *Am. Econ. Rev.* 122–137.
- Villas-Boas, J.M., 1999. Dynamic competition with customer recognition. *RAND J. Econ.* 30 (4), 604–631.