

Carroni, Elias; Ferrari, Luca; Righi, Simone

Working Paper

The price of discovering your needs online

Quaderni - Working Paper DSE, No. 1116

Provided in Cooperation with:

University of Bologna, Department of Economics

Suggested Citation: Carroni, Elias; Ferrari, Luca; Righi, Simone (2018) : The price of discovering your needs online, Quaderni - Working Paper DSE, No. 1116, Alma Mater Studiorum - Università di Bologna, Dipartimento di Scienze Economiche (DSE), Bologna, <https://doi.org/10.6092/unibo/amsacta/5761>

This Version is available at:

<https://hdl.handle.net/10419/177605>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/3.0/>



ISSN 2282-6483

Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

The Price of Discovering Your Needs Online

Elias Carroni
Luca Ferrari
Simone Righi

Quaderni - Working Paper DSE N°1116



The Price of Discovering Your Needs Online*

Elias Carroni[†]

Luca Ferrari[‡]

Simone Righi^{§¶}

Abstract

Thanks to new digital technologies, web users are continuously targeted by offers that potentially fit their interests even if they are not actively looking for a product. Does this matching always promote transactions with high social value? We consider a model in which web users with state-contingent preferences are targeted by relevant banners. We characterize the optimal strategy of a seller who, in addition to the price of the offered good, designs a banner. We show that, in equilibrium, there is a positive relationship between the price of the offered good and the accuracy of the banner sent to users. Then, we consider the strategic decision of a Platform that attracts sellers because of its targeting abilities and we underline that a reduction in seller's costs may translate into less informative banners and lower prices, fueling purchases of goods that rational individuals may regret due to the persuasive nature of banners.

Keywords: Bayesian Persuasion, Targeting, Platforms. *JEL codes:* D80, D82, D83, L10, M37.

*We wish to thank Ennio Bilancini, Leonardo Boncinelli, Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, 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).

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

[‡]University of Corsica, UMR CNRS 6240 LISA, Corte, France. 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

Ann is surfing the web, she just started to play the guitar and suddenly, from her Facebook feeds, new advertisements appear:¹ they are related to music, music equipment and mostly guitar lessons. Ann is a self-learner and she is captured by an advertisement of a software which promises her quick improvements. Importantly, Ann was not actively looking for such a product. On the contrary, because of her collected private information, the advertising platform exposed Ann to a banner promoting a product that might fit her needs. Nevertheless, the information she acquired through the banner does not answer the question Ann is struggling with: “do I need this product?”.

The software costs \$20 and Ann knows that, if she needs the product, she would enjoy a payoff of 30. She also knows that, if she does not need the product, she would receive a payoff of zero and suffer a utility loss due to the fact that she would still pay for a product that does not match her needs. Ann believes that she does not need the product with probability 0.7. Thus, she would never buy the product at this price since her expected payoff ($0.3 \times (30 - 20) - 0.7 \times 20 = -11$) would be lower than zero.

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? One obvious option would be the one of lowering the price of the sponsored product. For example, the seller may set a price of \$9, making Ann just indifferent between buying it or not. A different option would be to change her perception of needs, keeping the price of \$20. The seller could design an experiment which conveys some information about Ann’s payoff-relevant state, that is, her needs. To this purpose, the seller can commit to a signal structure over purchase recommendations.² In other terms, the seller designs a reliable experiment from which Ann may learn something new about the match between the product and her needs. 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

¹Realistically 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 that aim to use Ann’s collected information in order to show her content which is more likely to be close to her actual needs.

²Hence, the seller is a Sender with commitment power as in [Rayo and Segal \(2010\)](#); [Kamenica and Gentzkow \(2011\)](#). We discuss the related literature in the following section.

recommendations to “buy” and “not buy” are sent with probability 0.95 and 0.8 conditional on Ann’s true needs, respectively.³ Given this opportunity, Ann is willing to proceed with the free trial. When Ann likes 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.⁴

This simple example shows that, whenever an individual is uncertain about the benefits she would derive from consuming a good, a seller might be willing to reveal some information in order to increase the probability of selling the good. However, the example also shows that the experiment leaves Ann with positive surplus since her expected utility is larger than zero. This means that the seller, if allowed, could design an optimal price experiment pair in order to maximize his profits. In other words, the price of the offered good depends not only on how much the consumer values the product, but also on how much information the seller is able to convey in order to reduce user’s uncertainty.

Nevertheless, the most important concern that the seller has is finding Ann. To this end, the seller might be willing to employ the aid of a middleman who would be able to target Ann and all users that are similar to Ann on an Internet platform. In the reality of digital markets, advertisers outsource the job of reaching users interested in their products to Internet giants 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.”⁵

Throughout this work, we focus on those situations in which the users of a website are exposed to offers they are not actively looking for. Going back to our example, the seller does not know Ann and the match with Ann and other similar users, who potentially share the same interest in the product, is the result of the job of a middleman; the latter offers

³Formally: $\Pr(\text{not buy}|\text{no need}) = 0.8$ and $\Pr(\text{buy}|\text{need}) = 0.95$.

⁴One can verify that whenever Ann does not like the trial, she would prefer not to buy the product.

⁵See https://www.facebook.com/ads/about/?entry_product=ad_preferences

sellers the possibility to expose potentially-interested users to their offers in exchange of a per-click fee.

This paper has a two-fold aim. The first is to characterize the optimal strategy of a seller that consists of an information-design strategy and a price. More precisely, the seller decides the information structure of an experiment and the price of the offered good which can be interpreted as the price of the user's risky action "buy the product". We show that the reduction of user's uncertainty granted by the observation of the outcome of the experiment has to be compensated by an increase of the price of the offered good. That is to say, we derive an inverse relationship between the accuracy of the experiment and the price of the offered good. When the experiment is fully informative the seller asks for the maximum price he could ask for, on the contrary, when the experiment is not informative (i.e., when it does not change user's perception), then the seller asks for the minimum price. In particular, we show that there exists a multiplicity of optimal experiment-price strategies that maximize seller's profit and always leave the user with null expected payoff. Nevertheless, we also show that, when the production of the good as well as the design of the experiment are costly, then the set of equilibria shrinks to a unique optimal experiment-price pair characterized by a price that sustains a persuasive experiment. Thus, the equilibrium experiment may suggest the user to buy a product she does not need.

Having characterized the relationship between a generic seller and a generic user, we then switch our attention to the strategic decisions of an Internet platform maximizing profits by selling banner spaces in exchange for a per-click fee. Sellers are lured into joining the platform because of its targeting abilities which create potential transactions with the users they are matched with. Accordingly, our aim is to understand the role of the platform as a matchmaker and its effects on welfare. The first observation we make is that welfare is maximized whenever the platform has perfect targeting abilities, as each seller is matched with his preferred user. This allows sellers to tailor their price-experiment pair on the basis of users' characteristics. Users are left with null expected payoff and all surplus is extracted. In turn, sellers' profits are captured by the per-click fee which transfers the rent to the platform. Since in equilibrium sellers send persuasive experiments, there may emerge transactions generating negative social surplus. Finally, we show that a decrease of sellers' marginal cost of production might reduce the value created by banners. This is

because more efficient sellers design less informative experiments lowering the price of the good increasing sales, also in cases in which users would have preferred not to buy. Because a reduction of sellers' marginal cost of production increases its profits, the platform prefers to hire more efficient sellers whereas, in the interim, users' decisions become cheaper and less accurate. Therefore, the efficiency gain might be balanced by a decrease of accuracy, thus reducing welfare.

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 is devoted to the understanding of the impact of targeting on welfare and of the role of the platform in information provision. Finally, Section 5 concludes.

2 Related Literature

The emergence of online markets and 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.

Targeting is the natural way to respond to the huge mass of information accessible online. 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. 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, in our model targeting is always efficient, but it generates social waste, 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 finally, [Gomes \(2014\)](#) highlights a trade-off between rent extraction and clicking volumes, as advertiser with the highest willingness to pay does not necessarily offer the most relevant advertisement and long-run clicks depend on the relevance of sponsored links. Differently from this literature, 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 although 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, the 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 another 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. Therefore, our model interpretation of advertisement does not immediately match with the classification traditionally adopted by the literature, i.e., persuasive ([Robinson 1969](#); [Kaldor 1950](#)) and informative ([Stigler 1961](#); [Nelson 1974](#)) advertisement.⁶ 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 outcome of the experiment.

Our modeling choice is in the spirit of the recent stream of literature about information control, initiated by [Rayo and Segal \(2010\)](#) and [Kamenica and Gentzkow \(2011\)](#). In these

⁶ See [Bagwell \(2007\)](#) for a review on different classes of advertisement.

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 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-targeting⁷ may harm consumers depends on the ability of firms to price-discriminate, on the competition between senders and on consumers’ wariness.

3 The Model

An Internet platform runs a website and offers banner spaces to sellers in exchange for a per-click fee. Users who visit the website are exposed to banners, which can be tailored whenever possible. Whenever a user clicks on a banner, she 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. When the outcome of the experiment is positive, the user buys the product. In line with the advertising policy of Facebook and Google, the banners cannot contain contents that are manipulative or deceptive per se. In our setup, this is guaranteed by the commitment power we give to the advertisers, who cannot manipulate the outcome of the experiment.⁸ The timing of the model is as follows:

⁷[Hoffmann et al. \(2014\)](#) define hyper-targeting as “the collection and use of personally identifiable data by firms to tailor selective disclosure”.

⁸A possible source of commitment is reputation which sellers may build online thanks to the interactions with users who frequently comment on sponsored posts. For instance, [Best and Quigley \(2017\)](#) show that




Figure 1: Timeline

First, the platform sets a per-click fee that sellers have to pay in order to advertise their products through banners. Then, sellers enter the platform by paying the fee, and each of them reaches targeted users. Banners contain the price of the sponsored product and an experiment that may improve user’s knowledge of her needs. Given their prior beliefs, consumers each consumer decides whether to click on a banner and - conditional on the information received - whether to purchase the advertised product. Finally payoffs of all agents are realized.

For the sake of exposition, we proceed as follows. In the next section, we study a generic seller-user interaction. Subsequently, we study the optimal decisions of a platform that attracts sellers because of its targeting abilities. We say that the platform has perfect targeting abilities whenever it is informed about each user’s willingness to pay. When the platform has imperfect targeting abilities, the platform is just informed about the distribution of the willingness to pay of its users.

3.1 Offers through Banners

In this section, we formalize the example proposed in the introduction. We consider one seller targeting a generic user with willingness to pay $x \in [\underline{x}, \bar{x}]$. In terms of our example, the seller is targeting Ann and all the users who share the same characteristics 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

“the desire to persuade in the future can generate credibility, and hence persuasion, today”.

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. 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:

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

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 the maximum payoff which is 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, $U_{NB}(x, p) = 0$.

The seller reaches the user through his banner which contains the price of the offered good and an experiment whose outcomes are statistically correlated with user's state of necessity. In particular, since the experiment reveals something about the user's payoff-relevant state, the user might find it worth clicking on the banner. Formally, an experiment is characterized by a pair of conditional probability distributions π_0 and π_1 :

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

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

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

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

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

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.

Given that the seller designs the information content of the experiment but he does not manipulate directly the observed message s , the seller is a *Sender* with *commitment power* as in [Kamenica and Gentzkow \(2011\)](#). In other words, through the choice of the experiment $\pi \in [0, 1]^2$, the seller designs what the user learns from participating in the experiment. In our leading example, the free trial contains some information relevant to Ann, so that she is better informed about whether she needs the software after having tried its demo version. 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 *Sender* to choose the price of the sponsored product which, in our setup represents the cost of the risky action *buy*.

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 possible costs the seller may incur in, e.g., production or shipping costs, and information-design costs. We then relax the analysis by studying how these frictions affect seller’s optimal strategy.

Seller’s strategic variables are summarized by the banner through which he reaches the targeted user. Thus, the seller targeting user x is described by the banner $\beta_x = (p, \pi)$ where x is referred to the targeted user as explained above. The game 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 a shopping advice s and then reacts accordingly.

We solve the user-seller interaction problem by means of the Bayes Correlated Equilibrium (BCE) (Bergemann and Morris 2016) with the Sender-Preferred refinement. Thus, the equilibria are those BCE in which the *Sender* payoff is maximized, as in Kamenica and Gentzkow (2011). 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

$$\Pi = \Pr(s = 1)p. \quad (2)$$

Notice that the probability with which the outcome of the experiment results in a recommendation to buy the product is

$$\Pr(s = 1) = \mu(1 - \pi_0) + (1 - \mu)\pi_1, \quad (3)$$

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

$$\beta_x^* = \arg \max_{\beta_x} \Pr(s = 1)p \text{ s.t. IC0 and IC1.} \quad (4)$$

Proposition 1. *The Optimal Banner.*

Consider the seller targeting user x . The optimal banner $\beta_x^ = (p^*, \pi_0^*, \pi_1^*)$ is characterized by $\pi_1^* = 1$ whereas π_0^* and p^* satisfy*

$$\frac{1 - \mu}{1 - \mu\pi_0^*}x = p^*. \quad (5)$$

Therefore, there is a continuum of equilibria in which $\pi_0^(p^*)$ is strictly increasing in the price of the offered good and the equilibrium payoff is $(1 - \mu)x$ for any β_x^* .*

Proof. 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. Finally, for any $\pi_0 \in [0, 1]$, there exists a $p^* \in \mathbb{R}_+$ such that **IC1** is satisfied with equality, that is

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

From the relationship, it is immediate to note that p^* strictly increases in π_0^* . Substituting the optimal $\beta_x^* = (p^*, \pi_0^*, 1)$ into the profit equation of the seller each equilibrium combination yields $(1 - \mu)x$. ■

Proposition 1 describes the choice of the optimal banner as an information-design problem in which, in addition to the information content of the experiment, the seller also chooses the price of the product sold. At one extreme, we find the *fully informative* banner $\beta_x^* = (x, 1, 1)$ which always suggests the user not to buy when she doesn't need the product, and to buy in the good state of the world. As this equilibrium eliminates user's uncertainty about her state of necessity, the latter is willing to pay the maximal price. At the opposite extreme, the *non-informative* banner $\beta_x^* = ((1 - \mu)x, 0, 1)$ suggests the buyer to buy in all states of the world, thus effectively providing no information to the user. In turns, the latter will be willing to buy at a price equal to her expected payoff given her priors about the state of the world.⁹ Intuitively, the seller is facing an endogenous trade-off between low demand (selling with low probability) and high prices on one side, and high demand (selling with high probability) and lower prices on the other side. This reflects the fact that the experiment must be incentive compatible and, therefore, the price must accommodate the purchase recommendations. In particular, since $\pi_1 = 1$, the probability of sale resulting from a credible information structure is $\Pr(s = 1) = 1 - \mu\pi_0$, where π_0

⁹Observe that this would be the unique outcome of a Sender-Receiver game in which sender has no commitment power. Given that interests are opposed, there is no information the seller could share without commitment and her only possibility would be to decrease the price. Nevertheless, the commitment assumption underlines the fact that, thanks to the platform, the seller may build some reputation which allows her to increase the price of the good in exchange for a more accurate experiment.

has to satisfy the IC1. Thus, in order to increase the probability of sale, the seller has to decrease π_0 , and a percentage change of π_0 will give a percentage change in the probability of sale equal to $\frac{\partial \Pr(s=1)}{\partial \pi_0} \times \frac{\pi_0}{\Pr(s=1)} = \frac{-\mu\pi_0}{1-\mu\pi_0}$. Meanwhile, for the same percentage change to be incentive compatible, the price has to vary and a percentage change in π_0 gives $\frac{\partial p}{\partial \pi_0} \times \frac{\pi_0}{p} = \frac{\mu(1-\mu)x}{(1-\mu\pi_0)^2} \times \frac{(1-\mu\pi_0)\pi_0}{(1-\mu)x} = \frac{\mu\pi_0}{1-\mu\pi_0}$. Hence, the reduction in margins required to provide less information (and sell also in the bad state) is always perfectly offset by the increase in demand. As an example, consider a situation in which $x = 10$ and $\mu = 1/2$. If the seller sells only in the good state producing a *fully informative* experiment, $\pi_0 = 1$, the optimal price turns out to be $p = 10$ and the expected profit $\Pi = 5$. If the seller instead switches to the opposite extreme, proposing a *non-informative* experiment with $\pi_0 = 0$, the price cannot exceed $p = 5$. In relation to the fully informative solution, the seller doubles the probability of sale, but charges half the price: the expected profit would be precisely the same as in full information.

The positive relationship between the price and π_0 can be understood in terms of accuracy of the experiment. The accuracy of the experiment is defined as

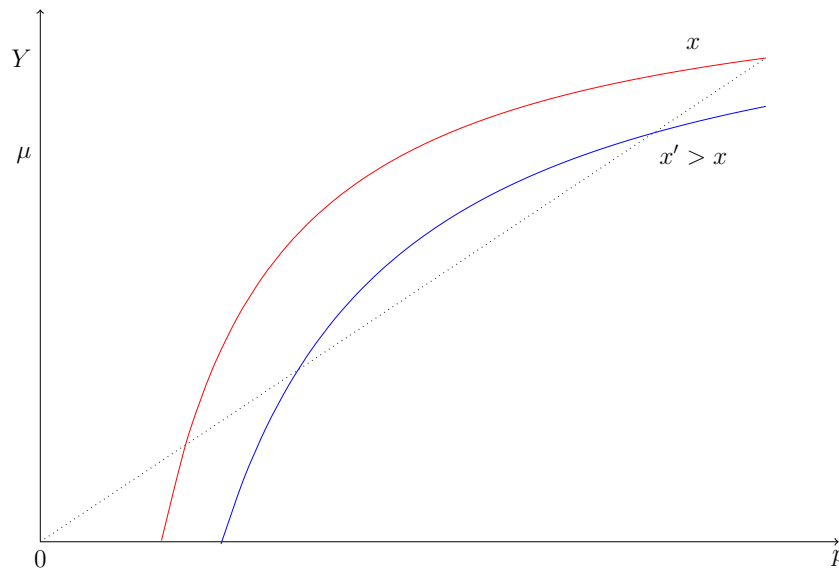


Figure 2: Accuracy of the experiment as a function of the price. The constraint is given by $\pi_0^*(p) \in [0, 1]$ and $\mu = 0.8$. Two users $x' > x$ receive different equilibrium banners.

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

that is the probability with which the purchase recommendation matches user's needs. Given that in equilibrium $\pi_1^* = 1$, from the relationship $\pi_0^*(p^*) \in [0, 1]$ 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.

Proposition 1 describes the abstract relationship linking the price of the offered good and the information content of the experiment neglecting the market frictions and transaction costs that affect seller's behavior. Indeed, the multiplicity of optimal banners targeting consumer x described in Proposition 1 vanishes when the seller faces a positive marginal cost of production $c > 0$, which can be thought as a shipping cost from the warehouse to the user. In particular, the seller has to pay a cost c whenever the good is purchased. Then, the seller's problem becomes

$$\max_{\beta_x} \Pr(s = 1)(p - c) \text{ s.t. IC0 and IC1.} \quad (6)$$

The cost the seller has to pay does not affect the user's incentive compatibility constraints which guarantee that, in equilibrium $\pi_1^* = 1$ and

$$\frac{1 - \mu}{1 - \mu\pi_0^*}x = p^*. \quad (7)$$

The next proposition shows how the set of equilibria identified in Proposition 1 shrinks in favor of fully informative experiments whenever the seller incurs a positive marginal cost $c > 0$.

Proposition 2. Marginal cost of production.

Consider the seller targeting user x . When the seller faces a marginal cost of production $c > 0$, the optimal banner $\beta_x^ = (p^*, \pi_0^*, \pi_1^*)$ is characterized by $\pi_1^* = \pi_0^* = 1$ and $p = x$.*

Proof. The relationship between p and π_0 needed to fulfil incentive compatibility, i.e., $\frac{1-\mu}{1-\mu\pi_0}x = p(\pi_0)$ reduces seller’s problem to the following:

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

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

The intuition behind this result is easily grasped by following up on our previous example. Consider again a switch from a price 10 (which is incentive compatible only if $\pi_0 = 1$) to a price of 5, but now assuming a marginal cost $c = 2$. The reduction in margins going from $10 - 2 = 8$ to $5 - 2 = 3$ is such that the new margin is $3/8 < 1/2$, whereas the probability of sale is again doubled. This means that providing full information - and the maximal price - is always the best solution for the seller. Therefore, a seller who wants to target user x will do so by designing a fully informative experiment, thus selling only in the good state at the maximal price.

Proposition 2 stresses the fact that positive costs of production provide an incentive for the seller to design a fully informative banner in order to maximize his profits. Nevertheless, experiments are rarely fully informative as gathering information is usually costly. Designing an experiment would require the seller to sustain relevant costs to study the drivers that can inform a user about her needs. In turn, 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 – the seller has less incentives to provide information given that it is costly to do so. The cost of designing an experiment can be related to the reduction of uncertainty the user experiences after having run the experiment (see for instance [Kamenica and Gentzkow 2014](#) and [Martin 2017](#)). Thus, the natural formulation of a cost function for such a problem is based on Shannon’s entropy function ([Shannon, 1948](#)). However, in order to not overly complicate the model, we make the following simplifying assumption about the cost of designing an experiment: for all sellers, the cost of the experiment $C(\pi)$ is given by $kC(\pi_0)$ where $C(\cdot)$ is a strictly convex function with $C(0) = 0$, $C'''(\cdot) = 0$ and $k \geq 0$ is a scale parameter which measures the ability of a seller to provide a more accurate experiment.¹⁰ Under this formulation, all sellers are freely able to set $\pi_1^* = 1$ whereas the accuracy of the experiment in state $\omega = 0$ is endogenously determined. This assumption

¹⁰The fact that $C'''(\cdot) = 0$ is a mathematical convenience that does not affect the results of the model as

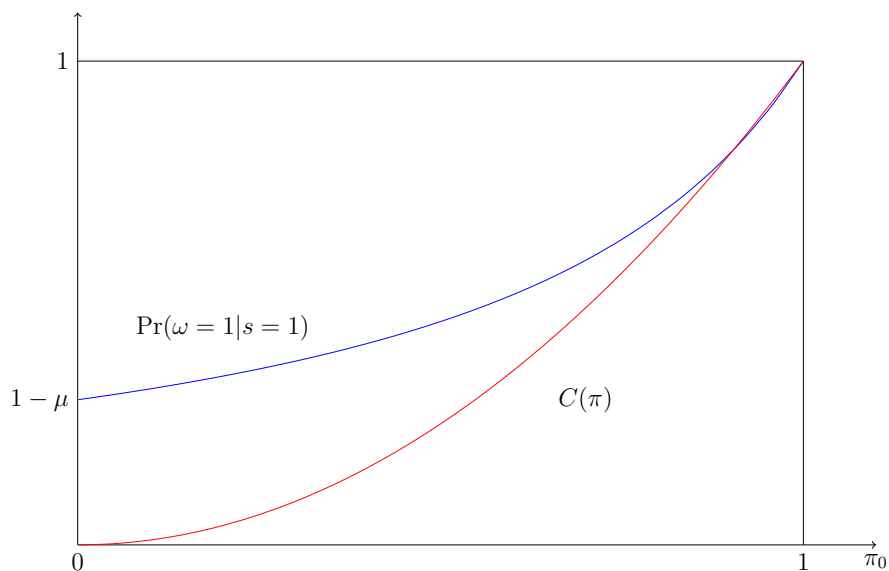


Figure 3: Posterior belief $\Pr(1|1)$ as a function of π_0 and cost function. The cost of designing the experiment strictly increases in π_0 . When $\pi_0 = 1$, the experiment becomes fully informative and the reduction of uncertainty is maximized, i.e., the user is informed about her necessities.

implies that the cost of designing an experiment increases in the posterior belief $1 - \mu_1$ the user reaches after observing message $s = 1$. In other words, given that $\pi_1^* = 1$, starting from $\pi_0 = 0$ one would have no reduction of uncertainty whereas ending up with $\pi_0^* = 1$ one would reach the fully informative experiment and therefore the maximum reduction of uncertainty (see Figure 3). Thus, in this final case, the seller's problem is

$$(p - c) \Pr(s = 1) - kC(\pi_0) \text{ s.t. IC0 and IC1.} \quad (8)$$

Proposition 3. Costly experiment.

Consider the seller targeting user x . When the seller faces marginal cost $c > 0$ and the cost of designing the experiment is $C(\pi) = kC(\pi_0)$ as described above, then the optimal banner satisfies $\beta_x^* = (p^*, \pi_0^*, \pi_1^*)$, where:

$$p^* = \frac{1 - \mu}{1 - \mu\pi_0} x, \pi_1^* = 1, \text{ and } \pi_0^* = C'^{-1} \left(\frac{\mu c}{k} \right)$$

Proof. From relationship (7), the seller's problem reduces to

$$\max_{\pi_0} = (1 - \mu)(x - c) - \mu(1 - \pi_0)c - kC(\pi_0).$$

From the convexity of $C(\pi_0)$, it follows that the FOC

$$\mu c = kC'(\pi_0^*)$$

are necessary and sufficient and π_0^* is implicitly defined by the cost function. ■

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 also 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. This trade-off is captured by the first order conditions, indeed:

$$\underbrace{\mu c}_{\text{MB of no sale in } \omega=0} = \underbrace{kC'(\pi_0^*)}_{\text{MC of info}}. \quad (9)$$

long is C is strictly convex.

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 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 9 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 of the cost of designing the experiment results in a higher equilibrium accuracy;*
- ii) A decrease of marginal costs results in a lower equilibrium accuracy;*
- iii) As $c \rightarrow 0$, the equilibrium experiment provides no additional accuracy, i.e., $\pi_0^* \rightarrow 0$.*

Proof. 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. ■

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 Figure 3. A change in $\frac{\mu c}{k}$ induces a change in the price according to 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^* .

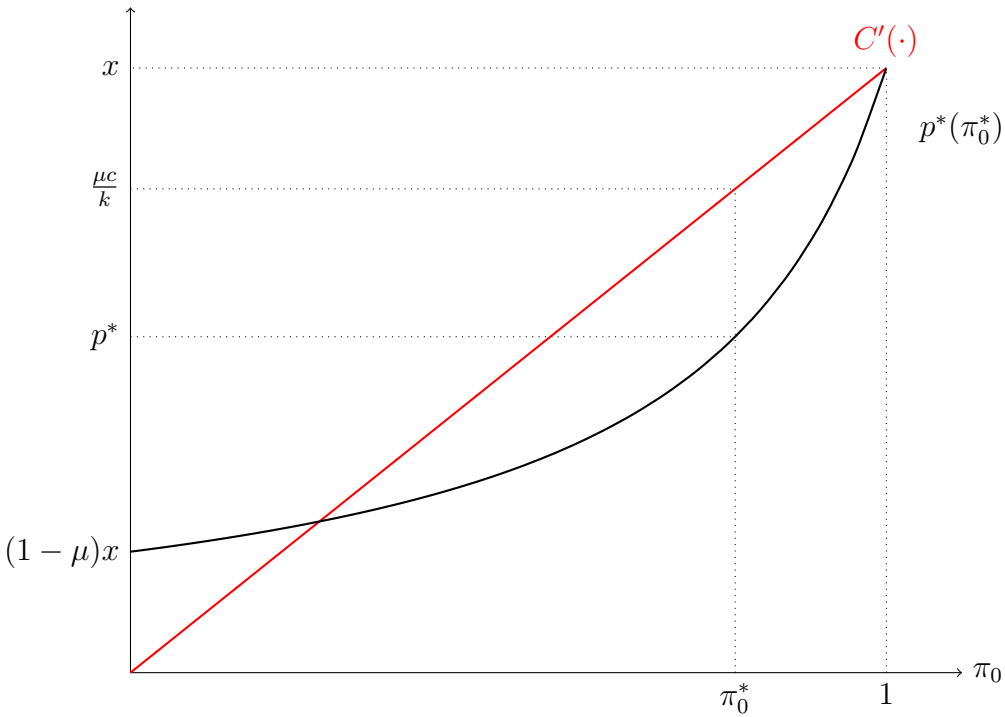


Figure 4: The equilibrium price $p^*(\pi_0^*)$ is determined by the optimal choice of how much information to convey in state 0, i.e., π_0 . The picture shows that the optimal level of π_0^* solves the trade-off between the cost of producing the good, c , and the marginal cost of increasing π_0 , $C'(\cdot)$. Then, the optimal price is determined by the equilibrium relationship $p^* = \frac{1-\mu}{1-\mu\pi_0^*}x$. Finally, 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.

In this section 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 user clicks on the banner and she receives a purchase recommendation behaving 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. The analysis has nevertheless been carried out without considering the role of the platform matching sellers and users. In the following section, we introduce the platform who matches sellers, characterized by their optimal banners, and users, characterized by their willingness to pay.

3.2 The platform: a perfect matchmaker

In the previous section, we characterized the optimal banner targeting user x . At this point, we are able to discuss the role of the platform as a matchmaker. The role of the platform is to target with a banner each user and it is precisely for this reason that sellers are willing to pay a fee for its services. We assume that the platform sets a per-click fee f , which means that whenever a user clicks on banner β_x , the correspondent seller has to pay the amount f . Since a banner β_x^* is incentive compatible for user x , then it would generate a click also for all users with $x' > x$. Moreover, the seller providing banner β_x^* is the one willing to pay the most to reach user x in the web.¹¹ The 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 low: $c > \frac{kC'(1)}{\mu}$.*

¹¹Notice that all banners $\beta_{x'}$ with $x' < x$ would not generate any profit when shown to user x , who does not click on it. 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.

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 the following, we are going to consider two different regimes: perfect and imperfect targeting. In the first, the platform knows each user's x and can thus perfectly tailor banners. In the second, the platform only knows the distribution.

Imperfect targeting. The platform chooses which banner to show to each user among the optimal banners as described in Proposition 3. The platform's problem can be reduced to the choice of which user x to target, since banners are optimally designed. If the platform decides to show banner β_x^* to all users, the maximal fee is $f = \Pi(\beta_x^*)$, as the fee cannot be higher than the seller's profit.¹² Hence, the platform's problem reduces to

$$\max_x \underbrace{[1 - G(x)]}_{\text{click volume}} \underbrace{\Pi(\beta_x^*)}_{\text{per-click fee}} . \quad (10)$$

Intuitively, the seller faces the following trade-off. On the one hand, by targeting a user with a lower x , the platform increases the number of clicks and reduces the fee. On the other hand, by targeting a user with a larger x , it is able to increase the fee at the expense of a reduction of click volume. The following lemma reports the optimal choice of the platform.

Lemma 2. *When targeting is imperfect, the platform shows a unique banner $\beta_{x^*}^*$, where:*

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

Proof. To see that a solution to platform's problem exists, notice that the function $[1 - G(x)]\Pi(\beta_x^*)$ is continuous in the closed and bounded interval $[\underline{x}, \bar{x}]$, thus, by the extreme value theorem, it attains a maximum at x^* . ■

¹²In other words, we implicitly assume that the advertising market is competitive, *vis-à-vis* the platform, because sellers compete for a limited number of banner spaces. Therefore, sellers' expected profits are zero. The only thing that would change if sellers had market power would be the surplus distribution.

Perfect Targeting. In this case, the platform is able to discriminate among users choosing for each $x \in [\underline{x}, \bar{x}]$ the optimal banner β_x^* which, as discussed in the previous case, maximizes the surplus of each seller in the interaction with user x and, in turn, the rent extractable by the platform.

Lemma 3. *When targeting is perfect, each user $x \in [\underline{x}, \bar{x}]$ is shown an optimal banner β_x^* .*

Proof. The proof is direct consequence of Proposition 3. ■

The only difference between Lemma 2 and Lemma 3 is in the number of different banners shown by the platform. When perfect targeting is not viable, the platform has no better option than placing the same banner β_{x^*} to all users. On the contrary, if individual x s are known by the platform, each user receives a personalized offer: a banner that contains a price and an experiment. The consequences of such a difference are better understood in terms of welfare as we will discuss in the following section.

4 Discussion

In this section, we discuss the welfare consequences of targeting and the platform's incentives to provide accurate experiments, showing also that there is an important relationship between accuracy of the experiments and social welfare.

In order to understand welfare, 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 the following table.

Table 1 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. Targeting makes more users click on a banner. In particular, when targeting is perfect, each

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	$-(1 - \pi_0^*)c$

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

of them 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. When targeting is imperfect, the clicking volumes are lower, because users with $x < x^*$ do not click on the unique banner $\beta_{x^*}^*$. The user with $x = x^*$ clicks on the banner and receives zero expected payoff, as the tailoring banner lets the seller extract her entire surplus. Finally, users with higher willingness to pay will receive a positive expected payoff, as $U(x > x^*, \beta_{x^*}^*) > 0$. Therefore, when targeting is perfect, users' surplus is completely extracted. Finally, the sellers are indifferent, as their surplus is always extracted by the platform through the fully discriminating fee. The following proposition summarizes the impact of targeting on surplus distribution.

Proposition 4. *Perfect targeting entails a reduction of users' surplus and an increase of the platform's profits. Sellers are always left with zero surplus.*

Proposition 4 shows that perfect targeting entails a shift of consumer surplus to the platform. This is a classic tale of two-sided price discrimination. When targeting is perfect, each seller is matched with his desired user whose surplus is totally extracted by the banner. In other words, perfect targeting makes all users click on a banner and possibly buy a product, but all of them become infra-marginal. Once the sellers extract the entire surplus of the other side of the market, the discriminating fee is used by the platform to fully take over the value created by banners.

The reasoning used in Table 1 applies to every user-seller interaction, thus the expected total surplus sums the surplus generated by each banner and takes into account also the cost of designing the experiments. The expected welfare is:

$$W = \underbrace{(1 - \mu) \int_{\underline{x}}^{\bar{x}} xg(x)B(x)dx}_{\text{surplus}} - \underbrace{(1 - \mu\pi_0^*)c \int_{\underline{x}}^{\bar{x}} g(x)B(x)dx}_{\text{production cost}} - \underbrace{C(\pi_0^*)}_{\text{information cost}} \quad (11)$$

where $B(x) = 1$ if agent x clicks on a banner and zero if she does not. With perfect targeting $B(x) = 1 \forall x$, while with imperfect targeting $B(x) = 1$ 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. Total welfare can be split into two components. The first one is the value generated within the platform, W_p , which represents the surplus component emerging thanks to the existence of the middleman who, through the banners, allows the interaction between sellers and users. The second one does not concern the platform, as it is the information cost that sellers sustain to provide potential consumers with experiments.

Comparing the welfare with perfect and imperfect targeting yields the following proposition.

Proposition 5. *Welfare is strictly larger under perfect targeting. The welfare gain due to perfect targeting drops as long as the banners become less accurate.*

Proof.

When individual x 's are not known, the total welfare is given by:

$$W^{IT} = (1 - \mu) \int_{x^*}^{\bar{x}} xg(x)dx - c(1 - \mu\pi_0^*) [1 - G(x^*)] - C(\pi_0^*) \quad (12)$$

With perfect targeting, the total welfare is given by:

$$W^{PT} = (1 - \mu) \int_{\underline{x}}^{\bar{x}} xg(x)dx - c(1 - \mu\pi_0^*) - C(\pi_0^*) \quad (13)$$

Comparing Equations (12) and (13) we find:

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

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

Moving from imperfect to perfect targeting entails a massive boost in the clicking volumes and purchases in the good as well as in the bad state. In particular, because experiments are costly, fully informative experiments are never sent, i.e., $\pi_0 < 1$ for all x s. Hence, although perfect targeting increases the overall supply of information conveyed and sellers would like to provide as much information as possible, the experiments are biased towards persuasion, thus generating undesirable transactions in the bad state.

As a consequence, accuracy becomes very important to understand the welfare generated by the platform by means of the banners. In particular, less accurate experiments result in a welfare loss both under imperfect and perfect targeting. Perfect targeting emphasizes this effect with respect to imperfect targeting, as all users, rather than some of them, click on non-accurate banners, giving us the second part of Proposition 5.

The accuracy of experiments is then a crucial aspect as long as welfare is concerned. Indeed, for given level of targeting, a decrease of accuracy has the effect of reducing welfare. Therefore, two questions naturally arise. Which advertisers has the platforms incentives to hire? Do they provide accurate experiments? Once answered these two questions, we are able to demonstrate that increasing productive efficiency may counterintuitively decrease the social value created by the platform, i.e., W_p in equation (11).

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 most profits. As a consequence, there may exist situations in which the platform hires sellers offering less accurate experiments, as expressed in the following proposition.

Lemma 4. *Consider the platform deciding which seller to hire for the banner targeting user x . The platform may prefer to hire sellers offering less accurate experiments.*

Proof. 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 always 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:

$$\Pi(\pi_0|c') \geq \Pi(\pi_0|c) \text{ for all } \pi_0,$$

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.

■

The intuition behind the result stated in Lemma 4 is that a more efficient seller can make higher profits when producing a banner tailored for user x . The discriminating power of the platform in setting the fees allows it to extract the entire surplus generated by the banner. Therefore, the optimal solution for the platform is just to hire the seller generating the highest possible surplus which is clearly reached by the most efficient seller. An important consequence is that the platform prefers a seller proposing less accurate experiments. Indeed, combining the results shown in Lemma 1 with the one in Lemma 4, 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.

Lemma 4 is particularly important when we consider the effect of increasing sellers' efficiency on the surplus that is generated within the platform. In particular, as shown in the following proposition, a positive productivity shock may lead to a drop of surplus brought about by banners.

Proposition 6. *Assume sellers to experience a positive productivity shock, so that c decreases. If $C'' < \frac{\mu^2 c}{k}$, this shock will result in a decrease of W_P , regardless the targeting regime.*

Proof. Let us consider the surplus created by the platform in equation (11):

$$W_p = (1 - \mu) \int_{\underline{x}}^{\bar{x}} xg(x)B(x)dx - (1 - \mu\pi_0^*)c \int_{\underline{x}}^{\bar{x}} g(x)B(x)dx. \quad (14)$$

From the sellers' first-order conditions we have that π_0^* is implicitly defined by $c\mu = kC'(\pi_0^*)$. Totally differentiating the first-order condition we get $\frac{d\pi_0^*}{dc} = \frac{\mu}{kC''}$. Let us consider a drop in the marginal cost from c to $c' < c$. We have:

$$W_p(c') - W_p(c) = [-(1 - \mu\pi_0^*(c'))c' + (1 - \mu\pi_0^*(c))c] \int_{\underline{x}}^{\bar{x}} g(x)B(x)dx. \quad (15)$$

The productivity shock does not affect clicking volumes, which will be determined by the targeting regime. Therefore, regardless the targeting regime, $W_p(c') - W_p(c) < 0$ when:

$$c - c' < \mu[c\pi_0^*(c) - c'\pi_0^*(c')]. \quad (16)$$

Given $\frac{d\pi_0^*}{dc} = \frac{\mu}{kC''}$, we have that $\pi_0^*(c) = \pi_0^*(c') + \frac{d\pi_0^*}{dc}(c - c') = \pi_0^*(c') + \frac{\mu}{kC''}(c - c')$. Plugging into equation (16), we get:

$$\begin{aligned} c - c' &< \mu \left[c \left(\pi_0^*(c') + \frac{\mu}{kC''}(c - c') \right) - c'\pi_0^*(c') \right] \\ c - c' &< \mu \left[(c - c')\pi_0^*(c') + \frac{\mu c}{kC''}(c - c') \right] \\ 1 &< \mu\pi_0^*(c') + \frac{\mu^2 c}{kC''} \end{aligned} \quad (17)$$

For all $\pi_0^*(c') \in [0, 1]$, the right-hand side is surely larger than $\frac{\mu^2 c}{kC''}$. Therefore, $C'' < \frac{\mu^2 c}{k}$ is a sufficient condition for the positive productivity shock to lower the welfare generated by the platform. ■

The sufficient condition found in Proposition 6 relies on the two effects triggered by a productivity shock. The first effect is related to the efficiency gain, since costs drop, i.e. $c' < c$. The second effect emerges indirectly from Lemma 4, which expresses the positive relationship between π_0^* and c . The two forces go in opposite directions. The direct effect pushes W_p to increase, but the second one goes towards a welfare deterioration:

$$W_p(c') - W_p(c) = \underbrace{(c - c')}_{\text{efficiency gain}} - \underbrace{\mu[c\pi_0^*(c) - c'\pi_0^*(c')]}_{\text{accuracy loss}} \underbrace{\int_{\underline{x}}^{\bar{x}} g(x)B(x)dx}_{\text{clicking volumes}}. \quad (18)$$

The decrease in welfare may only result from an accuracy loss. In particular, when the cost of designing the experiment changes quickly with π_0 then a positive productivity shock brings about a reduction of π_0^* which is small relatively to the efficiency gain of producing at a lower cost. In Figure 4, the line C' becomes steeper. Conversely, as long as the cost of designing the experiment increases slowly, the loss in terms of information overcomes the direct efficiency gain, with detrimental consequences on the surplus originated in the platform. In Figure 4, the line C' becomes less steep.

Notice that Proposition 6 is referred to the welfare that banners originate within the platform and abstracts from the reduced cost of information. Total welfare $W = W_p - C(\pi_0^*)$ must increase because of the enhanced efficiency. However, the efficiency introduced by the productivity shock is captured by a boost in persuasion but nevertheless does not insure an increase of socially valuable transactions. Indeed, $1 - \pi_0$ is just but the probability with which a banner “lies”. Therefore, a decrease of π_0 allows the experiments to be less reliable in exchange of lower prices.

5 Conclusion

In recent years, the evolution of Internet-market 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 created a market which platforms exploit to match sellers and users. Whereas the platform effectively increases 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. As the second effect might overwhelm the first the welfare might be reduced.

References

- Athey, S. and Ellison, G. (2011). Position Auctions with Consumer Search. *The Quarterly Journal of Economics*, 126(3):1213–1270.
- Bagwell, K. (2007). The economic analysis of advertising. *Handbook of industrial organization*, 3:1701–1844.
- 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 Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.
- Best, J. and Quigley, D. (2017). Persuasion for the long run.
- Chen, Y. (1997). Paying customers to switch. *Journal of Economics & Management Strategy*, 6(4):877–897.
- De Corniere, A. and De Nijs, R. (2016). Online advertising and privacy. *The RAND Journal of Economics*, 47(1):48–72.
- Fudenberg, D. and Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of Economics*, pages 634–657.

- Gomes, R. (2014). Optimal auction design in two-sided markets. *The RAND Journal of Economics*, 45(2):248–272.
- 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*.
- Johnson, J. P. and Myatt, D. P. (2006). On the simple economics of advertising, marketing, and product design. *The American Economic Review*, 96(3):756–784.
- 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.
- Kamenica, E. and Gentzkow, M. (2014). Costly persuasion. In *American Economic Review (Papers and Proceedings)*, volume 104, pages 457–462.
- Martin, D. (2017). Strategic pricing with rational inattention to quality. *Games and Economic Behavior*, 104:131–145.
- Nelson, P. (1974). Advertising as information. *Journal of political economy*, 82(4):729–754.
- 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.
- 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.

- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.
- Stigler, G. J. (1961). The economics of information. *Journal of political economy*, 69(3):213–225.
- 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.
- Villas-Boas, J. M. (1999). Dynamic competition with customer recognition. *The Rand Journal of Economics*, pages 604–631.



Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

Strada Maggiore 45
40125 Bologna - Italy
Tel. +39 051 2092604
Fax +39 051 2092664
<http://www.dse.unibo.it>